

Using Multi-Encoder Fusion Strategies to Improve Personalized Response Selection

Souvik Das Sougata Saha Rohini K. Srihari

{souvikda, sougatas, rohini}@buffalo.edu

Department of Computer Science and Engineering, University at Buffalo, NY.

Abstract

Personalized response selection systems are generally grounded on *persona*. However, a correlation exists between persona and empathy, which these systems do not explore well. Also, when a contradictory or off-topic response is selected, *faithfulness* to the conversation context plunges. This paper attempts to address these issues by proposing a suite of fusion strategies that capture the interaction between persona, emotion, and entailment information of the utterances. Ablation studies on the *Persona-Chat* dataset show that incorporating emotion and entailment improves the accuracy of response selection. We combine our fusion strategies and concept-flow encoding to train a BERT-based model which outperforms the previous methods by margins larger than 2.3% on original personas and 1.9% on revised personas in terms of **hits@1** (top-1 accuracy), achieving a new state-of-the-art performance on the *Persona-Chat* dataset.

1 Introduction

Currently, most response selection systems tend to perform well in most cases (Gu et al., 2021a; Zhang et al., 2021b; Gu et al., 2019a, 2020a). However, these re-ranking systems have the poor capability to detect and evade contradictory responses. Responses selected by these systems often contradict previous utterances, and any form of contradiction may disrupt the flow of conversation. Previous research has attempted to incorporate persona while selecting (Gu et al., 2021b; Zhang et al., 2021a) or generating (Wu et al., 2021) responses to maintain consistency. Additionally, a correlation exists between persona with personality (Leary and Allen, 2011), which influences empathy (Richendoller and Weaver III, 1994). Zhong et al. presented a multi-domain dataset collected from several empathetic Reddit threads contributing towards persona-based empathetic conversations. Nevertheless, no one has studied the emotion-persona interplay in

data presented in a more natural form. Figure 1 depicts situational emotion sometimes needs more preference than the chatbot’s persona in response selection.

On the contrary, different personality traits are related to emotion regulation difficulties (Pollock et al., 2016). Due to this, a person’s expected emotions can deviate based on his persona. Besides that, we also observe that concepts discussed in a conversational flow play an important role in response selection. However, no one has incorporated this in response selection.




Figure 1: For this conversation, the selected candidate response directly contradicts the context. Also, the bot’s persona influences the response selection, while the situational emotions and concepts get ignored. The underlines phrases/words denote the concepts.

Model	Emotion Inappropriate(%)	Contradictory(%)	Off-topic(%)
BERT-CRA	7.35	11.88	12.3

Table 1: Statistics of issues reported in the test split of *Persona-Chat* inferred by BERT-CRA (Gu et al., 2021b)¹

¹Insights drawn from the human evaluation done on 500 randomly selected data-points from self-persona original and partner-persona original sets of *Persona-Chat*

We can infer the significance of these problems from Table 1. So, to increase the usability of the personalized response selection systems, all these fundamental problems need to be addressed. We automatically annotate `Persona-Chat` (Zhang et al., 2018) dataset using a series of classifiers and rule-based modules. We model emotion-persona interaction, context-response entailment, and concept-flow using the annotations. To compare the ability of annotated features to enhance the emotion-persona interaction, contradiction avoidance, and adherence to the concept flow, we perform preliminary experiments by devising independent encoders based on BERT. Our baseline model extends BERT-CRA (Gu et al., 2021b) where we introduce an additional bot-encoder to represent the bot-utterances better. Subsequently, we propose three fusion strategies, emotion-aware(EmA), entailment-aware(EnA), persona-entailment-aware(P-EmA). These fusion strategies are designed based on emotion-persona interaction or persona-entailment information. We propose a concept-flow encoding technique that matches relevant concepts from the context and candidate responses with these fusion strategies.

We test our proposed methods on the `Persona-Chat` dataset with our automatic annotation. The results show that a model trained on a combination of our proposed fusion strategies outperforms the current state-of-the-art model by a margin of 2.3% in terms of top-1 accuracy **hits@1**.

In summary, the contributions of this paper are three-fold. (1) Automatically annotate `Persona-Chat` dataset with utterance level emotion, entailment, and concept information to provide extra supervision. (2) A suite of fusion strategies and a concept-flow encoder which are designed and implemented into a series of models, aiming to explore the impact of emotion, entailment, and concept-flow in the task of response selection. (3) Experimental results demonstrate that our proposed models outperform the existing state-of-the-art models by significant margins on the widely used `Persona-Chat` response selection benchmark.

2 Related Works

2.1 Personalized Response Selection

Chit-chat models typically trained over many dialogues with different speakers lack a consistent

personality and explicit long-term memory. These models produce an utterance given only a recent dialogue history. Li et al. proposed a persona-based neural conversation model to capture individual characteristics such as background information and speaking style. (Zhang et al., 2018) has constructed `Persona-Chat` dataset to build personalized dialog systems; this is by far the largest public dataset containing million-turn dialog conditioned on persona. Many benchmarks have been established for this dataset. For example, (Mazaré et al., 2018) proposed the fine-tuned `Persona-Chat` (FT-PC) model, which first pre-trained models using a large-scale corpus based on Reddit to extract valuable dialogues conditioned on personas and then fine-tuned these pre-trained models on the `Persona-Chat` dataset. (Wolf et al., 2019; Liu et al., 2020) also employed the pre-trained language model(GPT) for building personalized dialogue agents. (Gu et al., 2020c) proposed filtering before iteratively referring (FIRE) to ground the conversation on the given knowledge and then perform the deep and iterative matching. (Gu et al., 2021b) explored a new direction by proposing four persona fusion strategies, thereby incorporating partner persona in response selection.

2.2 Faithfulness to Context

Faithfulness in conversational systems to conversation context or knowledge is a very broad topic that can range from decreasing fact hallucination(Chen et al., 2021), reducing contradictory responses, staying on topic, etc. (Rashkin et al., 2021) has used additional inputs to act as stylistic controls that encourage the model to generate responses that are faithful to a provided evidence or knowledge. However, no one has studied the level of faithfulness the current personalized response selection systems exhibit to the conversation history. Thus, this paper thoroughly explores the impact of utilizing utterance-level emotions, entailment, and concepts on the performance of personalized response selection.

3 Dataset

In this work, we extend `Persona-Chat` (Zhang et al., 2018) and augment it with a series of annotators. The dataset consists of 8939 complete dialogues for training, 1000 for validation, and 968 for testing. Responses are selected at every turn of a conversation sequence, resulting in 65719 context-

responses pairs for training, 7801 for validation, and 7512 for testing. The positive and negative response ratio is 1:19 in the training, validation, and testing sets. There are 955 possible personas for training, 100 for validation, and 100 for testing, consisting of 3 to 5 profile sentences. A revised version of persona descriptions is also provided by rephrasing, generalizing, or specializing the original ones to make this task more challenging.

4 Automatic Dataset Annotation

We have annotated the Persona-Chat with the help of a series of automatic annotation schemes. Since we are studying the effect of emotions in personalized response selection, we assign emotion labels to the personas, context-utterances, and candidate responses using an emotion classifier. Personas and utterances were annotated using an entailment classifier to incorporate the entailment information while selecting responses. Finally, we follow a multi-layer keyword mining strategy to match meaningful concepts appearing in the context and response.

4.1 Emotion

We trained an emotion classifier on GoEmotions dataset (Demszky et al., 2020). This dataset contains 58k English Reddit comments, labeled for 27 emotion categories or Neutral. We fine-tuned RoBERTa using this dataset. We saved the checkpoint with the best Macro F1 of 49.4% (equal to the current state of the art) and used this for annotating each utterance. Since emotion classification is a challenging task and given the complexity of the affect lexicons in the corpus, we only consider the labels which can be predicted with more than 90% confidence (i.e., probability higher than 90%). The goal here is to study the effect of emotion in personalized response selection; developing a highly accurate emotion classifier is kept outside the scope of this work.

4.2 Entailment

For annotating entailment, we have used an ensemble of two models. The first one is RoBERTa based model trained on Stanford Natural Language Inference (SNLI) corpus (MacCartney and Manning, 2008) released by AllenAI². The second model is also a RoBERTa based model fine-tuned on DECODE (Nie et al., 2020). We take the two models’

²<https://github.com/allenai/allennlp-models>

weighted average of both probabilities during inference. The second model has a higher preference with 80% weightage as it is trained on conversational data. The entailment label is assigned to every persona-response and utterance-response pair.

4.3 Concept Mining

We mine keywords and key phrases from the persona sentences, utterances, and responses denoted as $\{pc_i\}_{i=1}^{N_{pc}}$, $\{uc_i\}_{i=1}^{N_{uc}}$, $\{rc_i\}_{i=1}^{N_{rc}}$ respectively. We follow the techniques proposed in (Tang et al., 2019) to extract the first level of keywords. Subsequently, we expand the concept lists by extracting key phrases using the RAKE (Rose et al., 2010). We hypothesize that concepts appearing in responses should adhere to the speaker’s persona. So, we prune some of the response/context keywords by calculating the average of Point-wise Mutual Information score between persona keywords and response/context keywords $\sum_{j=1}^{N_{pc}} PMI(pc_j, rc_i)/N_{pc}$ and rejecting the concepts which are below a threshold value(λ). Similarly, for response/concept key phrases extracted using RAKE, we only keep top N key phrases. Finally, we combine the persona and context keywords and treat them as context keywords(uc_i).

5 Methodology

5.1 Problem Definition

Given a dataset $D = \{(C_i, uc_i, p_i, r_i, rc_i, y_i)\}_{i=1}^N$ is a set of N tuples consisting context C_i , the persona of the speaker or the partner p_i , response to the context r_i , and the ground truth y_i . A set of concepts appearing in context and a response is denoted by uc_i and rc_i , respectively. The context can be represented as $C_i = \{(U_j, Emo_j, Entail_j)\}_{j=1}^L$ where U_j is an utterance, Emo_j is the dominant emotion present in U_j and $Entail_j$ is the entailment label of U_j with respect to r_i . The j^{th} utterance U_j is denoted by $U_j = \{u_{1j}, u_{2j}, \dots, u_{Mj}\}$ which consists of M tokens. Each response r_i contains single utterance, $y_i \in \{0, 1\}$, $Emo_j \in \{0, 1, \dots, P\}$, and $Entail_j \in \{\text{entailment, neutral, contradiction}\}$ where P are the total number of emotion types possible in the D . The task is to train a matching model for D , $g(C, uc, p, rc, r)$. Given a triple of context-persona-response the goal of the matching model $g(C, uc, p, rc, r)$ is to calculate the degree of match between (C, uc, p) and (rc, r) .

5.2 Bot Context Encoding

When two users communicate, many topics are often discussed in parallel, and sometimes a few utterances might not be relevant for response selection. To account for the model to be aware of the speaker change information, Gu et al. introduced a speaker disentanglement strategy in the form of *speaker embeddings* fused with the original token embeddings. This technique has proven to improve response selection performance (Gu et al., 2020b; Su et al., 2021). However, the problem of the maximum length of positional embeddings still exists. To circumvent this, we have created bot-context encoding, which captures the representation of the bot’s turns in the context while ignoring the user’s turns. The intent is to use the bot’s turns to maintain consistency during response selection. The input sequence that is sent to BERT to encode bot context is composed as follows:

$$x_{si} = [CLS]u_2[EOU]u_4[EOU]... \dots u_{n-1}[EOU][SEP]r_i[EOU] \quad (1)$$

Where u_1, u_4, \dots, u_{n-1} are bot’s utterances in the context, $[EOU]$ is a special token denoting the end of an utterance.

The resultant tokens x_{si} are passed through `bert-base-uncased`, the last hidden states of k layers i.e. $\{\mathbf{h}_{s1}^{(l)}, \mathbf{h}_{s2}^{(l)}, \dots, \mathbf{h}_{sT}^{(l)}\}$, for $l = 1, 2, \dots, k$ are used in downstream tasks.

5.3 Fusion Strategies

We use several fusion strategies to model the interdependencies of the persona, emotion, and entailment information. We use BERT (Devlin et al., 2019) as our base sentence encoder. Like the Bi-encoder (Humeau et al., 2020) we concatenate context utterances as a single context sentence before passing it into BERT.

5.3.1 Baseline

For the baseline, we have extended BERT-CRA (Gu et al., 2021b) where persona and context are concatenated to form sequence A and response form sequence B. Then, these two sequences are concatenated using $[SEP]$ token. We made two changes to this model; first, we added *speaker embeddings* with the original token representation. Secondly, we fuse bot-context encoding as described in the previous section with BERT-CRA encoding by doing multi-headed attention between

the hidden representation of the last k layers of both encoders. The token arrangement is as follows:

$$x_{CRAi} = [CLS]p_1p_2\dots p_i[EOP]u_1[EOU] \dots u_i[EOU][SEP]r_i[EOU] \quad (2)$$

Where $p_1p_2\dots p_i$ are the personalities of the speaker, $[EOP]$ token denotes end of personality representation, u_1, u_2, \dots, u_i are the utterances in the context. The resultant tokens x_{CRAi} are passed through `bert-base-uncased`, the hidden states of last k layers i.e. $\{\mathbf{h}_{c1}^{(l)}, \mathbf{h}_{c2}^{(l)}, \dots, \mathbf{h}_{cT}^{(l)}\}$, for $l = 1, 2, \dots, k$ are used in downstream tasks.

Interaction Layer : Since we are using a multi-encoder pipeline, it is crucial to capture the interaction between the encoders. For that, we use multi-head attention between hidden states of speaker context encoder and BERT-CRA. For ease of presentation, we denote the whole multi-headed attention layer as $f_{mha}(*, *)$. Then these attention outputs are passed through an aggregation layer, which basically concatenates then passes it through a two-layer feed-forward network and finally mean pools across all the layers to get \mathbf{h}_d . The output is passed through a *MLP* to get the matching degree with the response.

$$\{\tilde{\mathbf{h}}_{s1}^{(l)}, \tilde{\mathbf{h}}_{s2}^{(l)}, \dots, \tilde{\mathbf{h}}_{sT}^{(l)}\} = f_{mha}(\{\mathbf{h}_{s1}^{(l)}, \mathbf{h}_{s2}^{(l)}, \dots, \mathbf{h}_{sT}^{(l)}\}, \{\mathbf{h}_{c1}^{(l)}, \mathbf{h}_{c2}^{(l)}, \dots, \mathbf{h}_{cT}^{(l)}\}) \quad (3)$$

$$\{\tilde{\mathbf{h}}_{c1}^{(l)}, \tilde{\mathbf{h}}_{c2}^{(l)}, \dots, \tilde{\mathbf{h}}_{cT}^{(l)}\} = f_{mha}(\{\mathbf{h}_{c1}^{(l)}, \mathbf{h}_{c2}^{(l)}, \dots, \mathbf{h}_{cT}^{(l)}\}, \{\mathbf{h}_{s1}^{(l)}, \mathbf{h}_{s2}^{(l)}, \dots, \mathbf{h}_{sT}^{(l)}\}) \quad (4)$$


$$\mathbf{h}_d = \text{MeanPool}(\{\text{FFN}(\{\tilde{\mathbf{h}}_{s1}^{(l)}, \dots, \tilde{\mathbf{h}}_{sT}^{(l)}\}; \{\tilde{\mathbf{h}}_{c1}^{(l)}, \dots, \tilde{\mathbf{h}}_{cT}^{(l)}\})\}_{l=0}^k) \quad (5)$$

Loss Function: The MLP layer predicts whether a context-persona (C, p) pair matches with the corresponding response r based on the derived features. Subsequently, the output from the MLP layer is passed through a softmax output layer to return a probability distribution over all response candidates. All the models described in this paper are trained using MLP cross-entropy loss. Let Θ be the model parameters, then the loss function $\mathcal{L}(D, \Theta)$ for all the models can be formulated as follows:


$$\mathcal{L}(D, \Theta) = - \sum_{(C,p,r,y) \in D} y \log(g(C, p, r)) \quad (6)$$

5.3.2 BERT-EmA Emotion Aware Fusion:

In this strategy, an emotion incorporation framework is introduced. Similar to BERT-CRA a dual



(a) Dual encoder pipeline consisting of combination of all the encoding strategies.



(b) Concept-flow interaction network, the output of this network $h_{concept}$ can be concatenated with any of the BERT based dual encoder's output (h_d).

Figure 2: Overall Training Architecture. Though the BERT based dual encoder is independently depicted but it is trained along with concept-flow interaction(if included).




Figure 3: Interaction Layer

pipeline matching network is followed. The first pipeline encodes the speaker's and listener's emotional and personality characteristics in the context. While the other encodes the bot-context as described in section 5.2.

We attach the most probable emotion tag to each utterance to incorporate emotion features in the BERT contextual representation. The emotion-infused context representation is then concatenated with the original persona representation, as described in section 5.3.1. The main goal of representing the context in this way is to understand how the emotions of each utterance interact with the speaker's persona. The input to the emotion encoder is as follows:

$$x_{EmAi} = [CLS]p_1p_2...p_i[EOP][Emo_1]u_1[EOU] \dots [Emo_i]u_i[EOU][SEP]r_i[EOU] \quad (7)$$

Similar to baseline, the hidden states of last k layers i.e. $\{h_{e1}^{(1)}, h_{e2}^{(1)}, \dots, h_{eT}^{(1)}\}$, for $l = 1, 2, \dots, k$ are used in downstream tasks.

5.3.3 BERT-EnA-P: Entailment Aware Fusion

In this fusion strategy, the intention is to model the entailment information about each of the utterances and personas with the response. Like BERT-EnA we follow a dual encoder pipeline, the first encodes the entailment features, and the second encodes the bot-context. To incorporate entailment features into BERT contextual representation, we attach entailment tags i.e. $\langle contradiction \rangle$, $\langle entailment \rangle$ and $\langle neutral \rangle$ at the start of every utterance and persona. The response is concatenated with the context-entailment representation with a $[SEP]$ token. The input to the entailment encoder is as follows:

$$x_{EmA-Pi} = [CLS][Entail_{p1}]p_1\dots[EOP] [Entail_1]u_1[EOU] [Entail_2]u_2[EOU] \dots [Entail_i]u_i[EOU] [SEP]r_i[EOU] \quad (8)$$

The hidden states of last k layers i.e. $\{h_{en1}^{(1)}, h_{en2}^{(1)}, \dots, h_{enT}^{(1)}\}$, for $l = 1, 2, \dots, k$ are used in downstream tasks.

Finally, we experiment with a combined pipeline as depicted in Figure 2a.

5.4 Concept-Flow(CF) Interaction

In section 4.3, we describe how we extract relevant concepts from the context and the response. An appropriate response often has concepts most recently discussed in the context. So, to model that, we construct a concept-flow interaction network, where the interaction between the context-concepts and response-concepts are measured and used as a feature in response relevance classification.

Let us consider $\{CC_1, CC_2, \dots, CC_n\}$ are concepts extracted from context and $\{RC_1, RC_2, \dots, RC_n\}$ are concepts extracted from a response. Now, we pass each of these concepts through a transformer-based concept encoder f_c to get two sets of concept embeddings $\{\mathbf{ec}_1, \mathbf{ec}_2, \dots, \mathbf{ec}_n\}$, $\mathbf{ec}_i \in \mathbb{R}^{d_c}$ and $\{\mathbf{rc}_1, \mathbf{rc}_2, \dots, \mathbf{rc}_n\}$, $\mathbf{rc}_i \in \mathbb{R}^{d_c}$ for context and response concepts respectively. To learn the context flow representation for each set of concepts, we apply a bi-directional GRU network to capture sequential dependencies between subsequent concepts in a conversational situation. Context-concept and response-concept representation \mathbf{h}_i^{cc} , \mathbf{h}_i^{rc} can be formulated as:

$$\mathbf{c}_i^{\text{cc}}, \mathbf{h}_i^{\text{cc}} = \overleftarrow{\text{GRU}}(\mathbf{ec}_i, \mathbf{h}_{i-1}^{\text{cc}}) \quad (9)$$

$$\mathbf{c}_i^{\text{rc}}, \mathbf{h}_i^{\text{rc}} = \overleftarrow{\text{GRU}}(\mathbf{er}_i, \mathbf{h}_{i-1}^{\text{rc}}) \quad (10)$$

$$\mathbf{h}_{\text{cc}} = \tanh\left(\sum_{j \in 2 * N_l} \mathbf{W}_j \mathbf{h}_j^{\text{cc}}\right) \quad (11)$$

$$\mathbf{h}_{\text{rc}} = \tanh\left(\sum_{j \in 2 * N_l} \mathbf{W}_j \mathbf{h}_j^{\text{rc}}\right) \quad (12)$$

Where $\mathbf{h}_i^{\text{cc}} \in \mathbb{R}^{2d_c}$, $\mathbf{h}_i^{\text{rc}} \in \mathbb{R}^{2d_c}$ are the i -th hidden states and $\mathbf{c}_i^{\text{cc}} \in \mathbb{R}^{2d_c}$, $\mathbf{c}_i^{\text{rc}} \in \mathbb{R}^{2d_c}$ are the outputs of the respective GRU encoders, \mathbf{W}_j is a learn-able parameter and N_l is the number of layers in each GRUs. To model the interaction between \mathbf{h}_i^{cc} and \mathbf{h}_i^{rc} we follow the same interaction mechanism described in the earlier section. The output $\mathbf{h}_{\text{concept}}$ is concatenated with the dual encoder output \mathbf{h}_d before passing it through an MLP.

6 Experimental Setup

6.1 Training Details

The ratio of positive to negative samples in the training set is 1:19, so there is a high imbalance in training data. Taking inspiration from (Gu et al., 2021b) we adopted a dynamic negative sampling strategy in which the ratio of positive and negative responses is 1:1 in an epoch. We keep the

positive response constant and change the negative response for every epoch, generating data for 19 epochs. We use `bert-base-uncased` as the base for each of our pretraining-based fusion models. In concept mining strategy, we have taken the top 3 concepts extracted using RAKE, λ for PMI-based scoring was varied from 0.3 to 0.8 with 0.1 steps, and 0.5 was found optimum. The number of turns in the conversation history used for concept mining varied following this set: $\{2, 3, 4, 5, 6, 7\}$. We preserve the original parameters of `bert-base-uncased`. The number of k -last layers in the interaction layer varied following this set: $\{3, 4, 5, 6\}$; after some initial experimentation, 4 was found as the optimum value. The number of heads in the multi-head attention layer was kept 8. We use a 6-layered version MiniLM(Wang et al., 2020) to encode the concepts; the embedding dimension was 384. The number of layers in the bi-directional GRUs in the concept encoder is 2. A dropout rate of 0.7 is applied to the concept encoder hidden representation before we send it to the interaction layer. AdamW(Loshchilov and Hutter, 2019) optimizer was used for optimization. The initial learning rate was set to $2e-5$ and linearly decayed by L2 weight decay. The maximum sequence length was set to 320. The training batch size was 12. The relevance prediction head used a single feed-forward layer with sigmoid activation. All code was implemented using the PyTorch framework. Also, we used 2 NVIDIA RTX A5000 GPUs to train the models. The average training time for one epoch was 46 minutes, using all our fusion strategies and concept encoding.

6.2 Evaluation Metrics

We used the same evaluation metrics as the previous work to ensure comparable results. Each model aimed to select the best-matched response from available candidates for the given context and persona. We calculated the recall of the true positive replies, denoted as **hits@1**. In addition, the mean reciprocal rank (**MRR**) was also adopted to take the rank of the correct response overall candidates into consideration.

6.3 Comparison Methods

For comparison, we have only selected pretraining-based models.

- **FT-PC (Mazaré et al., 2018)**: employed the “pretrain and fine-tune” framework by first

Model	Self Persona				Partner Persona			
	Original		Revised		Original		Revised	
	hits@1	MRR	hits@1	MRR	hits@1	MRR	hits@1	MRR
FT-PC (Mazaré et al., 2018)	-	-	60.7	-	-	-	-	-
DIM (Gu et al., 2019b)	78.8	86.7	70.7	81.2	64.0	76.1	63.9	76.0
TransferTransfo (Wolf et al., 2019)	80.7	-	-	-	-	-	-	-
FIRE (Gu et al., 2020c)	81.6	-	74.8	-	-	-	-	-
BERT-CRA (Gu et al., 2021b)	84.3	90.3	79.4	86.9	71.2	80.9	71.8	81.5
Baseline	84.4	90.7	79.4	87.6	71.2	81.1	71.4	81.5
BERT-EmA	84.6	90.9	79.8	87.7	71.4	81.2	71.4	81.6
BERT-P-EnA	85.3	91.2	80.5	87.9	71.7	81.3	71.3	81.4
BERT-EmA+BERT-P-EnA	85.8	91.4	80.7	88.0	72.3	81.5	71.7	81.5
BERT-EmA+BERT-P-EnA+CF (All)	86.6*	91.6*	81.3*	88.6*	72.6*	81.9*	72.4*	81.9*

Table 2: Performance of the proposed and previous methods on the Persona-Chat dataset under different persona configurations. The meanings of "Self Persona", "Partner Persona", "Original", and "Revised" can be found in Section 3. "-" represents that the results were not reported in their papers. Numbers marked with * denote that the improvement over the best performing baseline is statistically significant (t-test with p-value < 0.05). Numbers in bold denote the combined fusion strategy that achieves the best performance.

pretraining on a domain-specific corpus, dialogues of which were extracted from Reddit, and then fine-tuning on the Persona-Chat.

- **DIM (Gu et al., 2019b)**: used a dually matching network (DIM) which performs interactive matching between responses and contexts and between responses and personas respectively for ranking response candidates.
- **TransferTransfo (Wolf et al., 2019)**: the paper fine-tunes a transformer model (GPT) using Persona-Chat dataset on a multi-task objective which combines several unsupervised tasks.
- **BERT-CRA (Gu et al., 2021b)**: This work presents four context-aware persona fusion strategies and the models are initialized and pretrained using BERT on Persona-Chat dataset.

6.4 Experimental Results

Table 2 reports the evaluation results of our proposed and previous methods on Persona-Chat under various persona configurations. We can see that incorporating the emotion and entailment knowledge of the utterances coupled with generic distributional semantics and external knowledge learned from pretraining rendered improvements on both **hits@1** and **MRR** conditioned on various personas. Compared to FT-PC (Mazaré et al., 2018), our best model outperformed it by 20.4 % in terms of hits@1 conditioned emotion, entailment and concepts. Compared to TransferTransfo (Wolf et al., 2019), which was also trained using pretrained transformer models, our combined model

outperformed it, which shows the effectiveness of fusion strategies and the concept-encoder. Lastly, our combined model outperformed the BERT-CRA (Gu et al., 2021b) in all the tasks. We see a 2.3 % and 1.9 % improvement in original and revised self-persona, and 1.4 % and 0.6 % improvement in original and revised partner-persona in terms of **hits@1**. The results bolster our hypothesis that emotion, entailment, and concepts play an important role in the task of response selection. Also, it is to be noted that Persona-Chat is a synthetic dataset, i.e., the data collection did not happen naturally. Therefore, the chances that the user will display this subtle interplay of persona and emotion is less. In addition to that, we observe the presence of contradictory distractor responses. We see a significant performance improvement from this information by introducing entailment-aware fusion and concept encoding.

6.5 Human Evaluation




Figure 4: Human evaluation results on Persona-Chat self-persona original test split

Since a qualitative study by humans is necessary to understand the effectiveness of the proposed methods, we further perform a human evaluation

Models	hits@1	MRR
Baseline	84.4	90.7
BERT-EmA(- Speaker Encoding)	84.5	90.8
BERT-EmA	84.6	90.9
BERT-EnA-P	85.3	91.2

Table 3: Ablation Study for Emotion and Entailment on self-original persona.

on a portion of the data. We randomly sampled 100 inferred examples from the test set by baseline model, combined model, combined model except for the emotion, and combined model except for the concept flow interaction. We combined all the samples and evaluated them using Amazon Mechanical Turk by 2 different turkers on three metrics: emotion inappropriate, contradictory, and off-topic. The turkers needed to select if any of the three issues were present in an example. The percentages of reported issues by both groups are shown in Figure 4. The results reveal that all our encoding pipelines effectively reduce contradictory responses and somewhat reduce off-topic and emotionally inappropriate responses. The agreement between the two groups was moderate (Krippendorff’s $\alpha = 0.713$).

7 Analysis

7.1 Ablation Study for Emotion and Entailment

We perform ablation studies (shown in Table 3) to validate the effectiveness of emotion and entailment fusion in our proposed models. We see a slight improvement in our baseline model that uses our proposed speaker embedding. Also, unsurprisingly that the effect of emotion is not significant. As the dataset is artificially created, and emotions exhibited by the annotators are not always true. However, some performance improvement is observed. Conditioning persona in entailment fusion improves performance considerably as responses may not entail the speaker’s persona.

7.2 Effect of Context Turns on Concept Representation

Concept matching boosts the evaluation performance further. However, the number of turns in the conversation history from which we mine the concepts influences the performances. It is evident from Figure 5 that the essential concepts in the most relevant response will be present in the recent conversation history.

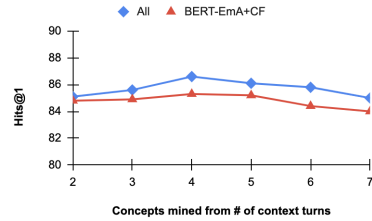


Figure 5: This graph shows how **hit@1** reaches an optimum value and then decreases with an increase in the number of turns used to mine concepts.

personas	my favorite color is blue. <ent: neutral> I enjoy reading mysteries. <ent: neutral> I have seven children. <ent: entail> I grew up on a large farm. <ent: neutral>
context	A: hello how are you today? <emo:curiosity> <ent: neutral> B: I am well. how are you? <emo:curiosity> <ent: neutral> A: I am doing great just got back from the beach <emo:excitement> <ent: neutral> B: that is great. I live far from the beach. <emo:caring> <ent: neutral> A: I am very lucky we live beside the beach. what do you do for a living? <emo:curiosity> <ent: neutral> B: I keep busy with my seven children. <emo:excitement> <ent: neutral> A: wow that much have taken some adjusting I teach kindergarten. <emo:surprise> <ent: neutral>
golden response	do you teach mysteries to your children ? they are my favorite type of novel . <emo:curiosity>
BERT-CRA	that must be a lot of work. but very rewarding. i bet <emo:realization>
All	do you teach mysteries to your children ? they are my favorite type of novel . <emo:curiosity>

Table 4: Case study showing concept flow.

7.3 Case Study

Table 4 shows the efficacy of concept-encoding. Sometimes models fine-tuned on pretrained transformer models like BERT-CRA tends to select more generic responses. These models pay less attention to the persona or specific keywords in the context while selecting responses. In this example, our proposed model performs better than BERT-CRA as it is conditioned on the concepts. Specifically, concepts in the correct response i.e. "mysteries", "novel" relates to "reading mysteries" concept in the persona and "your children" relates to "teach kindergarten" in the context.

8 Conclusion

This work proposes a suite of novel fusion strategies and concept-flow encoder, which leverages the utterances’ emotion, entailment, and concept information. These features help improve the performances of our models and provide critical insights into certain aspects of how humans communicate with each other. Though the techniques used in this paper are simple, it highlights the areas where response selection often falters, like detecting contraction, deviation from the concepts, etc. This work can be further extended by using a graphical

model to improve the concept representations.

Acknowledgements

We thank the anonymous reviewers for providing valuable feedback on our manuscript. This work is partly supported by NSF grant number IIS-2214070. The content in this paper is solely the responsibility of the authors and does not necessarily represent the official views of the funding entity.

References

- Sihao Chen, Fan Zhang, Kazuo Sone, and Dan Roth. 2021. [Improving faithfulness in abstractive summarization with contrast candidate generation and selection](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5935–5941, Online. Association for Computational Linguistics.
- Dorottya Demszky, Dana Movshovitz-Attias, Jeongwoo Ko, Alan Cowen, Gaurav Nemade, and Sujith Ravi. 2020. [Goemotions: A dataset of fine-grained emotions](#).
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [Bert: Pre-training of deep bidirectional transformers for language understanding](#).
- Jia-Chen Gu, Tianda Li, Zhen-Hua Ling, Quan Liu, Zhiming Su, Yu-Ping Ruan, and Xiaodan Zhu. 2021a. [Deep contextualized utterance representations for response selection and dialogue analysis](#). *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:2443–2455.
- 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](#).
- Jia-Chen Gu, Tianda Li, Quan Liu, Xiaodan Zhu, Zhen-Hua Ling, Zhiming Su, and Si Wei. 2020b. [Speaker-aware BERT for multi-turn response selection in retrieval-based chatbots](#). *CoRR*, abs/2004.03588.
- Jia-Chen Gu, Zhen-Hua Ling, and Quan Liu. 2019a. [Utterance-to-utterance interactive matching network for multi-turn response selection in retrieval-based chatbots](#).
- Jia-Chen Gu, Zhen-Hua Ling, Xiaodan Zhu, and Quan Liu. 2019b. [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. 2020c. [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. 2021b. [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.
- Samuel Humeau, Kurt Shuster, Marie-Anne Lachaux, and Jason Weston. 2020. [Poly-encoders: Transformer architectures and pre-training strategies for fast and accurate multi-sentence scoring](#).
- Mark R. Leary and Ashley Batts Allen. 2011. [Personality and persona: Personality processes in self-presentation](#). *Journal of Personality*, 79(6):1191–1218.
- Jiwei Li, Michel Galley, Chris Brockett, Georgios Spithourakis, Jianfeng Gao, and Bill Dolan. 2016. [A persona-based neural conversation model](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 994–1003, Berlin, Germany. Association for Computational Linguistics.
- 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](#).
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#).
- Bill MacCartney and Christopher D. Manning. 2008. [Modeling semantic containment and exclusion in natural language inference](#). In *Proceedings of the 22nd International Conference on Computational Linguistics (Coling 2008)*, pages 521–528, Manchester, UK. Coling 2008 Organizing Committee.
- 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*, pages 2775–2779, Brussels, Belgium. Association for Computational Linguistics.
- Yixin Nie, Mary Williamson, Mohit Bansal, Douwe Kiela, and Jason Weston. 2020. [I like fish, especially dolphins: Addressing contradictions in dialogue modeling](#).
- Noah C. Pollock, Gillian A. McCabe, Ashton C. Southard, and Virgil Zeigler-Hill. 2016. [Pathological personality traits and emotion regulation difficulties](#). *Personality and Individual Differences*, 95:168–177.

- Hannah Rashkin, David Reitter, Gaurav Singh Tomar, and Dipanjan Das. 2021. [Increasing faithfulness in knowledge-grounded dialogue with controllable features](#).
- Nadine R Richendoller and James B Weaver III. 1994. Exploring the links between personality and empathic response style. *Personality and individual Differences*, 17(3):303–311.
- Stuart Rose, Dave Engel, Nick Cramer, and Wendy Cowley. 2010. *Automatic Keyword Extraction from Individual Documents*, chapter 1. John Wiley & Sons, Ltd.
- Yixuan Su, Deng Cai, Qingyu Zhou, Zibo Lin, Simon Baker, Yunbo Cao, Shuming Shi, Nigel Collier, and Yan Wang. 2021. [Dialogue response selection with hierarchical curriculum learning](#). 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 1740–1751, Online. Association for Computational Linguistics.
- Jianheng Tang, Tiancheng Zhao, Chenyan Xiong, Xiaodan Liang, Eric P. Xing, and Zhiting Hu. 2019. [Target-guided open-domain conversation](#).
- Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. 2020. [Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers](#).
- Thomas Wolf, Victor Sanh, Julien Chaumond, and Clement Delangue. 2019. [Transfertransfo: A transfer learning approach for neural network based conversational agents](#).
- Yuwei Wu, Xuezhe Ma, and Diyi Yang. 2021. [Personalized response generation via generative split memory network](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1956–1970, Online. Association for Computational Linguistics.
- Chen Zhang, Hao Wang, Feijun Jiang, and Hongzhi Yin. 2021a. [Adapting to context-aware knowledge in natural conversation for multi-turn response selection](#). In *Proceedings of the Web Conference 2021, WWW '21*, page 1990–2001, New York, NY, USA. Association for Computing Machinery.
- 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?](#)
- Zhuo Zhang, Danyang Zheng, and Ping Gong. 2021b. [Multi-turn response selection in retrieval based chatbots with hierarchical residual matching network](#). *Journal of Physics: Conference Series*, 1757(1):012023.
- Peixiang Zhong, Chen Zhang, Hao Wang, Yong Liu, and Chunyan Miao. 2020. [Towards persona-based empathetic conversational models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6556–6566, Online. Association for Computational Linguistics.

A Appendix

A.1 Human Evaluation for Initial Study

Five hundred context response pairs randomly sampled from the Persona-Chat self-original test split inferred by BERT-CRA were evaluated by at least two AMT workers. The following questions were asked to the workers:

1. Is this response contain emotions that are consistent with the context? (Most definitely/ not at all)
2. Is this response contradicts the context? (Most definitely/ not at all)
3. Do the topics discussed in this response appropriate to the topics discussed in the context? (Most definitely/ not at all)

A.2 Final Human Evaluation

Same evaluation pattern is followed as A.1, average pay for each HIT was 0.07 \$.