

Put Chatbot into Its Interlocutor’s Shoes: New Framework to Learn Chatbot Responding with Intention

Hsuan Su^{1*}, Jiun-Hao Jhan^{2*}, Fan-yun Sun^{2†}, Saurav Sahay^{2◇} and Hung-yi Lee^{1*}

^{*}National Taiwan University, Taipei, Taiwan

[†]Stanford University, Stanford, CA, USA

[◇]Intel Labs, Santa Clara, CA, USA

{b04203058, r06942141, hungyilee}@ntu.edu.tw

fanyun@stanford.edu

saurav.sahay@intel.com

Abstract

Most chatbot literature that focuses on improving the fluency and coherence of a chatbot, is dedicated to making chatbots more human-like. However, very little work delves into what really separates humans from chatbots – humans intrinsically understand the effect their responses have on the interlocutor and often respond with an intention such as proposing an optimistic view to make the interlocutor feel better. This paper proposes an innovative framework to train chatbots to possess human-like intentions. Our framework includes a guiding chatbot and an interlocutor model that plays the role of humans. The guiding chatbot is assigned an intention and learns to induce the interlocutor to reply with responses matching the intention, for example, long responses, joyful responses, responses with specific words, etc. We examined our framework using three experimental setups and evaluated the guiding chatbot with four different metrics to demonstrate flexibility and performance advantages. Additionally, we performed trials with human interlocutors to substantiate the guiding chatbot’s effectiveness in influencing the responses of humans to a certain extent. Code will be made available to the public.

1 Introduction

Humans have evolved to become sensitive to their social interactions. The more they interact, the more they generally learn what to say and what not to say to light up people’s mood or to avoid upsetting others. In this paper, we aimed to train a chatbot to emulate these human-like qualities by making it learn from interactive conversation. A chatbot that understands the effect its utterances have on the interlocutor could be a significant step towards achieving human-level chatbots.

A chatbot that understands the effect of its utterances on the interlocutor is also critical in real-world applications. For instance, as shown in Figure 1, given a context from the interlocutor, both

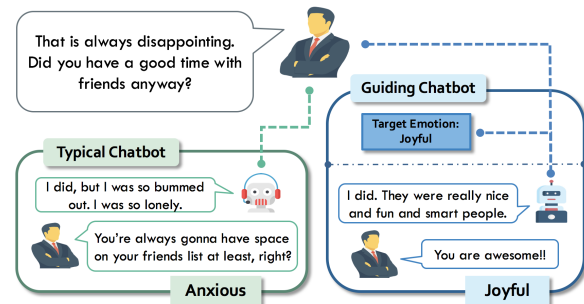


Figure 1: An example of the dialogue to show that how the chatbot interacts with the interlocutor, and how our chatbot affects the interlocutor’s response when assigned the intention, *making people respond joyful*.

responses "I did. They were really nice and fun and smart people." and "I did. I was so bummed out. I was so lonely." were relevant and reasonable responses, and were equally suitable for a typical chatbot. However, we could give an intention to the proposed chatbot (guiding chatbot), such as making the interlocutor feel joyful. In this way, the chatbot would respond in a positive way to induce joy in the interlocutor.

Much literature combine Reinforcement Learning(RL) (Kaelbling et al., 1996) with transformer-based (Vaswani et al., 2017) models to control the chatbot’s output. Gupta et al. (2019) proposed models to concentrate on crucial keyphrases presented in the context. Their models tended to generate outputs that were more coherent and specific to the conditionals, which led to more non-generic words. By training with a combination of the above criteria, their approach led to more diverse and interesting responses. However, these previous works focused on controlling the chatbot’s responses and completely neglected the interlocutor in their training.

In this paper, we made extensive use of the interlocutor’s responses as interactive experiences to train our guiding chatbot to influence the interlocutor with intentions. We introduce a novel train-

ing framework, in which there were two conversational models that simulated chatbot-interlocutor interaction. One model acted as the interlocutor model, while the other was the guiding chatbot to be trained. The interlocutor took the guiding chatbot's output as its input and generated corresponding responses. The learning framework was given a controllable factor, which represented the intention it had. We defined reward functions according to three different controllable factors, sentence length, emotion, and specific words, to make the guiding chatbot learn to induce the interlocutor model to generate desired responses using RL.

To evaluate our guiding chatbot, we designed several experiments to examine the rewards corresponding to three controllable factors, and empirical results demonstrate that our guiding chatbot can influence humans' responses. Moreover, we found that training with more interlocutor models together improved the guiding chatbot's performance on the human evaluation experiment. Furthermore, we analyzed recent off-the-shelf chatbots based on experimental results, aiming to find hidden tendencies these chatbot models had, such as cursing more or being more irritable.

2 Related Work

The most common chatbot model is sequence-to-sequence based (Sutskever et al., 2014). Recently, numerous researchers applied transformer to build coherent chatbots by retrieval-based (Zhou et al., 2016; Wu et al., 2019; Yang et al., 2018; Henderson et al., 2017; Yan et al., 2016) and generative-based (Ritter et al., 2011; Serban et al., 2016; Shang et al., 2015; Tammewar et al., 2018) approaches. Despite the decent fluency and coherence these chatbots achieved, they still hardly converse like a human. The reason might be that they are essentially devoid of emotions.

Furthermore, some used RL to improve their chatbot's performance (Serban et al., 2016; Williams, 1992) and others combined RL with GPT-2 (Radford et al., 2019) models to control the sentiment of their chatbot's response to make it more user-friendly (Han et al., 2019; Lee et al., 2018). Beyond the viewpoint of sentiment, the EmpatheticDialogues (ED) dataset was collected (Rashkin et al., 2019) to train a chatbot that could recognize the feeling of the interlocutor and know how to reply accordingly (Lin et al., 2020). However, these researchers neglected what really separated

humans from chatbots – humans understand the impact their responses have on the interlocutor and often responded with intentions and expectations. Note that it is not just about being empathetic as human's intentions could vary widely.

One previous work also considered interlocutor responses (Shin et al., 2019). It used a sentiment predictor to predict the interlocutor's sentiment given the chatbot's response, and also trained the chatbot with RL. Unlike this previous work, our proposed framework explicitly modeled the possible responses of interlocutors. Explicitly modeling interlocutor responses give the proposed framework more flexibility. For example, in addition to steering the interlocutor's sentiment as in this paper, the framework could be used to build a chatbot that induce the interlocutor to become more talkative by setting its learning target to be making the interlocutor generate longer sentences. Moreover, we also developed techniques to preserve the dialogue's coherence, so our chatbot could still generate fluent and appropriate responses in addition to having a particular intention.

Apart from influencing the interlocutor, the proposed framework also served as a way to analyze the underlying inclination of the off-the-shelf chatbots playing the role of interlocutor. Through the interaction, we could know what factors are apt to influence these off-the-shelf chatbots. Holzinger et al. (2017) claimed that the appealing performance of recent robust and SOTA models belied a potential problem of black-box models: these models lacked an explicit declarative knowledge representation. Hence, calling for a transparent representation, they dug into explaining trained models. In contrast to the previous contributions, we tried to explain the implied tendency of a chatbot, which was not obvious to recognize. According to the experiments, we were capable of telling whether the off-the-shelf black box chatbot possessed certain predispositions, such as tending to swear more or having a short temper.

3 Methodology

3.1 Framework

The proposed framework is shown in Figure 2. It consisted of two conversational models: the guiding chatbot and the interlocutor model. The interlocutor and guiding chatbot simulated the dialogue between a human and a chatbot. The guiding chatbot aimed to generate a response that maximize

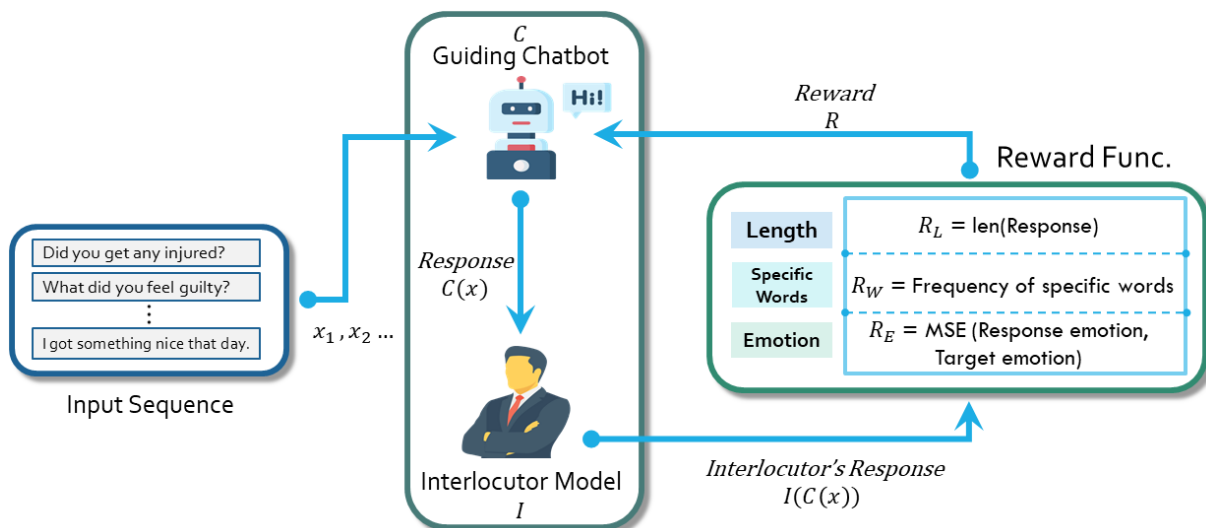


Figure 2: The framework that we proposed to teach the guiding chatbot how to achieve the intention assigned by the controllable factors.

rewards according to different controllable factors; the interlocutor models produced responses based on the guiding chatbot’s response in order to simulate a human’s response. Therefore, grounded in different controllable factors, we examined corresponding rewards to optimize the guiding chatbot to influence the interlocutor.

3.2 Conversational Models

Interlocutor The model I represented the interlocutor. I could be any off-the-shelf chatbot whose parameters were fixed during training; that is, it was unnecessary to know its parameters in the framework. I was only used during the training phase to train the guiding chatbot via interaction. In the testing phase, the guiding chatbot will interact with real human beings. The interlocutor models’ settings will be described in Section 5.3.

Guiding Chatbot The guiding chatbot model C was the chatbot we trained to induce desired responses in the interlocutor. We built the guiding chatbot model C based on DialoGPT (Zhang et al., 2020). To train model C , given the input sentence x , our chatbot C generated a sentence $C(x)$. The generated sentence $C(x)$ then became the input for I , and I output its response $I(C(x))$. We defined the reward R for C based on $C(x)$ and $I(C(x))$, and C was trained to maximize the value of R by the policy gradient. The definition of the reward R depended on the controllable factors, that is, the intention of the guiding chatbot (how the guiding chatbots wanted the interlocutor to respond). The

definition of our reward functions is in Section 3.3, and the controllable factors are in Section 4.

3.3 Rewards Functions

We introduce two kinds of reward functions: intention reward R_I and coherence reward R_C . The final reward that the guiding chatbot C learned to maximize will be a combination of R_I and R_C .

Intention To influence the interlocutor, the guiding chatbot C ought to learn from the interlocutor’s reaction. To be more specific, we collected responses $I(C(x))$ from the off-the-shelf chatbots when interacting with our guiding chatbot. Then the intention reward R_I was obtained by evaluating the interlocutor’s responses, that is, $I(C(x))$, based on the controllable factors of guiding chatbot C . Using the intention reward allowed the guiding chatbot to induce the interlocutor to perform specifically according to the controllable factors, namely our intentions. The formulation of R_I depended on the controllable factors. To observe the effectiveness of guiding these interlocutor models, in this paper, we had three controllable factors, which were equal to our intentions: to extend the sentence length, to make the interlocutor speak with a particular emotion, and to induce the interlocutor to speak specific words. Exact formulation of rewards for different controllable factors will be given in Section 4.

Coherence Using the intention reward as the only reward led to a drawback that the guiding chatbot ignored the coherence between the input

x and the generated response $C(x)$. To avoid this problem, an extra constraint on the guiding chatbot to maintain coherent responses was necessary: we applied another conversational model C' that served as a constraint maintaining coherence. Here we used the open-domain GPT-2 model as the C' . To be more specific, we estimated the difference in generated probability between C and C' and minimized the estimated difference. As a result, C would be less likely to produce responses unrelated to input x coherent to responses generated by C' . The additional reward R_C is defined as below.

$$R_C = P_{C'}(C(x) | x). \quad (1)$$

R_C was the likelihood that C' generated the sentence $C(x)$ given the input sentence x . This term served as a kind of regularization that avoids drift during training.

To sum up, the total reward is defined as:

$$R = \lambda R_I + (1 - \lambda) R_C, \quad (2)$$

where λ is the hyper-parameter.

4 Controllable Factors

Below are the three types of controllable factors studied in this paper. R_I in Section 3.3 could be either R_L for sentence length, R_E for emotion, or R_W for specific words, introduced below.

Sentence Length A chatbot that could inspire the interlocutor to become more talkative is desirable in many real world applications. We aimed to observe whether our chatbot was able to make the interlocutor more talkative, and extend the length of conversations. Hence, we counted the sentence length of interlocutor models' responses as R_L . By optimizing this reward, we anticipated that the guiding chatbot might extend sentence length from the interlocutor.

Emotion We studied whether our chatbot was capable of inducing the interlocutor to respond with different emotions. We selected eight emotions, including anger, anxiety, contentment, disgust, hope, joy, sadness, surprise. We selected the eight emotions such that two emotions are located in each of the four different quadrants of the Valence-Arousal coordinate (Russell, 1980).

To ascertain the emotion of sentences, we established an Emotion Detector. The Emotion Detector was an emotion classifier used to classify emotion

given an input sentence. We trained the Emotion Detector on the EnpatheticDialogue (ED) dataset (Rashkin et al., 2019). For each sentence, the Emotion Detector will employ a Valence-Arousal (VA) projection grounded on the Valence-Arousal coordinate (Russell, 1980; G. Paltoglou, 2013). Given an input sequence, the Emotion Detector would output a two-dimensional vector representing sequence's emotion, defined as emotional valence¹. More details related to the Valence-Arousal Coordinate will be discussed in Section 5.2. We utilized the BERT (Devlin et al., 2019) architecture, a pre-trained contextualized embedding model, to improve language understanding. Next, we fine-tuned the BERT model on an emotional classification task to enhance the model's capability of categorizing each emotion. The accuracy of our emotion detector was up to 82%, and, therefore, we could obtain a detected emotion and its emotional valence given an input sentence.

The Emotion Detector takes $I(C(x))$ as input and predicted its emotional valence according to the VA coordinate. Therefore, we could calculate the Mean Square Error (MSE) between the emotional valence of the interlocutor models' responses and the target emotion's emotional valence as the reward R_E .

Specific Words We aimed to induce the interlocutor to speak with words from specific groups. These word groups, including Bad Words, Sports, and Food, were collected from Google's team² and the Enchanted Learning website³. To provoke the interlocutor to respond to the sentence including the specific words we want, we calculated the frequency of the specific words in a sentence. We counted the frequency of interlocutor models' responses that contain words in these word groups as R_W . We anticipated that the interlocutor can generate a sentence that contains more words from the specific group and still be coherent as well as fluent.

5 Experimental Setup

5.1 Dataset

EmpatheticDialogues Dataset Rashkin et al. (2019) created an innovative dataset with around

¹e.g. fear=[-0.12, 0.79], joy=[0.85, 0.15]

²<https://gist.github.com/jamiew/1112488>

³<https://www.enchantedlearning.com/home.shtml>

25K conversations, each consisting of a speaker and a listener. The participants, acting as the speaker, initiated the talks, and the psychologists, serving as the listener, responded to the speaker empathetically. The dataset covers 32 different emotion labels including positive, negative, and neutral emotions. They firmly ensure that each emotion in the dataset was evenly distributed. Nonetheless, a few emotion classes were quite similar, such as "sentimental" and "nostalgic". Thus, we merged these equivalent emotion classes into one emotion class.

5.2 Valence-Arousal Coordinate Projection

In Valence-Arousal Coordinate study (Russell, 1980; G. Paltoglou, 2013), researchers assigned emotional values to nineteen kinds of emotions. We performed supervised training of the Emotion Detector based on these known emotions on the ED dataset. Each emotion could be represented as a two-dimensional vector. Therefore, we could map each emotion to the coordinate on the VA space.

5.3 Model Settings

RL Training Details We applied the Policy gradient (Sutton et al., 2000) as our RL algorithm. To implement an RL training chatbot, we applied the DialoGPT model, which fine-tuned the GPT-2 model on 147M multi-turn dialogues from Reddit discussion threads. The GPT-2 model was a transformer-based model with 36 layers, 20 attention heads in each layer, 345M parameters, and an embedding size was 1024. This model was trained on the WebText dataset and 50,257 tokens with invertible byte pair encoding to preserve capitalization and punctuation. In our training procedure, we fine-tuned the DialoGPT model on the ED dataset based on the reward function mentioned in Section 3.3.

Interlocutor Models The interlocutor models had three different setups:

- The Publicly available **Google bot** (Vinyals and Le, 2015)⁴ was trained on the dataset proposed by Danescu-Niculescu-Mizil and Lee (2011) with 220,579 conversational exchanges between 10,292 pairs. The whole corpus was split into training and testing sets.

⁴<https://github.com/Conchylicultor/DeepQA>

- The same **DialoGPT** model mentioned in Section 5.3 was used here to act as the interlocutor. The weights of the model were fixed.
- A BERT-based **Retrieval** chatbot trained on the ED dataset. Given input sentences, the chatbot chose the corresponding response from the candidate pool. The BERT encoder first embedded the sentences into sentence embedding and then computed cosine similarity between the input sentences and all candidates to select the most likely option. The candidate pool was comprised of all sentences in the ED dataset, which contained approximately 100K sentences.

5.4 Evaluation Metrics

Aside from the reward scores related to the intentions, we also reported the following three metrics in the experiments.

Conditional Perplexity The Conditional Perplexity here was to measure the dialogue coherence between the output sentence and input sentence x . The equation is shown below.

$$CPPL = \prod_{i=1}^T \frac{1}{(P(C(x)_i|x))^{1/T}} \quad (3)$$

$CPPL$ was the conditional perplexity, which was equal to the inverse of the product of each word's probability in the sentence $C(x)$ given the input sentence x . T was the length of the sentence $C(x)$.

Perplexity Here we employed the pretrained GPT-2 language model to judge if the output sentence $C(x)$ was an acceptable sentence. The computation of Perplexity (Chen et al., 1998) is shown below.

$$PPL = \prod_{i=1}^T \frac{1}{(P(C(x)_i))^{1/T}} \quad (4)$$

Self-BLEU While BLEU score (Papineni et al., 2002) is usually used to measure the correctness in machine translation, Self-BLEU (Zhu et al., 2018) was used here to measure the diversity of chatbot responses; we calculated the average BLEU score between sentences in our testing result as the Self-BLEU score.

Interlocutor while training	Interlocutor while testing	Sentence Length				Emotion (Anxiety)				Specific Words (Food)			
		$R_L \uparrow$	CPPL \downarrow	PPL \downarrow	SB-3 \downarrow	$R_E \downarrow$	CPPL \downarrow	PPL \downarrow	SB-3 \downarrow	$R_W \uparrow$	CPPL \downarrow	PPL \downarrow	SB-3 \downarrow
-	GPT-2	7.12	50.72	40.26	0.62	1.8	-	-	-	0.03	-	-	-
GPT-2	GPT-2	9.5	31.82	22.84	0.79	0.78	37.39	22.61	0.77	0.38	95.68	50.7	0.68
-	Google	3.74	49.89	39.81	0.62	1.82	-	-	-	0.02	-	-	-
Google	Google	10.14	110.59	41.67	0.91	0.7	26.57	10.99	0.8	0.002	97.84	48.27	0.8
-	Ret	12.77	50.3	39.47	0.62	1.77	-	-	-	0.08	-	-	-
Ret	Ret	19.79	76.55	18.46	0.94	0.75	26.57	10.98	0.8	1.29	69.0	35.7	0.81
GPT-2 + Google	GPT-2	8.52				0.52				0.51			
+ Ret	Google	4.31	39.68	30.2	0.75	0.5	39.95	34.05	0.71	0.51	72.4	40.55	0.8
	Ret	14.79				0.5				0.45			
Google + Ret	GPT-2	7.95	49.31	36.13	0.78	0.53	40.27	34.15	0.71	0.08	64.33	15.55	0.99
GPT-2 + Ret	Google	5.75	59.95	21.56	0.8	0.49	41.65	33.36	0.73	0.00	64.18	51.8	0.99
GPT-2 + Google	Ret	14.85	44.0	37.9	0.71	0.51	40.0	35.71	0.72	0.12	246.34	15.6	1

Table 1: Results of metrics and rewards according to different controllable factors. The metrics of Conditional Perplexity (CPPL), Perplexity (PPL), and Self-BLEU3(SB-3) are only examined on the guiding chatbot. Rewards are calculated on the interlocutor models during testing. The baseline performance is tested by the original guiding chatbot, the DialogPT pre-trained model that has not yet trained with any interlocutor model. Higher scores for R_L and R_W indicate better performance. Lower scores for R_E , CPPL, PPL, and SB-3 indicate better performance. The best results are boldfaced.

5.5 Human Evaluation Setups

For human evaluation, we recruited participants online. There were 19 participants; most of them were graduate or undergraduate students. Each participant was given several conversations, including an opening sentence and a corresponding response. They were asked to try to understand the conversation and provide a response to reply to the conversation. Therefore, we were able to collect numerous participants’ responses to calculate rewards. Moreover, participants were asked to score the relevance of the guiding chatbot’s response to the opening sentence. This task was rated on a Likert scale (Likert, 1932), ranging from 1 to 5: Score 1 means a firm disagreement, Score 3 meant neutral, and Score 5 meant an undoubted approval. Finally, we counted rewards from humans’ responses corresponding to the methods mentioned in Section 4.

6 Discussion and Analysis

6.1 Extending Sentence Length

The first controllable factor was sentence length. We aimed to guide the interlocutor to say more words in a single sentence. Table 1 reveals that our chatbot possessed the ability to encourage the interlocutor to be more talkative. The guiding chatbot interacted with the Google model while training could induce the interlocutor model to increase its sentence length from 3 to 10 words on average. However, as the sentence length increased, the conditional perplexity rose simultaneously. The result

reflected that the guiding chatbot trained with the Google model was forced to generate weird sentences so that the interlocutor model would produce a longer sentence. In contrast, although the guiding chatbot trained with the Retrieval model suffered from the same problem, the conditional perplexity increased only slightly, from 50.3 to 76.55, and the sentence length was much longer. Still, the high Self-BLEU3 score indicates that our chatbot might encounter a low-diversity problem. Therefore, the guiding chatbot trained with the GPT-2 model was the most desirable and stable chatbot to extend the interlocutor’s sentence length.

6.2 Guiding Emotion

The second task was to induce the interlocutor to speak with a particular emotion. These emotions included *anger*, *anxiety*, *contentment*, *disgust*, *hope*, *joy*, *sadness*, *surprise*. We examined the MSE loss between these emotions and the detected emotions of test sentences. Fig. 3a demonstrated that after training, all three interlocutors had similar performance in each emotion. Furthermore, Table 1 indicates that all guiding chatbots trained with any interlocutor model significantly decreased the MSE loss against baseline performance. As a result, independent of the choice of interlocutor model, our chatbot could successfully guide the interlocutor to speak with a specific emotion.

Positive Emotions Versus Negative Emotions

We investigated how positive/negative of the in-

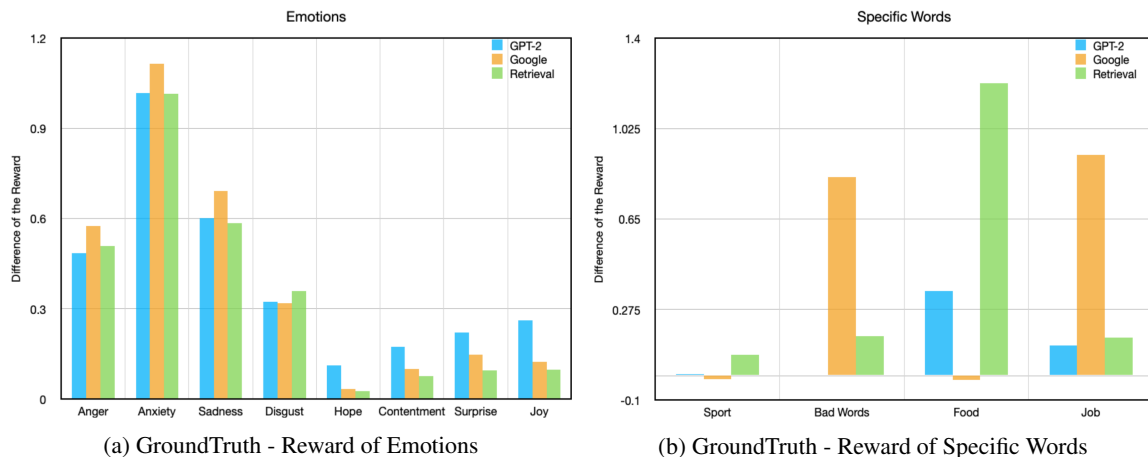


Figure 3: Experiments on controllable factors. The heights of the bars indicate the differences between the rewards of the interlocutor models before and after training.

	Positive	Negative
GPT-2	1.12	1.40
Google	1.05	1.41
Retrieve	1.08	1.42

Table 2: The MSE scores on the positive and negative emotions of the interlocutors without any fine-tuned.

terlocutors that interacted with our model without any fine-tuning. Table 2 shows that all three interlocutors responded more with positive emotions than with negative emotions.

Then, we evaluated how our chatbot realizes the way to influence the interlocutor. Figure 3a shows the difference between the MSE scores of the ground truth sentences and the MSE scores of the test sentences. We found that the improvements for negative emotions are greater than those of positive emotions. Table 2 shows that the average MSE scores of negative emotions is greater than positive emotions. According to the Fig. 3a, the Google model was easier to guided to reply with negative emotions, such as anxiety, sadness, and disgust. In comparison, the GPT-2 model was more easily encouraged to speak with positive emotion, such as joy, surprise, hope, and contentment. We attribute this phenomenon to the datasets underpinning each of these chatbots. The Google model was trained on the Cornell Movie Dialogue dataset, whereas the GPT-2 model was fine-tuned using the ED dataset. The movie dataset is full of simple, dramatic, exaggerated sentences. On the other hand, the ED dataset, designed to arouse the participants' sympathy tends be more positive. Furthermore, the Fig. 3a also displays that our chatbot performs ex-

ceptionally well on inducing the interlocutor speak with anxiety. The difference in the Google model's reward was up to 0.7, which means that we can significantly induce the interlocutor to speak with anxious emotion.

6.3 Inducing Specified Words

In another set of trials, our chatbot managed to make the interlocutor sentences contain certain groups of words, such as Food, Sports, Jobs, and Bad Words. We calculated the frequency of a word in a specific group. Table 1 shows that the ground truth's reward was close to 0, which suggests that the interlocutor models barely spoke words in the "Food" group before being exposure to by our guiding chatbot. Fig. 3b shows that our chatbot could successfully influence the interlocutor to talk about a sentence containing a word from the "Sports" group and "Food" group. On average, after interacting with the guiding chatbot, the Google model spoke 0.7 more words in the "Job" group, and the Retrieval model was induced to say 0.6 more words in the "Food" group. However, since the rewards of the ground truth are all near 0, Figure 3b indicates that fine-tuning the guiding chatbot using the RL approach can lead the interlocutor to say words they did not previously say.

We also found that the guiding chatbot trained with the GPT-2 model could only weakly induce the interlocutors to use words from the "Bad Word" group. This is almost certainly because bad words rarely appear in the ED dataset. The guiding chatbot trained with the Google model was more likely to induce the Google model interlocutor to say words in the "Bad Word" groups. We further an-

Interlocutor while training	Interlocutor while testing	Sentence Length		Emotion (Anxiety)		Specific Words (Food)	
		$R_L \uparrow$	Relevance \uparrow	$R_E \downarrow$	Relevance \uparrow	$R_W \uparrow$	Relevance \uparrow
-	Human	5.82	3.10	0.41	3.10	0.05	3.10
GPT-2 Google Ret	Human	6.05	2.10	0.27	3.89	0.16	2.63
		2.74	2.31	0.47	4.21	0.05	2.42
		5.90	1.52	0.46	3.68	0.21	1.47
GPT-2 + Google + Ret	Human	7.21	2.79	0.39	3.21	0.68	1.53

Table 3: Human Evaluation Results. Relevance represents the extent to which the guiding chatbot’s response is relevant. In contrast, the reward is based only on the interlocutor’s responses. We tested the baseline performance of the original guiding chatbot, the DialoGPT pre-trained model that has not yet trained with any interlocutor model. Top results were boldfaced.

alyzed the Cornell Movies dataset and found that, there are 24547 bad words out of 220579 sentences. We likewise concluded that dramatic utterances in the Cornell Movies dataset brought about the tendency for the interlocutor to say more bad words.

6.4 Cross Validation of Different Interlocutor Models while Training and Testing

Having proven that our guiding chatbot can significantly improve all three rewards against ground truth while training with a given interlocutor model, we experimented with the more formidable task of having the guiding chatbot consider all three interlocutor models at once. Table 1 demonstrates that the guiding chatbot could increase the performance, which indicates that the guiding chatbot could learn more experiences when interacting with more and different interlocutor models. While interacting with more models, the guiding chatbot can improve the "Emotion" and "Specific words" rewards against the guiding chatbot that was only trained with a single interlocutor model. Although the "Sentence Length" reward subtly decreased, the rewards still surpassed the ground truth reward, showing that the guiding chatbot could influence the interlocutor.

Moreover, since we could not assume that our interlocutor models are capable of representing all kinds of humans, we conducted an experiment to evaluate our guiding chatbot all-around. The detailed procedures are as follow: we tested our guiding chatbot on the interlocutor model that our guiding chatbot had seen before during training. For example, the guiding chatbot was trained with the GPT-2 and Google models but would be tested with the Retrieval model. Results in Table 1 shows that

all guiding chatbots trained with different interlocutor models could improve the rewards in three controllable factors. Also, we found that while testing on the Retrieval interlocutor model, this model was more likely to be induced to speak longer sentences than other interlocutor models. It is mainly because retrieving a longer response is easier than generating.

6.5 Human Evaluation Result

Human evaluation results sufficiently verify the guiding chatbot’s effectiveness of influencing humans’ responses to certain extents. Since the performances of the "anxiety" emotion and "Food" group were relatively well, shown in Table 1, we focused on these factors when conducting the human evaluation. Table 3 shows that the guiding chatbot could significantly induce humans to speak with anxiety, as well as maintain, or even enhance, the relevance within a conversation. This performance was consistent with the results in Table 1, in which the guiding chatbot acquired the ability to gain better rewards.

Nonetheless, the results of "Sentence Length" and "Specific Words" can hardly show a promising effect. Although the reward gained improvement slightly, humans generally felt the guiding chatbot’s response irrelevant: as the reward increased, the relevance decreased dramatically. This result demonstrates that the guiding chatbot might learn a tricky approach to gain higher rewards during training, but this method was not fully adaptive to humans. For instance, when training the guiding chatbot to influence the interlocutor to speak the sentence with the "Food" group, the guiding chatbot usually ended up with "What is your favorite

food?", ignoring the context. In contrast, the guiding chatbot could not only increase R_E reward but also improve the coherence between responses of the guiding chatbot and the interlocutor models.

6.6 Effects of R_C

We analyzed the effects brought by R_C . We trained a guiding chatbot model without R_C reward on aforementioned experimental settings in Section 5.3 and observed that the model was more prone to giving low diversity responses that were irrelevant to the context. In our experiments, the Self-BLEU3 score was near 0.99 and the CPPL was over 10000 without R_C reward.

7 Conclusion

This paper introduced a novel framework that aims to train a guiding chatbot to influence the interlocutor. We designed three different controllable factors for the guiding chatbot to induce the interlocutor to reply with responses matching the intention. We managed to prolong the length of the interlocutor's responses, influence the interlocutor to reflect with a particular emotion, and induce the interlocutor to use some specific words more frequently. Furthermore, we further enhanced the performance of the guiding chatbot by training it with more interlocutor models. Experiment results show that our proposed framework can successfully train chatbot with intentions.

Ethics

In this paper, we proposed a learning framework that trains chatbots to influence humans. We defined several rewards to reflect different behaviors that we want to induce to humans.

We undertook this work because we envisioned a future in which a chatbot can become a digital companion for humans. To that end, we need the chatbot to be able to understand a human's mental state and reply with appropriate responses. As a concrete example, chatbots could act as healthcare or relationship coaches for people who could not afford such services. Having a healthcare chatbot to talk to at anytime could alleviate the workload of nurses and therapists. Moreover, since our framework is reward-agnostic that could be optimized for any reward, we also expect that the experts could customize the profession reward definitions in their fields to bring the technique to higher level usage.

However, we also acknowledge the potential that this technique could be misused. Using our framework, ill-intentioned people could train chatbots with negative intentions and could threaten the stability of our society. For example, we have identified the following means by which a malicious actor could take advantage of our proposed technology:

- **Emotional Manipulation:** One could train chatbots with the intention of arousing negative emotions such as anxiety, sadness, or anger to influence human's mental state.
- **Social Antagonism:** One could train chatbots with the "Specific Words Intention Reward" to induce the interlocutors to exhibit gender biases or use racist terms to purposefully destabilize society.
- **Political Interference:** One could train chatbots with the malicious intentions of manipulating the public's political opinion.

To prevent the aforementioned abuse of our method, we propose the following methods to counter them.

- **Intention Classifier:** We could train a dialogue classifier that classifies whether a chatbot is purposefully influencing humans. We believe this is technically achievable as we could find many works that aim to distinguish whether a sentence is generated by humans or not (Gao et al., 2020). To further refine this work, we could easily collect training datasets for this classifier by interacting with chatbots trained by our framework and other general-purpose chatbots. By doing this, we could inform humans when we detect that the chatbot they are conversing with is being manipulative.
- **Special Token:** In the future, biomimetic technologies could blur the boundary between a living being and an artifact. We suggest that if the chatbot model generates the sentences, the sentence needs to be labeled with some special flag to tell people whether the chatbot generates the sentence with the intention. For instance, we can add "<chatbot | intention>" before any chatbot's response with the intention to inform people that a chatbot is trying to influence them. This will make users aware that they are interacting with a chatbot and can

undermine the effectiveness of a malevolent attack.

- **Safety Layer:** Inspired by (Adiwardana et al., 2020), we could use a safety layer (e.g., an additional classifier) to filter out sensitive or toxic responses from chatbots during inference.

Future Work To avoid malicious actors taking our framework and train their own chatbot. The development of the **Intention Classifier** become an essential research topic. In future work, we would set the development of the Intention Classifier as the top priority. The functions of the Intention Classifier are not only detect the intention of a dialogue system, it can also have an ability to generalize to any other dialogue systems. With the power of Meta-Learning (Finn et al., 2017) the classifier is expected to train on a dialogue system with few data and could have the ability to detect whether sentences generated by the dialogue system are with intention.

As developers of emerging technologies, we also take responsibility for defining the boundaries of these technologies. We will continue to refine the aforementioned methods to ensure that the proposed methodology improves public welfare as we intend it to.

References

- Daniel Adiwardana, Minh-Thang Luong, David R. So, Jamie Hall, Noah Fiedel, Romal Thoppilan, Zi Yang, Apoorv Kulshreshtha, Gaurav Nemade, Yifeng Lu, and Quoc V. Le. 2020. [Towards a human-like open-domain chatbot](#).
- Stanley F Chen, Douglas Beeferman, and Roni Rosenfeld. 1998. Evaluation metrics for language models.
- Cristian Danescu-Niculescu-Mizil and Lillian Lee. 2011. Chameleons in imagined conversations: A new approach to understanding coordination of linguistic style in dialogs. In *Proceedings of the Workshop on Cognitive Modeling and Computational Linguistics, ACL 2011*.
- 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.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. [Model-agnostic meta-learning for fast adaptation of deep networks](#). In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1126–1135, International Convention Centre, Sydney, Australia. PMLR.
- M. Thelwall G. Paltoglou. 2013. Seeing stars of valence and arousal in blog posts. *IEEE Transactions on Affective Computing*, 4:116–123.
- Xiang Gao, Yizhe Zhang, Michel Galley, Chris Brockett, and Bill Dolan. 2020. Dialogue response ranking with large-scale human feedback data. In *EMNLP*.
- Prakhar Gupta, Vinayshekhar Bannihatti Kumar, Mukul Bhutani, and Alan W Black. 2019. [Writer-Forcing: Generating more interesting story endings](#). In *Proceedings of the Second Workshop on Storytelling*, pages 117–126, Florence, Italy. Association for Computational Linguistics.
- Jing Han, Zixing Zhang, and Bjorn Schuller. 2019. Adversarial training in affective computing and sentiment analysis: Recent advances and perspectives. *IEEE Computational Intelligence Magazine*, 14(2):68–81.
- Matthew Henderson, Rami Al-Rfou, Brian Strope, Yunhsuan Sung, Laszlo Lukacs, Ruiqi Guo, Sanjiv Kumar, Balint Miklos, and Ray Kurzweil. 2017. [Efficient natural language response suggestion for smart reply](#).
- Andreas Holzinger, Chris Biemann, Constantinos S. Pattichis, and Douglas B. Kell. 2017. [What do we need to build explainable ai systems for the medical domain?](#)
- Leslie Pack Kaelbling, Michael L Littman, and Andrew W Moore. 1996. Reinforcement learning: A survey. *Journal of artificial intelligence research*, 4:237–285.
- Chih-Wei Lee, Yau-Shian Wang, Tsung-Yuan Hsu, Kuan-Yu Chen, Hung-yi Lee, and Lin-shan Lee. 2018. Scalable sentiment for sequence-to-sequence chatbot response with performance analysis. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6164–6168. IEEE.
- Rensis Likert. 1932. A technique for the measurement of attitudes. *Archives of psychology*.
- Zhaojiang Lin, Peng Xu, Genta Indra Winata, Farhad Bin Siddique, Zihan Liu, Jamin Shin, and Pascale Fung. 2020. [Caire: An end-to-end empathetic chatbot](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(09):13622–13623.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of*

- the 40th annual meeting on association for computational linguistics, pages 311–318. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Hannah Rashkin, Eric Michael Smith, Margaret Li, and Y-Lan Boureau. 2019. Towards empathetic open-domain conversation models: A new benchmark and dataset. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5370–5381, Florence, Italy. Association for Computational Linguistics.
- Alan Ritter, Colin Cherry, and William B. Dolan. 2011. Data-driven response generation in social media. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 583–593, Edinburgh, Scotland, UK. Association for Computational Linguistics.
- J. A. Russell. 1980. A circumplex model of affect. *Journal of Personality and Social Psychology*, 39:1161–1178.
- Iulian Vlad Serban, Alessandro Sordoni, Yoshua Bengio, Aaron C. Courville, and Joelle Pineau. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*, pages 3776–3784. AAAI Press.
- Lifeng Shang, Zhengdong Lu, and Hang Li. 2015. Neural responding machine for short-text conversation. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1577–1586, Beijing, China. Association for Computational Linguistics.
- Jamin Shin, Peng Xu, Andrea Madotto, and Pascale Fung. 2019. Happybot: Generating empathetic dialogue responses by improving user experience look-ahead.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112.
- Richard S Sutton, David A McAllester, Satinder P Singh, and Yishay Mansour. 2000. Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems*, pages 1057–1063.
- Aniruddha Tammewar, Monik Pamecha, Chirag Jain, Apurva Nagvenkar, and Krupal Modi. 2018. Production ready chatbots: Generate if not retrieve. In *The Workshops of the The Thirty-Second AAAI Conference on Artificial Intelligence, New Orleans, Louisiana, USA, February 2-7, 2018*, volume WS-18 of AAAI Workshops, pages 739–745. AAAI Press.
- 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.
- Oriol Vinyals and Quoc Le. 2015. A neural conversational model. *arXiv preprint arXiv:1506.05869*.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256.
- Yu Wu, Wei Wu, Chen Xing, Can Xu, Zhoujun Li, and Ming Zhou. 2019. A sequential matching framework for multi-turn response selection in retrieval-based chatbots. *Computational Linguistics*, 45(1):163–197.
- Rui Yan, Yiping Song, and Hua Wu. 2016. Learning to respond with deep neural networks for retrieval-based human-computer conversation system. In *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '16*, page 55–64, New York, NY, USA. Association for Computing Machinery.
- Yinfei Yang, Steve Yuan, Daniel Cer, Sheng-yi Kong, Noah Constant, Petr Pilar, Heming Ge, Yun-Hsuan Sung, Brian Strope, and Ray Kurzweil. 2018. Learning semantic textual similarity from conversations. In *Proceedings of The Third Workshop on Representation Learning for NLP*, pages 164–174, Melbourne, Australia. Association for Computational Linguistics.
- Yizhe Zhang, Siqu Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020. DIALOGPT : Large-scale generative pre-training for conversational response generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 270–278, Online. Association for Computational Linguistics.
- Xiangyang Zhou, Daxiang Dong, Hua Wu, Shiqi Zhao, Dianhai Yu, Hao Tian, Xuan Liu, and Rui Yan. 2016. Multi-view response selection for human-computer conversation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 372–381, Austin, Texas. Association for Computational Linguistics.
- Yaoming Zhu, Sidi Lu, Lei Zheng, Jiaxian Guo, Weinan Zhang, Jun Wang, and Yong Yu. 2018. Tegygen: A benchmarking platform for text generation models. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 1097–1100.