

GENSF: Simultaneous Adaptation of Generative Pre-trained Models and Slot Filling

Shikib Mehri and Maxine Eskenazi

Language Technologies Institute, Carnegie Mellon University

{amehri, max}@cs.cmu.edu

Abstract

In transfer learning, it is imperative to achieve strong alignment between a pre-trained model and a downstream task. Prior work has done this by proposing task-specific pre-training objectives, which sacrifices the inherent scalability of the transfer learning paradigm. We instead achieve strong alignment by simultaneously modifying both the pre-trained model and the formulation of the downstream task, which is more efficient and preserves the scalability of transfer learning. We present GENSF (**Generative Slot Filling**), which leverages a generative pre-trained open-domain dialog model for slot filling. GENSF (1) adapts the pre-trained model by incorporating inductive biases about the task and (2) adapts the downstream task by reformulating slot filling to better leverage the pre-trained model’s capabilities. GENSF achieves state-of-the-art results on two slot filling datasets with strong gains in few-shot and zero-shot settings. We achieve a **9 F₁ score** improvement in zero-shot slot filling. This highlights the value of strong alignment between the pre-trained model and the downstream task.

1 Introduction

The advent of pre-trained language models (Devlin et al., 2019; Radford et al., 2019) has transformed natural language processing. The dominant paradigm has shifted away from designing task-specific architectures towards transfer learning. Fine-tuning pre-trained models on downstream datasets achieves strong performance on a variety of natural language understanding tasks (Wang et al., 2018). Generally, prior to fine-tuning, the pre-trained models are adapted to the specifics of the downstream task through minor architectural modifications (e.g., adding a classification layer) (Chen et al., 2019; Mehri et al., 2020). By avoiding major task-specific changes to the models, it

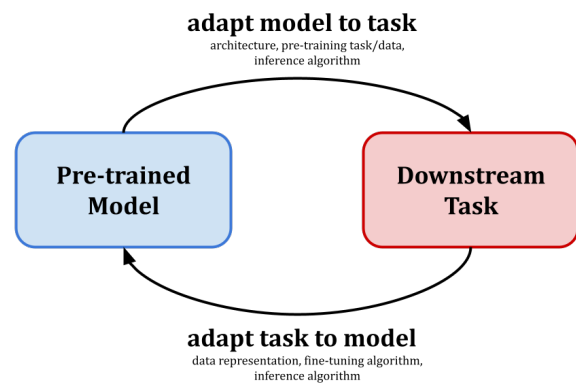


Figure 1: To achieve a stronger alignment, both the downstream task and the pre-trained models must be adapted. The downstream task can be adapted with knowledge of the properties and capabilities of the pre-trained models. Likewise, the pre-trained model can be adapted with knowledge of the downstream task/data.

is assumed that the underlying pre-trained models possess a degree of generality that allows transfer to a variety of tasks. We posit that this assumption is flawed. Consequently this paper demonstrates the importance of incorporating *inductive biases* that achieve *stronger alignment* between the pre-trained model and the downstream task.

Recent work has validated the idea that stronger alignment between pre-training and the downstream task results in improved performance. Rather than fine-tuning off-the-shelf models, it is more effective to first understand the downstream task and adapt the model’s architecture, pre-training and inference algorithm accordingly. Adapting pre-trained models in this manner is equivalent to **incorporating inductive biases about the downstream task**. For example, pre-training on open-domain dialog data results improves performance on downstream dialog tasks (Henderson et al., 2019; Mehri et al., 2020). Designing task-specific pre-training objectives has

yielded strong results in extractive question answering (Glass et al., 2019), paraphrase and translation (Lewis et al., 2020) and slot filling (Henderson and Vulić, 2020). This body of work attains stronger alignment by significantly modifying the pre-trained model through task-specific pre-training. This necessitates a new pre-trained model for every downstream task, and therefore relinquishes the inherent scalability of the transfer learning paradigm. Instead, we achieve stronger alignment by *simultaneously adapting* both the pre-trained model and the downstream task, such that both contain inductive biases about one another.

The downstream task can be adapted to achieve stronger alignment with the capabilities of the pre-trained model. To effectively leverage pre-trained models, it is important to first understand the properties and capabilities of the model derived from the model architecture, the pre-training data and task. Then the downstream task can be adapted to be better aligned with the model. Adapting the task to the model is equivalent to **incorporating inductive biases about the pre-trained model** into the downstream task. For example, given a pre-trained model that was trained with a ranking objective, it is likely to be more effective if the downstream fine-tuning and inference algorithms are modified to rank rather than to classify. By simultaneously adapting both the downstream task and the pre-trained model, we intend to achieve stronger alignment without sacrificing the inherent scalability of the transfer learning paradigm (i.e., avoiding task-specific pre-trained models).

We address the task of slot filling, a natural language understanding task with the goal of identifying values for pre-defined attributes (slots) in a natural language utterance. We leverage a DialoGPT (Zhang et al., 2020), a generative language model, pre-trained on open-domain dialog data. To achieve strong alignment between the slot filling task and DialoGPT, we (1) reformulate slot filling as a natural language response generation task, and (2) augment the DialoGPT architecture with a copy-mechanism, constrained decoding and a post-processing heuristic. The resulting model, GENSF (**Generative Slot Filling**), is shown to achieve state-of-the-art results on two slot filling datasets. GENSF achieves the strongest performance gains in few-shot and zero-shot settings, highlighting the importance of stronger alignment in the absence of abundant data. Our code is open-

sourced and can be found at https://github.com/shikib/generative_slot_filling.

2 Related Work

Slot filling is the task of identifying values for pre-defined attributes, or slots, in a natural language utterance (Tur and De Mori, 2011). Slot filling is a vital natural language understanding component of task-oriented dialog systems (Young, 2002, 2010). A variety of architectures have been explored for the task of slot filling, including CNNs (Vu, 2016), deep LSTMs (Yao et al., 2014), RNNs with external memory (Peng et al., 2015), encoder labeler LSTMs (Kurata et al., 2016) and joint pointer and attention seq2seq networks (Zhao and Feng, 2018). With the introduction of large-scale pre-trained language models (Devlin et al., 2019; Radford et al., 2019), strong slot filling results have been achieved with simple architectures (Chen et al., 2019).

Several approaches have been proposed for zero-shot slot filling. Bapna et al. (2017) leverage slot names and descriptions to align slots across domains. Shah et al. (2019) leverage examples for zero-shot slot filling. Liu et al. (2020) achieve strong results in zero-shot slot filling with a coarse-to-fine approach in combination with template regularization. We use the Coach+TR model (Liu et al., 2020) as a baseline in our zero-shot experiments.

Working on the hypothesis that pre-trained language models, such as BERT (Devlin et al., 2019), do not effectively capture the intricacies of dialog, recent work has attempted to mitigate this issue. Coope et al. (2020) use ConveRT (Henderson et al., 2019), a lightweight model pre-trained on dialog data, in combination with CNN and conditional random field (CRF) to outperform BERT. Mehri et al. (2020) achieves similar results with ConvBERT, a model that further pre-trains BERT on open-domain dialog data. Recently, Henderson and Vulić (2020) introduces a ‘*pairwise cloze*’ pre-training objective that uses open-domain dialog data to specifically pre-train for the task of slot filling. The resulting ConVEx model achieves significant improvements, particularly in few-shot settings. A common theme in recent work is achieving better alignment between the pre-trained models and the downstream task, either by pre-training on data that is closer to the domain of the downstream task (i.e., dialog data) (Henderson et al., 2019; Mehri et al., 2020) or by designing custom pre-training objectives that better model the down-

stream task (Henderson and Vulić, 2020). Our proposed approach shares the goal of achieving better alignment, but we simultaneously adapt both the pre-trained model and the downstream task, with the goal of leveraging a generative pre-trained dialog model, DialoGPT, for slot filling.

3 Methods

In order to effectively leverage a pre-trained generative dialog model, DialoGPT (Zhang et al., 2020), for the task of slot-filling, we introduce the GENSF model which achieves stronger alignment between the downstream task and the pre-trained model, by simultaneously adapting the task to the model and the model to the task. This paper first describes how the slot filling task is reformulated as a natural language response generation task to be better aligned with the DialoGPT model. Next, it describes several modifications to the DialoGPT architecture and inference algorithm that act as inductive biases for the slot filling task.

3.1 Slot Filling as Response Generation

Given an utterance $u = \{w_1, w_2, \dots, w_n\}$, a set of possible slot keys $s = \{s_1, s_2, \dots, s_k\}$, and a list of slots requested by the system $r = \{r_1, r_2, \dots, r_m\}$ (where $r_i \in s$ and $m \geq 0$), the task of slot filling is to assign a value to a subset of the slot keys. Concretely, for a given slot key s_i , the output will either be NULL or a contiguous span of words from the utterance: $s_i = \{w_i, \dots, w_{i+j}\}$.

In response generation, given a dialog context consisting of a sequence of utterances: $c = \{x_1, x_2, \dots, x_n\}$ wherein each utterance x_i is a sequence of words, the task is to generate a valid response $y = \{w_1, w_2, \dots, w_m\}$.

Many tasks can be represented as an *input to output* mapping (Raffel et al., 2019; Hosseini-Asl et al., 2020; Peng et al., 2020), making sequence-to-sequence a universal formulation. Trivially, slot filling can be represented as a sequence-to-sequence task by setting the context to be the concatenation of the utterance and the requested slots: $c = \{u, r\}$ and the target response to be the slot mappings $y = \{(s_1, w_{i:j}), (s_2, \text{NULL}), \dots, (s_k, (w_{j:n}))\}$. However, this does not leverage the natural language capabilities of pre-trained dialog models. While this trivial formulation may suffice with sufficient training, it will under-perform in few-shot and zero-shot settings. To this end, this paper presents a reformulation of slot filling that better aligns with

the natural language capabilities of DialoGPT.

We hypothesize that to some degree, large-scale dialog pre-training can result in a model implicitly learning to fill slots. For example, given the slot key ‘time’, such a model should understand what *time* is and should be able to generate a valid time (e.g., ‘4:15 pm’). An effective task formulation can leverage these implicitly learned slot filling capabilities. An off-the-shelf pre-trained model is likely to only be capable of filling generic slots (e.g., time, date, price, etc.). But by reformulating slot filling in a manner that is better aligned with the pre-training task, it should be easier for the model to adapt to novel slot keys.

Concretely, given a slot filling input (u, r) and a particular slot key s_i , we construct a natural language dialog context using a template-based approach: $c = \text{‘What is the } \{f(r)\} \text{? [eos] } \{u\} \text{ [eos] Ok, the } \{f(s_i)\} \text{ is’}$. Here, f denotes a manually constructed function that maps slot keys to a natural language phrase (e.g., *first_name: first name, departure_location: leaving from*). Given the constructed dialog context, the model is tasked with completing the partial response (i.e., *Ok, the } $\{f(s_i)\}$ is*) by auto-regressively generating the slot value. During training the model would be tasked with generating either the slot value or the phrase *not provided*. With this natural language reformulation, the slot filling task is being adapted to better leverage the capabilities of the pre-trained DialoGPT model. As this achieves better alignment between the pre-trained model and the downstream task, it should be more effective for zero-shot and few-slot filling. To better illustrate the conversion of the slot-filling input (utterance u and request slots r), several examples are shown in Table 1.

3.2 DialoGPT for Slot Filling

In order to adapt the pre-trained DialoGPT model to the slot filling task, we augment the architecture and modify the inference algorithm. These adaptations are motivated by the observation that if the slot value is provided, it will always be a contiguous span of tokens from the utterance. As such, the generative model can only produce: (1) ‘not provided’ if the slot does not appear in the utterance, (2) the end of sentence token, and (3) tokens from the input utterance.

A copy-mechanism is incorporated into the DialoGPT architecture to allow the model to explicitly generate tokens from the input utterance.

Utterance	Requested Slots	Slot Key	Natural Language Context
We will require an outside table to seat 9 people on August 23rd	None	date	We will require an outside table to seat 9 people on August 23rd [EOS] Ok, the date is
Laurice Hoisl	first_name, last_name	first_name	What is the first name, last name? [EOS] Laurice Hoisl [EOS] Ok, the first name is
My party will be 9 people. My name is Nancie Waltemeyer and the time is 7pm	None	people	My party will be 9 people. My name is Nancie Waltemeyer and the time is 7pm [EOS] Ok, the number of people is

Table 1: Examples of slot filling inputs reformulated as natural language dialog contexts

Given a context $c = \{x_1, x_2, \dots, x_n\}$, through its self-attention layers, the model will produce a hidden state representation for each token, $h = \{h_1, h_2, \dots, h_n\}$. A probability distribution over the vocabulary is then obtained by passing h_n through a classification layer:

$$P_{vocab} = \text{softmax}(Wh_n + b) \quad (1)$$

To explicitly generate tokens from the input, h_n is used to attend to $h_{1:n}$ to produce a probability distribution over $x_{1:n}$. The process for computing the probability for a specific word, $P_{copy}(w)$ is as follows:

$$\alpha = \text{softmax}(h_n^T h_{1:n}) \quad (2)$$

$$P_{copy}(w) = \sum_{i:x_i=w} \alpha_i \quad (3)$$

These two probability distributions are combined through a weighted sum. The weight assigned to each of the distributions is predicted using h_n :

$$p_{copy} = \sigma(W_{copy}h_n + b_{copy}) \quad (4)$$

The final probability distribution is therefore:

$$P_{final} = (1 - p_{copy})P_{vocab} + p_{copy}P_{copy} \quad (5)$$

The copy-mechanism requires training, as it introduces new weights (w_{copy} , b_{copy}) and the off-the-shelf DialoGPT model does not necessarily produce attention weights, α , that can be used to create an output probability distribution. As such, to attain strong zero-shot performance we must also modify the inference algorithm to account for

the aforementioned observation. This is done using both constrained decoding and a post-processing heuristic.

Constrained decoding is a modification of greedy decoding wherein the argmax sampling is modified to only generate (1) words that appear in the input utterance, (2) the end of sentence token and (3) the phrase ‘not provided’.

The slot values may consist of terms that the model has not frequently observed during pre-training (e.g., names, times). As such, because the DialoGPT model leverages a subword vocabulary, some subword tokens may be dropped during generation and therefore the slot values may be generated with typos (e.g., ‘Mocer’ vs ‘Mocher’). A simple post-processing heuristic is applied to mitigate this problem. If the slot value produced by the model is not present in the utterance, the Levenshtein distance to every contiguous span of tokens in the utterance is computed. If the best edit distance is within a certain threshold ($0.3 \times \text{len}(y)$), the corresponding span is returned as the slot value.

Through these modifications, the DialoGPT model is adapted to reflect the properties of the slot filling task. The copy-mechanism, constrained decoding and post-processing mechanism serve as an inductive bias to enable the pre-trained model to be better adapted for the downstream slot filling task.

4 Experiments

Experiments are performed to empirically validate the hypothesis that simultaneously adapting the downstream task and the pre-trained model results in stronger alignment and improved performance. We present experiments on two datasets and as-

Fraction	Span-Convert	Span-BERT	ConVEx	GenSF
1 (8198)	95.8	93.1	96.0	96.1
$1/2$ (4099)	94.1	91.4	94.1	94.3
$1/4$ (2049)	91.2	88.0	92.6	93.2
$1/8$ (1024)	88.5	85.3	90.6	91.8
$1/16$ (512)	81.1	76.6	86.4	89.7
$1/32$ (256)	63.8	53.6	81.8	82.1
$1/64$ (128)	57.6	42.2	76.0	76.1
$1/128$ (64)	40.5	30.6	71.7	72.2

Table 2: F_1 scores across all slots for the evaluation on the RESTAURANTS-8K test data with varying proportions of the training set. Numbers in brackets denote the training set sizes. The best scores (statistically significant by t-test to $p < 0.05$) are shown in boldface.

sess GENSF in full-data, few-shot and zero-shot settings. An ablation study is performed to characterize the source of the performance gains and demonstrate the importance of simultaneous adaptation.

4.1 Datasets

Experiments are carried out on RESTAURANTS-8K (Coope et al., 2020) and the DSTC8 datasets (Rastogi et al., 2020). RESTAURANTS-8K consists of 8,198 utterances from a commercial restaurant booking system and includes 5 slots (date, time, people, first name, last name). The DSTC8 datasets span four different domains (buses, events, homes, rental cars) for a total of 5,569 utterances with slot annotations extracted by Coope et al. (2020).

In both datasets, the value for a particular slot is always a contiguous span of the utterance. Some utterances consist of a set of slots requested by the system prior to the user utterance. This allows an otherwise ambiguous utterance like ‘four’ to be interpreted as either ‘four people’ or ‘four o’clock’.

4.2 Experimental Setup

We use the pre-processing and evaluation scripts provided by the DialoGLUE benchmark (Mehri et al., 2020). We follow the setup of Coope et al. (2020) and Henderson and Vulić (2020), wherein a validation set is not used and the experiments are therefore performed with fixed hyperparameters. Throughout all the experiments, the medium version of DialoGPT (Zhang et al., 2020) is used. We use the AdamW optimizer (Loshchilov and Hutter, 2017) with a learning rate of $5e-5$. On RESTAURANTS-8K, the models are trained for 10 epochs in the full-data setting, 20 epochs in the few-shot settings and 40 epochs in the extreme few-shot settings ($1/32$ - $1/128$; or less than 256 training examples). On the DSTC8 datasets, the models are

trained for 20 epochs in the full-data setting and 40 epochs in the few-shot setting.

The models are evaluated on the full test set, regardless of the amount of training data, using macro-averaged F_1 score (Coope et al., 2020).

To facilitate reproducibility, the code and the trained models will be released upon publication.

4.3 Slot Filling Results

Throughout the experiments we compare to several models from prior work. Span-Convert (Coope et al., 2020) and Span-BERT train a CNN and a CRF on top of contextual subword embeddings produced by Convert (Henderson et al., 2020) and BERT (Devlin et al., 2019), respectively. ConVEx (Henderson and Vulić, 2020) devises a *pairwise cloze* pre-training objective specifically for slot-filling. This task-specific pre-training objective is an example of significantly adapting the pre-trained model to the downstream task. In contrast to ConVEx, GENSF achieves strong alignment between the pre-trained model and the downstream task by simultaneously adapting both the task and the model. As such, GENSF does not need a task-specific pre-trained model and is inherently more scalable. The ConVEx pre-training takes 8 hours to train on 12 GPUs, while GENSF takes less than four hours to train on a single GTX 1080TI.

As shown in Table 2, GENSF achieves state-of-the-art results across all experimental settings on the RESTAURANTS-8K dataset. In the full-data setting, GENSF slightly outperforms ConVEx. Though the performance gain is small, this result signifies that our model can leverage an abundance of data. The value of strong alignment between the downstream task and the pre-trained model is better exemplified in the few-shot settings. Especially in the extreme few-shot settings (i.e., $1/32$ -

	Setting	Span-ConveRT	Span-BERT	ConVEx	GENSF
Buses_1	Full-Data (1133)	93.5	93.3	96.0	98.1
	Few-Shot (283)	84.0	77.8	86.7	90.5
Events_1	Full-Data (1498)	92.7	84.3	91.7	94.7
	Few-Shot (374)	82.2	78.6	87.2	91.2
Homes_1	Full-Data (2064)	94.8	96.3	98.3	96.9
	Few-Shot (516)	95.4	95.1	94.5	93.7
RentalCars_1	Full-Data (874)	94.0	92.8	92.0	93.5
	Few-Shot (218)	83.0	81.4	87.4	86.7

Table 3: F_1 scores across all slots for evaluation on the DSTC8 single-domain datasets in the full-data and few-shot settings. Numbers in brackets denote training set sizes. The best scores (statistically significant by t-test, to $p < 0.05$) are shown in boldface.

$1/128$ of the training set), GENSF strongly outperforms Span-ConveRT and Span-BERT, with greater than 30 F_1 score improvements. The few-shot performance of both ConVEx and GENSF in these few-shot settings underlies the value of effectively aligning the pre-trained model and the downstream task. However, GENSF achieves this alignment by simultaneously incorporating inductive biases about the model into the task rather than designing a complex pre-training objective. By incorporating inductive biases into both the task and the model, the approach outlined in this paper does not require task-specific pre-trained models and therefore preserves the inherent generality of the transfer learning paradigm. Furthermore, GENSF attains moderate improvements over ConVEx, especially in the few-shot settings, with a 3 F_1 score improvement in the $1/16$ th setting.

The results on the DSTC8 single-domain datasets is shown in Table 3. Here, we evaluate on both full-data and few-shot (25% of the training data) settings. On average, GENSF achieves strong performance improvements over prior work. In the full-data settings the best performance is observed on the buses and events domains, where GENSF achieves a 2.1 and 3.0 F_1 score improvement over ConVEx, respectively. In the few-shot settings, GENSF achieves a 4.0 F_1 score improvement over ConVEx on these domains and a 6.5 and 9.0 point improvement over Span-ConveRT. These strong improvements, over both Span-ConveRT and ConVEx, highlight the value of strong alignment between the pre-trained model and the downstream task, particularly in the few-shot experiments.

GENSF moderately underperforms on the homes and rental cars domains. On the homes domain, GENSF outperforms Span-ConveRT and Span-

BERT but scores 1.4 points below ConVEx. Similarly, on the rental cars domain, GENSF outperforms ConVEx and Span-BERT, but is 0.5 points below Span-ConveRT. Though GENSF is still competitive in these domains, these results nonetheless highlight a weakness of our model. Our use of a generative pre-trained dialog model, specifically DialoGPT (Zhang et al., 2020), was motivated by the hypothesis that such models can implicitly learn to identify certain slots through response generation pre-training. This hypothesis is empirically validated through improved performance on RESTAURANTS-8K and the buses/events domains of DSTC8. GENSF relies on the pre-trained model having an implicit understanding of the slots. This implicit understanding results in strong performance on slots like ‘time’ or ‘first name’, since such terms are likely to have been observed during pre-training. However, this is not the case for all slots and GENSF can underperform on slots that are ambiguous, ill-defined or are unlikely to have been observed during open-domain dialog pre-training. The homes domain consists of the slot, ‘area’, which has several definitions and is therefore challenging for the pre-trained model to understand and detect. The rental cars domain contains the slots ‘pickup date’ and ‘dropoff date’. While the DialoGPT model has learned to detect a ‘date’, the distinction between these two slots is more nuanced and therefore may cause some amount of confusion. As such, while GENSF is competitive in these domains and is only outperformed by one of the three models, these domains demonstrate that there are limitations at present to leveraging a generative pre-trained model. However, it is possible that by further adapting the downstream task to the pre-trained model, for example by renaming

these slots (e.g., ‘*area*’ may be renamed to ‘*city*’), the performance drops may be mitigated.

Overall, GENSF achieves impressive performance gains in both full-data and few-shot settings, underlying the value of achieving strong alignment between the pre-trained model and the downstream task. Furthermore, GENSF achieves this alignment by simultaneously adapting both the task and the model and without sacrificing the inherent scalability of the transfer learning paradigm or necessitating task-specific pre-training. In the RESTAURANTS-8K and the single-domain DSTC8 datasets, GenSF achieves state-of-the-art results and outperforms prior work. In few-shot settings, we achieve a 30 F_1 score improvement over SpanBERT and Span-ConveRT. On average, GenSF moderately outperforms ConVEx, with > 2.0 F_1 score improvements in the few-shot settings on RESTAURANT-8K, and both the full data and few-shot settings on two of the DSTC8 datasets. These experiments empirically validate (1) the importance of aligning the pre-trained model and the downstream task by simultaneously incorporating inductive biases into both the task and the model and (2) that through response generation pre-training, dialog models have implicitly learned to detect certain slots, which can be leveraged by effectively adapting the downstream task.

4.4 Zero-shot Slot Filling

For zero-shot slot filling, we must have strong alignment between the pre-trained model and the downstream task. Since the model is not fine-tuned on the task, it is necessary to effectively align the formulation of the downstream task to the capabilities of the model. As such, zero-shot experiments validate our proposed reformulation of slot filling as natural language response generation.

For these experiments, we compare to the published results of ConVEx (Henderson and Vulić, 2020). Furthermore, we run a Coach+TR model (Liu et al., 2020) on the RESTAURANT-8K dataset. Note that while ConVEx and GENSF have only been trained on open-domain dialog, Coach+TR trains on adjacent task-oriented domains (i.e., SNIPS), meaning that the zero-shot performance is higher on slots that are domain agnostic.

The experiments used the RESTAURANTS-8K dataset with the GENSF model. The copy-mechanism is removed from the model, as it adds additional weights to the model and therefore re-

Slot	Metric	Coach+TR	ConVEx	GenSF
First Name	P	1.7	2.3	13.7
	R	4.1	20.1	36.1
	F_1	2.5	4.1	19.8
Last Name	P	0	1.9	10.6
	R	0	16.2	19.7
	F_1	0	3.4	13.8
Date	P	10.2	2.2	10.7
	R	34.8	10.1	15.3
	F_1	15.7	3.6	12.6
Time	P	47.4	5.6	27.5
	R	27.9	23.6	46.9
	F_1	35.1	9.1	34.7
People	P	0	3.8	14.5
	R	0	13.9	18.9
	F_1	0	6.0	16.4
Average	F_1	10.7	5.2	19.5

Table 4: Zero-shot slot filling results on RESTAURANTS-8K. All models are evaluated on the test set without any training on the dataset.

quires training. However, the constrained decoding and the post-processing heuristic of GENSF, allow us to enforce that the slot values will always be a contiguous span from the input utterance. Table 4 demonstrates that GENSF significantly outperforms prior work on zero-shot slot filling with a **14 F_1 score improvement** over ConVEx and a **9 F_1 score improvement** over Coach+TR. These results further validate the hypothesis that pre-trained dialog models have implicitly learned to detect slots and that this ability can be leveraged through the proposed task reformulation.

Most noteworthy is the performance on the ‘*first name*’ and ‘*last name*’ slots. This suggests that, to some degree, DialoGPT (Zhang et al., 2020) can disambiguate between a first name and a last name when provided simultaneously (e.g., ‘*my name is Lakesha Mocher*’). It should be noted that the macro-averaged F_1 score used to evaluate the models considers a slot value to be incorrect unless it exactly predicts the ground-truth slot value. In many cases, the GENSF model produces appropriate slot values that differ from the ground-truth, e.g., ‘*wednesday*’ instead of ‘*next wednesday*’. It is possible that by incorporating additional inductive biases about the specific formulation of the slot values (e.g., slots should have maximal information) into the inference algorithm, the zero-shot performance can be further increased.

Model	Full-Data	Few-Shot ($1/16$)	Zero-Shot
GenSF	96.1	89.7	19.5
Removing Model Adaptation			
– Copy-mechanism	95.6	87.8	19.5
– Constrained Decoding	95.4	89.5	0.5
– Post-processing	96.1	89.7	18.1
– All model adaptation	95.4	87.8	0.5
Removing Task Adaptation			
– Natural Language Slot Names	95.3	86.6	12.2
– Natural Language Templates	94.8	88.5	0.0
– All Natural Language	95.5	88.9	0.0
Removing All Adaptation			
– All Adaptation	95.8	89.2	0.0

Table 5: Ablation experiments. We remove (1) adaptations to the model, (2) adaptations to the downstream task and (3) all adaptations proposed in this paper. The experiments are carried out on the full-data, few-shot ($1/16$ th of the training set) and zero-shot settings of RESTAURANTS-8K.

GENSF is shown to strongly outperform prior work on zero-shot slot filling. This impressive performance validates the proposed approach of simultaneously adapting both the downstream task and the pre-trained model. Furthermore, zero-shot performance also confirms the hypothesis that pre-trained response generation models have implicitly learned to understand and detect slots, thereby highlighting the potential of leveraging generative pre-trained models for language understanding tasks. Future work should explore mechanisms for reformulating other downstream tasks (e.g., intent prediction, dialog state tracking) in order to leverage generative pre-trained models. Furthermore, it is possible that these zero-shot results could be further improved through two-stage pre-training (e.g., further pre-train with the ‘*pairwise cloze*’ task).

4.5 Ablation

GENSF has been shown to outperform prior work in full-data, few-shot and zero-shot settings. To determine the source of the improvements, we perform an ablation study. The ablation experiments remove the adaptations used in GENSF and evaluate on RESTAURANTS-8K across full-data, few-shot ($1/16$ of the training set) and zero-shot settings. Removing all the ablation, is equivalent to training a DialoGPT model from scratch on the task, similar to the approach proposed by Madotto (2020).

As shown in Table 5, the various adaptations are vital to the strong performance of GENSF. Of the model adaptations, only the copy-mechanism is necessary in the full-data setting, since the model

effectively learns to copy tokens from the input utterance and therefore does not need constrained decoding and post-processing. However, constrained decoding is necessary for the zero-shot settings, as the zero-shot model does not leverage a copy-mechanism. Task adaptation, especially the use of natural language templates, is shown to be important across all of the experimental settings. This highlights the importance of formulating the downstream task in a manner that can effectively leverage the capabilities of the pre-trained models.

The results of the ablation study further validate this paper’s primary hypothesis. Pre-trained models work better for downstream tasks, when the task and the model are effectively aligned. As shown in the results of the ablation study, removing this adaptation results in a performance decrease.

5 Conclusion

This paper simultaneously adapts both the task and the pre-trained model in order to achieve strong alignment between a generative pre-trained dialog model and the downstream slot filling task. The resulting GENSF model achieves state-of-the-art results on two slot filling datasets, with particularly strong gains in few-shot and zero-shot settings. The empirical results underlie the importance of incorporating inductive bias into both the task and the pre-trained model. While this paper demonstrates the value of simultaneous adaptation for the task of slot filling, a similar paradigm could potentially be extended to alternate tasks. Future work should (1) explore improved mechanism for

achieving stronger alignment between the task and the model, (2) extend the simultaneous adaptation strategy to other problems and (3) explore the use of pre-trained generative models for language understanding tasks.

References

- Ankur Bapna, Gokhan Tur, Dilek Hakkani-Tur, and Larry Heck. 2017. Towards zero-shot frame semantic parsing for domain scaling. *arXiv preprint arXiv:1707.02363*.
- Qian Chen, Zhu Zhuo, and Wen Wang. 2019. Bert for joint intent classification and slot filling. *arXiv preprint arXiv:1902.10909*.
- Samuel Coope, Tyler Farghly, Daniela Gerz, Ivan Vulić, and Matthew Henderson. 2020. [Span-ConveRT: Few-shot span extraction for dialog with pretrained conversational representations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 107–121, Online. Association for Computational Linguistics.
- 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.
- Michael Glass, Alfio Gliozzo, Rishav Chakravarti, Anthony Ferritto, Lin Pan, GP Bhargav, Dinesh Garg, and Avirup Sil. 2019. Span selection pre-training for question answering. *arXiv preprint arXiv:1909.04120*.
- Matthew Henderson, Inigo Casanueva, Nikola Mrkšić, Pei-Hao Su, Tsung-Hsien Wen, and Ivan Vulić. 2019. Convert: Efficient and accurate conversational representations from transformers. *arXiv preprint arXiv:1911.03688*.
- Matthew Henderson, Iñigo Casanueva, Nikola Mrkšić, Pei-Hao Su, Tsung-Hsien Wen, and Ivan Vulić. 2020. [ConveRT: Efficient and accurate conversational representations from transformers](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2161–2174, Online. Association for Computational Linguistics.
- Matthew Henderson and Ivan Vulić. 2020. Convex: Data-efficient and few-shot slot labeling. *arXiv preprint arXiv:2010.11791*.
- Ehsan Hosseini-Asl, Bryan McCann, Chien-Sheng Wu, Semih Yavuz, and Richard Socher. 2020. A simple language model for task-oriented dialogue. *arXiv preprint arXiv:2005.00796*.
- Gakuto Kurata, Bing Xiang, Bowen Zhou, and Mo Yu. 2016. Leveraging sentence-level information with encoder lstm for semantic slot filling. *arXiv preprint arXiv:1601.01530*.
- Mike Lewis, Marjan Ghazvininejad, Gargi Ghosh, Armen Aghajanyan, Sida Wang, and Luke Zettlemoyer. 2020. Pre-training via paraphrasing. *arXiv preprint arXiv:2006.15020*.
- Zihan Liu, Genta Indra Winata, Peng Xu, and Pascale Fung. 2020. Coach: A coarse-to-fine approach for cross-domain slot filling. *arXiv preprint arXiv:2004.11727*.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Andrea Madotto. 2020. Language models as few-shot learner for task-oriented dialogue systems. *arXiv preprint arXiv:2008.06239*.
- Shikib Mehri, Mihail Eric, and Dilek Hakkani-Tur. 2020. Dialoglue: A natural language understanding benchmark for task-oriented dialogue. *arXiv preprint arXiv:2009.13570*.
- Baolin Peng, Chunyuan Li, Jinchao Li, Shahin Shayan-deh, Lars Liden, and Jianfeng Gao. 2020. Soloist: Few-shot task-oriented dialog with a single pre-trained auto-regressive model. *arXiv preprint arXiv:2005.05298*.
- Baolin Peng, Kaisheng Yao, Li Jing, and Kam-Fai Wong. 2015. Recurrent neural networks with external memory for spoken language understanding. In *Natural Language Processing and Chinese Computing*, pages 25–35. Springer.
- 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. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*.
- Abhinav Rastogi, Xiaoxue Zang, Srinivas Sunkara, Raghav Gupta, and Pranav Khaitan. 2020. Schema-guided dialogue state tracking task at dstc8. *arXiv preprint arXiv:2002.01359*.
- Darsh J Shah, Raghav Gupta, Amir A Fayazi, and Dilek Hakkani-Tur. 2019. Robust zero-shot cross-domain slot filling with example values. *arXiv preprint arXiv:1906.06870*.
- Gokhan Tur and Renato De Mori. 2011. *Spoken language understanding: Systems for extracting semantic information from speech*. John Wiley & Sons.

- Ngoc Thang Vu. 2016. Sequential convolutional neural networks for slot filling in spoken language understanding. *arXiv preprint arXiv:1606.07783*.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2018. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*.
- Kaisheng Yao, Baolin Peng, Yu Zhang, Dong Yu, Geoffrey Zweig, and Yangyang Shi. 2014. Spoken language understanding using long short-term memory neural networks. In *2014 IEEE Spoken Language Technology Workshop (SLT)*, pages 189–194. IEEE.
- Steve Young. 2002. Talking to machines (statistically speaking). In *Seventh International Conference on Spoken Language Processing*.
- Steve Young. 2010. Still talking to machines (cognitively speaking). In *Eleventh Annual Conference of the International Speech Communication Association*.
- Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020. [DIALOGPT : Large-scale generative pre-training for conversational response generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 270–278, Online. Association for Computational Linguistics.
- Lin Zhao and Zhe Feng. 2018. Improving slot filling in spoken language understanding with joint pointer and attention. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 426–431.