

Decoupling Representation and Knowledge for Few-Shot Intent Classification and Slot Filling

Jie Han¹, Yixiong Zou^{1*}, Haozhao Wang¹, Jun Wang², Wei Liu¹, Yao Wu³, Tao Zhang³, Ruixuan Li^{1*}

¹School of Computer Science and Technology, Huazhong University of Science and Technology

²iWudao Tech, ³Banma Network Technology

{jiejhan,yixiongz,hz_wang,idc_lw,rxli}@hust.edu.cn, jwang@iwudao.tech, {qifang.wy,billow.zhangt}@alibaba-inc.com

Abstract

Few-shot intent classification and slot filling are important but challenging tasks due to the scarcity of finely labeled data. Therefore, current works first train a model on source domains with sufficiently labeled data, and then transfer the model to target domains where only rarely labeled data is available. However, experience transferring as a whole usually suffers from gaps that exist among source domains and target domains. For instance, transferring domain-specific-knowledge-related experience is difficult. To tackle this problem, we propose a new method that explicitly decouples the transferring of general-semantic-representation-related experience and the domain-specific-knowledge-related experience. Specifically, for domain-specific-knowledge-related experience, we design two modules to capture intent-slot relation and slot-slot relation respectively. Extensive experiments on Snips and FewJoint datasets show that our method achieves state-of-the-art performance. The method improves the joint accuracy metric from 27.72% to 42.20% in the 1-shot setting, and from 46.54% to 60.79% in the 5-shot setting.

Introduction

Natural language understanding (NLU) is a critical component of conversational dialogue systems, such as Siri, Alexa, and Google Assistant. Specifically, NLU includes two sub-tasks: (1) intent classification, which classifies an utterance into an intent label, and (2) slot filling, which classifies each word in the utterance into a slot label. Both sub-tasks rely on large amounts of finely labeled data, which is difficult to obtain. Therefore, some work proposed to study NLU in the setting of few-shot, where only a few labeled data are available (Krone, Zhang, and Diab 2020; Gangadharaiyah and Narayanaswamy 2022). Recently, to boost NLU learning in the few-shot setting, some methods tried to utilize labeled data from other sufficiently labeled source domains to help NLU learning (in the rarely labeled target domains) (Rongali et al. 2023; Kwon et al. 2023). However, due to gaps that exist among different domains, it is hard to directly transfer prior experience from source domains to target domains.

In this paper, we hold that the prior experience in source domains consists of two parts: the **general semantic representation** that refers to utterance semantics, and the

*Corresponding author.

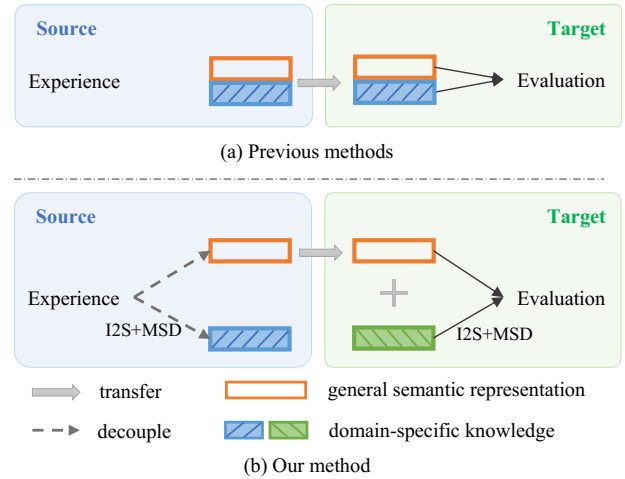


Figure 1: Framework of Joint Modeling with Relationship Masking (JMRM). Firstly, we explicitly decouple the general semantic representation and the domain-specific knowledge of source domains. Secondly, we only transfer the general semantic representation to target domains. Thirdly, we apply the knowledge specific to target domains to the transferred representations.

domain-specific knowledge that refers to the intent-slot relation and slot-slot relation. Intuitively, general semantics are shared across domains while the domain-specific intents/slots are not. Therefore, it is reasonable to transfer the general semantic representation from source domains to target domains, while it is difficult to transfer the domain-specific knowledge. Despite a few work has noticed the differences between the two parts, they still tried to conduct the transferring as a whole (Hou et al. 2021; Liu et al. 2021), which may be inefficient.

To this end, we propose **Joint Modeling with Relationship Masking (JMRM)**, which explicitly decouples the transferring of the general semantic representation and the domain-specific knowledge. Specifically, to achieve decoupling, we design two modules, **I2S-Mask (I2S)** and **Masked Slot Decoding (MSD)**, to capture the domain-specific knowledge (intent-slot relation and slot-slot relation). In the I2S module, we introduce an intent-slot correla-

tion score matrix to capture the intent-slot relation. Here, the matrix is automatically summarized from the labeled data. Furthermore, it regularizes the predicted intent and slot labels that are related. For example, as shown in Figure 2, in the predictions of an utterance, each intent label can only co-occur with the slot labels related to this intent, such as *PlayMusic* and *B-artist*, and an intent label should not co-occur with the slot labels unrelated to the intent, such as *PlayMusic* and *B-city*. In the MSD module, we introduce a slot-slot constraint score matrix to capture the slot-slot relation. Here, the matrix is automatically summarized from BIO annotation rules. Furthermore, it constrains the predicted slot label sequence is rational. For example, as shown in Figure 2, in the predicted slot label sequence, *I-artist* is allowed to follow *B-artist*, but not *O*.

Specifically, we implement the above design as two masking operations applied during both training and testing. During source-domain training, such operations will decouple the general semantic representation and the domain-specific knowledge. During target-domain testing, such operations will efficiently summarize target-domain information, and apply them to the transferred general representations, which therefore helps the target-domain recognition.

Experiments on two public datasets, the Snips dataset (Coucke et al. 2018) and the FewJoint dataset (Hou et al. 2020b), show that JMRM achieves state-of-the-art performance. Specifically, JMRM improves the joint accuracy from 27.72% to 42.20% in the 1-shot setting, and from 46.54% to 60.79% in the 5-shot setting. Furthermore, since JMRM is a plug-and-play method, we demonstrate that it can improve the effectiveness of other models. Moreover, extensive analysis suggests that decoupling the transfer of the general semantic representation and the domain-specific knowledge is more efficient than transferring as a whole.

In summary, the main contributions of this paper are summarized as follows:

1. To relieve the difficulty of transferring caused by the gaps between different domains, we try to explicitly decouple the general semantic representation and the domain-specific knowledge.
2. We propose the JMRM method, which contains I2S and MSD modules. The two modules explicitly utilize two relationship score matrices to capture the domain-specific knowledge. Furthermore, we jointly consider intent label and slot label sequence in the training process.
3. Our method achieves state-of-the-art performance. Furthermore, we validate the effectiveness of the decoupling by extensive experiments.

Method

In this section, we introduce the proposed Joint Modeling with Relationship Masking (JMRM) for few-shot intent classification and slot filling.

Background

Problem Definition Natural language understanding (NLU) contains two related tasks: intent classification, which is a sentence-level text classification task, and

slot filling, which is a word-level sequence labeling task. Formally, a labeled sample can be represented as $(\mathbf{x}, y, \mathbf{t})$, where $\mathbf{x} = \langle x_1, \dots, x_m \rangle$ is an utterance with m words, and y is the intent label of \mathbf{x} , and $\mathbf{t} = \langle t_1, \dots, t_m \rangle$ is the slot label sequence of \mathbf{x} . Here, t_i is the slot label of x_i . A few-shot NLU model aims to predict both y and \mathbf{t} for each \mathbf{x} according to the scarce labeled data in unseen target domains, with the experience transferred from source domains.

To match the source-domain setting with the target-domain setting, we construct multiple episodes in source domains. Each episode contains a support set \mathcal{S} with few labeled samples and a query set \mathcal{Q} with samples to be predicted. Formally, the support set is denoted as $\mathcal{S} = \{(\mathbf{x}^{(n)}, y^{(n)}, \mathbf{t}^{(n)})\}_{n=1}^{|\mathcal{S}|}$, and the query set is denoted as $\mathcal{Q} = \{(\mathbf{x}^{(n)}, y^{(n)}, \mathbf{t}^{(n)})\}_{n=1}^{|\mathcal{Q}|}$, where n denotes the index of the sample. During training, a model learns to predict labels of samples in \mathcal{Q} based on \mathcal{S} .

Preliminaries Before utilizing the two proposed modules, we obtain emission scores which are used to predict labels, based on Prototypical Networks (Snell, Swersky, and Zemel 2017). Firstly, we calculate the label representation as to the mean embedding of support samples belonging to the same class:

$$C_l = \frac{1}{|\mathcal{S}_l|} \sum_{(\mathbf{x}, y, \mathbf{t}) \in \mathcal{S}_l} E(\mathbf{x}), \quad (1)$$

$$C_o = \frac{1}{|\mathcal{S}_o|} \sum_{(\mathbf{x}, y, t_i) \in \mathcal{S}_o} E(x_i), \quad (2)$$

where $E(\cdot)$ is a BERT (Devlin et al. 2019) based encoding function, and $\mathcal{S}_l = \{(\mathbf{x}, y, \mathbf{t}) | y = l\}$ is the set of the samples with intent label l in \mathcal{S} , and $\mathcal{S}_o = \{(\mathbf{x}, y, t_i) | t_i = o\}$ is the set of the samples with slot label o in \mathcal{S} . Therefore, C_l denotes the label representation of intent label l , and C_o denotes the label representation of slot label o .

Then, for each query sample, we obtain the emission scores by calculating the similarity scores of the label representation and the sample embedding. Specifically, we calculate the intent emission score $f_l \in \mathbb{R}^{1 \times Y}$ and the original slot emission score $f_o \in \mathbb{R}^{m \times T}$ as follows:

$$f_l(y = l) = \text{SIM}(E(\mathbf{x}), C_l), \quad (3)$$

$$f_o(t_i = o) = \text{SIM}(E(x_i), C_o), \quad (4)$$

where Y and T denote the number of intent and slot classes, respectively, and $\text{SIM}(\cdot, \cdot)$ is a similarity function.

I2S-Mask

To capture the domain-specific knowledge of intent-slot relation, we propose I2S-Mask module. It utilizes an intent-slot correlation score matrix as the domain-specific information of intent-slot relation, which regularizes the predicted intent and slot labels are related.

Specifically, in this paper, we define that an intent label and a slot label are **related** if they appear in a support sample simultaneously, such as *PlayMusic* and *B-artist* in the support set of Figure 2. Otherwise, they are **unrelated**, such as *PlayMusic* and *B-city*.

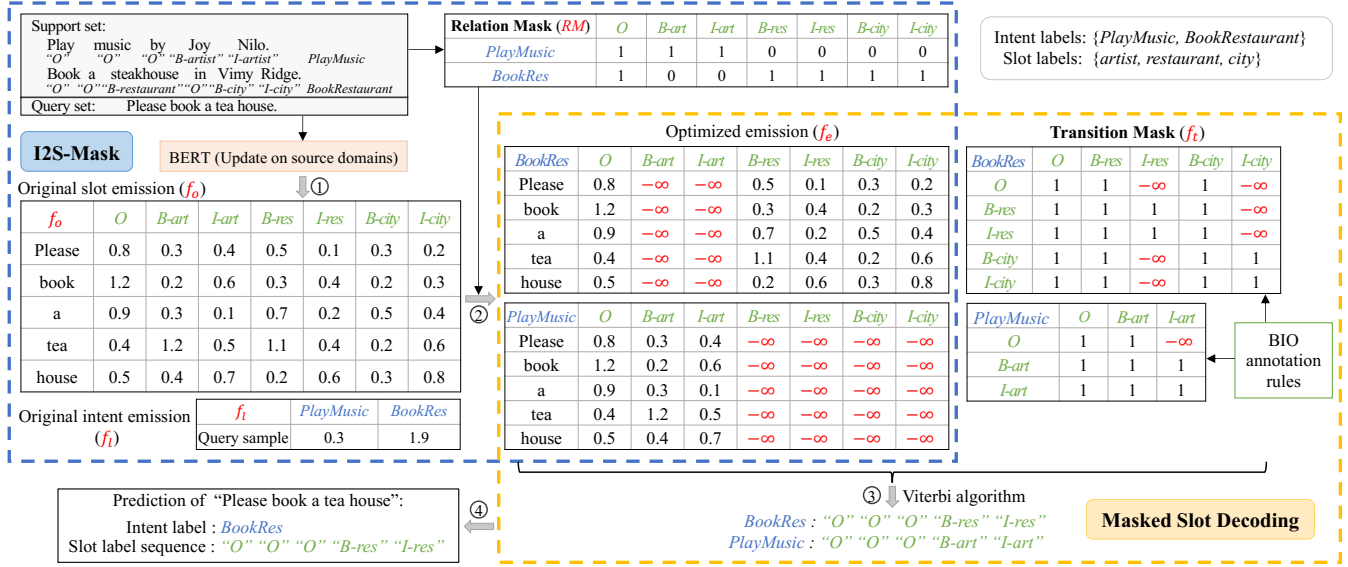


Figure 2: Illustration of two main components of our method: the I2S-Mask module, which captures the domain-specific knowledge of intent-slot relation, and the Masked Slot Decoding module, which captures the domain-specific knowledge of slot-slot relation. Specifically, *Bookres* denotes *BookRestaurant*, *art* denotes *artist*, and *res* denotes *restaurant*. We apply the two modules during both the training on source domains and the evaluation on target domains.

Firstly, we obtain an intent-slot relation mask, which is the intent-slot correlation score matrix summarized from the support set. Formally, we denote the relation mask as $RM \in \mathbb{R}^{Y \times T}$. It is shown in the I2S-Mask module of Figure 2, where 1 denotes the intent label in this row and the slot label in this column is related, and 0 denotes they are unrelated.

Then, we set the slot emission scores that are unrelated to the currently calculated intent label l to negative infinity according to the RM:

$$f_e(t_i = o | y = l) = \begin{cases} f_o(t_i = o), & RM_{l,o} = 1 \\ -\infty, & \text{otherwise} \end{cases} \quad (5)$$

where $f_e \in \mathbb{R}^{m \times T}$ and $RM_{l,o} = 1$ means that the intent label l and slot label o are related.

In summary, I2S-Mask calibrates the original slot emission (f_o) to the optimized emission (f_e) by the RM, which regularizes the predicted intent and slot labels are related. Thus, during source-domain training, the model could decouple the source-domain-specific knowledge of intent-slot relation and the general semantic representation. And during target-domain evaluation, without the source-domain-specific knowledge of intent-slot relation, the general semantic representation and the target-domain-specific knowledge of intent-slot relation make predictions more accurate.

Masked Slot Decoding

To capture the domain-specific knowledge of slot-slot relation, we propose Masked Slot Decoding module. It utilizes a slot-slot constraint score matrix as the domain-specific information of slot-slot relation, which constrains the predicted slot label sequence is rational.

Specifically, a **rational** slot label sequence follows BIO annotation rules. Formally, the *B-* should be the label for the first word of a slot phrase, and the *I-* should be the label for the other words of the slot phrase. *O* denotes *other*, indicating that the word is not important in the utterance. For example, slot labels of "Vimy Ridge" are "*B-city I-city*". And *I-city* is allowed to follow *B-city*, but not *O*.

Firstly, we obtain a slot-slot transition mask, which is the slot-slot constraint score matrix summarized according to the BIO annotation rules of the slot labels in the support set. Formally, we denote the transition mask as $f_t \in \mathbb{R}^{T \times T}$:

$$f_t(t_i = o_2 | t_{i-1} = o_1) = \begin{cases} 1, & \text{BIO}_{o_1, o_2} = 1 \\ -\infty, & \text{otherwise} \end{cases} \quad (6)$$

where $\text{BIO}_{o_1, o_2} = 1$ means that slot label o_2 is allowed to follow slot label o_1 .

Then, we utilize Viterbi algorithm (Lafferty, McCallum, and Pereira 2001) to predict the slot label sequence. The input of Viterbi contains the optimized emission (f_e) and the unlearned transition mask (f_t). The Viterbi algorithm utilizes the idea of dynamic programming to calculate the optimal slot label sequence among all possible predictions.

In summary, Masked Slot Decoding obtains the transition mask (f_t) according to BIO annotation rules, which regularizes the predicted slot label sequence is rational. Thus, during source-domain training, the model could decouple the source-domain-specific knowledge of the slot-slot relation and the general semantic representation. And during target-domain evaluation, without the source-domain-specific knowledge of slot-slot relation, the general semantic representation and the target-domain-specific knowledge of slot-slot relation make predictions more accurate.

Joint Modeling

Since the intent classification task and the slot filling task are strongly related, it is beneficial to learn them jointly (Weld et al. 2022). To this end, we jointly consider intent label and slot label sequence in the training process.

In the training process on source domains, firstly, given a query utterance \mathbf{x}^* and the support set \mathcal{S} , we calculate the score of the intent label of \mathbf{x}^* is y and the slot label sequence of \mathbf{x}^* is \mathbf{t} as follows, with the source-domain-specific knowledge:

$$R(y, \mathbf{t} | \mathbf{x}^*, \mathcal{S}) = \lambda \cdot f_I(y) + \sum_{i=1}^m (f_e(t_i | y) + f_t(t_i | t_{i-1})), \quad (7)$$

where λ is the weight of the intent score.

Secondly, we obtain the probability of the correct intent label y^* and slot label sequence \mathbf{t}^* :

$$p(y^*, \mathbf{t}^* | \mathbf{x}^*, \mathcal{S}) = \frac{\exp(R(y^*, \mathbf{t}^* | \mathbf{x}^*, \mathcal{S}))}{\sum_{y, \mathbf{t}} \exp(R(y, \mathbf{t} | \mathbf{x}^*, \mathcal{S}))}. \quad (8)$$

Thirdly, we utilize a single cross-entropy loss function to update the few-shot NLU model:

$$\mathcal{L} = -\log p(y^*, \mathbf{t}^* | \mathbf{x}^*, \mathcal{S}). \quad (9)$$

During the evaluation on target domains, the transferred model predicts the intent label and slot label sequence simultaneously as follows, with the target-domain-specific knowledge:

$$y^*, \mathbf{t}^* = \arg \max_{y, \mathbf{t}} p(y, \mathbf{t} | \mathbf{x}^*, \mathcal{S}). \quad (10)$$

The joint modeling considers both intent classification and slot filling tasks simultaneously. And as validated in Figure 4, the two related tasks guide each other to obtain better performance.

Experiments

Dataset and Domain Construction

Dataset. We conduct extensive experiments on two natural language understanding (NLU) benchmarks: the Snips dataset (Coucke et al. 2018) and the FewJoint dataset (Hou et al. 2020b). Specifically, Snips is an English dataset with 7 intent classes and 48 slot classes. FewJoint is a Chinese dataset with 141 intent classes and 193 slot classes.

Domain construction. We aim to provide a comprehensive description of domain construction across three levels of granularity: domain types, the quantity of episodes within each domain type, and episode construction.

In our experiments, we construct the training domains as source domains to update model parameters, the developing domains to select the best model, the testing domains as target domains to evaluate the performance. For Snips, we construct a training domain with 3 intent classes PlayMusic, AddToPlaylist, BookRestaurant, a developing domain with 2 intent classes RateBook, SearchScreeningEvent, and a testing domain with 2 intent classes GetWeather, SearchCreativeWork. For FewJoint, we construct 38 training domains,

5 developing domains, and 9 testing domains. Here, each domain has multiple intent labels and slot labels.

For the quantity of episodes in each domain, on Snips, we construct 200, 50, 10 episodes for the training, developing and testing domains, respectively. On FewJoint, we use the original few-shot episodes.

Formally, each episode contains a support set with few labeled samples and a query set with some samples to be predicted. For constructing episodes, we adopt the Mini-Including Algorithm (Hou et al. 2020a) that constructs a support set in the K -shot setting following two criteria: (1) Each slot class contains at least K examples. (2) Removing any utterance would cause the former be not hold.

Baselines and Evaluation Metrics

We compare our method with competitive few-shot NLU baselines:

- **SepProto** utilizes Prototypical Networks (Snell, Swersky, and Zemel 2017) for the intent classification and slot filling tasks with the separate BERT (Devlin et al. 2019) embedding. The model is trained on the training domains and then evaluated directly on the unseen testing domains without fine-tuning.
- **JointProto** (Krone, Zhang, and Diab 2020) jointly learns the intent and slot representations by sharing a single BERT encoder on source domains, without fine-tuning on target domains.
- **ConProm** (Hou et al. 2021) merges the intent and slot representations into one space and learns the representations by contrastive learning.
- **ConProm+TR+FT**, where FT denotes fine-tuning models, and TR denotes the transition rules of BIO annotation, which ban illegal slot predictions from left to right during target-domain evaluation.

There are 3 evaluation metrics in the NLU task: intent accuracy (Intent Acc), slot F1-score (Slot F1), and joint accuracy (Joint Acc). The Joint Acc means that both the predicted intent label and slot label sequence of a query utterance are correct, which is the most important metric. In our experiments, we select the model on the developing domains according to Joint Acc, and then report the performance of the model on testing domains as final results. To relieve the non-deterministic model training (Reimers and Gurevych 2017), we report the average score of 5 experiments with different random seeds for every setting.

Implementation Details

For a fair comparison, the hyperparameters in our experiments are set the same as in baselines. The batch size is 4 and the learning rate is 10^{-5} . We set a single BERT (Devlin et al. 2019) as the embedding function of intents and slots, where the intent representation of an utterance is the average of all word embedding in the utterance. The weight of the intent score λ in Eq. 7 is 1. We use ADAM (Kingma and Ba 2015) to update model parameters and transfer the learned general semantic representation from source domains to target domains without fine-tuning. To observe the robustness of our

Models	Snips			FewJoint		
	Intent Acc.	Slot F1	Joint Acc.	Intent Acc.	Slot F1	Joint Acc.
SepProto	98.23 ± 0.66	43.90 ± 1.98	9.47 ± 2.10	66.35 ± 0.51	27.24 ± 1.10	10.92 ± 0.89
JointProto	92.57 ± 0.57	42.63 ± 2.03	7.35 ± 1.70	58.52 ± 0.28	29.49 ± 1.01	9.64 ± 0.47
ConProm	96.67 ± 1.45	53.05 ± 0.81	21.72 ± 0.97	65.26 ± 0.23	33.09 ± 1.66	16.32 ± 0.75
ConProm+FT+TR	90.45 ± 0.52	56.04 ± 1.75	27.80 ± 2.33	63.67 ± 0.94	42.44 ± 0.51	27.72 ± 0.95
JMRM(Ours)	93.71 ± 0.74	66.55 ± 0.44	41.76 ± 1.25	65.97 ± 0.70	62.24 ± 1.17	42.20 ± 0.88

Table 1: Results of 1-shot natural language understanding task on Snips and FewJoint. TR denotes the BIO transition rule of adjacent slots, e.g. slot label “I” is not allowed to follow “O”. FT denotes fine-tuning model on target domains.

Models	Snips			FewJoint		
	Intent Acc.	Slot F1	Joint Acc.	Intent Acc.	Slot F1	Joint Acc.
SepProto	99.53 ± 0.11	53.28 ± 1.85	14.40 ± 3.00	75.64 ± 1.51	36.08 ± 0.65	15.93 ± 1.85
JointProto	99.17 ± 0.09	50.63 ± 2.01	13.40 ± 1.44	70.93 ± 2.45	39.47 ± 1.05	14.48 ± 1.11
ConProm	98.50 ± 0.42	61.03 ± 1.77	32.20 ± 2.06	78.05 ± 1.04	39.40 ± 1.75	24.18 ± 1.29
ConProm+FT+TR	98.40 ± 0.20	72.98 ± 0.41	52.95 ± 0.85	78.43 ± 1.86	69.44 ± 0.39	46.54 ± 0.72
JMRM(Ours)	97.84 ± 0.69	77.82 ± 1.47	59.24 ± 2.25	82.26 ± 0.64	76.15 ± 0.58	60.79 ± 0.30

Table 2: Results of 5-shot natural language understanding task on Snips and FewJoint.

method to the similarity function, we utilize 3 different similarity functions in our experiments, including cosine (Cos), euclidean distance (L2), and the vector projection with bias (VPB) (Zhu et al. 2020).

We conduct experiments of 1-shot on GeForce GTX 1080, and 5-shot on GeForce RTX 3090. Training on Snips takes 2.8 hours on average. Due to early convergence, on FewJoint, the 1-shot training takes 0.8 hours and the 5-shot training takes 2.2 hours.

Results and Analysis

To verify the proposed method, we conduct extensive experiments on the NLU task, in 1-shot and 5-shot settings. Moreover, we present some auxiliary experiments to analyze the effectiveness of the proposed decoupling and joint modeling for few-shot NLU.

Result of 1-shot The experimental results of the NLU task in the 1-shot setting are shown in Table 1. For Joint Acc, the most important metric, our method achieves state-of-the-art performance. Specifically, on Snips we improve Joint Acc from 27.80% to 41.76%, and on FewJoint we improve Joint Acc from 27.72% to 42.20%.

Interestingly, the Intent Acc results of SepProto are highest on both Snips and FewJoint. This is because SepProto has two BERT models, which learn intent embedding and slot embedding separately. In contrast, our method is not competitive on Intent Acc. This is because with sharing a single BERT encoder when we focus more on the slot filling task in the 1-shot setting, the performance of intent classification tends to decrease. The same phenomenon has also been observed in other works (Hou et al. 2021; Krone, Zhang, and Diab 2020).

Result of 5-shot The experimental results of the NLU task in the 5-shot setting are shown in Table 2. On the metric of Joint Acc, our method also achieves state-of-the-art performance. Specifically, on Snips we improve Joint Acc from 52.95% to 59.24%, and on FewJoint we improve Joint Acc from 46.54% to 60.79%.

Notably, our method achieves the best Intent Acc on FewJoint. This profits from the joint modeling, which makes the two related tasks guide each other. As an explanation, we explicitly add the intent guidance for slots and slot guidance for intents into training objectives. Here, an improvement in the performance of slot filling beyond a certain threshold can facilitate the learning process of intent classification.

Ablation study To analyze contributes of I2S-Mask (I2S) and Masked Slot Decoding (MSD), we conduct ablation experiments. Table 3 shows the Joint Acc results of 1-shot and 5-shot on FewJoint. JM is the proposed joint modeling, without I2S and MSD. +I2S denotes introducing the intent-slot relation mask in I2S into training objectives, and +MSD denotes introducing the slot-slot transition mask in MSD into training objectives. Furthermore, to observe the robustness of our method to the similarity function, we utilize three similarity functions, including Cos, L2, and VPB.

Remarkably, I2S and MSD improve performance in all settings. Specifically, compared with JointProto, JointProtoI2S performs better on all similarity functions, and the highest improvement of 1-shot is 7.35 points and that of 5-shot is 13.24 points, both on the L2 similarity function. And compared with JM, JMI2S and JMMSD perform better on all similarity functions. Interestingly, JMI2S is better than JMMSD. This is because, in I2S, the intents and slots calculated in Eq. 7 are all related, making the bi-directional guidance between intents and slots more helpful. And compared with JM, JMRM improves the performance by up to 17.06

Models	Training		1-shot			5-shot		
	+I2S	+MSD	Cos	L2	VPB	Cos	L2	VPB
(JointProto)			27.54	26.11	26.92	39.70	37.69	43.02
(JointProtoI2S)	✓		33.30 ↑5.76	33.46 ↑7.35	28.89 ↑1.97	48.37 ↑8.67	50.93 ↑13.24	52.25 ↑9.23
(JM)			24.29	25.79	24.98	38.93	37.24	45.37
(JMI2S)	✓		36.56 ↑12.27	37.09 ↑11.30	37.78 ↑12.80	49.31 ↑10.38	50.86 ↑13.62	54.57 ↑9.20
(JMMSD)		✓	32.99 ↑8.70	32.11 ↑6.32	32.90 ↑7.92	47.34 ↑8.41	46.54 ↑9.30	46.42 ↑1.05
JMRM	✓	✓	41.35 ↑17.06	41.75 ↑15.96	41.97 ↑16.99	59.54 ↑20.61	58.82 ↑21.58	60.79 ↑15.42

Table 3: Ablation study of the I2S-Mask (I2S) module and the Masked Slot Decoding (MSD) module on FewJoint, with three different similarity functions, including Cos, L2 and VPB. JM denotes joint modeling.

Models	Snips		FewJoint	
	1-shot	5-shot	1-shot	5-shot
JointProto	7.35	13.40	9.64	14.48
+RM	21.32	19.13	30.11	40.18
ConProm	21.72	32.20	16.32	24.18
+RM	30.65	47.30	34.21	37.72
JM	30.42	38.37	24.29	38.93
+RM	38.40	51.90	36.71	53.53

Table 4: Exploration of the effect of the target-domain-specific knowledge, with Joint Acc metrics.

points with Cos in the 1-shot setting and 21.58 points with L2 in the 5-shot setting.

Furthermore, in most instances, VPB is better than Cos and L2. The explanation is as follows. VPB is a metric function proposed by Zhu et al., and its formula is: $\text{SIM}(e, C) = e^T \frac{C}{\|C\|} - \frac{1}{2}\|C\|$, where e denotes the sample embedding and C denotes the label representation. The regularization of C in VPB prevents excessive influence when it becomes too large and reduces false positive errors, thereby improving performance (Zhu et al. 2020).

Analysis

- **Domain-specific knowledge helps predictions be more accurate.** To explore the effect of using the target-domain-specific knowledge, we conduct the experiments in Table 4. +RM denotes that a model utilizes I2S and MSD only during target-domain evaluation.

Remarkably, +RM improves the performance of all models. For example, in the 5-shot setting on FewJoint, +RM helps JointProto, ConProm, and JM improve the Joint Acc metrics up to 40.18%, 37.72%, and 53.53%, respectively. This is because +RM captures the target-domain-specific knowledge, which makes predictions satisfy the relationship constraints. Thus, domain-specific knowledge makes predictions more accurate.

Moreover, our method JM with RM achieves the best results in all settings compared to other methods with RM. This indicates that the proposed joint modeling is efficient for few-shot NLU. Furthermore, the JMRM is bet-

ter than JM+RM, which shows that the decoupling during source-domain training is efficient.

- **Decoupling makes transferring more efficient.** To investigate the effectiveness of decoupling the general semantic representation and the domain-specific knowledge, we conduct the experiment in Figure 3. Here, we list Joint Acc results of three different training methods, in the 5-shot setting on Snips. JMI2S uses only the I2S module during training, decoupling the domain-specific knowledge of intent-slot relation. JMRM uses both the I2S module and MSD module during training, decoupling the domain-specific knowledge of both intent-slot relation and slot-slot relation.

Notably, JM, JMI2S, and JMRM all perform similarly on source domains, which indicates that these three methods have learned source-domain experience of roughly equivalent magnitudes. During target-domain evaluation, JM transfers all source-domain experience as a whole, and JMI2S and JMRM transfer the decoupled general semantic representation. Since JMI2S is better than JM and JMRM is better than JMI2S, we conclude that, due to the gaps among different domains, the source-domain-specific knowledge will suffer the model’s performance on target domains. Thus, decoupling the general semantic representation and the domain-specific knowledge makes transferring more efficient.

- **JM achieves bi-directional guidance.** To explore the effect of the proposed joint modeling, we analyze the experiments in Figure 4. JointProtoI2S first predicts intent labels, and then predicts slot labels according to the relation mask. This means that JointProtoI2S uses intent information to guide the slot prediction. Significantly, JointProto compares with JointProtoI2S, the Intent Acc is about the same, the Slot F1 is a little better, but the Joint Acc is lower. It means that in JointProto, there are some samples whose slot predictions are correct, but their intent labels are wrong. Therefore, we conjecture that slot information can also guide the learning of intents.

To verify our conjecture, we further design SeqCEI2S which is different from JMI2S only in the joint modeling. Specifically, the loss of JointProtoI2S is: $\mathcal{L} = \mathcal{L}_{intent} + \sum_i \mathcal{L}_{slot_i}$, and the loss of SeqCEI2S is: $\mathcal{L} = \mathcal{L}_{intent} + \mathcal{L}_{slot.seq}$, and the loss of JMI2S is: $\mathcal{L} = \mathcal{L}_{intent+slot.seq}$, where $slot.seq$ denotes predicting the slot label sequence




Figure 3: Performance of different methods on source and target domains, in 5-shot setting on Snips.

using the Viterbi algorithm. The results show that the Intent Acc of JMI2S is higher than SeqCEI2S, which indicates that the intent task is indeed influenced by the guidance from the slot task. JMI2S also has higher Slot F1 and Joint Acc than others. Thus we conclude that joint modeling achieves bi-directional guidance of both intents and slots, resulting in better performance.

Related Work

Natural Language Understanding

Natural language understanding (NLU) is an important component of dialogue systems, including the intent classification task and the slot filling task. There are a lot of NLU methods (Goo et al. 2018; Wang, Shen, and Jin 2018; Liu et al. 2020; Ma et al. 2021b; Rosenbaum et al. 2022; Zheng et al. 2023; Ma et al. 2022). Qin et al. directly took the output of the intent task as the input to the slot task. Chen, Zhuo, and Wang utilized a shared BERT encoder to jointly learn intents and slots. Qin et al. utilized multiple label attention layers and co-interactive attention layers to jointly encode the intent and slot representations. Qin et al. listed some new areas and challenges related to NLU. However, these methods are limited in performance in specialist areas where data are highly variable and samples are difficult to collect.

Few-Shot Learning

Few-shot learning aims to learn models based on a few samples. Generally, the models are first trained on sufficient source domains, then evaluated on unseen target domains with few labeled data. Few-shot learning approaches in natural language processing mainly include three approaches. In the fine-tuning-based approaches (Sun et al. 2019; Shen et al. 2021), MAML (Finn, Abbeel, and Levine 2017) trained model parameters such that a small number of gradient updates will lead to fast learning on a new task with few labeled data, and ULMFiT (Howard and Ruder 2018) fine-tuned the language model and the classifier on target tasks. In the prompting-based approaches (Ma et al. 2021a; Gao, Fisch, and Chen 2021; Li and Liang 2021), PET (Schick and Schütze 2021) converted the text classification task into a masked language model task for few-shot learning. In the metric learning-based approaches (Snell, Swersky, and Zemel 2017; Hou et al. 2020a; Yuan et al. 2021), VPB (Zhu et al. 2020) utilized the projections of contextual word representations on each normalized label repre-




Figure 4: Analysis of the benefits of joint modeling.

sentation as the word-label similarity. However, these methods are not designed for intent and slot tasks in NLU. Our method specially solves the intent and slot problems by considering the relevance of these two tasks.

Recently, few-shot NLU attracts widespread attention due to the data scarcity problem (Rongali et al. 2023; Kwon et al. 2023; Gangadharaiah and Narayanaswamy 2022). ZEROTOP (Mekala, Wolfe, and Roy 2023) utilized large language models to complete NLU tasks and designed different prompts for different intent/slot labels. However, the large language model is worse at slot filling (He and Garner 2023; Qin et al. 2023). Thus, it is reasonable to use classical few-shot learning methods for few-shot NLU. Krone, Zhang, and Diab utilized classical few-shot learning methods MAML and Prototypical Networks to solve the few-shot intent classification and slot filling problem. However, these methods did not take into account the relationship between the two tasks. Therefore, few-shot joint learning of intents and slots becomes popular (Basu et al. 2021). Hou et al. merged the intent and slot representations into one space with the attention mechanism. Liu et al. utilized contrastive learning to learn the intent and slot representations jointly. However, the transferred experience remains difficult due to the gaps between different domains. Therefore, we consider decoupling the general semantic representation and the domain-specific knowledge to relieve the transfer disaster.

Conclusion

In this paper, we propose a new method, **JMRM**, for few-shot intent classification and slot filling. It explicitly decouples general semantic representation and domain-specific knowledge, and only transfers the general semantic representation to target domains. Specifically, to capture the domain-specific knowledge, we design the I2S-Mask and Masked Slot Decoding modules, which utilize two relationship score matrices to regularize predictions. Experiments validate that decoupling makes transferring more efficient.

Acknowledgments

We are grateful to the anonymous reviewers for their valuable comments. This work is supported by National Natural Science Foundation of China under grants 62376103, 62302184, 62206102, and Science and Technology Support Program of Hubei Province under grant 2022BAA046.

References

- Basu, S.; Sharaf, A.; Fischer, A.; Rohra, V.; Amoake, M.; El-Hammamy, H.; Nosakhare, E.; Ramani, V.; Han, B.; et al. 2021. Semi-Supervised Few-Shot Intent Classification and Slot Filling. *arXiv preprint arXiv:2109.08754*.
- Chen, Q.; Zhuo, Z.; and Wang, W. 2019. Bert for joint intent classification and slot filling. *arXiv preprint arXiv:1902.10909*.
- Coucke, A.; Saade, A.; Ball, A.; Bluche, T.; Caulier, A.; Leroy, D.; Doumouro, C.; Gisselbrecht, T.; Caltagirone, F.; Lavril, T.; et al. 2018. Snips voice platform: an embedded spoken language understanding system for private-by-design voice interfaces. *arXiv preprint arXiv:1805.10190*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 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)*, 4171–4186.
- Finn, C.; Abbeel, P.; and Levine, S. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, 1126–1135. PMLR.
- Gangadharaiah, R.; and Narayanaswamy, B. 2022. Zero-Shot Learning for Joint Intent and Slot Labeling. *arXiv preprint arXiv:2212.07922*.
- Gao, T.; Fisch, A.; and Chen, D. 2021. Making Pre-trained Language Models Better Few-shot Learners. 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)*, 3816–3830.
- Goo, C.-W.; Gao, G.; Hsu, Y.-K.; Huo, C.-L.; Chen, T.-C.; Hsu, K.-W.; and Chen, Y.-N. 2018. Slot-Gated Modeling for Joint Slot Filling and Intent Prediction. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, 753–757.
- He, M.; and Garner, P. N. 2023. Can ChatGPT Detect Intent? Evaluating Large Language Models for Spoken Language Understanding. *arXiv preprint arXiv:2305.13512*.
- Hou, Y.; Che, W.; Lai, Y.; Zhou, Z.; Liu, Y.; Liu, H.; and Liu, T. 2020a. Few-shot Slot Tagging with Collapsed Dependency Transfer and Label-enhanced Task-adaptive Projection Network. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 1381–1393.
- Hou, Y.; Lai, Y.; Chen, C.; Che, W.; and Liu, T. 2021. Learning to Bridge Metric Spaces: Few-shot Joint Learning of Intent Detection and Slot Filling. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, 3190–3200.
- Hou, Y.; Mao, J.; Lai, Y.; Chen, C.; Che, W.; Chen, Z.; and Liu, T. 2020b. Fewjoint: A few-shot learning benchmark for joint language understanding. *arXiv preprint arXiv:2009.08138*.
- Howard, J.; and Ruder, S. 2018. Universal Language Model Fine-tuning for Text Classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 328–339.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Krone, J.; Zhang, Y.; and Diab, M. 2020. Learning to Classify Intents and Slot Labels Given a Handful of Examples. In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI*, 96–108.
- Kwon, S. Y.; Bhatia, G.; Nagoudi, E. M. B.; Inciarte, A. A.; and Abdul-Mageed, M. 2023. Zero-Shot Slot and Intent Detection in Low-Resource Languages. *arXiv preprint arXiv:2304.13292*.
- Lafferty, J. D.; McCallum, A.; and Pereira, F. C. N. 2001. Conditional Random Fields: Probabilistic Models for Segmenting and Labeling Sequence Data. In *Proceedings of the Eighteenth International Conference on Machine Learning*, 282–289.
- Li, X. L.; and Liang, P. 2021. Prefix-Tuning: Optimizing Continuous Prompts for Generation. 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)*, 4582–4597.
- Liu, H.; Zhang, F.; Zhang, X.; Zhao, S.; and Zhang, X. 2021. An Explicit-Joint and Supervised-Contrastive Learning Framework for Few-Shot Intent Classification and Slot Filling. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, 1945–1955.
- Liu, Z.; Winata, G. I.; Xu, P.; and Fung, P. 2020. Coach: A Coarse-to-Fine Approach for Cross-domain Slot Filling. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 19–25.
- Ma, J.; Yan, Z.; Li, C.; and Zhang, Y. 2021a. Frustratingly Simple Few-Shot Slot Tagging. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, 1028–1033.
- Ma, Z.; Li, J.; Li, G.; and Cheng, Y. 2022. UniTranSeR: A Unified Transformer Semantic Representation Framework for Multimodal Task-Oriented Dialog System. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 103–114.
- Ma, Z.; Li, J.; Zhang, Z.; Li, G.; and Cheng, Y. 2021b. Intention Reasoning Network for Multi-Domain End-to-end Task-Oriented Dialogue. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 2273–2285.
- Mekala, D.; Wolfe, J.; and Roy, S. 2023. ZEROTOP: Zero-Shot Task-Oriented Semantic Parsing using Large Language Models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 5792–5799.
- Qin, C.; Zhang, A.; Zhang, Z.; Chen, J.; Yasunaga, M.; and Yang, D. 2023. Is ChatGPT a General-Purpose Natural Language Processing Task Solver? In *Proceedings of*

- the 2023 Conference on Empirical Methods in Natural Language Processing, 1339–1384.
- Qin, L.; Che, W.; Li, Y.; Wen, H.; and Liu, T. 2019. A Stack-Propagation Framework with Token-Level Intent Detection for Spoken Language Understanding. 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)*, 2078–2087.
- Qin, L.; Liu, T.; Che, W.; Kang, B.; Zhao, S.; and Liu, T. 2021a. A Co-Interactive Transformer for Joint Slot Filling and Intent Detection. In *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 8193–8197.
- Qin, L.; Xie, T.; Che, W.; and Liu, T. 2021b. A Survey on Spoken Language Understanding: Recent Advances and New Frontiers. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021*, 4577–4584.
- Reimers, N.; and Gurevych, I. 2017. Reporting Score Distributions Makes a Difference: Performance Study of LSTM-networks for Sequence Tagging. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 338–348.
- Rongali, S.; Sridhar, M.; Khan, H.; Arkoudas, K.; Hamza, W.; and McCallum, A. 2023. Low-Resource Compositional Semantic Parsing with Concept Pretraining. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, 1410–1419.
- Rosenbaum, A.; Soltan, S.; Hamza, W.; Versley, Y.; and Boese, M. 2022. LINGUIST: Language Model Instruction Tuning to Generate Annotated Utterances for Intent Classification and Slot Tagging. In *Proceedings of the 29th International Conference on Computational Linguistics*, 218–241.
- Schick, T.; and Schütze, H. 2021. It’s Not Just Size That Matters: Small Language Models Are Also Few-Shot Learners. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2339–2352.
- Shen, Z.; Liu, Z.; Qin, J.; Savvides, M.; and Cheng, K.-T. 2021. Partial is better than all: revisiting fine-tuning strategy for few-shot learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 9594–9602.
- Snell, J.; Swersky, K.; and Zemel, R. S. 2017. Prototypical Networks for Few-shot Learning. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017*, 4077–4087.
- Sun, Q.; Liu, Y.; Chua, T.-S.; and Schiele, B. 2019. Meta-transfer learning for few-shot learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 403–412.
- Wang, Y.; Shen, Y.; and Jin, H. 2018. A Bi-Model Based RNN Semantic Frame Parsing Model for Intent Detection and Slot Filling. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, 309–314.
- Weld, H.; Huang, X.; Long, S.; Poon, J.; and Han, S. C. 2022. A survey of joint intent detection and slot filling models in natural language understanding. *ACM Computing Surveys*, 55(8): 1–38.
- Yuan, Y.; Pan, J.; Jia, X.; Liu, L.; and Peng, M. 2021. DCEN: A Decoupled Context Enhanced Network For Few-shot Slot Tagging. In *2021 International Joint Conference on Neural Networks (IJCNN)*, 1–7.
- Zheng, B.; Li, Z.; Wei, F.; Chen, Q.; Qin, L.; and Che, W. 2023. Hit-scir at mmnl-22: Consistency regularization for multilingual spoken language understanding. *arXiv preprint arXiv:2301.02010*.
- Zhu, S.; Cao, R.; Chen, L.; and Yu, K. 2020. Vector projection network for few-shot slot tagging in natural language understanding. *arXiv preprint arXiv:2009.09568*.