

Different Strokes for Different Folks: Investigating Appropriate Further Pre-training Approaches for Diverse Dialogue Tasks

Yao Qiu, Jinchao Zhang*, Jie Zhou

Pattern Recognition Center, WeChat AI, Tencent Inc, China
{yasinqiu, dayerzhang, withtomzhou}@tencent.com

Abstract

Loading models pre-trained on the large-scale corpus in the general domain and fine-tuning them on specific downstream tasks is gradually becoming a paradigm in Natural Language Processing. Previous investigations prove that introducing a **further pre-training** phase between pre-training and fine-tuning phases to adapt the model on the domain-specific unlabeled data can bring positive effects. However, most of these further pre-training works just keep running the conventional pre-training task, e.g., masked language model, which can be regarded as the domain adaptation to bridge the data distribution gap. After observing diverse downstream tasks, we suggest that different tasks may also need a further pre-training phase with appropriate training tasks to bridge the task formulation gap. To investigate this, we carry out a study for improving multiple task-oriented dialogue downstream tasks through designing various tasks at the further pre-training phase. The experiment shows that different downstream tasks prefer different further pre-training tasks, which have intrinsic correlation and most further pre-training tasks significantly improve certain target tasks rather than all. Our investigation indicates that it is of great importance and effectiveness to design appropriate further pre-training tasks modeling specific information that benefit downstream tasks. Besides, we present multiple constructive empirical conclusions for enhancing task-oriented dialogues.

1 Introduction

Pre-trained models, e.g., BERT (Devlin et al., 2019a), RoBERTa (Liu et al., 2019) and GPT2 (Radford et al., 2019), have been widely used in many NLP tasks. These models are pre-trained on the large-scale general text corpus, such as Wikipedia or books, with self-supervised train-

ing objectives. Fine-tuning these models on downstream tasks can achieve excellent performance.

Recently, Gururangan et al. (2020) proposed a domain-adaptive pre-training method, they further pre-training the RoBERTa on a large corpus of unlabeled domain-specific text, e.g., biomedical papers and computer science papers, before fine-tuning on downstream tasks and achieved strong performance. Besides, they proved that it is also helpful to continue pre-training on the task-specific text. Wu et al. (2020) applied this method to task-oriented dialogue and proposed a new self-supervised pre-training objective on dialogue corpus. Despite they achieved performance improvements, the improvements on different downstream tasks vary a lot, some tasks even obtain no improvement, which indicates that different downstream tasks may need different further pre-training tasks.

To investigate this issue, we carry out experiments in the area of task-oriented dialogue. We choose one popular pre-training language model, BERT (Devlin et al., 2019a) as our base model, and construct a large scale domain-specific dialogue corpus which consists of nine task-oriented datasets for further pre-training (Wu et al., 2020). We also select four core task-oriented dialogue tasks, intent recognition, dialogue action prediction, response selection, and dialog state tracking as the downstream tasks used in fine-tuning phase. We aim to explore the following questions: 1) In the area of task-oriented dialogue, can further pre-training using the masked language model improve the performance of all downstream tasks? 2) Do different further pre-training tasks have different effects on different downstream tasks? 3) Which factors affect whether a further pre-training task can achieve improvement on a certain downstream task? 4) Does combining different further pre-training tasks benefits more downstream tasks?

To answer these questions, we design five self-supervised pre-training tasks according to different

*Jinchao Zhang is the corresponding author.

characteristics of the downstream tasks. Specifically, we first use specially designed pre-training tasks to further pre-training BERT on the domain-specific corpus, obtaining multiple new pre-trained models, denoted as BERT’s variants. Then, we fine-tune these variants on all downstream tasks and observe the effect of different pre-training tasks on different downstream tasks. From experiment results, we figure out that: 1) Further pre-training with masked language model does not achieve improvements for all downstream tasks, it is necessary to design special further pre-training tasks according to the characteristics of dialogue data. 2) Different pre-training tasks do have different effects on different downstream tasks, and there is a need to design a specific pre-training task for a certain downstream task. 3) Model’s ability and structure are two key factors influencing effectiveness of the further pre-training on a certain downstream task. 4) Training two further pre-training tasks in a multi-task paradigm does not lead to incremental performance improvements on downstream tasks.

The main contribution of our work is to give a set of empirical principles about how to design effective further pre-training tasks for enhancing the task-oriented dialogue. The key points of the design are to make the model structures of the pre-training task and the downstream task similar and let the model learn the abilities required by downstream tasks in the pre-training phase while maintaining the masked language model’s training. We release the source code at the GitHub repo.¹

2 Background

2.1 Pre-trained Models

Two-stage Training. Large pre-training models, such as BERT (Devlin et al., 2019b), RoBERTa (Liu et al., 2019), GPT2 (Radford et al., 2019), XLNet (Yang et al., 2019) T5 (Raffel et al., 2020), are trained on massive general domain text with self-supervised training objectives, like masked language model (Devlin et al., 2019b) and permutation language model (Yang et al., 2019). These models learned strong and general word representations, fine-tuning these pre-trained models on downstream tasks is proved to be effective.

Three-stage Training. Recently, further pre-training large language models on domain-specific corpus before fine-tuning on downstream tasks has

become a popular and effective paradigm. Gururangan et al. (2020) proposed domain-adaptive pre-training and task-adaptive pre-training methods, and they proved that such a second phase of pre-training in a specific domain leads to performance gains. Wu et al. (2020) applied the second phase pre-training on task-oriented dialogue, in addition to masked language modeling objective, they also proposed a new self-supervised objective according to the characteristic of dialogue corpus. However, the performance improvement gained from their proposed methods varies a lot across different downstream tasks, which indicates different downstream tasks may need different further pre-training tasks rather than the conventional one, such as MLM.

2.2 Task-oriented Dialogue

A task-oriented dialog system aims to assist the user in completing certain tasks in one or several specific domains, such as restaurant booking, weather query, and flight booking. The entire system usually consists of four modules, including natural language understanding (NLU), dialog state tracking (DST), dialog policy, and natural language generation (NLG). In this work, we focus on four core tasks:

- **Intent recognition:** The model is required to predict the intent type given the user utterance. Intent type is a high-level classification label of the user utterance, such as *Query* and *Inform*, which indicates the function of the user utterance.
- **Dialog act prediction:** The model is required to predict the dialog act (e.g., *Question*, *Statement*) of the next response given the whole dialog history.
- **Response selection:** The model is required to select the proper response from many candidate responses given the whole dialog history. The negative candidate responses are randomly sampled.
- **Dialog state tracking:** The dialog state tracker estimates the user’s goal in each time step by taking the entire dialog context as input. The dialog state at time t can be regarded as an abstracted representation of the previous turns until t .

¹<https://github.com/FFYYang/DSDF.git>

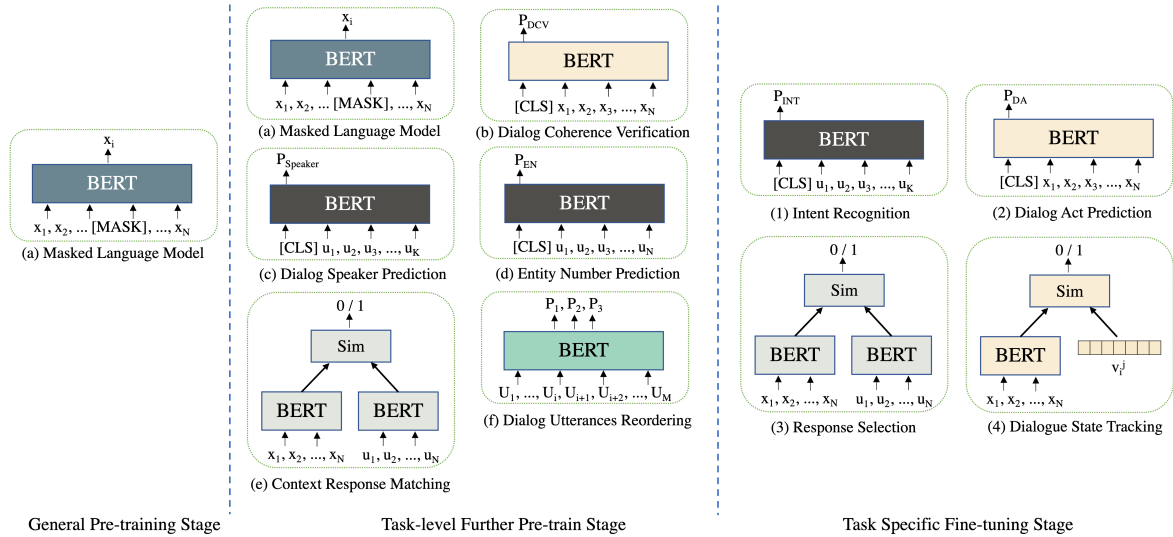


Figure 1: The pipeline of our training process, started from a pre-trained BERT, then further pre-train BERT on dialogue corpus using multiple proposed tasks, the resulting models are fine-tuned on downstream tasks.

3 Approaches

In this section, we firstly present the three-stage training framework, then introduce five expressly designed further pre-training tasks and downstream tasks. At last, we present some heuristic analysis on the relations between the tasks in the further pre-training and the fine-tuning stage.

3.1 Three-stage Training for the Task-oriented Dialogue

We design a three-stage training framework, includes the **general pre-training stage**, **task-level further pre-training stage** and the **task-specific fine-tuning stage** for enhancing the various tasks in the task-oriented dialogue, as shown in Figure 1. The general pre-training stage aims to learn general word representation. The task-level further pre-training stage contains multiple optional tasks trained on the un-labeled dialogue corpus. The task-specific fine-tuning stage is to train specific models for solving the downstream task such as intent recognition. To be emphasized, our further pre-training stage attempts to bridge the task-level gap between the pre-training and the fine-tuning stage rather than the domain adaptation on the data-level (Gururangan et al., 2020).

3.2 Task-level Further Pre-training

To enhance the task-oriented dialogue through bridging the task-level gap between pre-training and fine-tuning, we design multiple optional tasks

which can be trained on dialogue corpus without any human annotation.

Dialog Speaker Prediction (DSP). The model is required to predict the speaker (user or agent) of a given utterance. The model can learn a better single utterance representation from this task. The input of the model is a single utterance $U = u_1, u_2, \dots, u_K$, where K is the utterance length. The model outputs a binary result indicating the speaker is a user or agent.

$$P_{\text{Speaker}} = \text{Softmax}(W_{\text{DSP}} \cdot F_{\text{bert}}(U)) \quad (1)$$

Where F_{bert} is the forward function of BERT, we use its [CLS] representations as the utterance representation. W_{DSP} is a trainable linear mapping matrix. The task is trained with the cross-entropy loss.

Context Response Matching (CRM). Given a dialog context, the model selects the proper response from many randomly sampled candidate responses. This task is in the same as the response contrastive loss proposed by Wu et al. (2020). The model can learn the dialogue coherence information from this task.

Dialogue Coherence Verification (DCV). This task asks the model to predict whether a dialog is coherent. The incoherent dialog is constructed by randomly replacing some utterances in the dialog. The model can learn a better multi-turn dialog

representation from this task. We first randomly select half of the training data and randomly replace some utterances in the dialogue to destroy their coherence. The input of the model is the whole dialog, concatenating all utterances together, denoted as $S = x_1, x_2, \dots, x_N$, where N is the sequence length. The model outputs a binary prediction.

$$P_{DCV} = \text{Softmax}(W_{DCV} \cdot F_{bert}(S)) \quad (2)$$

Where F_{bert} is the forward function of BERT, we use its [CLS] representations as the dialog representation. W_{DCV} is a trainable linear mapping matrix. The task is trained with the cross-entropy loss.

Entity Number Prediction (ENP). The model predicts the number of entities contained in an utterance. Entities are extracted using the open-source tool *stanza*². The model can learn a better single utterance representation and entity information. This task is formulated as a multi-class classification problem, where we input a single utterance U , and the model predicts one single class indicating how many entities are contained in the utterance.

$$P_{EN} = \text{Softmax}(W_{ENP} \cdot F_{bert}(U)) \quad (3)$$

Where F_{bert} is the forward function of BERT. W_{ENP} is a trainable linear mapping matrix. The task is trained with the cross-entropy loss.

Dialog Utterances Reordering (DUR). The model reorders a group of shuffled utterances. The model can learn dialog coherence information from this task. The input of the model is the whole dialog, but some utterances' positions are shuffled. We put special tokens [USR] and [SYS] at the front of each utterance indicating it is spoken by a user or agent. We concatenate all utterances together, feed them to BERT, and take the representation of [USR] and [SYS] as the representation of each utterance. The model predicts the correct relative position of the shuffled utterances. For example, utterances U_i, U_{i+1}, U_{i+2} are shuffled, we first use BERT to get their representations R_i, R_{i+1}, R_{i+2} , and use a FFN and softmax function to get the probability distribution of their relative positions, $y_p = [y_p^1, y_p^2, y_p^3]$. The loss is calculated as:

$$L_{DUR} = \text{Avg}(-\text{SUM}(y_t * \log(y_p + eps))) \quad (4)$$

Where y_p is the correct probability distribution of these utterances relative positions, for example, suppose the correct relative position is $[2, 1, 3]$, then $y_p = \text{Softmax}([2, 1, 3])$.

²<https://github.com/stanfordnlp/stanza>

3.3 Task Specific Fine-tuning

After further pre-training, we fine-tuning our models on each downstream task individually. These downstream tasks are modeled in different forms following (Wu et al., 2020).

Intent Recognition (INT). The task is a multi-class classification problem, the input of the model is a single utterance U and model predicts one single intent type.

$$P_{INT} = \text{Softmax}(W_{INT} \cdot F_{bert}(U)), \quad (5)$$

The task is trained with the cross-entropy loss.

Dialogue Act Prediction (DA). The task is modeled as a multi-label classification problem, since a system response may contain multiple dialogue acts. The model's input is the whole dialogue history S , and the model outputs a binary prediction for each possible dialogue act.

$$P_{DA} = \text{Sigmoid}(W_{DA} \cdot F_{bert}(S)), \quad (6)$$

It is trained with the binary cross-entropy loss.

Response Selection (RS). The model selects the most proper system response from multiple candidates. We utilize a siamese structure and compute similarity scores between dialogue history H and a candidate response R_i .

$$s_i = \text{Sim}(F_{bert}(H), F_{bert}(R_i)), \quad (7)$$

Where s_i is the cosine similarity. The negative candidates are randomly sampled from the corpus.

Dialogue State Tracking (DST) is modeled as a multi-class classification task based on a predefined ontology. The model's input is the whole dialogue history S , and the model predicts the value of the slot for each (domain, slot) pair. We define v_i^j as the i -th value for j -th (domain, slot) pair, we use BERT to obtain its representation which is fixed during the whole fine-tuning stage.

$$S_i^j = \text{Sim}(G_j(F_{bert}(X)), F_{bert}(v_i^j)), \quad (8)$$

Where Sim is the cosine similarity function, and S^j is the probability distribution of the j -th (domain, slot) pair over its possible values. G_j is the slot projection layer of the j -th (domain, slot) pair, and the number of layers $|G|$ is equal to the number of (domain, slot) pairs. The task is trained with the cross-entropy loss summed over all the pairs.

	Abilities			Entity Information	Structures			Rank Loss
	Single Turn Representation	Multi Turn Representation	Coherence		Single Turn Classifier	Multi Turn Classifier	Siamese Model	
INT	※			※	※			
DA		※				※		
RS			※				※	
DST		※		※		※		
DSP	※				※			
CRM			※				※	
DCV		※				※		
ENP	※			※	※			
DUR			※				※	

Table 1: This table shows the comparison between further pre-training and downstream tasks from the ability and structure perspective. The above four tasks are downstream tasks, the below five tasks are further pre-training tasks. ※ indicates the task has the ability or belongs to the model structure.

All of the proposed tasks are trained with the masked language model in a multi-task paradigm. In addition, these tasks are optional, we focus on investigating their relations with each downstream task.

3.4 Heuristic Analysis on Task Relations between Further Pre-training and Fine-tuning

We analyse the task relations from two perspectives: model *ability* and *structure*. Ability refers to the information or knowledge the model learns, for example, the ability of single turn representation, the knowledge about the entity. Structure refers to the model’s network structure and its objective function, for example, the siamese structure and list-wise ranking loss function. The details are shown in Table 1. We suggest that if a further pre-training task learns similar abilities or has a similar model structure the with the downstream task, then the further pre-training will be more effective for fine-tuning.

4 Experimental Setup

4.1 Dialogue Datasets for Further Pre-training

Following Wu et al. (2020), we construct the further pre-training dataset by combining nine different multi-turn goal-oriented datasets (Frames (El Asri et al., 2017), MetaLWOZ (Lee et al., 2019), WOZ (Mrkšić et al., 2017), CamRest676 (Wen et al., 2017), MSR-E2E (Li et al., 2018), MWOZ (Budzianowski et al., 2018), Schema (Rastogi et al., 2020), SMD (Eric et al., 2017) and Taskmaster (Byrne et al., 2019)). In total, there are 100,707 dialogues containing 1,388,152 utterances over 60 domains.

4.2 Evaluation Datasets

We select four datasets, OOS, DSTC2, GSIM, and MWOZ, for downstream evaluation. Details of each evaluation dataset are discussed below.

OOS. (Larson et al., 2019) It contains 151 intent types across ten domains, including 150 in-scope and one out-of-scope intent.

DSTC2. (Henderson et al., 2014) It is a machine-human task-oriented dataset, We follow Wu et al. (2020) to map the original dialogue act labels to universal dialogue acts, resulting in 19 acts.

MWOZ. (Budzianowski et al., 2018) It is a popular benchmark for task-oriented dialogues. It has 30 (domain, slot) pairs across seven different domains. We use the revised version MWOZ 2.1.

GSIM. (Shah et al., 2018) It is a human-rewrote task-oriented dataset. Following Wu et al. (2020) we combine movie and restaurant domains into one single corpus, and map its dialogue act labels to universal dialogue acts, resulting in 13 acts.

4.3 Training Setting

For further pre-training, we set the learning rate equal to $5e-5$, batch size to 32, and maximum sequence length to 512. For fine-tuning, we set the learning rate to $5e-5$ (except dialog state tracking task, which is $3e-5$). We use the batch size that maximizes the GPU usage. We train our models using the Adam optimizer. Models are early-stopped using the loss of a validation set. We train each downstream task three times with different seeds. We use 4 NVIDIA V100 GPUs for further pre-training and one for fine-tuning. Our code is based on Transformers³

³<https://github.com/huggingface/transformers>

	DA						INT			
	MWOZ		DSTC2		GSIM		Acc	Acc	Acc	Recall
	$f1_{micro}$	$f1_{macro}$	$f1_{micro}$	$f1_{macro}$	$f1_{micro}$	$f1_{macro}$	(all)	(in)	(out)	(out)
$BERT^2$	90.90	81.31	91.16	37.67	99.07	45.49	84.96	95.20	88.70	38.87
MLM^3	91.51	79.77	87.76	36.99	99.35	45.70	85.10	95.90	88.30	36.70

	RS						DST	
	MWOZ		DSTC2		GSIM		acc_{joint}	acc_{slot}
	$R_{100}@1$	$R_{100}@3$	$R_{100}@1$	$R_{100}@3$	$R_{100}@1$	$R_{100}@3$		
$BERT^2$	47.39	74.46	48.33	63.30	19.21	40.02	47.96	96.86
MLM^3	59.55	82.63	54.25	68.91	35.71	58.73	48.72	96.94

Table 2: The results of the experiment investigating the effect of data-level further pre-train. $BERT^2$ does not contain a further pre-training stage, MLM^3 utilizes masked language model to further pre-train BERT on unlabeled dialogue corpus before fine-tuning. MLM^3 does not surpass $BERT^2$ in all metrics and datasets.

	DA						INT			
	MWOZ		DSTC2		GSIM		Acc	Acc	Acc	Recall
	$f1_{micro}$	$f1_{macro}$	$f1_{micro}$	$f1_{macro}$	$f1_{micro}$	$f1_{macro}$	(all)	(in)	(out)	(out)
MLM^3	91.51	79.77	87.76	36.99	99.35	45.70	85.10	95.90	88.30	36.70
DSP	91.46	80.87	93.26	43.54	99.35	45.72	85.78	95.52	89.28	41.97
CRM	91.52	79.87	92.59	40.48	99.31	45.68	85.02	94.91	88.92	40.53
DCV	91.67	80.47	93.84	42.44	99.41	45.77	85.00	95.76	88.32	36.60

	RS						DST	
	MWOZ		DSTC2		GSIM		acc_{joint}	acc_{slot}
	$R_{100}@1$	$R_{100}@3$	$R_{100}@1$	$R_{100}@3$	$R_{100}@1$	$R_{100}@3$		
MLM^3	59.55	82.63	54.25	68.91	35.71	58.73	48.72	96.94
DSP	61.79	84.04	53.06	66.94	36.59	59.37	49.02	96.96
CRM	64.27	86.15	58.12	71.86	41.94	66.47	48.89	96.97
DCV	60.14	83.51	54.57	68.45	29.35	53.12	51.18	97.15

Table 3: Results of the experiment investigating the effect of further pre-training tasks. These three tasks outperform MLM on most metrics, and different further pre-training tasks benefit to different downstream tasks.

5 Results and Discussion

In this section, we collect experimental results and analyse the effects of different further pre-training tasks on different downstream tasks.

5.1 Effect of the Data-level Further Pre-training

To investigate the effect of the data-level further pre-training, we firstly further pre-train BERT with masked language model (MLM) on the un-labeled task-oriented dialogue corpus, then fine-tune the model on each downstream task, we denote this experiment as MLM^3 . In contrast, we also directly fine-tune BERT on downstream tasks, the experiment is denoted as $BERT^2$. The experiment results are shown in Table 2, MLM^3 outperforms $BERT^2$ on response selection and dialog state tracking task, as for dialog act prediction and intent recognition task, MLM^3 does not surpass $BERT^2$ in all metrics and datasets. From the result, we can conclude that *further pre-training using MLM objective does not bring performance improvement for all downstream tasks*, so it is necessary to design special further pre-training tasks ac-

ording to the characteristics of the dialogue data.

5.2 Effect of Various Further Pre-training Tasks

To investigate the effects of different further pre-training tasks on different downstream tasks, we compare three further pre-training tasks, dialogue speaker prediction (DSP), context response matching (CRM), and dialogue coherence verification (DCV), each of which has special characteristics. From the experiment results shown in Table 3, DSP, CRM, and DCV are better than MLM^3 on most of the metrics, this indicates the effectiveness of these auxiliary pre-training tasks. In addition, we also observe that different pre-training tasks are more beneficial to different downstream tasks, for example, DSP is more beneficial to downstream intent recognition task than others, CRM is mainly beneficial to response selection, DCV is beneficial to dialogue act prediction and dialogue state tracking. We can conclude that *different pre-training tasks do have different effects on different downstream tasks*, so there is a need to design a specific pre-training task for a downstream task.

	DA						INT			
	MWOZ		DSTC2		GSIM		Acc	Acc	Acc	Recall
	$f1_{micro}$	$f1_{macro}$	$f1_{micro}$	$f1_{macro}$	$f1_{micro}$	$f1_{macro}$	(all)	(in)	(out)	(out)
DSP	91.46	80.87	93.26	43.54	99.35	45.72	85.78	95.52	89.28	41.97
ENP	91.38	80.31	92.47	40.38	99.57	45.85	86.27	95.67	89.60	44.00

	RS						DST	
	MWOZ		DSTC2		GSIM		acc_{joint}	acc_{slot}
	$R_{100}@1$	$R_{100}@3$	$R_{100}@1$	$R_{100}@3$	$R_{100}@1$	$R_{100}@3$		
DSP	61.79	84.04	53.06	66.94	36.59	59.37	49.02	96.96
ENP	59.36	82.96	54.97	68.83	33.33	56.17	49.65	97.07

Table 4: The experiments investigating the effect of the ability. ENP are designed to learn more abilities which are needed by downstream INT and DST task, and its performance on these two tasks is completely higher than DSP.

	DA						INT			
	MWOZ		DSTC2		GSIM		Acc	Acc	Acc	Recall
	$f1_{micro}$	$f1_{macro}$	$f1_{micro}$	$f1_{macro}$	$f1_{micro}$	$f1_{macro}$	(all)	(in)	(out)	(out)
DUR	91.56	80.16	94.99	45.36	99.48	45.79	85.80	95.67	89.11	41.37
CRM	91.52	79.87	92.59	40.48	99.31	45.68	85.02	94.91	88.92	40.53

	RS						DST	
	MWOZ		DSTC2		GSIM		acc_{joint}	acc_{slot}
	$R_{100}@1$	$R_{100}@3$	$R_{100}@1$	$R_{100}@3$	$R_{100}@1$	$R_{100}@3$		
DUR	59.87	83.28	55.48	69.27	29.37	53.76	49.36	97.10
CRM	64.27	86.15	58.12	71.86	41.94	66.47	48.89	96.97

Table 5: The experiments investigating the effect of the structure. The model structure of CRM is more similar to downstream task RS, and its performance on this task is completely higher than DUR.

5.3 Empirical Analysis on Task Relations between Further Pre-training and Fine-tuning

In session 3.4, we provide a heuristic analysis on task relations between further pre-training and fine-tuning. We suggest ability and structure are two key factors that influence the effectiveness of further pre-training to fine-tuning.

We define **nice pair** meaning that a further pre-training task is effective to a downstream task. From Table 3 we can find DSP is more beneficial for INT, CRM is for RS, while DCV is for DA and DST. So there are four nice pairs, (DSP, INT), (CRM, RS), (DCV, DA), and (DCV, DST). Among these four nice pairs, we can find there is one thing in common, *the further pre-training task and downstream task in the same nice pair almost share the same ability and the model structure*. Take (CRM, RS) pair as an example, both CRM and RS mainly learn the ability of dialogue coherence and belong to the siamese structure.

To further investigate the effect of the ability, we compare dialogue speaker prediction (DSP) and entity number prediction (ENP). Their structures are the same, that is, single turn classification, but the abilities they learn are different, DSP mainly learns the ability of single turn representation, while ENP also learns entity information. Experiment results

are shown in Table 4, ENP outperforms DSP on intent recognition and dialogue state tracking tasks across all metrics because these two tasks also need the ability about entity information. This indicates ability is important for further pre-training.

To further investigate the effect of the structure, we compare context response matching (CRM) and dialogue utterances reordering (DUR). Both of them mainly learn the ability about dialogue coherence, but their structures are different. Results in Table 5 show that CRM surpasses DUR on the response selection task because the CRM model is a siamese structure which is the same as the response selection task. This indicates the structure is also a crucial factor for the effectiveness of further pre-training.

5.4 Effect of Combining Further Pre-training Tasks

We jointly further pre-train entity number prediction (ENP) and context response matching (CRM) in the multi-task paradigm, the experiment is denoted as Joint. We expect the joint model can combine the advantages of ENP and CRM, and bring improvement on downstream INT, RS, and DST. The results in Table 6 are not fully consistent with our expectation, specifically, on intent recognition, joint model’s performance drops significantly, on the other three downstream tasks, joint model’s

	DA						INT			
	MWOZ		DSTC2		GSIM		Acc	Acc	Acc	Recall
	$f1_{micro}$	$f1_{macro}$	$f1_{micro}$	$f1_{macro}$	$f1_{micro}$	$f1_{macro}$	(all)	(in)	(out)	(out)
ENP	91.38	80.31	92.47	40.38	99.57	45.85	86.27	95.67	89.60	44.00
CRM	91.52	79.87	92.59	40.48	99.31	45.68	85.02	94.91	88.92	40.53
Joint	91.65	80.55	92.42	40.13	99.46	45.78	84.31	94.96	88.26	36.40

	RS						DST	
	MWOZ		DSTC2		GSIM		acc_{joint}	acc_{slot}
	$R_{100}@1$	$R_{100}@3$	$R_{100}@1$	$R_{100}@3$	$R_{100}@1$	$R_{100}@3$		
ENP	59.36	82.96	54.97	68.83	33.33	56.17	49.65	97.07
CRM	64.27	86.15	58.12	71.86	41.94	66.47	48.89	96.97
Joint	62.33	85.55	57.98	72.30	39.48	64.42	49.23	97.04

Table 6: Results of the experiment investigating the effect of combining multiple further pre-training tasks. The joint model does not improve all the downstream tasks that ENP and CRM beneficial to.

	DA						INT			
	MWOZ		DSTC2		GSIM		Acc	Acc	Acc	Recall
	$f1_{micro}$	$f1_{macro}$	$f1_{micro}$	$f1_{macro}$	$f1_{micro}$	$f1_{macro}$	(all)	(in)	(out)	(out)
ENP	91.38	80.31	92.47	40.38	99.57	45.85	86.27	95.67	89.60	44.00
w.o. mlm	91.06	80.27	90.75	38.16	99.34	45.72	85.79	95.57	89.18	41.77

	RS						DST	
	MWOZ		DSTC2		GSIM		acc_{joint}	acc_{slot}
	$R_{100}@1$	$R_{100}@3$	$R_{100}@1$	$R_{100}@3$	$R_{100}@1$	$R_{100}@3$		
ENP	59.36	82.96	54.97	68.83	33.33	56.17	49.65	97.07
w.o. mlm	58.09	82.29	37.12	52.98	16.92	38.70	50.30	97.21

Table 7: The experiment investigating the effect of combining data-level and task-level further pre-training. Removing masked language model objective can cause performance drop on almost all the downstream tasks.

performance is between ENP and CRM.

5.5 Effect of Combining Data-level and Task-level Further Pre-training

In the former experiments, each proposed further pre-trained task is trained with masked language model (MLM), we suppose MLM is for data-level adaptation while the proposed task is for task-level adaptation. In this section, we investigate the effect of MLM by removing MLM objective from further pre-training stage, this experiment is denoted as w.o. mlm. Experiment results are shown in Table 7. Removing MLM leads to performance drop across almost all downstream tasks, indicating MLM is important to further pre-training stage.

5.6 Experiment Summary

Through all the experiments, we can conclude that: In the area of task-oriented dialogue, 1) Masked language model alone is not enough for further pre-training, but it still plays an important role for enhancing fine-tuning. And there is a need to design special further pre-training tasks according to the characteristics of dialogue data. 2) Different pre-training tasks do have different effects on different downstream tasks, and it is necessary to design a specific pre-training task for a specific

downstream task. 3) Ability and structure of a further pre-training task are key factors influencing the performance of fine-tuning on a downstream task. 4) Training two further pre-training tasks in the multi-task paradigm does not lead to incremental performance improvement.

From these conclusions, we can obtain multiple empirical principles to design further pre-training tasks: 1) The ability learned by pre-training task should be similar to the ability required by the downstream task. 2) the modeling structure should also be similar, 3) the masked language model training objective should be kept.

6 Conclusion

In this work, we study how to make further pre-training more effective to downstream tasks in the area of the task-oriented dialog. Firstly, we notice that further pre-training using MLM objective does not improve all downstream tasks, then we designed multiple pre-training tasks for dialog data, finding that different pre-training tasks benefit different downstream tasks. Further, we observe that ability and structure are key factors influencing whether a pre-training task is helpful to a downstream task. These finds can be used as empirical principles to design pre-training tasks.

Acknowledgments

We would like to thank all the reviewers for their insightful and valuable comments and suggestions.

References

- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gasic. 2018. Multiwoz-a large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5016–5026.
- Bill Byrne, Karthik Krishnamoorthi, Chinnadhurai Sankar, Arvind Neelakantan, Ben Goodrich, Daniel Duckworth, Semih Yavuz, Amit Dubey, Kyu-Young Kim, and Andy Cedilnik. 2019. [Taskmaster-1: Toward a realistic and diverse dialog dataset](#). 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 2019, Hong Kong, China, November 3-7, 2019*, pages 4515–4524. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019a. 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.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019b. [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, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- Layla El Asri, Hannes Schulz, Shikhar Kr Sarma, Jeremie Zumer, Justin Harris, Emery Fine, Rahul Mehrotra, and Kaheer Suleman. 2017. Frames: a corpus for adding memory to goal-oriented dialogue systems. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 207–219.
- Mihail Eric, Lakshmi Krishnan, Francois Charette, and Christopher D Manning. 2017. Key-value retrieval networks for task-oriented dialogue. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 37–49.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A Smith. 2020. Don’t stop pretraining: Adapt language models to domains and tasks. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8342–8360.
- Matthew Henderson, Blaise Thomson, and Jason D Williams. 2014. The second dialog state tracking challenge. In *Proceedings of the 15th annual meeting of the special interest group on discourse and dialogue (SIGDIAL)*, pages 263–272.
- Stefan Larson, Anish Mahendran, Joseph J Peper, Christopher Clarke, Andrew Lee, Parker Hill, Jonathan K Kummerfeld, Kevin Leach, Michael A Laurenzano, Lingjia Tang, et al. 2019. An evaluation dataset for intent classification and out-of-scope prediction. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1311–1316.
- S Lee, H Schulz, A Atkinson, J Gao, K Suleman, L El Asri, M Adada, M Huang, S Sharma, W Tay, et al. 2019. Multi-domain task-completion dialog challenge. *Dialog System Technology Challenges*, 8.
- Xiujun Li, Sarah Panda, Jingjing Liu, and Jianfeng Gao. 2018. [Microsoft dialogue challenge: Building end-to-end task-completion dialogue systems](#). *CoRR*, abs/1807.11125.
- 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](#). *CoRR*, abs/1907.11692.
- Nikola Mrkšić, Diarmuid Ó Séaghdha, Tsung-Hsien Wen, Blaise Thomson, and Steve Young. 2017. Neural belief tracker: Data-driven dialogue state tracking. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1777–1788.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *J. Mach. Learn. Res.*, 21:140:1–140:67.
- Abhinav Rastogi, Xiaoxue Zang, Srinivas Sunkara, Raghav Gupta, and Pranav Khaitan. 2020. Towards scalable multi-domain conversational agents: The schema-guided dialogue dataset. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8689–8696.

- Pararth Shah, Dilek Hakkani-Tur, Bing Liu, and Gokhan Tur. 2018. Bootstrapping a neural conversational agent with dialogue self-play, crowdsourcing and on-line reinforcement learning. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 3 (Industry Papers)*, pages 41–51.
- Tsung-Hsien Wen, David Vandyke, Nikola Mrkšić, Milica Gasic, Lina M Rojas Barahona, Pei-Hao Su, Stefan Ultes, and Steve Young. 2017. A network-based end-to-end trainable task-oriented dialogue system. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 438–449.
- Chien-Sheng Wu, Steven CH Hoi, Richard Socher, and Caiming Xiong. 2020. Tod-bert: Pre-trained natural language understanding for task-oriented dialogue. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 917–929.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2019. [Xlnet: Generalized autoregressive pretraining for language understanding](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 5754–5764.