

Intent-calibrated Self-training for Answer Selection in Open-domain Dialogues

Wentao Deng¹, Jiahuan Pei², Zhaochun Ren¹, Zhumin Chen¹, Pengjie Ren^{1*}

¹Shandong University, Qingdao, China

wentao.deng@mail.sdu.edu.cn

{zhaochun.ren, chenzhumin, renpengjie}@sdu.edu.cn

²Centrum Wiskunde & Informatica, Amsterdam, Netherlands

Jiahuan.Pei@cwi.nl

Abstract

Answer selection in open-domain dialogues aims to select an accurate answer from candidates. The recent success of answer selection models hinges on training with large amounts of labeled data. However, collecting large-scale labeled data is labor-intensive and time-consuming. In this paper, we introduce the predicted intent labels to calibrate answer labels in a self-training paradigm. Specifically, we propose intent-calibrated self-training (ICAST) to improve the quality of pseudo answer labels through the intent-calibrated answer selection paradigm, in which we employ pseudo intent labels to help improve pseudo answer labels. We carry out extensive experiments on two benchmark datasets with open-domain dialogues. The experimental results show that ICAST outperforms baselines consistently with 1%, 5%, and 10% labeled data. Specifically, it improves 2.06% and 1.00% of F1 score on the two datasets, compared with the strongest baseline with only 5% labeled data.

1 Introduction

Open-domain dialogue systems (ODSs) interact with users by dialogues in open-ended domains (Huang et al., 2020). The responses in ODS can be divided into different types, such as answer, gratitude, greeting, and junk (Qu et al., 2018). In this paper, we focus on the selection of answers, which aims to identify the correct answer from a pool of candidates given a dialogue context. Typically, there are two main branches of approaches to produce answers, i.e., generation-based methods and selection-based methods (Park et al., 2022). The former generate a response token by token; and the

latter select a response from a pool of candidates. Currently, pure generation methods such as ChatGPT still face challenges: (1) They may generate incorrect content. (2) They cannot generate timely answers. Thus, it still needs selection-based methods to improve the correctness and timeliness of generation-based method.

Figure 1 illustrates our idea by comparing the answer selection paradigms of (a) context-aware methods, (b) intent-aware methods, and (c) intent-calibrated methods. Context-aware methods (see Figure 1(a)) capture the context of the ongoing dialogue for understanding users' information needs to select the most relevant responses from answer candidates (Jeong et al., 2021). Unlike task-oriented dialogue systems, it is much more challenging for ODSs to infer users' information needs due to their open-ended goals (Huang et al., 2020).

To this end, user intents, i.e., a taxonomy of utterances, are introduced to guide the information-seeking process (Qu et al., 2018, 2019a; Yang et al., 2020). If the intent of the previous (OQ) is not satisfied by the (PA) provided by a system, then the users' next intent is more likely to be information request (IR). For example, if the user asks: "Can you send me a website, so I can read more information?", the user's intent is IR. If the system does not consider the intent label IR, then it may provide an answer which does not satisfy the user's request.

Intent-aware methods (see Figure 1(b)) adopt intents as an extra input to better understand users' information needs in an utterance (Yang et al., 2020). However, they require sufficient human-annotated intent labels for training, the construction of which is time-consuming and labor-intensive.

*Corresponding author.

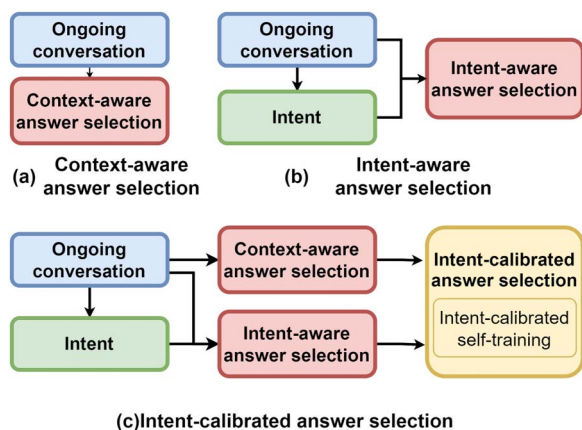


Figure 1: Comparison between previous answer selection models and our proposed framework. (a) Context-aware answer selection. (b) Intent-aware answer selection. (c) Intent-calibrated answer selection.

Self-training has been widely used to mitigate label scarcity problem (Liu et al., 2022; Yang et al., 2022; Zhang et al., 2022a). But it is still under-explored for answer selection in ODSs. The principle of self-training is to iteratively learn a model by assigning pseudo-labels for large-scaled unlabeled data to extend the training set (Amini et al., 2022). The teacher-student self-training framework has been widely used in much recent work, where the teacher generates pseudo-labels and the student makes predictions (Xie et al., 2020; Ghiasi et al., 2021; Li et al., 2021; Karamanolakis et al., 2021). However, noisy pseudo labels incur error propagation across iterations, so the key challenge is how to assure both quality and quantity of pseudo labels (Karamanolakis et al., 2021).

In this paper, we introduce an intent-calibrated answer selection paradigm, as in Figure 1(c). It first conducts both context-aware and intent-aware answer selection to predict pseudo intent and answer labels, and then it selects high-quality intent labels to calibrate final answer labels. To be more specific, we develop an intent-calibrated self-training (ICAST) algorithm based on the teacher-student self-training and intent-calibrated answer selection paradigm.

The core procedure is: First, we train a teacher model on the labeled data and predict pseudo intent labels for the unlabeled data. Second, we select high-quality intent labels by estimating intent confidence gain and then add selected intents to the input of the answer selection model. The intent confidence gain measures how much information a candidate intent label can bring to the

model. Third, we re-train a student model on both the labeled and pseudo-labeled data. Intuitively, ICAST synthesizes pseudo intent and answer labels and integrates them into teacher-student self-training, which can assure synthetic answer quality by high-quality intents.

We conduct experiments on two datasets: MS-DIALOG¹ (Qu et al., 2018) and MANTIS² (Penha et al., 2019). The experimental results show that ICAST outperforms the state-of-the-art baseline by 2.51%/0.63% of F1 score on the MSDIALOG/MANTIS dataset, with 1% labeled data. The results demonstrate the effectiveness of ICAST which selects accurate answers with incorporating high-quality predicted intent labels.

2 Related Work

In this section, we summarize related work in terms of three categories, i.e., traditional answer selection models, intent-aware answer selection models, and self-training for data argumentation.

2.1 Traditional Answer Selection Models

The dominant work focuses on modeling the representation of dialogue contexts, responses, and their relevance to select appropriate answers (Zhou et al., 2016, 2018; Chaudhuri et al., 2018). Wang et al. (2019) propose a sequential matching network to model the relation between the contextual utterances and the response by a cross-attention matrix. Yang and Choi (2019) encode dialogue contexts and responses for answer utterance selection and answer span selection using multiple self-attention models, e.g., R-Net (Wang et al., 2017) based on RNN and QANet (Yu et al., 2018) based on CNN. Many researchers also explore to enhance the dialogue contexts or candidate responses. Medveđ et al. (2020) extend the input candidate sentence with selected information from preceding sentence context. Fu et al. (2020) extend the contexts of the responses and integrate the context-to-context matching with context-to-response matching. Several studies (Ohmura and Eskenazi, 2018; Barz and Sonntag, 2021) also propose to improve the quality of answers by re-ranking answer candidates.

More recently, transformer-based pre-trained models have been the state-of-the-art paradigms

¹<https://citr.cs.umass.edu/downloads/msdialog/>.

²<https://github.com/guzpenha/mantis>.

(Kim et al., 2019; Henderson et al., 2019a; Tao et al., 2021). Researchers (Henderson et al., 2019b; Yang and Choi, 2019) apply a BERT encoder (Devlin et al., 2019) pre-trained on large-scaled open-domain dialogue corpus and fine-tune the model on small-scale in-domain dataset to capture the nuances. Likewise, Whang et al. (2020) also use a BERT encoder and perform context-response matching, but they also introduce the next utterance prediction and masked language modeling tasks during the post-training. Gu et al. (2020) incorporate speaker-aware embeddings into BERT to help with context understanding in multi-turn dialogues. Liu et al. (2021a) conduct utterance-aware and speaker-aware representations for dialogue contexts based on masking mechanisms in transformer-based pre-trained models, including BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019b), and ELECTRA (Clark et al., 2019).

There are several studies which use auxiliary tasks to enhance answer selection. Wu et al. (2020) incorporate a BERT-based response selection model with a contrastive learning objective and multiple auxiliary learning tasks, i.e., intention recognition, dialogue state tracking, and dialogue act prediction. Xu et al. (2021) enhance the response selection task with several auxiliary tasks, which can bring in extra supervised signals in multi-task learning manner. Pei et al. (2021) jointly learn missing user profiles with personalized response selection, which can improve response equality gradually based on enriched user profiles and neighboring dialogues.

2.2 Intent-aware Answer Selection Models

Intent detection is a key prior to understand users' intent for answer selection, especially in multiple turn dialogues (Gu et al., 2020; Park et al., 2022). Various deep NLP models have been adopted to classify intents (Chen et al., 2017; Liu et al., 2019a; Weld et al., 2021; Wang et al., 2021a). Chen et al. (2016) generate new intents to bridge the semantic relation across domains for intent expansion and classification. Wu et al. (2020) improve pre-trained BERT with an extra contrastive objective for intention recognition. The key challenge is natural language understanding with the state-of-the-art NLP models, e.g., CNNs (Chen et al., 2016), RNNs (Firdaus et al., 2021), transformers (Zhao et al., 2020), and pretrained language models (PLMs) (Wu et al., 2020; Yan et al., 2022).

Intent calibration research attempts to predict additional information to resolve users' ambiguous or uncertain intents. Lin and Xu (2019) calibrate the confidence of the softmax outputs for unknown intent detection. Gong et al. (2022) represent labels in hyperspherical space uniformly and calibrate confidence to trade-off accuracy and uncertainty. However, none of the above research has adapted the detected intents to answer selection. The most related work is IART (Yang et al., 2020), which weights the context by attending predicted intents for response selection.

Unlike the above methods, we propose to improve the performance of answer selection by using a large amount of unlabeled data. We devise the intent-calibrated self-training to improve the quality of pseudo answer labels by considering user intents.

2.3 Self-training for Answer Selection

Self-training has received remarkable attention in natural language processing (Luo, 2022) and machine learning (Karamanolakis et al., 2021; Amini et al., 2022). In general, the core idea is to augment the model training with pseudo supervision signals (Wu et al., 2020; Yan et al., 2022).

Sachan and Xing (2018) introduce a self-training algorithm for jointly learning to answer and generate questions, which augments labeled question-answer pairs with unlabeled text. Wu et al. (2018) introduce a pre-trained sequence-to-sequence model as an annotator to generate pseudo labels for unlabeled data to supervise the training process. Deng et al. (2021) propose to use the fine-tuned question generator and answer generator to generate pseudo question-answer pairs. Lin et al. (2020) introduce a fine-tuned generation-base model to generate gray-scale data.

Differently, the proposed ICAST in this work seeks to improve the quality of pseudo answer labels by introducing the intent-calibrated pseudo labeling mechanism which uses high-quality pseudo intent label to calibrate pseudo answer labels.

3 Preliminary

3.1 Answer Selection Task

We form answer selection as a binary classification task (Yang et al., 2020). We denote the

labeled dataset $D^l = \{([\mathbf{x}_i, \mathbf{e}_i], y_i)\}_{i=1}^{|D^l|}$ and unlabeled dataset $D^u = \{\mathbf{x}_i\}_{i=1}^{|D^u|}$. For the i -th sample, $\mathbf{x}_i = (\mathbf{u}_i, a_i)$ is a context-candidate pair, which consists of a sequence of utterances as the context $\mathbf{u}_i = [u_1, \dots, u_{|\mathbf{u}_i|}]$ and a candidate answer $a_i \in A$ (a set of all candidate answers). $\mathbf{e}_i = [e_1, \dots, e_{|\mathbf{u}_i|}]$ is a sequence of user intent labels. $y_i \in \{0, 1\}$ is the answer label, $y_i = 1$ denotes a_i is a correct answer, otherwise $y_i = 0$.

Our task is to learn a model $f = [f^\alpha, f^\beta]$. The intent generation module f^α predicts a set of intents $\tilde{\mathbf{e}}_i$ given a context-candidate pair, parameterized by α ; The answer selection module f^β predicts an answer label given context and the predicted intents, parameterized by β . Formally, we estimate the following probabilities:

$$f^\alpha := P(\mathbf{e}_i | \mathbf{u}_i, a_i, \alpha), \quad (1)$$

$$f^\beta := P(y_i | \mathbf{u}_i, \tilde{\mathbf{e}}_i, a_i, \beta). \quad (2)$$

3.2 BERT for Answer Selection

BERT (Devlin et al., 2019) is widely used to model the semantic dependency between context and candidate answers in recent research (Qu et al., 2019b; Li et al., 2019; Matsubara et al., 2020; Yang et al., 2020). First, we format the input of BERT as $\mathbf{x}_i = [[CLS]; \mathbf{u}_i; [SEP]; a_i]$, where the special token $[CLS]$ indicates the beginning of a context-candidate pair, and $[SEP]$ is a separator. Then, we use BERT to encode \mathbf{x}_i and get the representation of $[CLS]$ token \mathbf{h}_i^{CLS} . Next, let \mathbf{h}_i^{CLS} pass through a linear layer followed by an activation function to compute the probability p_i of a candidate answer. Formally,

$$p_i = \sigma(\mathbf{W}\mathbf{h}_i^{CLS} + \mathbf{b}), \quad (3)$$

$$\mathbf{h}_i^{CLS} = BERT(\mathbf{x}_i), \quad (4)$$

where \mathbf{W} and \mathbf{b} are trainable parameters, σ is sigmoid function.

3.3 Teacher-Student Self-training Framework

The teacher-student self-training framework (Li et al., 2021) is shown in Figure 2(a). It first trains the teacher model with the labeled data D^l to predict correct answer probabilities. Then at each iteration, the pseudo labeling module selects samples by using teacher's predictions to assign pseudo answer labels. Finally, the student model

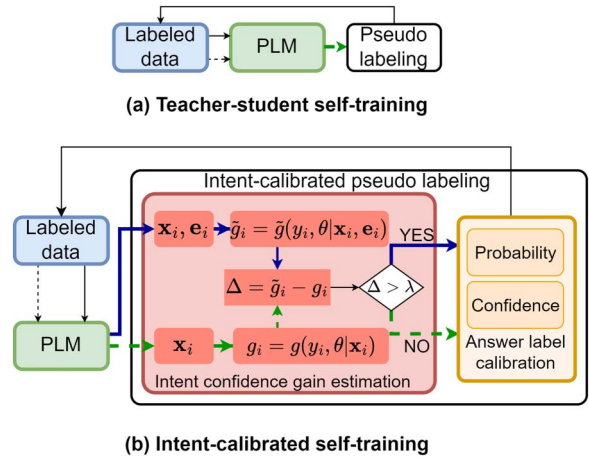


Figure 2: Comparison of self-training frameworks. (a) Teacher-student self-training framework. (b) Intent-calibrated self-training framework. The dashed thin line and solid thin line represent the workflow of teacher model and student model, respectively. The dashed thick line and solid thick line represent intent-aware and context-aware workflow.

is trained with the labeled data and pseudo-labeled data. At the next iteration, the student model is used as a new teacher model.

Pseudo Labeling Module. The principle is to determine a subset of samples and assign the unlabeled samples with pseudo answer labels. Following Tur et al. (2005) and Amini et al. (2022), we introduce thresholds λ^+ and λ^- for the positive and the negative classes to select a subset of unlabeled data with which the classifier is the most confident. For each unlabeled data, the selection criterion is defined as:

$$d = (\exists! p_i > \lambda^+) ? 1 : 0, \quad (5)$$

where $\exists!$ means exists one and only one. If there exists one and only one candidate answer, the probability p_i (see Eq. 3) of which is larger than the positive threshold λ^+ , then $d = 1$ and we add the current sample to the subset for pseudo labeling. Then, the pseudo answer label y_i of each sample $\mathbf{x}_i \in X$ is assigned by:

$$y_i = \Phi(\lambda^+, \lambda^-, p_i) = t \begin{cases} 1, & p_i > \lambda^+ \\ 0, & p_i < \lambda^- \end{cases}. \quad (6)$$

If the probability p_i is sufficiently high ($p_i > \lambda^+$), then positive label ‘‘1’’ is assigned to y_i ; if the probability p_i is sufficiently low ($p_i < \lambda^-$), then negative label ‘‘0’’ is assigned to y_i ; otherwise

Algorithm 1: Intent-calibrated self-training

input : Dataset w/ and wo/ labels $[D^l, D^u]$,
teacher model $f = [f^\alpha, f^\beta]$,
threshold λ of intent confidence gain.
output: Student model \tilde{f} .

- 1 Train teacher model f with D^l ;
- 2 **for** each $\mathbf{x}_i \in D^u$ **do**
- 3 $\mathbf{e}_i \leftarrow f^\alpha(\mathbf{x}_i)$; // Predict intents (Eq.7)
- 4 $g_i \leftarrow g(y_i, \beta|\mathbf{x}_i)$; // Context-aware confidence
 (Eq.10)
- 5 $\tilde{g}_i \leftarrow \tilde{g}(y_i, \beta|\mathbf{x}_i, \mathbf{e}_i)$; // Intent-aware confidence
 (Eq.11)
- 6 $\Delta = \tilde{g}_i - g_i$; // Intent confidence gain
- 7 **if** $\Delta > \lambda$ **then**
- 8 $\tilde{\mathbf{x}}_i \leftarrow (\mathbf{x}_i, \mathbf{e}_i)$; // Update input with intents
- 9 **else**
- 10 $\tilde{\mathbf{x}}_i \leftarrow \mathbf{x}_i$;
- 11 **end**
- 12 **end**
- 13 Select samples by Eq.5 and Eq.12;
- 14 Assign a label for each sample by Eq.14 and
 collect a dataset $D^p = \{(\tilde{\mathbf{x}}_i, y_i)\}_{i=1}^{|D^p|}$;
- 15 Re-train \tilde{f} on $D^l \cup D^p$ with 5 epochs;
- 16 $f \leftarrow \tilde{f}$; // Update teacher model with student model
- 17 Back to line 2 to iterate line 2-16 until the
 maximum epoch.

$p_i \in [\lambda^-, \lambda^+]$, y_i cannot be assigned a pseudo answer label, and this sample will not be used to train the student model.

4 Intent-calibrated Self-training

4.1 Overview

We illustrate the proposed (ICAST), as shown in Figure 2(b). First, we train a teacher model on labeled data D^l to predict pseudo intent labels for unlabeled data D^u (see §4.2). Second, we conduct intent-calibrated pseudo labeling (See §4.3). Specifically, we estimate intent confidence gain to select samples with high-quality intent labels, and we calibrate the answer labels by incorporating selected intent labels as an extra input for answer selection. Third, we train the student model with labeled and pseudo-labeled data (See §4.4). We summarize the proposed intent-calibrated self-training in Algorithm 1.

4.2 Teacher Model Training

We first train a teacher model $f = [f^\alpha, f^\beta]$ with the labeled dataset D^l . The intent generation module f^α constructs its input as a sequence of tokens, i.e., $\mathbf{x}_i = [u_1; e_1; \dots; u_{|u_i|}; e_{|u_i|}; [SEP]]$. It generates an intent label e_j by computing the probability of a candidate intent label, where $j \in [1, |u_i|]$:

$$p_i^e = \sigma(W\mathbf{h}_i^{SEP} + b), \quad (7)$$

$$\mathbf{h}_i^{SEP} = BERT(\mathbf{x}_i). \quad (8)$$

The answer selection module f^β constructs its input as a sequence of tokens, i.e., $\mathbf{x}_i = [[CLS]; u_1; e_1; \dots; u_{|u_i|}; e_{|u_i|}; [SEP]; a_i]$ and computes the probability of a candidate answer as Eq. 3 to decide if it is the correct answer.

4.3 Intent-calibrated Pseudo Labeling

4.3.1 Intent Confidence Gain Estimation

The intent-aware calibrator selects high-quality intent labels by estimating intent confidence gain. Intent confidence gain refers to the increase in confidence score after considering the predicted intents. A larger intent confidence gain indicates that the predicted intent can bring a greater increase in confidence score. We define the intent confidence gain as:

$$\Delta = \tilde{g}(y_i, \beta|\mathbf{x}_i, \mathbf{e}_i) - g(y_i, \beta|\mathbf{x}_i), \quad (9)$$

where β is the model parameters sampled by (MC dropout). Eq. 9 is formulated as the difference of two terms. The first term $\tilde{g}(y_i, \beta|\mathbf{x}_i, \mathbf{e}_i)$ is the confidence score of MC dropout with predicted intents, while the second term $g(y_i, \beta|\mathbf{x}_i)$ is the confidence score of MC dropout alone. The difference of two terms refers to the increase in confidence score after considering the predicted intents.

The $g(y_i, \beta|\mathbf{x})$ is the confidence score of MC dropout (Gal et al., 2017), which measures the decrease in Shannon entropy of answer prediction after using MC dropout sampling, i.e., the difference between the entropy of posterior and the expectation of the entropy of posteriors with

MC dropout. Formally, it can be defined and approximated as:

$$\begin{aligned}
& g(y_i, \beta | \mathbf{x}_i) \\
&= \mathbf{H}[y_i | \mathbf{x}_i] - \mathbf{E}_{P(\beta)}[\mathbf{H}[y_i | \mathbf{x}_i; \beta]] \\
&\approx \mathbf{H}[\mathbf{E}_{P(\beta)}[y_i | \mathbf{x}_i; \beta]] - \mathbf{E}_{P(\beta)}[\mathbf{H}[y_i | \mathbf{x}_i; \beta]] \\
&\approx -\frac{1}{T} \sum_{t=1}^T p_t \cdot \log \frac{1}{T} \sum_{t=1}^T p_t \\
&\quad + \frac{1}{T} \sum_{t=1}^T p_t \cdot \log p_t,
\end{aligned} \tag{10}$$

where the $\mathbf{H}[\cdot]$ is the Shannon entropy. The confidence score of MC dropout is calculated by the difference of two terms: the first term is the Shannon entropy with MC dropout, and the second term is the mean value of Shannon entropy with multiple MC dropout samplings.

Similarly, the confidence score of MC dropout with predicted intents $\tilde{g}(y_i, \beta | \mathbf{x}_i, \mathbf{e}_i)$ can be defined and approximated as:

$$\begin{aligned}
& \tilde{g}(y_i, \beta | \mathbf{x}_i, \mathbf{e}_i) \\
&= \mathbf{H}[y_i | \mathbf{x}_i, \mathbf{e}_i] - \mathbf{E}_{P(\beta)}[\mathbf{H}[y_i | \mathbf{x}_i, \mathbf{e}_i; \beta]] \\
&\approx \mathbf{H}[\mathbf{E}_{P(\beta)}[y_i | \mathbf{x}_i, \mathbf{e}_i; \beta]] - \mathbf{E}_{P(\beta)}[\mathbf{H}[y_i | \mathbf{x}_i, \mathbf{e}_i; \beta]] \\
&\approx -\frac{1}{T} \sum_{t=1}^T p_i^e \tilde{p}_t \cdot \log \frac{1}{T} \sum_{t=1}^T p_i^e \tilde{p}_t \\
&\quad + \frac{1}{T} \sum_{t=1}^T p_i^e \tilde{p}_t \cdot \log p_i^e \tilde{p}_t,
\end{aligned} \tag{11}$$

the confidence score of MC dropout with predicted intents indicates the decrease of Shannon entropy of answer prediction after using MC dropout sampling with considering predicted intents. It includes the predicted intents as inputs to the model, which is different from Eq. 10.

The first term of intent confidence gain is the confidence score of MC dropout after considering pseudo intents and the second term of intent confidence gain is the confidence score of MC dropout. The intent confidence gain is to measure how much confidence can pseudo intents can bring to the model with MC dropout. The higher the intent confidence gain, the more improvement that predicted intents can bring to the confidence score. We set a threshold λ to determine if the predicted intents can bring enough improvement to the confidence score. If the intent confidence

gain is larger than λ , we conclude that the predicted intents can improve the confidence score sufficiently and add them to the model's inputs. Specifically, if $\Delta > \lambda$, then we update the input with extra predicted intent labels \mathbf{e}_i , i.e., $\tilde{\mathbf{x}}_i = [\mathbf{x}_i, \mathbf{e}_i]$, which is expected to bring higher confidence score to the model, otherwise $\tilde{\mathbf{x}}_i = \mathbf{x}_i$.

4.3.2 Answer Label Calibration

To make use of more unlabeled samples, we introduce and extra three thresholds $\tilde{\lambda}^+$, $\tilde{\lambda}^-$, and λ^h to revise Eq. 5 as:

$$d = (\bar{p}_i > \tilde{\lambda}^+) \wedge (g_i > \lambda^h) ? 1 : 0, \tag{12}$$

where $\lambda^- < \tilde{\lambda}^- \leq \tilde{\lambda}^+ < \lambda^+$ and therefore we can consider extra samples with probabilities $p_i \in [\lambda^-, \lambda^+]$. The probability \bar{p}_i is approximated by T times MC dropout. The threshold λ^h is to select samples with high confidence g_i . Formally, \bar{p}_i and g_i are defined as:

$$\begin{cases} \bar{p}_i = \frac{1}{T} \sum_{t=1}^T \tilde{p}_t p_i^e, g_i = \tilde{g}(y_i, \beta | \mathbf{x}_i, \mathbf{e}_i), \Delta > \lambda \\ \bar{p}_i = \frac{1}{T} \sum_{t=1}^T p_t, g_i = g(y_i, \beta | \mathbf{x}_i), \Delta \leq \lambda \end{cases} \tag{13}$$

To calibrate an answer label for each sample $\tilde{\mathbf{x}}_i$, we revise Eq. 6 as:

$$y_i = \begin{cases} \Phi(\lambda^+, \lambda^-, \bar{p}_i), \bar{p}_i \in [0, \lambda^-) \cup (\lambda^+, 1] \\ \Phi(\tilde{\lambda}^+, \tilde{\lambda}^-, \bar{p}_i), \bar{p}_i \in [\lambda^-, \lambda^+] \end{cases} \tag{14}$$

Afterward, we can get a pseudo labeled dataset $D^p = \{\tilde{\mathbf{x}}_i, y_i\}_{i=1}^{|D^p|}$.

Note that line 13 in Algorithm 1 shows the process of selecting pseudo answer labels for re-training the answer selection module of the student model: Eq. 5 is used to determine if we add a sample to the subset for pseudo labeling. Eq. 12 is a revision of Eq. 5, which aims to make use of more unlabeled samples by introducing extra three thresholds. The goal of selecting samples by criteria of Eq. 5 and Eq. 12 (line 13 of Algorithm 1) is to prepare a set of candidates in primaries for high-quality pseudo labeling.

4.4 Student Model Re-training

We re-train the student model $\tilde{f} = [\tilde{f}^\alpha, \tilde{f}^\beta]$ with the extended dataset $D^l \cup D^p$. We minimize

three types of binary cross entropy losses, i.e., intent generation loss \mathcal{L}_i^e , answer selection loss without intent labels \mathcal{L}_i , and answer selection loss with intent labels $\tilde{\mathcal{L}}_i$, which are calculated as follows:

$$\begin{aligned}\mathcal{L}_i^e &= -\mathbf{e}_i \log \tilde{f}^\alpha \mathbf{x}_i + (1 - \mathbf{e}_i) \log(1 - \tilde{f}^\alpha(\mathbf{x}_i)), \\ \mathcal{L}_i &= -y_i \log \tilde{f}^\beta \mathbf{x}_i + (1 - y_i) \log(1 - \tilde{f}^\beta(\mathbf{x}_i)), \\ \tilde{\mathcal{L}}_i &= -y_i \log \tilde{f}^\beta(\mathbf{x}_i, \mathbf{e}_i) \\ &+ (1 - y_i) \log(1 - \tilde{f}^\beta(\mathbf{x}_i, \mathbf{e}_i)).\end{aligned}\tag{15}$$

The intent generation loss \mathcal{L}_i^e calculates the cross entropy loss between predicted intents and ground-truth intents. It can be used to optimize the intent generation module \tilde{f}^α .

The answer selection loss without intent labels \mathcal{L}_i calculates the cross entropy loss between predicted answers and ground-truth answers. It can be used to optimize the answer selection module \tilde{f}^β when the intent confidence gain is lower than the threshold.

The answer selection loss with intent labels $\tilde{\mathcal{L}}_i$ calculates the cross entropy loss between predicted answers and ground-truth answers. It can be used to optimize the answer selection module \tilde{f}^β when the intent confidence gain is larger than the threshold.

5 Experimental Setup

5.1 Datasets and Evaluation Metrics

We test all methods on our extension of two benchmark datasets: MSDIALOG (Qu et al., 2018) and MANTIS (Penha et al., 2019). The MSDIALOG dataset contains multi-turn question answering across 4 topics collected from the Microsoft community.³ It has 12 different types of intents. The MANTIS dataset provides multi-turn dialogs with user intent labels across 14 domains crawled from Stack Exchange.⁴ It has 10 different types of intents. Note that we require a small number of data with intent labels in our experiments. There are other response selection datasets (e.g., UDC [Lowe et al., 2015]), however, they do not contain dialogues with intent labels. To this end, we select the MSDIALOG and MANTIS datasets which contain a small amount of data with

³<https://answers.microsoft.com>.

⁴<https://stackexchange.com/>.

		Train		Validation	Test
		Labeled	Unlabeled		
MSDIALOG	1%	1,410	140,420	5,000	21,280
	5%	7,050	134,780	5,000	21,280
	10%	14,100	127,730	5,000	21,280
MANTIS	1%	2,640	260,990	12,000	50,000
	5%	13,200	250,430	12,000	50,000
	10%	26,400	237,230	12,000	50,000

Table 1: The statistics of experimental datasets, where labeled proportion denotes the proportion of labeled data in the training set.

intent labels, which can satisfy our experimental requirements.

In particular, we extend both datasets with unlabeled data. For MSDIALOG, we treat data without intent labels as unlabeled data; for MANTIS, we crawl unlabeled data from Stack Exchange⁴ from 2021 to 2022. For fair comparison with baselines, we follow previous work (Zhang et al., 2022b; Yang et al., 2020; Han et al., 2021): We use the ground-truth label as positive sample and use BM25 algorithm (Robertson and Zaragoza, 2009) to retrieve 9 relevant samples from different dialogues as negative samples. There could be a small number of negatives that are false-negatives, because there are cases which have same answers in different dialogues, so false-negatives may exist, but the number is very small. Besides, we use a different data partitioning strategy: First, we only extract conversations containing accurate answers, and the ground truth labels of all data are accurate answers. This is because we focus on the answer selection task, while the prior works focus on the response selection task. Note that not all the responses can serve as answers to users’ questions. Second, we put the data with intent labels into the training set. This is because the amount the data with intent labels is small, and we want to fully utilize the intent labels. In order to compare different methods in low-resource settings, we design three different low-resource simulation experiments: including 1%, 5%, and 10% labeled data and a large amount of unlabeled data. The statistics of the extended datasets is shown in Table 1.

We use 2 types of metrics to evaluate the models: classification metrics, i.e., Precision (P), Recall (R), and F1 score, and ranking metrics (Yang et al., 2020; Pan et al., 2021), i.e., mean average precision (MAP) and Recall@k (R@k).

5.2 Baselines

We compare the proposed ICAST with recent state-of-the-art methods that have reported results on the MSDIALOG and MANTIS datasets, respectively.

- **IART** (Yang et al., 2020) proposes the intent-aware attention mechanism to weight the utterances in context.
- **SAM** (Zhang et al., 2022b) captures semantic and similarity features to enhance answer selection.
- **JM** (Zhang et al., 2021) concatenates the context and all candidate responses as input to select the most proper response.
- **BIG** (Deng et al., 2021) uses the bilateral generation method to augment data and designs a contrastive loss function for training.
- **GRN** (Liu et al., 2021b) uses NUP and UOP pre-training tasks, and combines the graph network and sequence network to model the reasoning process of multi-turn response selection.
- **GRAY** (Lin et al., 2020) generates grayscale data by a fine-tuned generation model and proposes a multi-level ranking loss function for training.
- **BERT_FP** (Whang et al., 2020) learns the interactions between utterances in context to enhance answer selection.
- **BERT** (Devlin et al., 2019) is a general classification framework, which predicts answer labels on the vector of $[CLS]$ token.
- **Teacher-student self-training (TSST)** (Li et al., 2021) is a semi-supervised method, a teacher model is first trained with small, labeled data to generate pseudo labels on a large unlabeled dataset and then train a student model with pseudo labels.

5.3 Implementation Details

All models are implemented based on PyTorch⁵ and HuggingFace.⁶ We conduct hyper-parameter tuning on the validation dataset and report results on the test dataset. We use BERT-base-uncased model (Devlin et al., 2019) as an encoder in both

⁵<https://pytorch.org/>.

⁶<https://huggingface.co/>.

f^α and f^β , where the parameters are shared. We use AdamW (Loshchilov and Hutter, 2017) as optimizer. The batch size is 16, initial learning rate is $5e-5$, and weight decay is 0.01. The maximum number of context turn is set to 4. The maximum length of context and answer are set to 400 and 100. The dropout ratios is 0.1. ICAST generates pseudo labels every 5 epochs. MC dropout conducts sampling by $T = 5$ times. For thresholds of pseudo labeling, we set $\lambda^+ = 0.8$, $\lambda^- = 0.1$, $\tilde{\lambda}^+ = 0.5$, $\tilde{\lambda}^- = 0.5$ and $\lambda^h = 0.2$. For the threshold of intent confidence gain, we set $\lambda = 0.0$ for MANTIS dataset with 5% and 10% labeled data, otherwise, $\lambda = 0.02$.

For each parameter, we fix other hyper-parameters and select a specific value for the best performance on validation datasets. λ^- , λ^+ , and λ^h are selected in $(0, 1)$, the grid is 0.1. $\tilde{\lambda}^-$ and $\tilde{\lambda}^+$ are selected in $(0.1, 0.8)$, the grid is 0.1. $\tilde{\lambda}$ is selected in $(0, 0.05)$, the grid is 0.01. The number of ICAST’s parameters is 109,493,005. We train ICAST on 2 2080Ti GPUs with random seed 42, and the time cost is 48 hours.

6 Results

6.1 Overall Performance

We compare the overall performance of ICAST against the baseline methods. We also report the results of ICAST(Teacher). ICAST(Teacher) uses intent labels but BERT and BERT_FP do not, which is not a fair comparison. We conduct these experiments to see whether our method can outperform baselines without using unlabeled data. The results of overall performance are shown in Table 2.

First, in terms of all classification metrics, ICAST and ICAST (Teacher) outperform the baselines in each setting, excluding only one setting: The R score of ICAST (Teacher) is 0.42% lower than BERT trained on 10% labeled MSDIALOG dataset. Specifically, on the MSDIALOG dataset, ICAST with 1%, 5%, and 10% labeled data improves their corresponding strongest baselines by 2.51%, 2.06%, and 2.12% of F1 scores. On the MANTIS dataset, ICAST with 1%, 5%, and 10% labeled data improves their corresponding strongest baselines by 0.63%, 1.00%, and 0.62% of F1 scores. This demonstrates the effectiveness of ICAST on the performance of classifying correct answers. We believe there are two reasons: (i) the predicted intent labels can provide more

Setting	Model	MSDIALOG							MANTIS						
		P	R	F1	R@1	R@2	R@5	MAP	P	R	F1	R@1	R@2	R@5	MAP
1% labeled	IART [†]	22.18	46.75	30.08	25.65	46.28	77.58	47.74	48.29	52.22	50.18	50.40	68.34	86.22	66.12
	SAM [†]	44.17	44.36	44.26	46.89	59.06	77.02	60.72	57.75	58.62	58.18	65.10	76.32	88.54	75.60
	JM [‡]	44.80	44.59	44.70	44.54	60.76	84.30	61.26	62.95	62.62	62.78	62.64	77.32	92.30	75.25
	BIG [‡]	44.07	44.78	44.42	50.93	66.30	87.50	66.15	57.91	57.42	57.66	70.22	83.04	95.12	80.78
	GRAY [‡]	41.68	42.15	41.91	51.26	66.40	85.62	66.10	61.30	60.72	61.01	64.67	77.34	88.32	75.57
	GRN [‡]	43.41	43.37	43.39	43.28	61.60	86.46	61.19	61.75	61.10	61.42	61.06	76.64	93.66	74.56
	BERT_FP [†]	44.32	42.95	43.62	56.76	72.08	91.25	70.90	66.26	62.86	64.51	75.62	86.14	95.22	84.11
	BERT [‡]	48.56	45.34	46.90	54.79	68.32	85.80	68.04	67.28	65.62	66.44	74.82	83.00	92.16	82.41
	ICAST (Teacher)	49.82	46.33	48.01	56.86	67.81	85.38	69.03	68.48	66.12	67.28	77.28	86.12	94.98	82.98
1% labeled +all unlabeled	TSST [‡]	53.72	52.58	53.14	61.04	73.91	89.70	73.04	73.73	72.60	73.16	82.94	91.08	97.88	89.18
	ICAST	57.05	54.32	55.65	62.21	76.31	91.07	73.77	74.89	72.72	73.79	83.68	90.68	96.42	88.31
5% labeled	IART [†]	23.52	49.38	31.86	28.80	48.02	79.93	49.97	50.24	53.60	51.86	51.56	70.66	89.52	67.75
	SAM [†]	49.52	51.45	50.47	54.27	67.66	83.03	67.32	59.16	57.82	58.48	66.52	76.88	89.28	76.51
	JM [‡]	50.98	49.81	50.39	50.37	67.62	89.47	66.56	67.16	66.82	66.99	66.92	80.83	94.94	78.47
	BIG [‡]	50.82	50.93	50.88	58.12	73.07	89.80	71.53	61.34	60.88	61.11	74.22	87.58	96.94	84.02
	GRAY [‡]	48.99	48.26	48.62	55.16	69.54	86.23	68.75	62.53	66.50	64.45	70.24	80.92	90.62	79.48
	GRN [‡]	49.28	50.04	49.66	49.76	66.77	89.52	66.04	64.27	63.00	63.62	63.78	78.60	93.38	76.27
	BERT_FP [†]	49.74	50.93	50.33	62.96	77.16	92.76	75.41	70.04	68.32	69.17	80.22	89.36	97.30	87.37
	BERT [‡]	52.01	49.67	50.81	61.23	72.60	85.19	72.13	71.17	67.56	69.32	77.70	86.82	95.48	85.20
	ICAST (Teacher)	54.22	51.83	53.00	62.59	74.38	90.36	74.16	73.13	69.20	71.11	80.82	88.12	95.86	84.27
5% labeled +all unlabeled	TSST [‡]	58.34	58.78	58.56	64.89	74.62	86.41	74.61	74.33	72.92	73.61	81.62	89.32	96.10	87.83
	ICAST	61.54	59.72	60.62	69.54	80.35	93.09	77.29	74.60	74.62	74.61	84.38	90.76	97.06	89.78
10% labeled	IART [†]	34.38	47.22	39.79	39.05	58.31	84.77	58.00	50.77	53.04	51.88	51.80	71.20	89.28	68.04
	SAM [†]	55.63	54.27	54.94	59.53	70.62	85.05	71.00	61.39	60.00	60.69	66.88	77.92	90.84	77.08
	JM [‡]	57.64	57.56	57.60	57.61	73.12	90.97	71.70	68.06	68.98	68.52	68.22	80.46	94.04	79.02
	BIG [‡]	56.15	55.96	56.06	62.96	76.78	90.08	74.92	62.74	62.34	62.54	76.60	87.62	96.28	85.08
	GRAY [‡]	54.46	53.05	53.75	62.45	76.08	90.60	74.51	65.20	65.26	65.23	74.80	85.74	94.66	83.51
	GRN [‡]	54.06	53.43	53.74	53.52	70.67	90.08	68.96	66.01	64.92	65.46	66.10	80.60	93.34	77.83
	BERT_FP [†]	57.95	56.81	57.38	67.48	80.16	94.07	78.56	71.04	68.20	69.59	80.72	89.38	96.82	87.59
	BERT [‡]	61.94	60.19	61.05	64.38	73.77	85.99	74.12	70.33	69.56	69.94	82.12	91.00	97.70	88.72
	ICAST (Teacher)	62.41	<u>59.77</u>	61.06	66.54	76.55	89.09	76.43	71.89	70.24	71.05	81.92	90.02	97.12	88.29
10% labeled +all unlabeled	TSST [‡]	63.28	63.34	63.31	70.91	81.95	93.18	80.37	76.17	73.34	74.73	83.70	91.18	97.50	89.43
	ICAST	65.98	64.89	65.43	72.27	81.95	91.63	79.63	77.43	73.36	75.35	84.60	91.52	97.36	88.59

Table 2: Overall performance of answer selection. Bold and underlined fonts indicate leading and compared results in each setting. 1%, 5%, and 10% are the proportion of labeled data in training dataset. The symbol [†] indicates the baselines reproduced by the released source codes and [‡] indicates the baselines we implemented based on the papers. Note that we cannot fairly compare with the reported results in the IART paper, because we use a different data partitioning for a different task (see Section 5.1).

information that are useful for selecting correct answers; and (ii) the self-training paradigm can calibrate answer labels for continuous improvement. For example, with self-training on 10% labeled MSDIALOG dataset and all unlabeled data, the R score of ICAST is 4.70%/1.55% higher than BERT/TSST, respectively.

Second, in terms of ranking metrics, we have the following observations: (i) ICAST outperforms all baselines in terms of R@1 score in each setting. Specifically, on the MSDIALOG dataset, ICAST with 1%, 5%, and 10% labeled data achieve 1.17%, 4.65%, and 1.36% higher of R@1 scores than their corresponding strongest baselines, respectively. On the MANTIS dataset, ICAST with 1%, 5%, and 10% labeled data achieve 0.74%, 2.76%, and 0.90% higher of R@1 scores than their corresponding strongest baselines, respectively. It indicates that ICAST can rank an accurate answer on top. (ii) For R@2, R@5, and MAP, ICAST achieves the highest scores in most of the set-

tings, excluding: On the MSDIALOG dataset, with 10% labeled data and all unlabeled data, R@5 and MAP scores decrease 1.55% and 0.74%. On the MANTIS dataset, with 1% labeled data and all unlabeled data, R@2, R@5, and MAP scores decrease 0.40%, 1.46%, and 0.87%; with 10% labeled data and all unlabeled data, R@5 and MAP scores decrease 0.14% and 0.84%. Our method does not possess a significant advantage in terms of R@2, R@5, and MAP, as the primary objective of answer selection is to identify the answer rather than generating a ranking list. Hence, the fundamental performance measurements are precision, recall, and the F1-score (Wang et al., 2021b). Accordingly, we evaluate the models using standard evaluation metrics for a fair comparison. Additionally, we present supplementary ranking metrics (i.e., R@2, R@5, MAP) to assess whether improvements in selection metrics result in a noteworthy decline in ranking metrics. The results demonstrate that

Setting	Model	MSDIALOG							MANTIS						
		P	R	F1	R@1	R@2	R@5	MAP	P	R	F1	R@1	R@2	R@5	MAP
1% labeled +all unlabeled	ICAST	57.05	54.32	55.65	62.21	76.31	91.07	73.77	74.89	72.72	73.79	83.68	90.68	96.42	88.31
	-ICGE	54.81	54.04	54.42	60.24	72.85	86.46	71.83	74.44	71.78	73.08	81.68	89.02	94.86	87.55
	-ICGE-ALC	54.13	53.52	53.82	60.80	73.77	89.38	72.86	74.00	71.96	72.96	82.66	90.82	97.72	88.94
	-IG	53.72	52.58	53.14	61.04	73.91	89.70	73.04	73.73	72.60	73.16	82.94	91.08	97.88	89.18
5% labeled +all unlabeled	ICAST	61.54	59.72	60.62	69.54	80.35	93.09	77.29	74.60	74.62	74.61	84.38	90.76	97.06	89.78
	-ICGE	60.66	58.55	59.58	64.94	76.31	90.32	75.61	74.35	74.12	74.23	83.18	91.14	97.50	89.29
	-ICGE-ALC	59.94	58.92	59.43	66.30	77.82	91.63	76.86	74.19	73.56	73.87	82.80	90.72	96.66	88.81
	-IG	58.34	58.78	58.56	64.89	74.62	86.41	74.61	74.33	72.92	73.61	81.62	89.32	96.10	87.83
10% labeled +all unlabeled	ICAST	65.98	64.89	65.43	72.27	81.95	91.63	79.63	77.46	73.36	75.35	84.60	91.52	97.36	88.59
	-ICGE	65.71	62.78	64.21	71.42	81.53	93.70	80.51	76.19	73.82	74.98	83.82	91.26	96.98	89.46
	-ICGE-ALC	64.27	64.09	64.18	70.63	81.95	92.95	80.15	75.60	74.26	74.92	84.06	91.38	97.32	89.69
	-IG	63.28	63.34	63.31	70.91	81.95	93.18	80.37	76.17	73.34	74.73	83.70	91.18	97.50	89.43

Table 3: Ablation study. Impact of different modules in our proposed framework. \uparrow and \uparrow indicate an increase of the performance compared with ICAST and ICAST-ICGE, respectively.

our method exhibits no considerable decrease in ranking metrics.

Third, using self-training with unlabeled data has the largest impact on all settings in terms of both classification and ranking metrics. Specifically, on the MSDIALOG dataset with 1%, 5%, and 10% labeled data, F1 scores increase 7.64%, 7.62%, and 4.37%; MAP scores increase 2.87%, 1.88%, and 1.81%. On the MANTIS dataset with 1%, 5%, and 10% labeled data, F1 scores increase 6.51%, 3.50%, and 4.30%; MAP scores increase 5.07%, 2.41%, and 0.71%. This reveals that ICAST benefits from making good use of unlabeled data with self-training. Besides, the influence of classification performance is larger than the ranking performance in each setting.

Last but not least, we do not require too much data with intent labels. This also motivates us to conduct experiments on only a small amount of data with labels (1%, 5%, 10%). For example, with 1% labeled data, our method outperforms the baselines with only 141 and 264 intent labels on the MSDIALOG and MANTIS datasets, respectively. Thus, it is possible to apply our method in practice, even without a large amount of intent labels.

6.2 Ablation Study

To better understand the contribution of each functional component of ICAST, i.e., intent confidence gain estimation (ICGE), answer label calibration (ALC), and intent generation (IG), we conduct the following ablation studies: After removing the IG (denoted as “-IG”), ICGE and ALC do not work, so ICAST degenerates to the TSST model. The ICGE estimates intent confidence gain according to the predicted intents which are the outputs of the IG module. The

ALC uses the intent confidence scores to select unlabeled samples, the computation of intent confidence scores also needs the predicted intents. After removing the ICGE (denoted as “-ICGE”), the model does not select the predicted intents according to the intent confidence gain and adds all of predicted intent labels into the inputs. After removing the ICGE and ALC (denoted as “-ICGE-ALC”), the model does not select unlabeled data with intent confidence score and degenerates into the TSST with adding all predicted intents into inputs. Table 3 reports the results of the ablation studies.

First, intent confidence gain estimation (ICGE), answer label calibration (ALC), and intent generation (IG) have positive influence on overall performance of classification in all settings on both the MSDIALOG and MANTIS datasets with 1%, 5%, and 10% labeled data, respectively. Removing IG from ICAST, F1 scores decrease 2.51%/2.06%/2.12% on the MSDIALOG dataset and 0.63%/1.00%/0.62% on the MANTIS dataset. This proves our hypothesis that the generated intents can provide more useful information for selecting correct answers. Removing ICGE from ICAST, F1 scores decrease 1.23%/1.04%/1.22% on the MSDIALOG dataset and 0.71%/0.38%/0.37% on the MANTIS dataset. This reveals that intent confidence gain can select high-quality intent labels that are helpful to select correct answers. Removing ALC from ICAST without ICGE, F1 scores decrease 0.60%/0.15%/0.03%. This shows that ALC can bring extra improvement even though ICGE is absent. Meanwhile, it works better together with the other two components.

Second, in terms of ranking performance, R@1 decreases when removing ICGE, ALC, and IG

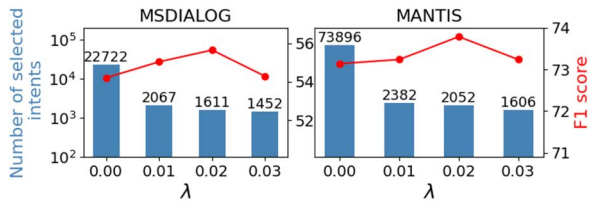


Figure 3: F1 scores (w.r.t. the line) and average number of selected intents for answer label calibration (w.r.t. the bar) with different values of $\lambda = [0.00, 0.01, 0.02, 0.03]$ on MSDIALOG (left) and MANTIS (right) with 1% labeled data.

from ICAST in all settings on the MSDIALOG and MANTIS datasets. Removing ICGE/ALC/IG with 1%, 5%, and 10% labeled data, R@1 drops 1.97%/1.41%/1.17%, 4.60%/3.24%/4.65%, and 0.85%/1.64%/1.36% on the MSDIALOG dataset; R@1 drops 2.00%/1.02%/0.74%, 1.20/1.58/2.76%, and 0.78%/0.54%/0.90% on the MANTIS dataset. This shows that the three functional components are helpful to rank correct answers on top.

6.3 Analysis

Figure 3 shows the impact of threshold of intent confidence gain λ on classification performance of ICAST. We can see that as λ increases, the average number of selected intents decreases. Meanwhile, F1 scores increase first, achieve top at $\lambda = 0.02$, and then descend, because with a larger λ , more intents are selected to calibrate the answer labels, which leads to an increase of F1 scores. Then, adding more generated intents might also introduce noise for answer selection, which is the possible reason for the decrease of F1. Thus, λ can balance between more predicted intents and less noisy intents.

6.4 Case Study

Table 4 shows a case study of how ICAST and TSST select different answers for the same given context.

In general, a model chooses the candidate answer with the highest probability among all candidate answers as the correct answer. In this case, the strongest baseline TSST incorrectly chooses the second candidate answer (A2) with the probability of 0.96, instead of the first candidate answer (A1) which has a probability of 0.00. It shows that selecting answers based solely on their

Context Utterances	Intents		
User: How does a photon picture make the pattern?	OQ		
Agent: Photons in mainstream physics, are quantum mechanical entities which in great numbers build up the classical electromagnetic radiation...	PA		
User: Do you know why the photon which is hitting forward is causing an electron to move up-down?	IR		
Candidate Answers	Model	ICG	Probability
A1: The theories of quantum mechanics for electron photon interactions can be found in https://www.website.com .	TSST	/	0.00
	ICAST	0.14	0.99
A2: The energy of a photon is equal to the level spacing of a two-level system. It is a result of energy conservation...	TSST	/	0.96
	ICAST	-0.13	0.71

Table 4: answers by the ICAST and TSST models. Each model chooses the candidate answer with the highest probability among all candidate answers as the correct answer. If ICG is larger than λ , then ICAST combines the intents and context to select answer. Here, $\lambda = 0.00$. Note that the first candidate answer (A1) is the correct answer.

probabilities can result in significant bias. ICAST calibrates the probabilities based on ICG, and it correctly chooses A1 with the probability of 0.99, while skipping A2 with the probability of 0.71. ICAST computes ICGs by combining context and its predicted intents, and each candidate answer. The ICG of the correct answer is larger than λ , which indicates that ICAST can capture the intent information from the correct answer, so ICAST increases the probability of correct answer from 0.00 to 0.99. Meanwhile, the ICG of the incorrect answer is less than λ , which indicates ICAST cannot capture the intent information from and incorrect answer, so ICAST decreases the probability of incorrect answer from 0.96 to 0.71. Furthermore, we explain the intuition. In context utterances, the user asks the (OQ), and then the agent gives a (PA) which can explain the original question, but the user still raises the (IR) to ask the agent for more detailed information. Next, the user anticipates an answer that includes a link or document providing more detailed information, instead of a continuous explanation in text. Intuitively, the predicted intents can aid in monitoring changes in the user’s expectations throughout the utterances.

7 Conclusion and Future Work

In this paper, we propose intent-calibrated self-training (ICAST) based on the teacher-student

self-training and intent-calibrated answer selection: We train a teacher model on labeled data to predict intent labels on unlabeled data; select high-quality intents by intent confidence gain to enrich inputs and predict pseudo answer labels; and retrain a student model on both the labeled and pseudo-labeled data. We conduct massive experiments on two benchmark datasets and the results show that ICAST outperforms baselines even with small but similar proportions (i.e., 1%, 5%, and 10%) of labeled data, respectively. Note that we understand that a greater proportion of labeled data may lead to an increase of performance, e.g., BERT-FP with 10% labeled data beat ICAST with 1% labeled data for across all metrics for MSDIALOG. However, we focus on verifying if the proposed ICAST outperforms other methods given a very few amounts of labeled data. In some cases, ICAST can outperform baselines with fewer labeled data. In the future work, we will explore more predictable dialogue context (e.g., user profiles) than intents.

8 Reproducibility

To facilitate reproducibility of the results reported in this paper, the code and data used are available at <https://github.com/dengwentao99/ICAST>.

Limitations

Our proposed ICAST also has the following limitations. First, ICAST only considers user intents to enhance answer selection. It is limited because we only capture the user's expectations from the predicted intent labels, without considering other user-centered factors, such as user profiles and user feedback. Second, like retrieval-based methods that have been shown to have a good effect in professional question-answering fields, ICAST also has limitations when it comes to diversity. For example, ICAST cannot retrieve multiple correct answers with different expressions given the same context. Third, since our model needs to predict the intent labels, to complete this task, the model needs a few additional parameters.

Ethics Considerations

We realize that there are risks in developing the dialogue system, so it is necessary to pay attention

to the ethics issues of the dialogue system. It is crucial for a dialogue system to give correct answers to the users while avoiding ethical problems such as privacy preservation problems. Due to the fact that we have used public datasets to train our model, these datasets are carefully processed by publishers to ensure that there are no ethical problems. Specifically, the dataset publishers performed user ID anonymization on all datasets, and only the tokens "user" and "agent" are used to represent the roles in the conversation. The utterances do not contain any user privacy information (e.g., names, phone numbers, addresses) to prevent privacy disclosure.

Acknowledgments

We would like to thank the editors and reviewers for their helpful comments. This research was supported by the National Key R&D Program of China (grants No.2022YFC3303004, No.2020YFB1406704), the Natural Science Foundation of China (62102234, 62272274, 62202271, 61902219, 61972234, 62072279), the Key Scientific and Technological Innovation Program of Shandong Province (2019JZZY010129), the Natural Science Foundation of Shandong Province (ZR2021QF129), the Fundamental Research Funds of Shandong University, and VOXReality (European Union grant 101070521). All content represents the opinion of the authors, which is not necessarily shared or endorsed by their respective employers and/or sponsors.

References

- Massih-Reza Amini, Vasilii Feofanov, Loic Pauletto, Emilie Devijver, and Yury Maximov. 2022. Self-training: A survey. *ArXiv:2202.12040v1*. <https://doi.org/10.48550/arXiv.2202.12040>
- Michael Barz and Daniel Sonntag. 2021. Incremental improvement of a question answering system by re-ranking answer candidates using machine learning. In *Increasing Naturalness and Flexibility in Spoken Dialogue Interaction*, pages 367–379, Springer, Singapore. https://doi.org/10.1007/978-981-15-9323-9_34

- Debanjan Chaudhuri, Agustinus Kristiadi, Jens Lehmann, and Asja Fischer. 2018. Improving response selection in multi-turn dialogue systems by incorporating domain knowledge. In *Proceedings of the 22nd Conference on Computational Natural Language Learning*, pages 497–507, Brussels, Belgium. Association for Computational Linguistics. <https://doi.org/10.18653/v1/K18-1048>
- Hongshen Chen, Xiaorui Liu, Dawei Yin, and Jiliang Tang. 2017. A survey on dialogue systems: Recent advances and new frontiers. *ACM SIGKDD Explorations Newsletter*, 19(2):25–35. <https://doi.org/10.1145/3166054.3166058>
- Yun-Nung Chen, Dilek Hakkani-Tür, and Xiaodong He. 2016. Zero-shot learning of intent embeddings for expansion by convolutional deep structured semantic models. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6045–6049, Shanghai, China. <https://doi.org/10.1109/ICASSP.2016.7472838>
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2019. ELECTRA: Pre-training text encoders as discriminators rather than generators. In *Proceedings of International Conference on Learning Representations*. <https://doi.org/10.48550/arXiv.2003.10555>
- Yang Deng, Wenxuan Zhang, and Wai Lam. 2021. Learning to rank question answer pairs with bilateral contrastive data augmentation. In *Proceedings of the Seventh Workshop on Noisy User-generated Text (W-NUT 2021)*, pages 175–181, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.wnut-1.20>
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of 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. <https://doi.org/10.18653/v1/N19-1423>
- Mauajama Firdaus, Hitesh Golchha, Asif Ekbal, and Pushpak Bhattacharyya. 2021. A deep multi-task model for dialogue act classification, intent detection and slot filling. *Cognitive Computation*. <https://doi.org/10.1007/s12559-020-09718-4>
- Zhenxin Fu, Shaobo Cui, Mingyue Shang, Feng Ji, Dongyan Zhao, Haiqing Chen, and Rui Yan. 2020. Context-to-session matching: Utilizing whole session for response selection in information-seeking dialogue systems. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1605–1613. <https://doi.org/10.1145/3394486.3403211>
- Yarin Gal, Riashat Islam, and Zoubin Ghahramani. 2017. Deep Bayesian active learning with image data. *ArXiv:1703.02910v1*. <https://doi.org/10.48550/arXiv.1703.02910>
- Golnaz Ghiasi, Barret Zoph, Ekin D. Cubuk, Quoc V. Le, and Tsung-Yi Lin. 2021. Multi-task self-training for learning general representations. In *Proceedings of IEEE/CVF International Conference on Computer Vision*. <https://doi.org/10.1109/ICCV48922.2021.00873>
- Yantao Gong, Cao Liu, Fan Yang, Xunliang Cai, Guanglu Wan, Jiansong Chen, Weipeng Zhang, and Houfeng Wang. 2022. Confidence calibration for intent detection via hyperspherical space and rebalanced accuracy-uncertainty loss. In *Proceedings of AAAI Conference on Artificial Intelligence*. <https://doi.org/10.1609/aaai.v36i10.21314>
- Jia-Chen Gu, Tianda Li, Quan Liu, Zhen-Hua Ling, Zhiming Su, Si Wei, and Xiaodan Zhu. 2020. Speaker-aware BERT for multi-turn response selection in retrieval-based chatbots. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management October*, pages 2041–2044. <https://doi.org/10.1145/3340531.3412330>
- Janghoon Han, Taesuk Hong, Byoungjae Kim, Youngjoong Ko, and Jungyun Seo. 2021. Fine-grained post-training for improving retrieval-based dialogue systems. In *Proceedings of North American Chapter of the Association for Computational Linguistics: Human Lan-*

- guage Technologies*, pages 1549–1558, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.naacl-main.122>
- Matthew Henderson, Ivan Vulić, Iñigo Casanueva, Paweł Budzianowski, Daniela Gerz, Sam Coope, Georgios Spithourakis, Tsung-Hsien Wen, Nikola Mrkšić, and Pei-Hao Su. 2019a. PolyResponse: A rank-based approach to task-oriented dialogue with application in restaurant search and booking. 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): System Demonstrations*, pages 181–186, Hong Kong, China. <https://doi.org/10.18653/v1/D19-3031>
- Matthew Henderson, Ivan Vulić, Daniela Gerz, Iñigo Casanueva, Paweł Budzianowski, Sam Coope, Georgios Spithourakis, Tsung-Hsien Wen, Nikola Mrkšić, and Pei-Hao Su. 2019b. Training neural response selection for task-oriented dialogue systems. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5392–5404, Florence, Italy. <https://doi.org/10.18653/v1/P19-1536>
- Minlie Huang, Xiaoyan Zhu, and Jianfeng Gao. 2020. Challenges in building intelligent open-domain dialog systems. *ACM Transactions on Information Systems*, 38(3):1–32. <https://doi.org/10.1145/3383123>
- Myeongho Jeong, Seungtaek Choi, Jinyoung Yeo, and Seung-won Hwang. 2021. Label and context augmentation for response selection at DSTC8. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:2541–2550. <https://doi.org/10.1109/TASLP.2021.3076876>
- Giannis Karamanolakis, Subhabrata Mukherjee, Guoqing Zheng, and Ahmed Hassan Awadallah. 2021. Self-training with weak supervision. In *Proceedings of North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 845–863, Online. Association for Computational Linguistic. <https://doi.org/10.18653/v1/2021.naacl-main.66>
- Seokhwan Kim, Michel Galley, Chulaka Gunasekara, Sungjin Lee, Adam Atkinson, Baolin Peng, Hannes Schulz, Jianfeng Gao, Jinchao Li, Mahmoud Adada, Minlie Huang, Luis Lastras, Jonathan K. Kummerfeld, Walter S. Lasecki, Chiori Hori, Anoop Cherian, Tim K. Marks, Abhinav Rastogi, Xiaoxue Zang, Srinivas Sunkara, and Raghav Gupta. 2019. The eighth dialog system technology challenge. *ArXiv:1911.06394v1*. <https://doi.org/10.48550/arXiv.1911.06394>
- Dongfang Li, Yifei Yu, Qingcai Chen, and Xinyu Li. 2019. BERTSel: Answer selection with pre-trained models. *ArXiv:1905.07588v1*. <https://doi.org/10.48550/arXiv.1905.07588>
- Zheng Li, Danqing Zhang, Tianyu Cao, Ying Wei, Yiwei Song, and Bing Yin. 2021. MetaTS: Meta teacher-student network for multilingual sequence labeling with minimal supervision. In *Proceedings of Empirical Methods in Natural Language Processing*, pages 3183–3196, Online and Punta Cana, Dominican Republic. <https://doi.org/10.18653/v1/2021.emnlp-main.255>
- Ting-En Lin and Hua Xu. 2019. A post-processing method for detecting unknown intent of dialogue system via pre-trained deep neural network classifier. *Knowledge-Based Systems*. <https://doi.org/10.1016/j.knosys.2019.104979>
- Zibo Lin, Deng Cai, Yan Wang, Xiaojiang Liu, Hai-Tao Zheng, and Shuming Shi. 2020. The world is not binary: Learning to rank with grayscale data for dialogue response selection. In *Proceedings of Empirical Methods in Natural Language Processing (EMNLP)*, pages 9220–9229, Online. <https://doi.org/10.18653/v1/2020.emnlp-main.741>
- Hongrui Liu, Binbin Hu, Xiao Wang, Chuan Shi, Zhiqiang Zhang, and Jun Zhou. 2022. Confidence may cheat: Self-training on graph neural networks under distribution shift. In *Proceedings of The ACM Web Conference*, pages 1248–1258. <https://doi.org/10.1145/3485447.3512172>
- Jiao Liu, Yanling Li, and Min Lin. 2019a. Review of intent detection methods in the human-machine dialogue system. *Journal of Physics:*

- Conference Series*, 1267:25–27. <https://doi.org/10.1088/1742-6596/1267/1/012059>
- Longxiang Liu, Zhuosheng Zhang, Hai Zhao, Xi Zhou, and Xiang Zhou. 2021a. Filling the gap of utterance-aware and speaker-aware representation for multi-turn dialogue. In *Proceedings of AAAI Conference on Artificial Intelligence*. <https://doi.org/10.1609.aaai.v35i15.17582>
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019b. RoBERTa: A robustly optimized BERT pretraining approach. *ArXiv:1907.11692v1*. <https://doi.org/10.48550/arXiv.1907.11692>
- Yongkang Liu, Shi Feng, Daling Wang, Kaisong Song, Feiliang Ren, and Yifei Zhang. 2021b. A graph reasoning network for multi-turn response selection via customized pre-training. In *Proceedings of the AAAI Conference on Artificial Intelligence*. <https://doi.org/10.1609/aaai.v35i15.17585>
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *ArXiv:1711.05101v3*. <https://doi.org/10.48550/arXiv.1711.05101>
- 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. *ArXiv:1506.08909v3*. <https://doi.org/10.18653/v1/W15-4640>
- Hongyin Luo. 2022. *Self-Training for Natural Language Processing*. Ph.D. thesis, Massachusetts Institute of Technology.
- Yoshitomo Matsubara, Thuy Vu, and Alessandro Moschitti. 2020. Reranking for efficient transformer-based answer selection. In *Proceedings of International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1577–1580. <https://doi.org/10.1145/3397271.3401266>
- Marek Medveď, Radoslav Sabol, and Aleš Horák. 2020. Employing sentence context in Czech answer selection. In *Proceedings of International Conference on Text, Speech, and Dialogue*, pages 112–121, Springer, Cham. https://doi.org/10.1007/978-3-030-58323-1_12
- Junki Ohmura and Maxine Eskenazi. 2018. Context-aware dialog re-ranking for task-oriented dialog systems. In *IEEE Spoken Language Technology Workshop (SLT)*, pages 846–853, Athens, Greece. <https://doi.org/10.1109/SLT.2018.8639596>
- Haojie Pan, Cen Chen, Chengyu Wang, Minghui Qiu, Liu Yang, Feng Ji, and Jun Huang. 2021. Learning to expand: Reinforced response expansion for information-seeking conversations. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, pages 4055–4064, New York, NY. Association for Computing Machinery. <https://doi.org/10.1145/3459637.3481932>
- Yeongjoon Park, Youngjoong Ko, and Jungyun Seo. 2022. BERT-based response selection in dialogue systems using utterance attention mechanisms. *Expert Systems with Applications*. <https://doi.org/10.1016/j.eswa.2022.118277>
- Jiahuan Pei, Pengjie Ren, and Maarten de Rijke. 2021. A cooperative memory network for personalized task-oriented dialogue systems with incomplete user profiles. In *Proceedings of The Web Conference*, pages 1552–1561, New York, NY. Association for Computing Machinery. <https://doi.org/10.1145/3442381.3449843>
- Gustavo Penha, Alexandru Balan, and Claudia Hauff. 2019. Introducing MANtIS: A novel multi-domain information seeking dialogues dataset. *ArXiv:1912.04639v1*. <https://doi.org/10.48550/arXiv.1912.04639>
- Chen Qu, Liu Yang, W. Bruce Croft, Johanne R. Trippas, Yongfeng Zhang, and Minghui Qiu. 2018. Analyzing and characterizing user intent in information-seeking conversations. In *Proceedings of the 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 989–992, New York, NY. Association for Computing Machinery. <https://doi.org/10.1145/3209978.3210124>
- Chen Qu, Liu Yang, W. Bruce Croft, Yongfeng Zhang, Johanne R. Trippas, and Minghui Qiu.

- 2019a. User intent prediction in information-seeking conversations. In *Human Information Interaction and Retrieval*, pages 25–33, New York, NY. Association for Computing Machinery.
- Chen Qu, Liu Yang, Minghui Qiu, W. Bruce Croft, Yongfeng Zhang, and Mohit Iyyer. 2019b. BERT with history answer embedding for conversational question answering. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information*, pages 1133–1136, New York, NY. Association for Computing Machinery. <https://doi.org/10.1145/3331184.3331341>
- Stephen Robertson and Hugo Zaragoza. 2009. The probabilistic relevance framework: BM25 and beyond. *Foundations and Trends® in Information Retrieval*. <https://doi.org/10.1561/1500000019>
- Mrinmaya Sachan and Eric Xing. 2018. Self-training for jointly learning to ask and answer questions. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 629–640, New Orleans, Louisiana. Association for Computational Linguistics. <https://doi.org/10.18653/v1/N18-1058>
- Chongyang Tao, Jiazhan Feng, Rui Yan, Wei Wu, and Daxin Jiang. 2021. A survey on response selection for retrieval-based dialogues. In *Proceedings of International Joint Conference on Artificial Intelligence*, pages 4619–462, International Joint Conferences on Artificial Intelligence Organization. <https://doi.org/10.24963/ijcai.2021/627>
- Gokhan Tur, Dilek Hakkani-Tür, and Robert E. Schapire. 2005. Combining active and semi-supervised learning for spoken language understanding. *Speech Communication*. <https://doi.org/10.1016/j.specom.2004.08.002>
- Benyou Wang, Qianqian Xie, Jiahuan Pei, Zhihong Chen, Prayag Tiwari, Zhao Li, and Jie Fu. 2021a. Pre-trained language models in biomedical domain: A systematic survey. *ArXiv:2110.05006v2*. <https://doi.org/10.48550/arXiv.2110.05006>
- Bingning Wang, Ting Yao, Weipeng Chen, Jingfang Xu, and Xiaochuan Wang. 2021b. ComQA: Compositional question answering via hierarchical graph neural networks. In *Proceedings of the Web Conference 2021*, pages 2601–2612, New York, NY. Association for Computing Machinery. <https://doi.org/10.1145/3442381.3449993>
- Heyuan Wang, Ziyi Wu, and Junyu Chen. 2019. Multi-turn response selection in retrieval-based chatbots with iterated attentive convolution matching network. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, pages 1081–1090, New York, NY. Association for Computing Machinery. <https://doi.org/10.1145/3357384.3357928>
- Wenhui Wang, Nan Yang, Furu Wei, Baobao Chang, and Ming Zhou. 2017. Gated self-matching networks for reading comprehension and question answering. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 189–198, Vancouver, Canada. Association for Computational Linguistics. <https://doi.org/10.18653/v1/P17-1018>
- Henry Weld, Xiaoqi Huang, Siqu Long, Josiah Poon, and Soyeon Caren Han. 2021. A survey of joint intent detection and slot filling models in natural language understanding. *ACM Computing Surveys*, 55(8):1–38. <https://doi.org/10.1145/3547138>
- Taesun Whang, Dongyub Lee, Chanhee Lee, Kisu Yang, Dongsuk Oh, and Heuseok Lim. 2020. An effective domain adaptive post-training method for BERT in response selection. *Annual Conference of the International Speech Communication Association*. <https://doi.org/10.21437/Interspeech.2020-2153>
- Chien-Sheng Wu, Steven C. H. Hoi, Richard Socher, and Caiming Xiong. 2020. TOD-BERT: Pre-trained natural language understanding for task-oriented dialogue. In *Proceedings of Empirical Methods in Natural Language Processing (EMNLP)*, pages 917–929, Online.

- Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.emnlp-main.66>
- Yu Wu, Wei Wu, Zhoujun Li, and Ming Zhou. 2018. Learning matching models with weak supervision for response selection in retrieval-based chatbots. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 420–425, Melbourne, Australia. Association for Computational Linguistics. <https://doi.org/10.18653/v1/P18-2067>
- Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V. Le. 2020. Self-training with noisy student improves imagenet classification. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10684–10695, Seattle, WA. <https://doi.org/10.1109/CVPR42600.2020.01070>
- Ruijian Xu, Chongyang Tao, Daxin Jiang, Xueliang Zhao, Dongyan Zhao, and Rui Yan. 2021. Learning an effective context-response matching model with self-supervised tasks for retrieval-based dialogues. In *AAAI Conference on Artificial Intelligence*. <https://doi.org/10.1609/aaai.v35i16.17666>
- Guojun Yan, Jiahuan Pei, Pengjie Ren, Zhaochun Ren, Xin Xin, Huasheng Liang, Maarten de Rijke, and Zhumin Chen. 2022. ReMeDi: Resources for multi-domain, multi-service, medical dialogues. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information*, pages 3013–3024, New York, NY. <https://doi.org/10.1145/3477495.3531809>
- Lihe Yang, Wei Zhuo, Lei Qi, Yinghuan Shi, and Yang Gao. 2022. ST++: Make self-training work better for semi-supervised semantic segmentation. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*. <https://doi.org/10.1109/CVPR52688.2022.00423>
- Liu Yang, Minghui Qiu, Chen Qu, Cen Chen, Jiafeng Guo, Yongfeng Zhang, W. Bruce Croft, and Haiqing Chen. 2020. IART: Intent-aware response ranking with transformers in information-seeking conversation systems. In *Proceedings of the Web Conference*, pages 2592–2598, New York, NY. Association for Computing Machinery. <https://doi.org/10.1145/3366423.3380011>
- Zhengzhe Yang and Jinho D. Choi. 2019. FriendsQA: Open-domain question answering on tv show transcripts. In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 188–197, Stockholm, Sweden. Association for Computational Linguistics. <https://doi.org/10.18653/v1/W19-5923>
- Adams Wei Yu, David Dohan, Minh-Thang Luong, Rui Zhao, Kai Chen, Mohammad Norouzi, and Quoc V. Le. 2018. QANet: Combining local convolution with global self-attention for reading comprehension. In *Learning Representations*. <https://doi.org/10.48550/arXiv.1804.09541>
- Chen Zhang, Luis Fernando D’Haro, Thomas Friedrichs, and Haizhou Li. 2022a. MDD-Eval: Self-training on augmented data for multi-domain dialogue evaluation. In *Proceedings of AAI Conference on Artificial Intelligence*. <https://doi.org/10.1609/aaai.v36i10.21420>
- Linhao Zhang, Dehong Ma, Sujian Li, and Houfeng Wang. 2021. Do it once: An embarrassingly simple joint matching approach to response selection. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4872–4877, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.findings-acl.430>
- Rongjunchen Zhang, Tingmin Wu, Sheng Wen, Surya Nepal, Cecile Paris, and Yang Xiang. 2022b. SAM: Multi-turn response selection based on semantic awareness matching. *ACM Transactions on Internet Technology*, 2(1):1–18. <https://doi.org/10.1145/3545570>
- Xinyan Zhao, Feng Xiao, Haoming Zhong, Jun Yao, and Huanhuan Chen. 2020. Condition aware and revise transformer for question answering. In *Proceedings of The Web Conference*, pages 2377–2387, New York, NY. Association for Computing Machinery.

<https://doi.org/10.1145/3366423.3380301>

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. <https://doi.org/10.18653/v1/D16-1036>

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 (Volume 1: Long Papers)*, pages 1118–1127, Melbourne, Australia. <https://doi.org/10.18653/v1/P18-1103>