

C2C-GenDA: Cluster-to-Cluster Generation for Data Augmentation of Slot Filling

Yutai Hou*, Sanyuan Chen*, Wanxiang Che†, Cheng Chen, Ting Liu

Research Center for Social Computing and Information Retrieval, Harbin Institute of Technology, China
{ythou, sychen, car, tliu}@ir.hit.edu.cn, 170400202@stu.hit.edu.cn

Abstract

Slot filling, a fundamental module of spoken language understanding, often suffers from insufficient quantity and diversity of training data. To remedy this, we propose a novel Cluster-to-Cluster generation framework for Data Augmentation (DA), named C2C-GenDA. It enlarges the training set by reconstructing existing utterances into alternative expressions while keeping semantic. Different from previous DA works that reconstruct utterances one by one independently, C2C-GenDA jointly encodes multiple existing utterances of the same semantics and simultaneously decodes multiple unseen expressions. Jointly generating multiple new utterances allows to consider the relations between generated instances and encourages diversity. Besides, encoding multiple existing utterances endows C2C with a wider view of existing expressions, helping to reduce generation that duplicates existing data. Experiments on ATIS and Snips datasets show that instances augmented by C2C-GenDA improve slot filling by 7.99 (11.9%↑) and 5.76 (13.6%↑) F-scores respectively, when there are only hundreds of training utterances. Code: <https://github.com/Sanyuan-Chen/C2C-DA>.

Introduction

Slot filling is a fundamental module of the Spoken Language Understanding (SLU) in the task-oriented dialogue system (Young et al. 2013). The “inputs” in Figure 1 shows examples of slot filling, where key entities within user utterances are tagged with slot labels. Due to the high cost of manual annotation and the rapidly changing nature of dialogue domain, slot filling often faces the lack of quantity and diversity of training data. Such insufficiency in training data poses serious challenges for slot-filling models to handle myriad ways in which users express their demands.

Data augmentation (DA) technique, which improves diversity and quantity of training data with synthetic instances, offers an appealing solution to the data scarcity problem of SLU. Success has been achieved with data augmentation on a wide range of problems, including computer vision (Krizhevsky, Sutskever, and Hinton 2012), speech recognition (Hannun et al. 2014), text classification (Zhang, Zhao,

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

*Equal contributions.

†Corresponding author.

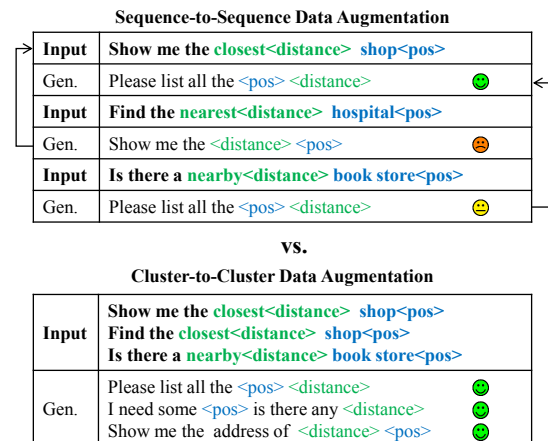


Figure 1: Examples of sequence-to-sequence data augmentation and cluster-to-cluster data augmentation. 🟢 denotes novel utterance. 🟡 denotes duplication to existing utterance. 🟠 denotes duplication to other generated utterances.

and LeCun 2015a), and question answering (Fader, Zettlemoyer, and Etzioni 2013).

For slot filling, state-of-the-art data augmentation works focus on generative methods (Shin, Yoo, and Lee 2019). One of their typical ideas is generating new utterances by reconstructing existing utterances into alternative expressions while keeping the semantics. Previous works learn a Sequence-to-Sequence (Seq2Seq) model to reconstruct each existing utterance one-by-one (Yoo 2020; Hou et al. 2018; Kurata, Xiang, and Zhou 2016). However, these methods tend to generate duplicated utterances, because they can only consider the expression variance between one input-output pair at a time. For example in Figure 1, each new utterance is only generated to be different from the corresponding input utterance, and thus often unconsciously duplicates other generated utterances (🟡) or other input utterances (🟡). Such duplication will hinder the effectiveness of data augmentation. We argue these defects can be easily avoided by breaking the shackles of current one-by-one augmentation paradigm and considering the extensive instance relations during generation.

In this paper, we propose a novel Cluster-to-Cluster Gen-

eration framework for Data Augmentation of slot filling, named C2C-GenDA. As shown in Figure 1, different from previous works that augment each utterance one-by-one independently, we jointly generate multiple new instances by reconstructing a cluster of existing utterances with the same semantics. Such cluster-to-cluster generation allows model to consider the duplication between generated utterances and aware of more existing expressions in original data. These advantages of C2C-GenDA remedy the aforementioned defects of Seq2Seq DA and help to improve generation diversity. To encourage diversity and quality of generation, we propose the Duplication-aware Attention and Diverse-Oriented Regularization mechanisms, both of which promote diverse decoding. To learn to generate diverse new utterances, we train the C2C-GenDA model with cluster-to-cluster ‘paraphrasing’ pairs, and introduce a Dispersed Cluster Pairing algorithm to extract these cluster pairs from existing data.

Experiments on ATIS and Snips datasets show that the proposed method significantly improves the performance of slot-filling systems. Case studies and analysis of augmented data also confirm that our method generates diverse utterances. Our contributions can be summarized as follow (1) We propose a novel Cluster-to-Cluster generation framework for data augmentation of slot filling, which can remedy the duplication problem of existing one-by-one generation methods. (2) We propose the Duplication-aware Attention and Diverse-Oriented Regularization mechanism to improve diversity of the augmented utterances. (3) We introduce a Dispersed Cluster Pairing algorithm to extract cluster-to-cluster ‘paraphrasing’ pairs for data augmentation model training.

Problem Description

In this paper, we study the data augmentation for slot filling task that maps utterances into semantic frames (slot type and slot value pairs). Slot filling is commonly treated as a sequence labeling problem, where *slot type* labels are assigned to contiguous sequences of words indicating these sequences are the corresponding *slot values*.

We specify the data augmentation (DA) for slot filling as exploiting existing training instances to generate new expressions for each semantic frame. Suppose existing slot filling training data is $D = \{(\mathbf{u}_i, \mathbf{s}_i)\}_{i=1}^N$. Given a semantic frame \mathbf{s}_j and the corresponding existing utterances $C = \{\mathbf{u} | (\mathbf{u}, \mathbf{s}) \in D \wedge \mathbf{s} = \mathbf{s}_j\}$, DA generates a set of new utterances $C' = \{\mathbf{u}'_i\}_{i=1}^{M'}$ with unseen expressions. Then DA constructs new training instances $D'_{\mathbf{s}_j} = \{(\mathbf{u}'_i, \mathbf{s}_j)\}_{i=1}^{M'}$ by associating new utterances with the semantic frame. Finally, DA takes the union of all new instances $D' = \bigcup_{\mathbf{s}_j} D'_{\mathbf{s}_j}$ as the additional data to reinforce the model training.

Proposed Framework

In this section, we present an overview of our data augmentation framework, and introduce the Cluster2Cluster generation model. Then, we discuss how to extract cluster-to-cluster paraphrasing data for generation model training.

Overview

Here, we introduce the overview of the proposed cluster-to-cluster data augmentation framework for slot filling. For each semantic frame, we use a Cluster2Cluster (C2C) model to generate new expressions from existing utterances. The input of our framework is a cluster of existing instances for a certain semantic frame, and the output is a cluster of generated new instances with unseen expressions.

Following Hou et al. (2018), we perform delexicalized generation. Specifically, both the inputs and outputs of C2C generation model are delexicalized utterances, where slot values tokens are replaced by slot label tokens. For the example in the Figure 1, C2C takes in “show me the <distance> <pos>” and reconstruct the expression as “please list all the <pos> <distance>”. The delexicalization focuses the model on generating diverse expressions rather than slot values and reduces the vocabulary size. Then after generation, we recover the delexicalized utterances by filling the slots with context-suitable slot values. Such delexicalization is important since it allow us to generate both the utterance and accurate slot annotations simultaneously.

To learn the ability of generating diverse and new expressions, we construct cluster-to-cluster paraphrasing pairs from original training data with the Dispersed Cluster Pairing algorithm, which simulates the data augmentation process of generating novel expressions from existing expressions for a specific semantic frame.

Cluster2Cluster Generation Model

Cluster2Cluster (C2C) model is a generation model that lies at the core of our C2C-GenDA framework and aims to reconstruct input utterances into alternative expressions while keeping semantic. As Figure 2 shown, the C2C model first encodes the input cluster of utterances $C = \{\mathbf{u}_i\}_{i=1}^M$ for a certain semantic frame, then jointly decodes a new cluster of utterances $C' = \{\mathbf{u}'_i\}_{i=1}^{M'}$ with different expressions, where M and M' are size of input and output cluster respectively.

To further encourage the diversity of the generated utterances, we propose two novel mechanisms: (1) **Duplication-aware Attention** that attends to the existing expressions to avoid duplicated generation for each decoding step. (2) **Diverse-Oriented Regularization** that guides the synchronized decoding of multiple utterances to improve the internal diversity of the generated cluster.

Cluster Encoder We jointly encode multiple input utterances by concatenating them, and representing the whole sequence with an L -layer transformer (Vaswani et al. 2017):¹

$$\begin{aligned} \mathbf{H}_0 &= \{e_{i,1}, e_{i,2}, \dots, e_{i,n_i}\}_i^M \\ \mathbf{A}_l &= \text{LN}(\mathbf{H}_{l-1} + \text{MultiHead}(\mathbf{H}_{l-1}, \mathbf{H}_{l-1})) \\ \mathbf{H}_l &= \text{LN}(\mathbf{A}_l + \text{FFN}(\mathbf{A}_l)) \\ \mathbf{R} &= \mathbf{H}_L \end{aligned}$$

where \mathbf{R} is the final representations of the input tokens, $e_{i,j}$ is the embedding of j th token in the i th input utterance, \mathbf{H}_0 is the package of all input token embeddings, and \mathbf{H}_l

¹We separate input utterances with special tokens <SEP>.

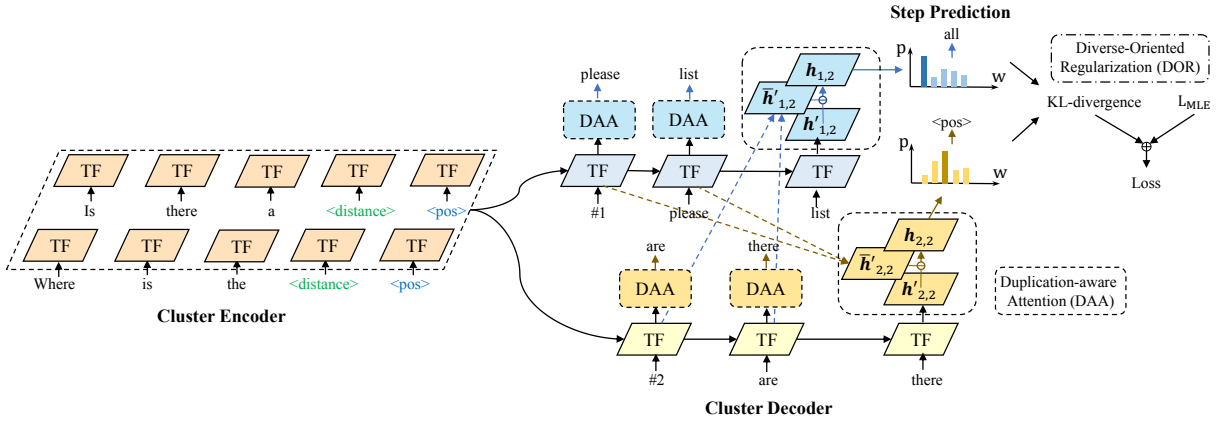


Figure 2: Cluster2Cluster generation model

is the outputs of l th layer. LN is the layer normalization. MultiHead(\mathbf{Q}, \mathbf{V}) is the multi-head self-attention function operating on vector packages of queries \mathbf{Q} , values \mathbf{V} (also used as keys). FFN is a position-wise feed-forward network.

Cluster Decoder with Duplication-aware Attention To reduce duplication and encourage diversity, we propose a cluster decoder with the Duplication-aware Attention (DAA) mechanisms. It decodes each new utterance while being aware of the existing expressions in both the input cluster and other generated utterances.

Intuitively, we decode the r th target utterance \mathbf{u}'_r depending on both input cluster C and other output utterances $C' \setminus \{\mathbf{u}'_r\}$. We also incorporate the diversity rank token $\#r$ (Hou et al. 2018) as generation conditions to encourage diversity and distinguish different output utterances. Details of the diverse rank in C2C will be introduced in a later section. Then the C2C model is formalized as:

$$p(C' | C) = p(\mathbf{u}'_1, \dots, \mathbf{u}'_{M'} | C) = \prod_r p(\mathbf{u}'_r | C, \#r, C' \setminus \{\mathbf{u}'_r\}).$$

However, it is unrealistic to decode a target utterance depending on all the other target utterances, because we jointly decode all the target utterances and the generation of other target utterances has not finished. Therefore, we approximate the dependence between target utterances and depend the decoding on already generated tokens of all the target utterances. For each step, we simultaneously decode one token for all the target utterances which depends on all the previously decoded tokens $\{\mathbf{u}'_{i,1:t-1}\}_i^{M'}$:

$$p(C' | C) = \prod_t \prod_r p(u'_{r,t} | C, \#r, \{\mathbf{u}'_{i,1:t-1}\}_i^{M'}),$$

where T is the number of decoding steps.

We calculate the decoding possibility for the t th step of r th utterance as $p_{r,t} = \text{Softmax}(\text{MLP}(\mathbf{h}_{r,t}))$, where $\mathbf{h}_{r,t}$ is a hidden state that combines feature representations of C , $\#r$ and $\{\mathbf{u}'_{i,1:t-1}\}_i^{M'}$.

Here, we obtain the hidden state $\mathbf{h}_{r,t}$ with DAA which contains two terms: $\mathbf{h}'_{r,t}$ and $\bar{\mathbf{h}}'_{r,t}$. The first term $\mathbf{h}'_{r,t}$ mainly

records the information of what token should be generated. To achieve this, $\mathbf{h}'_{r,t}$ encodes previously decoded tokens of current utterance and semantic information from the input cluster. Since $\mathbf{h}'_{r,t}$ encodes existing expressions in the input cluster, it also allows to reduce generation duplicated to existing expressions. For the r th target utterance, we compute $\mathbf{h}'_{r,t} = \mathbf{R}'_{r,t}$ with an L -layer transformer as decoder:

$$\begin{aligned} \mathbf{H}'_0 &= \{\#r, e'_{r,1}, e'_{r,2}, \dots, e'_{r,t-1}\} \\ \mathbf{A}'_l &= \text{LN}(\mathbf{H}'_{l-1} + \text{MultiHead}(\mathbf{H}'_{l-1}, \mathbf{H}'_{l-1} \cup \mathbf{H}_{l-1})) \\ \mathbf{H}'_l &= \text{LN}(\mathbf{A}'_l + \text{FFN}(\mathbf{A}'_l)) \\ \mathbf{R}'_r &= \mathbf{H}'_L \end{aligned}$$

where \mathbf{R}'_r is a package of hidden states for all t decoding steps. \mathbf{H}'_0 is the package of all decoded token embeddings and \mathbf{H}'_l is the outputs of l th decoding layer. \mathbf{H}_l is the input cluster representation from l th encoding layer.

The second term $\bar{\mathbf{h}}'_{r,t}$ mainly records the duplicated expressions that should not be generated, it encodes expressions generated by other target utterances as $\bar{\mathbf{h}}'_{r,t} = \text{MultiHead}(\mathbf{h}'_{r,t}, \{\mathbf{h}'_{i,1:t-1}\}_{i \neq r}^{M'})$.

Finally, the hidden-state for decoding is $\mathbf{h}_{r,t} = \mathbf{h}'_{r,t} - \lambda \cdot \bar{\mathbf{h}}'_{r,t}$, where λ is a balance factor. Subtraction makes $\mathbf{h}_{r,t}$ different for each target utterances, and $-\bar{\mathbf{h}}'_{r,t}$ can implicitly punish decoding of commonly shared words.

Model Training with Diverse-Oriented Regularization

We train the C2C model with a Diverse-Oriented Regularization (DOR) to encourage internal diversity within the generated utterance cluster.

To achieve this, we propose to enlarge the distance between distributions of utterances in the output cluster. However, the distribution of an utterance is hard to estimate during the decoding process. Thus, we approximately enlarge two utterances' distribution by encouraging the divergence of token distributions. As shown in Figure 2, we train the model to enlarge the Kullback-Leibler Divergence (KL) between decoding distribution of different output utterances at each step. Formally, we define the distance between two output utterances \mathbf{u}'_{r_i} and \mathbf{u}'_{r_j} as:

$$\text{Dist}(\mathbf{u}'_{r_i}, \mathbf{u}'_{r_j}) = \sum_t \text{KL}(p_{r_i,t} || p_{r_j,t}),$$

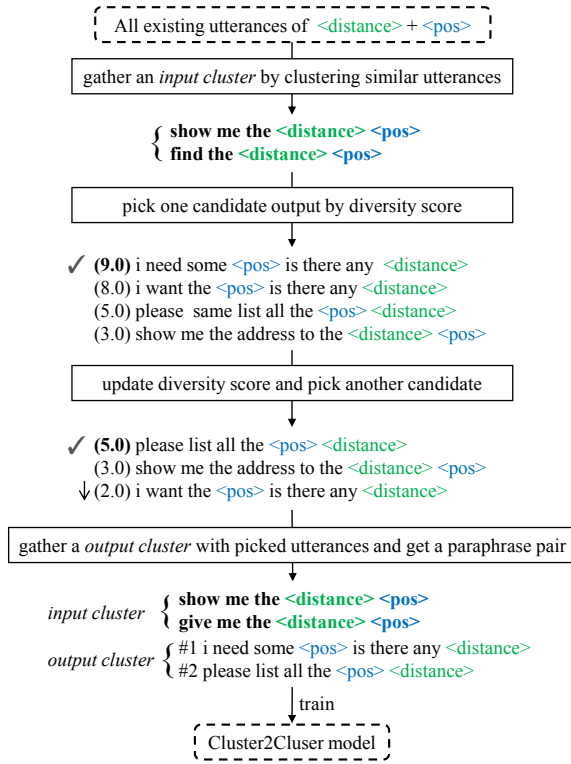


Figure 3: Construct training instance for our Cluster2Cluster model with Dispersed Cluster Pairing.

where $p_{r,t}$ denotes token distribution of r_{th} output utterance at t_{th} decoding step. Then we define Diverse-Oriented Regularization of generation as:

$$R_{\text{diverse}} = -\sum_{r_i, r_j (r_i \neq r_j)} \text{Dist}(\mathbf{u}'_{r_i}, \mathbf{u}'_{r_j}).$$

Overall, we train C2C model to minimize:

$$\text{Loss} = -\log(p(C' | C)) + \gamma \cdot R_{\text{diverse}},$$

where γ is a balancing factor.

Generation Pre-training When training data is insufficient, the data augmentation model itself is often poorly trained due to limited expression in the training data. To remedy this, we initialize the transformer encoder/decoder with pre-trained language model GPT-2 (Radford et al. 2019).

Cluster-to-Cluster Data Construction

To learn to generate diverse new utterances, we train the C2C model with cluster-to-cluster ‘paraphrasing’ pairs extracted from existing training data, and propose a **Dispersed Cluster Pairing** algorithm to construct these pairs.

We hope the cluster-to-cluster generation pairs simulate the data augmentation process, where we generate diverse new utterances from limited expressions. Therefore, given all utterances with same semantic, we gather similar utterances as an input cluster and pick the utterances with the

Algorithm 1: Dispersed Cluster Pairing

Input: Original training data $D = \{(\mathbf{u}_i, \mathbf{s}_i)\}_{i=1}^N$

- 1: Initialize cluster-to-cluster pairs $\mathcal{P} = \{\}$
- 2: All semantic frame $\mathcal{S} = \{\mathbf{s} \mid \exists(\mathbf{u}, \mathbf{s}) \in D\}$
- 3: Delexicalize all training data $D' = \text{Delexicalize}(D)$
- 4: **for** \mathbf{s} **in** \mathcal{S} **do**
 - $C_s = \{(\tilde{\mathbf{u}}, \tilde{\mathbf{s}}) \mid (\tilde{\mathbf{u}}, \tilde{\mathbf{s}}) \in D \wedge \tilde{\mathbf{s}} = \mathbf{s}\}$
 - $C_{\text{lexical}} = \text{K_Medoids}(C_s)$
 - for** $C = \{(\mathbf{u}_i, \mathbf{s}_i)\}_{i=1}^{M'}$ **in** C_{lexical} **do**
 - Initialize target cluster $C' = \{\}$
 - while** $|C'| < M'$ **do**
 - $\mathbf{u}' = \arg \max_{\mathbf{u} \in C_s \setminus C} \text{DS}(\mathbf{u}, C \cup C')$
 - Update $C' = \{\mathbf{u}'\} \cup C'$
 - Update $\mathcal{P} = \mathcal{P} \cup \{(C, C')\}$
- 5: **return** \mathcal{P}

most different expressions as the output cluster. For each semantic frame \mathbf{s} , we construct the input cluster C with *lexical clustering* and the construct output cluster C' with *furthest including* mechanism.

Figure 3 and Algorithm 1 present the workflow of the cluster-to-cluster data construction. Firstly, we perform *lexical clustering* on the utterances with K-Medoids clustering method (Park and Jun 2009). Each lexical cluster contains similar utterances and is used as an input cluster C .

Then, for each source cluster C , we sample target utterance according to a *furthest including* principle. Each time, we pick the \mathbf{u}' that has the highest *diversity score* and include it in target cluster C' .

We compute *diversity score* between a candidate utterance \mathbf{u}' and the union of source cluster C and current target cluster C' as $\text{DS} = \min_{\mathbf{u} \in C \cup C'} \text{EDITDISTANCE}(\mathbf{u}, \mathbf{u}')$. Notice that maximizing the diverse score between \mathbf{u}' and source cluster C increases the target cluster’s novelty against the source cluster. The diverse score between \mathbf{u}' and target cluster C' helps to avoid duplication within the target cluster.

Diversity Rank As mentioned in the decoder section, we adopt the diversity rank to encourage diversity and distinguish sentences in the output cluster. Consequently, we incorporate the diverse rank in training data of C2C model by associating each output utterance with a diverse rank token $\#r$ (See the examples in Figure 3). Since the output cluster utterances are greedily picked by diversity score, we naturally use this greedy picking order as the diversity rank, which models the novelty of output utterance. When augmenting new data, we generate the new utterances at rank from 1 to M' , where M' is a preset size of output cluster.

Cross Expansion After training of the C2C model, we generate unseen new utterances from the constructed input clusters. To avoid the new utterances to overfit to the original output utterances seen in C2C training, we perform data augmentation with a *Cross Expansion* mechanism. We partition all the cluster-to-cluster pairs \mathcal{P} into training ones $\mathcal{P}_{\text{train}}$ and reserved ones $\mathcal{P}_{\text{reserved}}$. Then we train the C2C model only

Model	ATIS			Snips		
	Full	Medium	Small	Full	Medium	Small
Baseline	94.93	85.85	67.33	89.30	64.84	42.33
+ NoiseSeq2Seq (Kurata, Xiang, and Zhou 2016)	94.61	87.34	67.93	-	-	-
+ Slot Expansion (Shin, Yoo, and Lee 2019)	94.67	87.58	74.83	-	-	-
+ Rel-Seq2Seq (Hou et al. 2018)	94.82	88.72	73.71	-	-	-
+ C-VAE (Shin, Yoo, and Lee 2019)	95.04	88.82	71.97	90.93	65.13	38.46
+ Ours w/o pre-train	95.06	90.87*	75.21*	90.33	67.49*	46.94*
+ Ours	95.29	90.95*	75.32*	91.01	67.90*	48.09*

Table 1: Comparison of data augmentation methods for slot filling on ATIS and Snips datasets. Results marked with + are results of the same Bi-LSTM trained with different data augmentation methods. *w/o pre-train* initializes C2C with random parameters rather than pretrained GPT. The Snips results are re-implemented. * indicates that the result is statistically significant over the strongest data augmentation baseline under t-test (p-value < 0.05).

with $\mathcal{P}_{\text{train}}$ and generate new utterances from the input clusters of reserved pairs $\mathcal{P}_{\text{reserved}}$. To make full use of existing utterances, we repeat such partition in a crossing manner.

Experiment

We evaluate the proposed data augmentation method on two slot filling datasets.²

Data We conduct experiments on ATIS and Snips datasets. ATIS (Hemphill, Godfrey, and Doddington 1990) is extensively used for slot filling and provides a well-founded comparison for data augmentation methods. It contains 4,978 training utterances and 893 testing utterances. To simulate the data insufficient situations, we follow Chen et al. (2016a); Hou et al. (2018); Shin, Yoo, and Lee (2019), and evaluate our model on two small proportions of the training data which is *small* proportion (1/40 of the original training set with 129 instances) and *medium* proportion (1/10 of the original training set with 515 instances). We use a development set of 500 instances.

Snips (Coucke et al. 2018) dataset is collected from the Snips personal voice assistant. There are 13,084 training utterances and 700 testing utterances. We use another 700 utterances as the development set. We also split the snips training set into *small* proportion (1/100 of the original training set with 130 instances) and *medium* proportion (1/20 of the original training set with 654 instances).

Evaluation Following previous works (Shin, Yoo, and Lee 2019; Hou et al. 2018), we compute F1-score as evaluation metric with the *conllev* script.³

Implementation We built our Cluster2Cluster model with the transformer implemented by Wolf et al. (2019). For pre-trained parameters, we used the GPT-2, which has 12 layers,

²We only focus on DA for the sequence-labeling problem of slot-filling. So the results may be lower than some joint SLU models, which perform slot-filling using additional information from intents (Peng et al. 2020; Louvan and Magnini 2020).

³www.clips.uantwerpen.be/conll2000/chunking/conllev.txt

110M parameters and the hidden state dimension of 768. We used AdamW (Loshchilov and Hutter 2019) optimizer with initial learning rate 6.25e-5 or 5e-5 for training. We varied λ in {0.1, 0.02, 0.01, 0.002, 0.001} and set γ as 1.0.

Following previous works (Shin, Yoo, and Lee 2019; Hou et al. 2018), we conduct experiments with Bi-LSTM as slot-filling model and train it with both original training data and data augmented by different data augmentation methods. We use the same Bi-LSTM implements as previous work.⁴ The dimension of word embeddings and hidden states was set to 300 and 128, respectively. We used GloVe (Pennington, Socher, and Manning 2014) to initialize word embedding. We varied training batch size in {16, 128}, set dropout rate to 0.5, and trained the model with Adam as suggested by Kingma and Ba (2015).

For all models, best hyperparameter settings are determined on the development set. We report the average of 5 differently-seeded runs for each result.

Main Results for Data Augmentation

Table 1 shows the evaluation results of data augmentation methods on two slot filling datasets: ATIS and Snips. To simulate data insufficient situations, we compare the proposed method with previous data augmentation methods with different proportions following previous works (Chen et al. 2016a; Hou et al. 2018; Shin, Yoo, and Lee 2019). Baseline results are obtained with a Bi-LSTM slot-filling model trained on original training data. And results of each data augmentation methods are obtained with Bi-LSTM models that have the same architecture as the baseline but are trained with both original data and generated data.

On ATIS dataset, our model significantly outperforms the baseline model by 5.10 and 7.99 F-scores on medium and small proportion respectively. There are similar improvements on Snips dataset. These improvements show the effectiveness of our augmentation method in the data-insufficient scenarios. When tested with data sufficient scenarios on full proportions, our model also brings improvements over baselines models. The improvements are narrowed comparing to

⁴github.com/AtmaHou/Bi-LSTM_PosTagger

Model	Full	Medium	Small
Ours	91.01	67.90	48.09
- cluster-wise gen.	90.28	66.23	45.93
- diverse reg.	90.32	66.11	44.32
- dup. attention	90.16	66.43	44.37

Table 2: The results of the ablation test.

those in data scarcity settings. We address this to the fact that full ATIS and Snips are large enough for slot-fillings, which limit the effects of additional synthetic data. When we augment new data without generation pre-training, our performance drops but still achieves significant improvements in most settings, which shows the effectiveness of pre-training and C2C structure respectively. We will discuss pre-training in detail later.

We compare our methods to two kinds of popular data augmentation methods for slot filling: rephrasing-based and sampling-based methods. Similar to our methods, rephrasing-based data augmentation methods reconstruct existing data into alternative expressions.⁵ For this kind of method, NoiseSeq2Seq (Kurata, Xiang, and Zhou 2016) and Rel-Seq2Seq (Hou et al. 2018) learn seq2seq models to reconstruct the existing utterances. To generate unseen expression, NoiseSeq2Seq Introduce noise to decoding, and Rel-Seq2Seq considering the relation between expression alternatives. Slot Expansion (Shin, Yoo, and Lee 2019) generates the new data by randomly replacing the slot values of existing utterances. These methods argument each new utterance independently, thus often generate duplicated expressions that are helpless to improve slot-filling training. Our C2C model mitigates this by jointly encoding and decoding multiple utterances and considering the extensive relation between instances. Such advantages result in higher diversity and help to achieve better performance.

For the second type of data augmentation, we compare with the sampling-based data augmentation method of C-VAE (Shin, Yoo, and Lee 2019). C-VAE leverages a conditioned VAE model to sample new utterances and generates corresponding annotations at the same time. It also faces the diversity problem, since it samples each new data independently. Our methods outperform this strong baseline on all the six slot-filling settings. The improvements come from the better diversity and fluency of the proposed Cluster2Cluster generation. Notably, we gain significant improvements of 9.63 and 3.35 F1-scores on Snips-small and ATIS-small. It shows that our methods are more effective in data scarcity situations.

Analysis

Ablation Test We perform an ablation study to evaluate the importance of each component in C2C framework. Table 2 shows the results on Snips. For the model without cluster-wise generation, we directly fine-tune GPT to gen-

⁵Traditional paraphrasing and back-translation methods are not compared here, because they are not capable to generate token-level annotation for sequence labeling problem.

Model	TF Layers	Full	Medium	Small
w/ pre-train	12	91.01	67.90	48.09
w/o pre-train	12	90.42	67.32	44.96
w/o pre-train	2	90.33	67.49	46.94
w/o pre-train	1	90.10	66.45	46.59

Table 3: Effect analysis of generation pre-training.

Model	ATIS			Snips		
	Full	Medium	Small	Full	Medium	Small
Baseline	94.93	85.85	67.33	89.30	64.84	42.33
+ BERT	95.53	91.18	82.56	96.63	86.34	66.43
+ BERT + Ours	95.45	92.13	85.92	96.12	88.63	73.37

Table 4: Analysis of data augmentation effects with deep pre-trained embeddings

erate new data in a seq-to-seq manner. The drops of F1-score demonstrate the superiority of the cluster-wise generation. If removing either Diverse-Oriented Regularization or Duplication-aware attention from the model, performance drops are witnessed. This shows that both of the two mechanisms help to improve slot-filling by encouraging diversity.

Effects of Generation Pre-training We analyze the impact of initializing C2C model with pre-trained language model. We randomly initialize C2C model and vary the model sizes to avoid overfitting caused by large model sizes. As shown in Table 3, the pre-training helps to improve the effects of data augmentation on all settings. We attribute this to the fact that pre-training can improve generation fluency. However, as revealed in both Table 3 and Table 1, the drops are limited compared to the overall improvements, which shows the inherent effectiveness of C2C model.

Effects over Deep Pre-trained Embeddings For data scarcity problem, deep pre-trained embeddings, such as BERT (Devlin et al. 2019), are also demonstrated as an effective solution (Wang et al. 2020). To see whether DA is still effective when using deep pre-trained embeddings, we conduct DA experiments over a BERT-based slot-filling model.⁶ As shown in Table 4, although BERT greatly improves the performance of slot-filling, our model still achieved improvements on Medium and Small proportion data. This shows the effectiveness of our DA methods for data scarcity problems. Our augmentation method slightly lags the BERT-only model on Full proportion. We address this to the fact that full data is large enough for slot-filling and BERT can be misled by the noise within generated data.

Evaluation for Generation Diversity

Increasing the diversity of generation is one of the essential goals of the data augmentation methods. Following Shin,

⁶We fine-tune the BERT-uncased-base model with AdamW optimizer and initial learning rate of 5e-5.

Model	Inter		Intra	
	Ratio	MED	Ratio	MED
NoiseSeq2Seq	74%	1.20	86%	1.03
Rel-Seq2Seq	96%	3.16	90%	1.81
C-VAE	23%	0.62	11%	0.42
Ours (delexicalized)	96%	5.88	92%	5.55
Ours	100%	9.03	95%	4.85

Table 5: Diversity evaluation of utterance generation.

Yoo, and Lee (2019), we evaluate the diversity of generated data from two aspects: Inter and Intra. *Inter*: ratio of utterances that did not appear in the original training set. *Intra*: ratio of unique utterances among all generated new data.

Such metrics only measure the whole-sentence level diversity, but fail to measure expression diversity at token level. To remedy this, we introduce a token-level diversity metric: Minimum Edit Distance (MED). For each generated utterance \mathbf{u}' , we calculate its MED to a set of utterances C as $MED(\mathbf{u}', C) = \min_{\mathbf{u} \in C} \text{EDITDISTANCE}(\mathbf{u}, \mathbf{u}')$. MED measures novelty of a sentence comparing to a set of existing sentences at token level. We report the averaged MED of each generated utterance to the original training set (*Inter*) and to the other generated utterances (*Intra*).

Table 5 shows the evaluation of the generation diversity on the ATIS-Full. For *Inter* Diversity, our method significantly outperforms all previous methods on both Ratio and average MED metrics. We note that we can achieve the best diversity even evaluating the generated delexicalized utterances. It shows the great ability of the C2C model in generating unseen expressions. This is mainly due to that cluster-wise encoding mechanism allows model to be aware of more existing expression during generation.

For *Intra* Diversity, our method also achieves the best performances over the previous works. These improvements show that considering relations between generated utterances can significantly reduce duplication.

Diversity Analysis To understand how the proposed method enhances expression diversity, we investigate the diversity distribution of generated delexicalized utterances on the ATIS-full. We measure the diversity with *Inter* MED. As shown in Figure 4, Seq2Seq generation yields more existing expressions, and the MED scores are mostly distributed in low-value areas. Comparing to Seq2Seq, Cluster2Cluster model generally has higher MED scores. This demonstrates the intrinsic advantage of the cluster-wise generation to generate new expressions.

When training the Cluster2Cluster model with Diverse-Oriented Regularization and Duplication-aware Attention, there is much fewer existing expressions within generated utterances, and we can see a continuous drifting of distribution towards higher diversity. This shows that the proposed mechanisms help to generate more diverse utterances.




Figure 4: Diversity distribution of generated expressions. Edit Distance is the minimum edit distance to existing expressions.

Case Study of Generation

We conduct case study to see how the proposed Cluster2Cluster model generates unseen expressions. Specifically, we randomly picking a generated utterance \mathbf{u}' , and searching in the original training set D for the most similar utterance $\mathbf{u} = \arg \min_{\hat{\mathbf{u}} \in D} \text{EDITDISTANCE}(\hat{\mathbf{u}}, \mathbf{u}')$. To focus on expressions change, we perform experiments on delexicalized utterances.

By comparing the difference between each pair of \mathbf{u}' and \mathbf{u} , we find the new expressions in several interesting manners: Replace Phrases, Enrich Info, Change Syntax, Change Semantics. The examples are listed in Table 6. We find that the most common new expressions are from replacing phrases and enriching information. As shown from the given examples, Cluster2Cluster model can replace semantically similar phrases and bring additional information. The phrase replacing ability is mainly learned from cluster-to-cluster training data, and the enriched information often comes from the pretraining process. As for Change Syntax, it is interesting to see that Cluster2Cluster can yield new expression by using alternative syntax, which is very close to humans in paraphrasing.

Related Work

Data augmentation (DA) solves data scarcity problems by enlarging the size of training data (Fader, Zettlemoyer, and Etzioni 2013; Zhang, Zhao, and LeCun 2015a; Zhao, Zhu, and Yu 2019; Kim, Roh, and Kim 2019; Yin et al. 2020; Li et al. 2019). Previous DA works propose back-translation methods (Sennrich, Haddow, and Birch 2016; Wang et al. 2019) and paraphrasing methods (Zhang, Zhao, and LeCun 2015b; Iyyer et al. 2018; Hu et al. 2019; Gao et al. 2020) to generate semantically similar sentences. However, these DA methods are not applicable to the sequence labeling problem of slot-filling. Because slot filling requires token-level annotations of semantic frame, while these methods can only provide sentence-level labels.

Spoken Language understanding, including slot filling and intent detection tasks, has drawn a lot of research attention recently (Yao et al. 2013, 2014; Mesnil et al. 2013, 2015; Chen et al. 2016a,b; Goo et al. 2018; Haihong et al. 2019; Liu et al. 2019; Zhu, Cao, and Yu 2020). In this pa-

Type	u and u'
Replace Phrases	u: show me the flights from <from_city> to <to_city> with stop in <stop_city> u': give me flights from <from_city> to <to_city> with stopover in <stop_city>
Enrich Info	u: i 'd like information on all the flights from <from_city> to <to_city> on <depart_date> u': i 'm sorry to see all the flights that i take from <from_city> to <to_city> on <depart_date>
Change Syntax	u: how much is a flight from <from_city> to <to_city> u': how much does a flight cost from <from_city> to <to_city>
Change Semantics	u: show all airlines with flights between <from_city> and <to_city> u': show me more airlines with seats between <from_city> and <to_city>

Table 6: Case study of new expression. For each generated example u' , we find the most similar existing utterance u to it and compare the differences.

per, we only focus on the slot filling task. For data augmentation of slot filling, previous works focus on generation-based methods. Kurata, Xiang, and Zhou (2016); Hou et al. (2018); Peng et al. (2020) augment the training data with a Sequence-to-Sequence model. Shin, Yoo, and Lee (2019); Yoo, Shin, and Lee (2019) introduced Variational Auto-Encoder (Kingma and Welling 2014) and jointly generate new utterances and predict the labels. Louvan and Magnini (2020) introduce simple rules to generate new utterances. Different from our C2C framework, these methods augment each instance independently and often unconsciously generate duplicated expressions.

Conclusion

In this paper, we study the data augmentation problem for slot filling and propose a novel data augmentation framework C2C-GenDA, which generates new instances from existing training data in a cluster-to-cluster manner. C2C-GenDA improves generation diversity by considering the relation between generated utterances and capturing more existing expressions. To further encourage diversity, we propose Duplication-aware Attention and Diverse-Oriented Regularization mechanism. We introduce a Dispersed Cluster Pairing algorithm to construct cluster-to-cluster paraphrasing pairs for C2C-DA training. Experiments show that the proposed framework can improve slot-filling by generating diverse new training data and outperform existing data augmentation systems of slot-filling.

Acknowledgments

We are grateful for the helpful comments and suggestions from the anonymous reviewers. This work was supported by the National Key R&D Program of China via grant 2020AAA0106501 and the National Natural Science Foundation of China (NSFC) via grant 61976072 and 61772153.

References

Chen, Y.-N.; Hakkani-Tür, D.; Tur, G.; Celikyilmaz, A.; Guo, J.; and Deng, L. 2016a. Syntax or semantics? knowledge-guided joint semantic frame parsing. In *Proc. of SLT*, 348–355.

Chen, Y.-N. V.; Hakkani-Tür, D.; Tur, G.; Gao, J.; and Deng, L. 2016b. End-to-End Memory Networks with Knowledge Carryover for Multi-Turn Spoken Language Understanding. In *Proc. of INTERSPEECH*.

Coucke, A.; Saade, A.; Ball, A.; Bluche, T.; Caulier, A.; Leroy, D.; Doumouro, C.; Gisselbrecht, T.; Caltagirone, F.; Lavril, T.; Prinet, M.; and Dureau, J. 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 *Proc. of NAACL*, 4171–4186.

Fader, A.; Zettlemoyer, L.; and Etzioni, O. 2013. Paraphrase-Driven Learning for Open Question Answering. In *Proc. of ACL*.

Gao, S.; Zhang, Y.; Ou, Z.; and Yu, Z. 2020. Paraphrase Augmented Task-Oriented Dialog Generation. In *Proc. of ACL*.

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 *Proc. of NAACL*, 753–757.

Haihong, E.; Niu, P.; Chen, Z.; and Song, M. 2019. A novel bi-directional interrelated model for joint intent detection and slot filling. In *Proc. of ACL*, 5467–5471.

Hannun, A. Y.; Case, C.; Casper, J.; Catanzaro, B.; Diamos, G.; Elsen, E.; Prenger, R.; Satheesh, S.; Sengupta, S.; Coates, A.; and Ng, A. Y. 2014. Deep Speech: Scaling up end-to-end speech recognition. *arXiv preprint arXiv:1412.5567*.

Hemphill, C. T.; Godfrey, J. J.; and Doddington, G. R. 1990. The ATIS spoken language systems pilot corpus. In *Proc. of Speech and Natural Language Workshop*.

Hou, Y.; Liu, Y.; Che, W.; and Liu, T. 2018. Sequence-to-Sequence Data Augmentation for Dialogue Language Understanding. In *Proc. of COLING*, 1234–1245.

Hu, J. E.; Khayrallah, H.; Culkin, R.; Xia, P.; Chen, T.; Post, M.; and Durme, B. V. 2019. Improved Lexically Constrained Decoding for Translation and Monolingual Rewriting. In *Proc. of NAACL*, 839–850.

- Iyyer, M.; Wieting, J.; Gimpel, K.; and Zettlemoyer, L. 2018. Adversarial Example Generation with Syntactically Controlled Paraphrase Networks. In *Proc. of NAACL*, 1875–1885.
- Kim, H.-Y.; Roh, Y.-H.; and Kim, Y.-G. 2019. Data Augmentation by Data Noising for Open-vocabulary Slots in Spoken Language Understanding. In *Proc. of NAACL*, 97–102.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *Proc. of ICLR*.
- Kingma, D. P.; and Welling, M. 2014. Auto-Encoding Variational Bayes. In *Proc. of ICLR*.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In Pereira, F.; Burges, C. J. C.; Bottou, L.; and Weinberger, K. Q., eds., *Proc. of NeurIPS*, 1097–1105.
- Kurata, G.; Xiang, B.; and Zhou, B. 2016. Labeled Data Generation with Encoder-Decoder LSTM for Semantic Slot Filling. In *Proc. of INTERSPEECH*, 725–729.
- Li, J.; Qiu, L.; Tang, B.; Chen, D.; Zhao, D.; and Yan, R. 2019. Insufficient Data Can Also Rock! Learning to Converse Using Smaller Data with Augmentation. In *Proc. of AAAI*, volume 33, 6698–6705.
- Liu, Y.; Meng, F.; Zhang, J.; Zhou, J.; Chen, Y.; and Xu, J. 2019. CM-Net: A Novel Collaborative Memory Network for Spoken Language Understanding. In *Proc. of EMNLP-IJCNLP*, 1050–1059.
- Loshchilov, I.; and Hutter, F. 2019. Decoupled Weight Decay Regularization. In *Proc. of ICLR*. OpenReview.net.
- Louvan, S.; and Magnini, B. 2020. Simple is Better! Lightweight Data Augmentation for Low Resource Slot Filling and Intent Classification. *arXiv preprint arXiv:2009.03695*.
- Mesnil, G.; Dauphin, Y.; Yao, K.; Bengio, Y.; Deng, L.; Hakkani-Tür, D. Z.; He, X.; Heck, L. P.; Tür, G.; Yu, D.; and Zweig, G. 2015. Using Recurrent Neural Networks for Slot Filling in Spoken Language Understanding. *Proc. of TASLP* 23(3): 530–539.
- Mesnil, G.; He, X.; Deng, L.; and Bengio, Y. 2013. Investigation of recurrent-neural-network architectures and learning methods for spoken language understanding. In *Proc. of INTERSPEECH*.
- Park, H.-S.; and Jun, C.-H. 2009. A simple and fast algorithm for K-medoids clustering. *Expert systems with applications* 36(2): 3336–3341.
- Peng, B.; Zhu, C.; Zeng, M.; and Gao, J. 2020. Data Augmentation for Spoken Language Understanding via Pre-trained Models. *arXiv preprint arXiv:2004.13952*.
- Pennington, J.; Socher, R.; and Manning, C. D. 2014. GloVe: Global Vectors for Word Representation. In *Proc. of EMNLP*, 1532–1543.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; and Sutskever, I. 2019. Language models are unsupervised multitask learners. *OpenAI Blog* 1(8): 9.
- Sennrich, R.; Haddow, B.; and Birch, A. 2016. Improving Neural Machine Translation Models with Monolingual Data. In *Proc. of ACL*.
- Shin, Y.; Yoo, K. M.; and Lee, S.-G. 2019. Utterance Generation With Variational Auto-Encoder for Slot Filling in Spoken Language Understanding. *IEEE Signal Processing Letters* 26(3): 505–509.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Proc. of NIPS*, 5998–6008.
- Wang, S.; Liu, Y.; Wang, C.; Luan, H.; and Sun, M. 2019. Improving Back-Translation with Uncertainty-based Confidence Estimation. In *Proc. of EMNLP-IJCNLP*, 791–802.
- Wang, Y.; Hou, Y.; Che, W.; and Liu, T. 2020. From static to dynamic word representations: a survey. *International Journal of Machine Learning and Cybernetics* 1–20.
- Wolf, T.; Debut, L.; Sanh, V.; Chaumond, J.; Delangue, C.; Moi, A.; Cistac, P.; Rault, T.; Louf, R.; Funtowicz, M.; and Brew, J. 2019. HuggingFace’s Transformers: State-of-the-art Natural Language Processing. *ArXiv abs/1910.03771*.
- Yao, K.; Peng, B.; Zhang, Y.; Yu, D.; Zweig, G.; and Shi, Y. 2014. Spoken language understanding using long short-term memory neural networks. In *Proc. of the IEEE Spoken Language Technology Workshop*, 189–194. IEEE.
- Yao, K.; Zweig, G.; Hwang, M.-Y.; Shi, Y.; and Yu, D. 2013. Recurrent neural networks for language understanding. In *Proc. of Interspeech*, 2524–2528.
- Yin, Y.; Shang, L.; Jiang, X.; Chen, X.; and Liu, Q. 2020. Dialog State Tracking with Reinforced Data Augmentation. In *Proc. of AAAI*.
- Yoo, K. M. 2020. *Deep Generative Data Augmentation for Natural Language Processing*. Ph.D. thesis, Seoul National University.
- Yoo, K. M.; Shin, Y.; and Lee, S.-g. 2019. Data augmentation for spoken language understanding via joint variational generation. In *Proc. of AAAI*, volume 33, 7402–7409.
- Young, S.; Gašić, M.; Thomson, B.; and Williams, J. D. 2013. Pomdp-based statistical spoken dialog systems: A review. *Proc. of the IEEE* 101(5): 1160–1179.
- Zhang, X.; Zhao, J.; and LeCun, Y. 2015a. Character-level convolutional networks for text classification. In *Proc. of NIPS*, 649–657.
- Zhang, X.; Zhao, J. J.; and LeCun, Y. 2015b. Character-level Convolutional Networks for Text Classification. In *Proc. of NIPS*, 649–657.
- Zhao, Z.; Zhu, S.; and Yu, K. 2019. Data Augmentation with Atomic Templates for Spoken Language Understanding. In *Proc. of EMNLP-IJCNLP*, 3628–3634.
- Zhu, S.; Cao, R.; and Yu, K. 2020. Dual learning for semi-supervised natural language understanding. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 28: 1936–1947.