

# GRASP: Guiding model with Relational Semantics using Prompt for Dialogue Relation Extraction

Junyoung Son\*, Jinsung Kim\*, Jungwoo Lim\*, Heuseok Lim†

Computer Science and Engineering, Korea University

Republic of Korea

{s0ny, jin62304, wjddn803, limhseok}@korea.ac.kr

## Abstract

The dialogue-based relation extraction (DialogRE) task aims to predict the relations between argument pairs that appear in dialogue. Most previous studies utilize fine-tuning pre-trained language models (PLMs) only with extensive features to supplement the low information density of the dialogue by multiple speakers. To effectively exploit inherent knowledge of PLMs without extra layers and consider scattered semantic cues on the relation between the arguments, we propose a Guiding model with Relational Semantics using Prompt (GRASP). We adopt a prompt-based fine-tuning approach and capture relational semantic clues of a given dialogue with 1) an argument-aware prompt marker strategy and 2) the relational clue detection task. In the experiments, GRASP achieves state-of-the-art performance in terms of both F1 and F1<sub>c</sub> scores on a DialogRE dataset even though our method only leverages PLMs without adding any extra layers.

## 1 Introduction

The relation extraction (RE) task aims to extract semantic relations from unstructured text such as a sentence, a document, or even a dialogue. RE plays a critical role in information extraction and knowledge base construction as it can extract structured relational information (Ji et al., 2010; Swampilal and Stevenson, 2010). However, the utilization of sentence-level RE in a conversational setting is limited because numerous relational facts appear across multiple sentences with more than one speaker in a dialogue (Yao et al., 2021). Thus, the dialogue-based relation extraction (DialogRE) task, which includes argument pairs and their corresponding relations, has been proposed to encourage building a model that captures the underlying se-

\*These authors have equally contributed to this work

† Corresponding author

Dialogue		
1	S1: Hey guys! Hey!	
2	S2: Hey <b>Pheeb</b> s, guess who we saw today.	
3	S3: Ooh, ooh, fun! Okay... um, Liam Neeson.	
...		
7	S3: The woman who cuts my hair!	
8	S4: Okay, look, <u>this</u> could be a really long game.	
9	S5: Your <u>sister</u> Ursula.	
10	S3: Oh, really.	
11	S5: Yeah, yeah, she <u>works over at</u> that place, uh...	
12	S3: <b>Rift</b> 's. Yeah, I know.	
...		
17	S6: Um, <b>Pheeb</b> s, so, you guys just don't get along?	
18	S3: It's mostly just dumb <u>sister</u> stuff, you know, I mean, like, everyone always thought of her as the pretty one, you know... ..	
Argument pair	Trigger	Relation Type
<b>R1</b> (Pheeb's, PER) (Ursula, PER)	<u>sister</u>	per:siblings
<b>R2</b> (S3, PER) (Pheeb's, PER)	none	per:alternate_names
<b>R3</b> (Rift's, ORG) (Ursula, PER)	<u>works over at</u>	per:employees_or_members

Table 1: Example of DialogRE data. The arguments are bold, and the triggers are underlined. The arguments and triggers are scattered throughout the dialogue, which leads to low information density. The triggers determine the direction of the proper relation indirectly.

mantic spread in the dialogue, as presented in Table 1 (Yu et al., 2020).

In previous studies where state-of-the-art (SoTA) performance is achieved on DialogRE benchmarks, fine-tuning is employed on pre-trained language models (PLMs) (Lee and Choi, 2021; Long et al., 2021), such as BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019). As fine-tuning requires the addition of extra layers on top of the PLMs and the training objectives are different from those used in the pre-training phase, PLMs cannot effectively exploit their learned knowledge in the downstream task, resulting in less generalized capability (Chen et al., 2021).

To effectively utilize knowledge from PLMs, several studies involving prompt-based fine-tuning have been conducted. They employ the PLM directly as a predictor and completing a cloze task



Figure 1: An illustration of (1) standard fine-tuning approach, and (2) prompt-based fine-tuning approach. In prompt-based fine-tuning approach, a set of *label words* are mapped into a *class set* by a certain mapping function.

to bridge the gap between pre-training and fine-tuning (Gao et al., 2020; Han et al., 2021). As presented in Figure 1, prompt-based fine-tuning treats the downstream task as a masked language modeling (MLM) problem by directly generating the textual response to a given template. In specific, prompt-based fine-tuning updates the original input on the basis of the template and predicts the *label words* with the [MASK] token. Afterwards, the model maps predicted *label words* to corresponding task-specific class sets.

However, the prompt-based fine-tuning approach is still not sufficient in terms of performance compared with the fine-tuning-based approach (Han et al., 2021; Chen et al., 2021). We attribute this phenomenon to the following properties of conversation: higher person-pronoun frequency (Wang and Liu, 2011) and lower information density (Biber, 1991) by multiple speakers<sup>1</sup>. Therefore, a prompt-based fine-tuning approach that collects the sparse semantics in the dialogue is required to understand relation between the arguments.

We propose a method named Guiding model with Relational Semantics using Prompt (GRASP) for DialogRE. To maximize the advantages of the prompt-based fine-tuning approach for the DialogRE task, we suggest an argument-aware prompt marking (APM) strategy and a relational clue detection (RCD) task. The APM strategy guides the model to the significant arguments scattered in the dialogue by carefully considering arguments. For our APM strategy, we conduct empirical study based on the diverse marker types to validate our APM strategy. Along with the APM strategy, the suggested RCD task with a training objective leads the model to pay

<sup>1</sup>In the DialogRE dataset, 65.9% of relational triples involve arguments that never appear in the same utterance, demonstrating that multi-turn reasoning plays an important role.

attention to significant relational clues. Specifically, the model is trained to determine whether each token in a dialogue belongs to a subject, object, or trigger. As a result, the PLM is trained on RCD and MLM jointly. In the experiments, our method achieves SoTA performance at a significant level in the DialogRE task for both the full-shot and few-shot settings. Only PLMs are employed without the addition of an extra layer as a predictor, and GRASP exhibits a higher performance than other baselines. The significant performance improvement indicates that attending to significant semantic clues guides the PLMs to predict the correct relation with its inherent knowledge in both full-shot and few-shot settings. Moreover, we provide ablation studies and qualitative analysis on the robustness of GRASP.

Our contributions are as follows:

- We adopt a prompt-based fine-tuning approach to utilize a PLM’s inherent knowledge directly for dialogues with relatively low information density.
- We introduce an APM strategy and a RCD task that guide PLMs on the significant relational clues, which are semantic information to predict relations.
- We demonstrate that our proposed method achieves SoTA performance on the DialogRE task in both full-shot and few-show settings.
- We conduct ablation studies and qualitative analysis to validate the robustness of GRASP.

The remainder of this paper is organized as follows. In Section 3, we present the entire process of our method in detail. The experimental setup and the results are explained in Section 4. The further analyses is provided in Section 5, and Section 6

presents the conclusions. Appendix 2 provides related works including the prompt-based learning, and the DialogRE.

## 2 Related Works

**Prompt-based learning** Prompt-based learning is a method of reducing the gap between the pre-training objective and that of fine-tuning. For example, language models such as BERT (Devlin et al., 2019) use masked language modeling (MLM) objective in pre-training phase where the model fills the [MASK] token whereas the model trains without [MASK] token in fine-tuning phase by adding an extra classifier layer. As a result, the discrepancy between training objectives prevents PLM from leveraging knowledge acquired from the pre-training enormous corpus (Chen et al., 2021). Also, prompt-based learning shows better performance than fine-tuning especially in the few-shot setting (Gao et al., 2020; Schick and Schütze, 2021; Li and Liang, 2021; Liu et al., 2021).

**DialogRE** Recent studies on the DialogRE dataset show a tendency to fine-tune the PLM with task-specific objectives and use the model with high-complexity (Xue et al., 2021; Long et al., 2021; Lee and Choi, 2021). In detail, Lee and Choi (2021) shows considerable performance with the contextualized turn representations from the diverse type of nodes and edges. Moreover, task-specific objectives of fine-tuning lead to a gap between pre-training and fine-tuning.

To overcome the limitation, Han et al. (2021) utilizes multiple [MASK] tokens for each argument and the relation with logical rules by concentrating subject and object in DialogRE. There also exists an approach that incorporates potential knowledge contained in relation labels into prompt construction with trainable virtual type words and answers words. This approach also carefully initializes the virtual tokens with implicit semantic words and employs prior distributions estimated from the data (Chen et al., 2021).

Despite the prompt-based approach’s high potential, few prompting studies sufficiently consider low information density and difficulty of capturing intrinsic relational information of the data between the argument pair of dialogue relation extraction task. We focus on building a light model with prompt-based fine-tuning with implicit semantic information of the relation which alleviates the

sparsity problem.

## 3 Methodology

An overview of GRASP is illustrated in Figure 2. First, an input with a prompt template is constructed by using the APM strategy. Then, the PLM receives the constructed input for prompt-based fine-tuning and estimates probability distributions of the model’s vocabulary by using contextualized representations of the PLM. The RCD task lets the model predict the relational clue type of each token, and the model takes the [MASK] representation to predict a final relation in the MLM task. In other words, our model is trained through multi-task learning to encourage mutual communication between relational clues and the final relation for the argument pair.

### 3.1 Problem Formulation

Each example  $\mathcal{X}$  includes dialogue  $D$ , subject  $a_1$ , and object  $a_2$ . Note that  $D$  denotes  $\{s_1 : u_1, s_2 : u_2, \dots, s_N : u_N\}$ , where  $s_n$  is the  $n$ -th speaker and  $u_n$  is the corresponding utterance. Given  $\mathcal{X} = \{D, a_1, a_2\}$ , the goal of DialogRE is to predict relation  $y$  between arguments  $a_1$  and  $a_2$  by leveraging  $D$ . To describe the DialogRE task in terms of prompt-based fine-tuning, a template function,  $\mathcal{T}(\cdot)$ , is defined to map each example to  $\mathcal{X}_{prompt} = \mathcal{T}(\mathcal{X})$ . A [MASK] is inserted into  $\mathcal{X}_{prompt}$ , and used to predicting the label words of relation  $y$ . The formulation of  $\mathcal{T}(\mathcal{X})$  is as follows:

$$\mathcal{T}(\mathcal{X}) = \text{“}[\text{CLS}] D [\text{SEP}] [\text{subj}] a_1 [\text{subj}] [\text{MASK}] [\text{obj}] a_2 [\text{obj}] [\text{SEP}] \text{”} \quad (1)$$

Based on the structure of  $\mathcal{T}(\cdot)$ , we construct our template function,  $\mathcal{T}'(\cdot)$ , by applying two steps: transformation of  $D$  to  $D'$  with an argument-aware prompt marker, described in 3.2, and prompt initialization for [subj] and [obj], explained in 3.3. Subsequently, we introduce the RCD task in 3.4 to train the model by utilizing  $\mathcal{T}'(\cdot)$  with a multitask training strategy on MLM, as detailed in 3.5.

### 3.2 Argument-aware Prompt Marker

We propose an APM strategy that considers both speaker and non-speaker arguments. In previous studies, a dialogue is encoded by focusing on speaker information (Lee and Choi, 2021; Yu et al., 2020; Chen et al., 2021) without focusing non-speaker arguments. However, in the DialogRE dataset, approximately 77.4% of relation triples

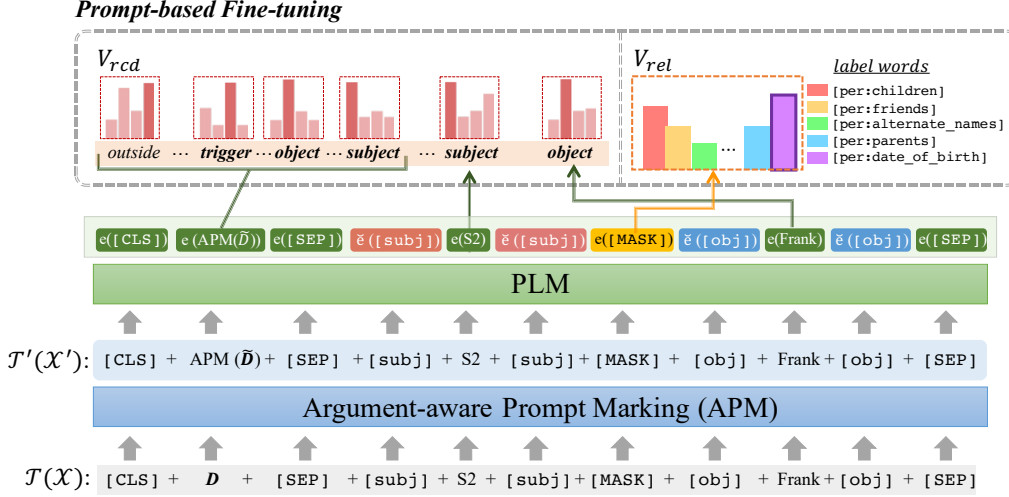


Figure 2: The overall model architecture of GRASP. By formalizing specific tasks as MLM tasks, the model predicts answers with the tokens from the model’s vocabulary; for example, the *label words* of  $V_{rel}$  and  $V_{rcd}$  are from the model’s vocabulary.

include at least one non-speaker argument, implying that consideration of non-speaker arguments is also inevitable to enhance the model’s argument-awareness.

Our methods is inspired by the previous works of Soares et al. (2019) and Han et al. (2021); the former addresses the importance of entity markers regarding recognition of the entity position, and the latter improves the model performance by using specific additional punctuation in the model’s original vocabulary as the entity marker. Accordingly, we insert the argument-aware prompt marker token,  $[p]$ , as an entity marker. The argument-aware prompt markers allow our model to obtain informative signs to determine which token is the indicative component for relation prediction. We initialize the feature of  $[p]$  as the embedding of the space token in the vocabulary of the model. Our empirical experiments reveal that the space token can perceive the start position of the arguments. Consequently, our prompt marker enhances the model to discriminate which part of the dialogue plays a critical role in predicting a relation.

Using the proposed argument-aware prompt marker, we strengthen the token replacement method of BERT<sub>s</sub> (Yu et al., 2020). Given example  $\mathcal{X}$ , BERT<sub>s</sub> constructs  $\tilde{\mathcal{X}} = \{\tilde{D}, \tilde{a}_1, \tilde{a}_2\}$ , where  $\tilde{D} = \{\tilde{s}_1 : u_1, \tilde{s}_2 : u_2, \dots, \tilde{s}_N : u_N\}$  and  $\tilde{s}_n$  is

$$\tilde{s}_n = \begin{cases} [S_1] & \text{if } s_n = a_1 \\ [S_2] & \text{if } s_n = a_2 \\ s_n & \text{otherwise.} \end{cases} \quad (2)$$

$[S_1]$  and  $[S_2]$  are special tokens for speakers. In addition,  $\tilde{a}_m$  ( $m \in \{1, 2\}$ ) is defined as  $[S_m]$  if  $\exists n (s_n = a_m)$  and  $a_m$  otherwise.

Even though BERT<sub>s</sub> prevents the model from overfitting and demonstrates higher generalization capacity (Yu et al., 2020), BERT<sub>s</sub> does not consider non-speaker arguments. To explore the disregarded arguments, the APM strategy expands BERT<sub>s</sub> by considering both types of arguments in the DialogRE task. We define function  $\text{APM}(\cdot)$  that encodes an utterance by inserting  $[p]$  in front of each argument token. Given dialogue  $\tilde{D}$  constructed by using BERT<sub>s</sub>, we construct  $D' = \{\tilde{s}_1 : u'_1, \tilde{s}_2 : u'_2, \dots, \tilde{s}_N : u'_N\}$  by applying  $\text{APM}(u_n)$  for each utterance in  $\tilde{D}$ . Consequently, the APM strategy constructs  $\mathcal{X}' = \{D', \tilde{a}_1, \tilde{a}_2\}$  based on  $\tilde{\mathcal{X}}$ .

$$u'_n = \begin{cases} \text{APM}(u_n) & \text{if } \exists m (\tilde{a}_m \in u_n) (m \in \{1, 2\}) \\ u_n & \text{otherwise} \end{cases} \quad (3)$$

For instance,  $\text{APM}(\cdot)$  encodes the text “I am Tom Gordon” to  $[I, \text{am}, [p], \text{Tom}, \text{Gordon}]$ , inserting  $[p]$  in front of “Tom Gordon.”

### 3.3 Prompt Construction

We update the constructed input  $\mathcal{X}'$  with a template function and conduct a deliberate initialization. We utilize the prior distribution of argument types for initialization inspired by the study conducted by Chen et al. (2021). Prompt tokens  $[subj]$  and  $[obj]$  are used to inject argument-type information. We define the argument-type set,  $\mathcal{AT} = \{\text{“PER,” “ORG,” “GPE,” “VALUE,”$

“STRING”}, using the types pre-defined in the dataset as depicted in Table 1. We calculate the distributions of argument types  $\phi^{[\text{subj}]}$  and  $\phi^{[\text{obj}]}$  over  $\mathcal{AT}$  by using frequency statistics. We aggregate each argument type,  $at \in \mathcal{AT}$ , with the corresponding prior distribution to initialize prompt tokens  $[\text{subj}]$  and  $[\text{obj}]$ . The specific initialization equations are as follows:

$$\begin{aligned}\check{e}([\text{subj}]) &= \sum_{at \in \mathcal{AT}} \phi_{at}^{[\text{subj}]} \cdot e(at) \\ \check{e}([\text{obj}]) &= \sum_{at \in \mathcal{AT}} \phi_{at}^{[\text{obj}]} \cdot e(at),\end{aligned}\quad (4)$$

where  $e(\cdot)$  is the embedding from the PLM of an input token and  $\check{e}(\cdot)$  is the initialized embedding of the prompt token. Suppose the subject has prior distribution,  $\phi^{[\text{subj}]} = \{\text{“PER”}: 0.5, \text{“ORG”}: 0.5, \text{“GPE”}: 0.0, \text{“VALUE”}: 0.0, \text{“STRING”}: 0.0\}$ . The initial embedding of the  $[\text{subj}]$  token can be calculated as a weighted average, i.e.,  $\check{e}([\text{subj}]) = 0.5 \cdot e(\text{“PER”}) + 0.5 \cdot e(\text{“ORG”})$ .

Consequently, we can formalize  $\mathcal{T}'(\cdot)$ , which converts  $\mathcal{X}'$  to  $\mathcal{X}'_{prompt}$ , where  $\mathcal{X}'_{prompt}$  is an argument-enhanced input example, by using the APM strategy and prompt construction with deliberate initializations, i.e.,  $\mathcal{X}'_{prompt} = \mathcal{T}'(\mathcal{X}')$ . Then, the final input structure for prompt-based fine-tuning is as follows:

$$\mathcal{T}'(\mathcal{X}') = \text{“}[\text{CLS}] D' [\text{SEP}] [\text{subj}] \tilde{a}_1 [\text{subj}] [\text{MASK}] [\text{obj}] \tilde{a}_2 [\text{obj}] [\text{SEP}] \text{”}.\quad (5)$$

In addition, for relation prediction by applying MLM, we also define  $\mathcal{V}_{rel}$  as a set of *label words* in the model’s vocabulary as illustrated in Figure 2. In detail, we utilize its metadata for the initialization of each relation representation to inject its semantics. For instance, we add a special token to the model’s vocabulary,  $[\text{per}:\text{date\_of\_birth}]$  as a *label word*, and initialize this token by aggregating the embeddings of the words in the metadata, i.e.,  $\{\text{“person,” “date,” “of,” “birth”}\}$  for a class  $\text{“per}:\text{date\_of\_birth.”}$

### 3.4 Relational Clue Detection task

To improve the understanding capability on relational clues by employing a prompt-based fine-tuning approach, we introduce a RCD task. We define a set of *label words*,  $\mathcal{V}_{rcd} = \{[\text{subject}], [\text{object}], [\text{trigger}], [\text{outside}]\}$ , and add to the model’s vocabulary. Then, we construct a sequence of *label words* for RCD,  $\mathcal{C}_{rcd}$ , by assigning each token in  $\mathcal{X}'_{prompt}$  to the corresponding

clue type word from  $\mathcal{V}_{rcd}$ . For instance,  $\mathcal{C}_{rcd}$  is constructed as follows:  $\{[\text{subject}], [\text{trigger}], [\text{object}]\}$  when the token sequence is given by  $\{\text{“Pheeb’s”}, \text{“lives in”}, \text{“LA”}\}$ , where  $\text{“Pheeb’s”}$  is a subject and  $\text{“LA”}$  is an object argument, and  $\text{“lives in”}$  is a trigger.

The RCD task exploits the MLM head that is used to predict the  $[\text{MASK}]$  token. Except for the  $[\text{MASK}]$  token, each token is sequentially tagged with  $\mathcal{V}_{rcd}$  using the meta-data provided in the dataset. In other words, the RCD task allows the model to identify which non- $[\text{MASK}]$  tokens correspond to certain relational clue types. RCD supports the model in collecting scattered information from the entire dialogue by indicating where to focus in the dialogue to predict a relation. In this respect, the model pays considerably more attention to the semantic clues of the relation, such as triggers. Moreover, our model maintains a lightweight complexity by conducting the RCD task without an additional classifier. The loss for the RCD task over each token  $x \in \mathcal{X}'_{prompt}$  is aggregated as follows:

$$\mathcal{L}_{RCD} = - \sum_{x \in \mathcal{X}'_{prompt}} \log P(x = \mathcal{C}_{rcd}(x) | \mathcal{X}'_{prompt}).\quad (6)$$

### 3.5 Model Training

Before the final relation-prediction, we mark the position of the trigger using the  $[\text{p}]$  token, i.e., the argument-aware prompt marker, based on the results of the RCD task. Given example  $\mathcal{X}'_{prompt}$ , the model predicts the *label words* of  $\mathcal{V}_{rcd}$  for all non- $[\text{MASK}]$  tokens. Subsequently,  $[\text{p}]$  is appended in front of the words predicted as a trigger to train the model to distinguish essential clues that encourage determining the relation.

**Multitask Learning**  $\mathcal{M}_{rel} : \mathcal{Y} \rightarrow \mathcal{V}_{rel}$  is a mapping function that converts a class set,  $\mathcal{Y}$ , into a set of *label words*,  $\mathcal{V}_{rel}$ . For each input  $\mathcal{X}'_{prompt}$ , the purpose of MLM is to fill  $[\text{MASK}]$  with the relation *label words* in the model’s vocabulary. As the model predicts the correct *label word* at the position of  $[\text{MASK}]$ , we can formulate  $p(y|x) = P([\text{MASK}] = \mathcal{M}_{rel}(y) | \mathcal{X}'_{prompt})$ . The training objective of the relation prediction is to minimize

$$\mathcal{L}_{REL} = - \log P([\text{MASK}] = \mathcal{M}_{rel}(y) | \mathcal{X}'_{prompt}).\quad (7)$$

To improve the model’s ability to capture relational clues through the interaction between the MLM task for relation prediction and the RCD

<i>Full-shot Setting</i>				
Method	V1		V2	
	F1	F1 <sub>c</sub>	F1	F1 <sub>c</sub>
Fine-tuning based approach				
BERT <sub>s</sub> (Yu et al., 2020)	61.2	55.4	-	-
RoBERTa <sub>s</sub> (Lee and Choi, 2021)	-	-	71.3	63.7
Dual (Bai et al., 2021)	67.3	<u>61.4</u>	67.1	61.1
CoIn (Long et al., 2021)	<u>72.3</u>	-	-	-
TUCORE-GCN (Lee and Choi, 2021)	-	-	<u>73.1</u>	<u>65.9</u>
Prompt-based fine-tuning approach				
PTR (Han et al., 2021)	63.2	-	-	-
KnowPrompt (Chen et al., 2021)	68.6	-	-	-
<b>GRASP<sub>base</sub></b> (Our model)	69.2	62.4	69.0	61.7
<b>GRASP<sub>large</sub></b> (Our model)	<b>75.1 (+2.8)</b>	<b>66.7 (+5.3)</b>	<b>75.5 (+2.4)</b>	<b>67.8 (+1.9)</b>

Table 2: Performances of **GRASP** on test set of DialogRE. V1 and V2 represent the version of the dataset. The underlined scores are the previous SoTA performances. Subscript in parentheses represents advantages of **GRASP** over the best results of baselines (the underlined). Best results are bold.

task, we employ a multitask learning to train the **GRASP** model by using the joint loss expressed in Equations (6) and (7). Therefore, the final learning objective is to minimize the joint loss, where  $\lambda_1$  and  $\lambda_2$  are hyperparameters.

$$\mathcal{L}_{GRASP} = \lambda_1 \cdot \mathcal{L}_{RCD} + \lambda_2 \cdot \mathcal{L}_{REL} \quad (8)$$

## 4 Experiments

### 4.1 Experimental Setup

For the base PLMs, RoBERTa-base and RoBERTa-large are adopted, as denoted by **GRASP<sub>base</sub>** and **GRASP<sub>large</sub>**, respectively. The evaluation metrics are the F1 and F1<sub>c</sub> (Yu et al., 2020) scores. F1<sub>c</sub> is an evaluation metric for supplementing the F1 score in a conversational setting and is computed by employing part of the dialogue necessary for predicting the relation between given arguments as input instead of the entire dialogue. The detailed settings for **GRASP** can be found in Appendix A.

In a full-shot setting, **GRASP** is compared with both fine-tuning-based and prompt-based fine-tuning approaches. TUCORE-GCN (Lee and Choi, 2021) is a typical fine-tuning-based model using turn-level features with a graph convolution network (Kipf and Welling, 2016). CoIn (Long et al., 2021) employs utterance-aware and speaker-aware representations, and Dual (Bai et al., 2021) models the relational semantics using abstract meaning representations (Banarescu et al., 2013). Moreover, PTR (Han et al., 2021) and KnowPrompt (Chen

et al., 2021) are the prompt-based fine-tuning baseline models. In the few-shot setting, 8-, 16-, and 32-shot experiments were conducted based on LM-BFF (Gao et al., 2020) by using three different randomly sampled data.

### 4.2 Experimental Results

**Full-shot setting** As presented in Table 2, it is shown that **GRASP<sub>large</sub>** surpasses all of the baseline models, including the current SoTA models, that is, CoIn and TUCORE-GCN. From the result in the full-shot setting, the baselines of the fine-tuning-based approach, such as TUCORE-GCN or CoIn, show better performance than those of the prompt-based fine-tuning approach. In particular, CoIn outperforms all of the other baselines, including PTR and KnowPrompt on V1, and TUCORE-GCN exhibits the best performance on V2. Interestingly, even though the performance of the prompt-based fine-tuning baselines is much lower than fine-tuning based models, our **GRASP<sub>large</sub>** outperforms regardless of the way of training approach.

In addition, **GRASP<sub>large</sub>** shows its efficiency in conversational settings by exceeding all the baselines in terms of F1<sub>c</sub>, thereby indicating that our method effectively overcomes the low information density of dialogues. These results imply that guiding the model to pay attention to relational clues with a prompt-based fine-tuning approach can be more effective than adding additional features and layers. The slightly low performance of **GRASP<sub>base</sub>** is attributed to the gap in the model size; for example, TUCORE-GCN has 401M

<i>Few-shot Setting</i>			
Method	Shot		
	$K=8$	$K=16$	$K=32$
Fine-tuning based approach			
RoBERTa (Chen et al., 2021)	29.8	40.8	49.7
TUCORE-GCN (Lee and Choi, 2021)	24.6	40.0	53.8
Prompt-based fine-tuning approach			
PTR (Han et al., 2021)	35.5	43.5	49.5
KnowPrompt (Chen et al., 2021)	43.8	50.8	55.3
<b>GRASP<sub>base</sub></b> (Our model)	<b>45.4</b>	<b>52.0</b>	<b>56.0</b>
<b>GRASP<sub>large</sub></b> (Our model)	36.0	<b>55.3</b>	<b>62.6</b>

Table 3: Low-resource RE performance of F1 scores (%) on different test sets. We use  $K = 8, 16, 32$  (# of examples per class) for few-shot experiments. Best results are bold and the second place results are underlined.

weight parameters with RoBERTa-large model, and **GRASP<sub>base</sub>** has 125M weight parameters with RoBERTa-base. Nevertheless, **GRASP<sub>base</sub>** demonstrates 6.0%p and 0.6%p improvements compared with the other prompt-based fine-tuning baselines, i.e., PTR and KnowPrompt, respectively.

**Few-shot setting** As presented in Table 3, **GRASP** still exhibits robust performance in few-shot settings. **GRASP<sub>base</sub>** outperforms the baselines of both the fine-tuning and prompt-based fine-tuning methods regardless of the number of shots, demonstrating 20%p or higher performance in the 8-shot setting compared with TUCORE-GCN and indicating that our method is more efficient than the fine-tuning method in a low-resource setting. **GRASP<sub>base</sub>** also demonstrates improved performance compared with KnowPrompt in all-shot settings, indicating the effectiveness of the considerations on the properties of the dialogue with prompt-based fine-tuning. Except for the 8-shot setting, **GRASP<sub>large</sub>** presents outstanding performance, achieving up to 15.3%p of absolute improvement in the 16-shot setting. Although **GRASP<sub>large</sub>** outperforms the fine-tuning-based models and PTR, the limited performance of **GRASP<sub>large</sub>** in the 8-shot setting can be attributed to an insufficient number of examples.

We also observe that the fine-tuning-based models, such as TUCORE-GCN, perform at least 5.7%p worse than the prompt-based fine-tuning models, such as PTR, in the 8-shot setting, indicating that the fine-tuning-based models may have difficulty in sufficiently capturing relational semantics when the data are extremely scarce. In particular, TUCORE-GCN indicates a 5.2%p lower perfor-

Method	Dev		Test	
	F1	F1 <sub>c</sub>	F1	F1 <sub>c</sub>
<b>GRASP<sub>base</sub></b>	<b>70.3</b>	<b>63.3</b>	<b>69.0</b>	<b>61.7</b>
-APM	69.5	62.9	66.9	60.7
-RCD	69.2	62.8	67.7	61.0
-Prompt manual init.	68.1	61.9	65.9	59.8

Table 4: Ablation study on DialogRE dataset.

mance than the fine-tuned RoBERTa, indicating that the high complexity requires a larger amount of data than the other models.

**Ablation Study** We conduct an ablation study to validate the effectiveness of the proposed modules. As shown in Table 4, each of the proposed modules improves the overall performance for both F1 and F1<sub>c</sub> settings. Without the APM strategy, the performance of **GRASP<sub>base</sub>** decreases the F1 score by 0.8%p and the F1<sub>c</sub> score by 0.4%p on the development set, and the F1 score drops sharply by 2.1%p and F1<sub>c</sub> by 1.0%p on the test set. This result indicates that argument-awareness can be obtained through both speaker and non-speaker argument information. When the RCD task is excluded, the performance of **GRASP<sub>base</sub>** decreases the F1 score by 1.1%p and F1<sub>c</sub> by 0.5%p on the development set and the F1 score by 1.3%p and F1<sub>c</sub> by 0.7%p for the test set. These results demonstrate that the RCD task alleviates the low information density of the dialogue by guiding the model to focus on relational clues.

In addition, the performance without the manual initialization of prompt construction of **GRASP<sub>base</sub>** is reduced by 2.2%p for the F1 score and 1.4%p for the F1<sub>c</sub> score on the development set and by 3.1%p for the F1 score and 1.9%p for the F1<sub>c</sub> score on the test set. This result suggests that prompt construction is a basic step in training the model in the prompt-based fine-tuning case. The deliberate initialization of prompts is critical for modeling the tasks in an appropriate direction.

## 5 Analysis

### 5.1 Analysis on marker type for APM

To analyze argument-awareness regarding the types of markers, we conduct experiments on diverse prompt markers, as shown in Table 5. The result reveals that considering argument types leads to performance improvement in the model. Punctuation marker “;” shows comparable performance among other punctuation markers, and “!” and “@”

Marker Type	Input Example	F1
Entity marker	[CLS] [E1] <b>Frank</b> [/E1] lives in [E2] <b>Montauk</b> [/E2] . [SEP]	65.9
Type marker	[CLS] [E1:PER] <b>Frank</b> [/E1:PER] lives in [E2:GPE] <b>Montauk</b> [/E2:GPE] . [SEP]	66.7
Punctuation marker (!)	[CLS] ! <b>Frank</b> ! lives in ! <b>Montauk</b> ! . [SEP]	66.7
Punctuation marker (@)	[CLS] @ <b>Frank</b> @ lives in @ <b>Montauk</b> @ . [SEP]	65.3
Punctuation marker (:)	[CLS] ; <b>Frank</b> ; lives in ; <b>Montauk</b> ; . [SEP]	67.1
APM marker (front)	[CLS] [p] <b>Frank</b> lives in [p] <b>Montauk</b> . [SEP]	<b>69.0</b>
APM marker (surrounding)	[CLS] [p] <b>Frank</b> [p] lives in [p] <b>Montauk</b> [p] . [SEP]	65.9

Table 5: The performance based on the marker type. The arguments are bold. The embedding of [p] is initialized with the space token from the model’s vocabulary. In type marker, [E1:PER] represents a start position of subject which has a person type and [/E1:PER] represents an end position of object that has the same type.

Dialogue				
S1: Hey!!				
S2: Hey!				
S1: Guess what. <b>Frank Jr.</b> , and <b>Alice</b> <u>got married!</u>				
S2: Oh my God!!				
S1: And! And, they’re gonna have a baby! And! And, they want me to grow it for them in my uterus.				
S3: My God!				
S4: Are you serious?				
S1: Yeah				
S5: You’re really thinking about having sex with your brother?!				
S1: Ewww! And "Oh no!" It’s—they just want me to be the surrogate. It’s her-it’s her egg and her sperm, and I’m-I’m just the oven, it’s totally their bun.				
S5: Huh.				
Argument pair	Ground Truth	RoBERTa Predicted Relation	GRASP Predicted ( <i>subject, object, trigger</i> )	Predicted Relation
(Frank Jr., Alice)	<i>per:spouse</i>	unanswerable	(Frank Jr, Alice, <u>got married</u> )	<i>per:spouse</i>
(Alice, Frank Jr.)	<i>per:spouse</i>	<i>per:siblings</i>	(Alice, Frank Jr, <u>got married</u> )	<i>per:spouse</i>

Table 6: The qualitative analysis on the prediction of **GRASP** based on the comparison with the RoBERTa model. Predicted (*subject, object, trigger*) is that **GRASP** predicted on RCD task.

achieve a limited score. We presume that the higher frequency of “;” acts as a delimiter, which results in decent performance.

We also observe that the APM marker (front) performs the best, with a 69.0% F1 score among all other marker types. In addition, we conduct an experiment based on the position of the prompt marker, [p], by comparing two versions of the APM marker: APM marker (front) and APM marker (surrounding). The APM marker (surrounding) display 3.1%p lower performance than the APM marker (front). Based on these results, we empirically adopt the embedding initialization of our [p] prompt marker using the space token and located it in front of the arguments.

## 5.2 Qualitative Analysis on GRASP

Since we train **GRASP** attending on relational clues through the APM strategy and the RCD task, we further conduct analysis to investigate that the relational clues such as triggers contribute to predicting a relation between the arguments in a

prompt-based manner, as shown in Table 6. Specifically, we compare **GRASP**<sub>large</sub> with a fine-tuned RoBERTa-large model to validate our method.

We observe that the fine-tuned RoBERTa model struggles to capture the symmetrical relations including the trigger. The fine-tuned RoBERTa model fail to capture relational clues, such as “got married”, misleading the model into predicting an inappropriate relations for both symmetrical relations between “Frank Jr.” and “Alice.” “got married” is a critical cue to distinguish the relations between “per:spouse” and “per:siblings” because this phrase implies a romantic relationship in a dictionary definition. In contrast, **GRASP**, which is trained using the RCD task and the APM strategy in a prompt-based manner, predicted the correct relations, capturing the correct relational clues including the trigger “got married” for both symmetrical argument pairs. This result presents the effectiveness of **GRASP** designed to guide the model on the relational clues, alleviating the difficulties of low information density in dialogues. Additional



Method	MELD	EmoryNLP
RoBERTa (Liu et al., 2019)	62.0	37.3
COSMIC (Ghosal et al., 2020)	65.2	38.1
TUCORE-GCN (Lee and Choi, 2021)	65.4	39.2
<b>GRASP<sub>large</sub> (Ours)</b>	<b>65.6</b>	<b>40.0</b>

Table 7: Experimental results of **GRASP<sub>base</sub>** on MELD and EmoryNLP tasks.

examples demonstrating similar phenomena for the symmetrical relations are provided in the Appendix B.

### 5.3 Analysis on the applicability of GRASP

To demonstrate the robustness of our APM strategy and RCD task, we evaluated **GRASP** on MELD (Poria et al., 2019) and EmoryNLP (Zahiri and Choi, 2018) datasets, which are designed for emotion recognition in conversations (ERC). MELD (Poria et al., 2019) is a multimodal dataset collected from a TV show named Friends and consists of seven emotion labels and 2,458 dialogues with only textual modality. EmoryNLP (Zahiri and Choi, 2018) is also collected from Friends and comprises seven emotion labels and 897 dialogues. Each utterance in these datasets is annotated with one of the seven emotion labels. The weighted-F1 is calculated to evaluate the MELD and EmoryNLP datasets.

For baselines, we employ the fine-tuned RoBERTa (Liu et al., 2019), COSMIC (Ghosal et al., 2020), and TUCORE-GCN models (Lee and Choi, 2021). COSMIC (Ghosal et al., 2020) uses RoBERTa-large as the encoder. It is a framework that models various aspects of commonsense knowledge by considering mental states, events, actions, and cause-effect relations for emotional recognition in conversations.

As presented in Table 7, **GRASP** is applied to other dialogue-based tasks by alleviating the low information density of the given dialogue. In particular, the performance of **GRASP** surpasses that of TUCORE-GCN, which is the current SoTA model in DialogRE, and those of the baseline specialized on ERC tasks, such as COSMIC in both MELD and EmoryNLP.

## 6 Conclusion

In this paper, we proposed **GRASP**, which is a method for guiding PLMs to relational semantics using prompt-based fine-tuning for the DialogRE

task. We focused on alleviating the critical challenge in dialogues, that is, low information density, by effectively capturing relational clues. In **GRASP**, we constructed prompts with deliberate initialization and suggested 1) an APM strategy considering both speaker and non-speaker arguments and 2) the RCD task, which guides the model to determine which token belongs to the relational clues. Experimental results on the DialogRE dataset revealed that **GRASP** achieved SoTA performance in terms of the F1 and F1<sub>c</sub> scores, even though our method only leveraged a PLM without adding any extra layers.

## Acknowledgements

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. 2020-0-00368, A Neural-Symbolic Model for Knowledge Acquisition and Inference Techniques), the MSIT (Ministry of Science and ICT), Korea, under the ITRC (Information Technology Research Center) support program (IITP-2022-2018-0-01405) supervised by the IITP, and the MSIT, Korea, under the ICT Creative Con-silience program (IITP-2022-2020-0-01819) supervised by the IITP.

## References

- Xuefeng Bai, Yulong Chen, Linfeng Song, and Yue Zhang. 2021. [Semantic representation for dialogue modeling](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4430–4445, Online. Association for Computational Linguistics.
- Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. [Abstract Meaning Representation for sembanking](#). In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 178–186, Sofia, Bulgaria. Association for Computational Linguistics.
- Douglas Biber. 1991. *Variation across speech and writing*. Cambridge University Press.
- Xiang Chen, Ningyu Zhang, Xin Xie, Shumin Deng, Yunzhi Yao, Chuanqi Tan, Fei Huang, Luo Si, and Huajun Chen. 2021. [Knowprompt: Knowledge-aware prompt-tuning with synergistic optimization for relation extraction](#). *arXiv preprint arXiv:2104.07650*.

- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Tianyu Gao, Adam Fisch, and Danqi Chen. 2020. Making pre-trained language models better few-shot learners. *arXiv preprint arXiv:2012.15723*.
- Deepanway Ghosal, Navonil Majumder, Alexander Gelbukh, Rada Mihalcea, and Soujanya Poria. 2020. [COSMIC: COMmonSense knowledge for eMotion identification in conversations](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2470–2481, Online. Association for Computational Linguistics.
- Xu Han, Weilin Zhao, Ning Ding, Zhiyuan Liu, and Maosong Sun. 2021. Ptr: Prompt tuning with rules for text classification. *arXiv preprint arXiv:2105.11259*.
- Heng Ji, Ralph Grishman, Hoa Trang Dang, Kira Griffitt, and Joe Ellis. 2010. Overview of the tac 2010 knowledge base population track. In *Third text analysis conference (TAC 2010)*, volume 3, pages 3–3.
- Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- Bongseok Lee and Yong Suk Choi. 2021. [Graph based network with contextualized representations of turns in dialogue](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 443–455, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Xiang Lisa Li and Percy Liang. 2021. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*.
- Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding, Yujie Qian, Zhilin Yang, and Jie Tang. 2021. Gpt understands, too. *arXiv preprint arXiv:2103.10385*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Xinwei Long, Shuzi Niu, and Yucheng Li. 2021. [Consistent inference for dialogue relation extraction](#). In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 3885–3891. International Joint Conferences on Artificial Intelligence Organization. Main Track.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Soujanya Poria, Devamanyu Hazarika, Navonil Majumder, Gautam Naik, Erik Cambria, and Rada Mihalcea. 2019. [MELD: A multimodal multi-party dataset for emotion recognition in conversations](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 527–536, Florence, Italy. Association for Computational Linguistics.
- Timo Schick and Hinrich Schütze. 2021. [Exploiting cloze-questions for few-shot text classification and natural language inference](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 255–269, Online. Association for Computational Linguistics.
- Livio Baldini Soares, Nicholas Fitzgerald, Jeffrey Ling, and Tom Kwiatkowski. 2019. Matching the blanks: Distributional similarity for relation learning. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2895–2905.
- Kumutha Swampillai and Mark Stevenson. 2010. Inter-sentential relations in information extraction corpora. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC’10)*.
- Dong Wang and Yang Liu. 2011. A pilot study of opinion summarization in conversations. In *Proceedings of the 49th annual meeting of the Association for Computational Linguistics: Human language technologies*, pages 331–339.
- Fuzhao Xue, Aixin Sun, Hao Zhang, and Eng Siong Chng. 2021. Gdpnet: Refining latent multi-view graph for relation extraction. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI*, pages 2–9.
- Yuan Yao, Jiaju Du, Yankai Lin, Peng Li, Zhiyuan Liu, Jie Zhou, and Maosong Sun. 2021. Codred: A cross-document relation extraction dataset for acquiring knowledge in the wild. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4452–4472.
- Dian Yu, Kai Sun, Claire Cardie, and Dong Yu. 2020. [Dialogue-based relation extraction](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4927–4940, Online. Association for Computational Linguistics.
- Sayed M Zahiri and Jinho D Choi. 2018. Emotion detection on tv show transcripts with sequence-based convolutional neural networks. In *Workshops at the thirty-second aai conference on artificial intelligence*.

## A Detailed experimental settings

Hyper-parameters	GRASP <sub>base</sub>	GRASP <sub>large</sub>
Learning rate	$5e-5$	$5e-6$
Max seq. len	512	
Batch size	8	
Num. epochs	30	
Joint ratio ( $\lambda_1$ & $\lambda_2$ )	0.7 / 0.3	

Table 8: Hyper-parameter values used in prompt-tuning process on test set.

GRASP is trained using AdamW (Loshchilov and Hutter, 2017) as an optimizer with no weight decay. The number of training epochs is set to 30 with early stopping, and the ratio of  $\lambda_1$  and  $\lambda_2$  for the joint loss is 0.7 to 0.3. A learning rate of  $5e-5$ , batch size of 8, and maximum sequence length of 512 are adopted for RoBERTa-base with identical parameters for RoBERTa-large, except for the learning rate of  $5e-6$ .

## B Qualitative Analysis Examples

Table 9 shows additional examples to demonstrate the prediction tendency of GRASP and the fine-tuned RoBERTa models on the symmetrical relations described in Section 5.2.

---

**Dialogue**


---

S1: Hey! Hi!  
S2: Hey!  
S1: What's up?  
S2: Well umm, **Chandler** and I are moving in together.  
S1: Oh my God. Ohh, my little sister and my best friend... shaking up. Oh, that's great. That's great.  
S3: Guys, I'm happy too.  
S2: Okay, come here!  
S3: Wow! Big day huh? People moving in, people getting annulled...  
S2: Okay, I gotta go find Rachel but umm, if you guys see her could you please try to give her some really bad news so that mine doesn't seem so bad?  
...

Argument pair	Ground Truth	RoBERTa	GRASP	
		Predicted Relation	Predicted ( <i>subject, object, trigger</i> )	Predicted Relation
(S2, <b>Chandler</b> )	per:girl/boyfriend	per:roommate	(S2, Chandler, <u>moving in together</u> )	per:girl/boyfriend
( <b>Chandler</b> , S2)	per:girl/boyfriend	per:roommate	(Chandler, S2, <u>moving in together</u> )	per:girl/boyfriend

---

**Dialogue**


---

S1: Rach, Rach, I just remembered. I had a dream about **Mr. Geller** last night.  
S2: Really?!

S1: Yeah, I dreamt that he saved me from a burning building and he was so brave and so strong! And it's making me look at him totally differently. Y'know, I mean he used to be just, y'know "Jack Geller **Monica** and **Ross**'s dad" and now he's he's "Jack Geller, dream hunk."  
S2: I dunno. Y'know to me he'll always be "Jack Geller, walks in while you're changing."

Argument pair	Ground Truth	RoBERTa	GRASP	
		Predicted Relation	Predicted ( <i>subject, object, trigger</i> )	Predicted Relation
(Mr. Geller, Monica)	<i>per:children</i>	unanswerable	(Mr Geller, Monica, <u>dad</u> )	per:children
(Mr. Geller, Ross)	<i>per:children</i>	unanswerable	(Mr Geller, Ross, <u>dad</u> )	per:children
(Ross, Mr. Geller)	<i>per:parents</i>	per:alternate_names	(Ross, Mr Geller, <u>dad</u> )	per:parents

Table 9: The additional examples for qualitative analysis on the prediction of **GRASP** based on the comparison with the RoBERTa model. Predicted (*subject, object, trigger*) is that **GRASP** predicted on RCD task.