

# Data Augmentation with Atomic Templates for Spoken Language Understanding

Zijian Zhao\*, Su Zhu\* and Kai Yu†

MoE Key Lab of Artificial Intelligence  
SpeechLab, Department of Computer Science and Engineering  
Shanghai Jiao Tong University, Shanghai, China  
{1248uu, paul2204, kai.yu}@sjtu.edu.cn

## Abstract

Spoken Language Understanding (SLU) converts user utterances into structured semantic representations. Data sparsity is one of the main obstacles of SLU due to the high cost of human annotation, especially when domain changes or a new domain comes. In this work, we propose a data augmentation method with atomic templates for SLU, which involves minimum human efforts. The atomic templates produce exemplars for fine-grained constituents of semantic representations. We propose an encoder-decoder model to generate the whole utterance from atomic exemplars. Moreover, the generator could be transferred from source domains to help a new domain which has little data. Experimental results show that our method achieves significant improvements on DSTC 2&3 dataset which is a domain adaptation setting of SLU.

## 1 Introduction

The SLU module is a key component of spoken dialogue system, parsing user utterances into corresponding semantic representations in a narrow domain. The typical semantic representation for SLU could be semantic frame (Tur and De Mori, 2011) or dialogue act (Young, 2007). In this paper, we focus on SLU with the dialogue act that a sentence is labelled as a set of *act-slot-value* triples. For example, the utterance “*Not Chinese but I want Thai food please*” has an annotation of “*deny(food=Chinese), inform(food=Thai)*”.

Deep learning has achieved great success in the SLU field (Mesnil et al., 2015; Liu and Lane, 2016; Zhao et al., 2019). However, it is notorious for requiring large labelled data, which limits the scalability of SLU models. Despite re-

\* Zijian Zhao and Su Zhu are co-first authors and contribute equally to this work.

† The corresponding author is Kai Yu.

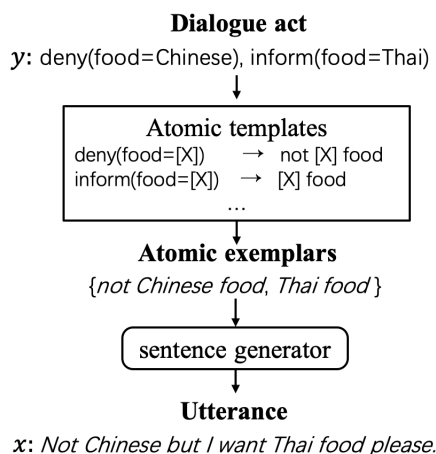


Figure 1: Workflow of the data augmentation with atomic templates for SLU.

cent advancements and tremendous research activity in semi-supervised learning and domain adaptation, the deep SLU models still require massive amounts of labelled data to train.

Therefore, data augmentation for SLU becomes appealing, and it needs three kinds of capabilities:

- **Expression diversity:** There are always various expressions for the same semantic meaning. Lack of expression diversity is an obstacle of SLU.
- **Semantic diversity:** The data augmentation method should generate data samples with various semantic meanings.
- **Domain adaptation:** The method should be able to utilize data from other domains and adapt itself to a new domain rapidly.

In this work, we propose a new data augmentation method for SLU, which consists of atomic templates and a sentence generator (as shown in Figure 1). The method starts from dialogue acts

which are well-structured and predefined by domain experts. Rich and varied dialogue acts can be created automatically, which is guided by a domain ontology.

To enhance the capability of domain adaptation for the sentence generator, we propose to first interpret dialogue acts in natural language so that new dialogue acts can be understood by the generator. However, interpreting a dialogue act at sentence level costs a lot, atomic templates are exploited to alleviate human efforts. The atomic templates are created at phrase level, which produce exemplars for fine-grained constituents (i.e. *act-slot-value* triple) of semantic representations (i.e. dialogue act). Thus, the sentence generator can be an encoder-decoder model, paraphrasing a set of atomic exemplars to an utterance.

We evaluate our approach on DSTC 2&3 dataset (Henderson et al., 2014a,b), a benchmark of SLU including domain adaptation setting. The results show that our method can obtain significant improvements over strong baselines with limited training data.

## 2 Related Work

**Data augmentation for SLU** Hou et al. (2018) use a sequence-to-sequence model to translate between sentence pairs with the same meaning. However, it cannot generate data for new semantic representations. Yoo et al. (2018) use a variational autoencoder to generate labelled language. It still lacks variety of semantic meanings. Let  $x$  and  $y$  denote an input sentence and the corresponding semantic representation respectively. These two works estimate  $p(x, y)$  with generation models to produce more labelled samples, while they have two drawbacks. 1) The models cannot control which kinds of semantic meaning should be generated for supply. 2) The models cannot generate data for new semantic representations which may contain out-of-vocabulary (OOV) labels. We propose to generate utterance based on semantic representations by estimating  $p(x|y)$ , because  $y$  is well-structured and easy to be synthesized. To overcome the OOV problem of semantic labels, we first map them to atomic exemplars with a little of human effort.

**Zero-shot learning of SLU** Besides data augmentation, zero-shot learning of SLU (Ferreira et al., 2015; Yazdani and Henderson, 2015) is also related, which can adapt to unseen semantic labels.

Yazdani and Henderson (2015) exploit a binary classifier for each possible *act-slot-value* triple to predict whether it exists in the input sentence. Zhao et al. (2019) propose a hierarchical decoding model for SLU. However, they still have problem with new *act* and *slot*. Bapna et al. (2017); Lee and Jha (2018); Zhu and Yu (2018) try to solve it with textual slot descriptions. In this paper, we propose atomic templates to describe *act-slot-value* triples but not separate slots or acts.

## 3 Method

In this section, our data augmentation method with atomic templates will be introduced, which generates additional data samples for training SLU model.

### 3.1 SLU Model

For dialogue act based SLU, we adopt the hierarchical decoding (HD) model (Zhao et al., 2019) which dynamically parses *act*, *slot* and *value* in a structured way. The model consists of four parts:

- A shared utterance encoder, bidirectional LSTM recurrent neural network (BLSTM) (Graves, 2012);
- An *act* type classifier on the utterance;
- A *slot* type classifier with the utterance and an act type as inputs;
- A *value* decoder with the utterance and an *act-slot* type pair as inputs.

The *value* decoder generates word sequence of the value by utilizing a sequence-to-sequence model with attention (Luong et al., 2015) and pointer network (Vinyals et al., 2015a) which helps handling out-of-vocabulary (OOV) values.

### 3.2 Data Generation with Atomic Templates

As illustrated in Figure 1, the workflow of the data augmentation is broken down into: 1) mapping *act-slot-value* triples to exemplars with atomic templates, and 2) generating the corresponding utterance depending on the atomic exemplars.

Let  $x = x_1 \cdots x_{|x|}$  denote the utterance (word sequence), and  $y = \{y_1, \cdots, y_{|y|}\}$  denote the dialogue act (a set of *act-slot-value* triples). We wish to estimate  $p(x|y)$ , the conditional probability of utterance  $x$  given dialogue act  $y$ .

However, there are also some disadvantages to directly using dialogue acts as inputs:

Triple	Templates
bye()	goodbye bye
request(addr)	the address what’s the address
inform(food=[food])	[food] [food] food
inform(hastv=true)	television with a television

Table 1: Examples of atomic templates in DSTC 2&3 dataset. “[food]” is an arbitrary value of slot “food”.

- The dialogue acts from different narrow domains may conflict, e.g. different slot names for the same meaning.
- Act and slot types may not be defined in a literal way, e.g. using arbitrary symbols like “city\_1” and “city\_2” to represent city names in different contexts.

Thus, it is hard to adapt the model  $p(x|y)$  to new act types, slot types, and domains.

We propose to interpret the dialogue act in short natural language, and then rephrase it to the corresponding user utterance. While interpreting the dialogue act at sentence level costs as much as building a rule-based SLU system, we choose to interpret *act-slot-value* triples with atomic templates which involve minimum human efforts.

### 3.2.1 Atomic templates

Table 1 gives some examples of atomic templates used in DSTC 2&3 dataset. Atomic templates produce a simple description (atomic exemplar  $e_i$ ) in natural language for each *act-slot-value* triple<sup>1</sup>  $y_i$ . If there are multiple templates for triple  $y_i$ , which generate multiple atomic exemplars  $E(y_i)$ , we choose the most similar one  $e_i = \operatorname{argmax}_{e \in E(y_i)} \operatorname{sim}(e, x)$  in the training stage, and randomly select one  $e_i$  from  $E(y_i)$  in the data augmentation stage. The similarity function  $\operatorname{sim}(e, x)$  we used is Ratcliff-Obershelp algorithm (Black, 2004). Therefore,

$$p(x|y) = p(x|\{y_1 \cdots y_{|y|}\}) = p(x|\{e_1 \cdots e_{|y|}\})$$

### 3.2.2 Sentence generator

An encoder-decoder model is exploited to generate the utterance based on the set of atomic exemplars by estimating  $p(x|\{e_1 \cdots e_{|y|}\})$ .

<sup>1</sup>It should be noted that an *act-slot-value* triple may have an empty value, e.g. “request(phone)” refers to asking for phone number. The triple may also need no slot, e.g. “bye()”.

As the set of atomic exemplars is unordered, we encode them independently. For each atomic exemplar  $e_i = w_{i1} \cdots w_{iT_i}$  (a sequence of words with length  $T_i$ ), we use a BLSTM to encode it. The hidden vectors are recursively computed at the  $j$ -th time step via:

$$\begin{aligned} \vec{\mathbf{h}}_{ij} &= \mathbf{f}_{\text{LSTM}}(\psi(w_{ij}), \vec{\mathbf{h}}_{ij-1}), j = 1, \dots, |T_i| \\ \overleftarrow{\mathbf{h}}_{ij} &= \mathbf{b}_{\text{LSTM}}(\psi(w_{ij}), \overleftarrow{\mathbf{h}}_{ij+1}), j = |T_i|, \dots, 1 \\ \mathbf{h}_{ij} &= [\vec{\mathbf{h}}_{ij}; \overleftarrow{\mathbf{h}}_{ij}] \end{aligned}$$

where  $[\cdot; \cdot]$  denotes vector concatenation,  $\psi(\cdot)$  is a word embedding function,  $\mathbf{f}_{\text{LSTM}}$  is the forward LSTM function and  $\mathbf{b}_{\text{LSTM}}$  is the backward one. The final hidden vectors of the forward and backward passes are utilized to represent each atomic exemplar  $e_i$ , i.e.  $\mathbf{c}_i = [\vec{\mathbf{h}}_{iT_i}; \overleftarrow{\mathbf{h}}_{i1}]$ .

After encoding all the atomic exemplars, we have a list of hidden vectors  $H(y) = [\mathbf{h}_{11} \cdots \mathbf{h}_{1T_1}; \cdots; \mathbf{h}_{|y|1} \cdots \mathbf{h}_{|y|T_{|y|}}]$ . A LSTM model serves as the decoder (Vinyals et al., 2015b), generating the utterance  $x$  word-by-word:

$$p(x|y) = p(x|H(y)) = \prod_{t=1}^{|x|} p(x_t|x_{<t}, H(y))$$

Before the generation, the hidden vector of the decoder is initialized as  $1/|y| \sum_{i=1}^{|y|} \mathbf{c}_i$ , the mean of representations of all the atomic exemplars. The pointer softmax (Gulcehre et al., 2016) enhanced by a trick of targeted feature dropout (Xu and Hu, 2018) is adopted to tackle OOV words, which will switch between generation and copying from the input source dynamically.

## 4 Experiment

### 4.1 Data

In our experiments, we use the dataset provided for the second and third Dialog State Tracking Challenge (DSTC 2&3) (Henderson et al., 2014a,b). DSTC2 contains a large number of training dialogues ( $\sim 15.6\text{k}$  utterances) related to restaurant search while DSTC3 is designed to address the problem of adaptation to a new domain (tourist information) with only a small amount of seed data (*dstc3\_seed*, 109 utterances). The manual transcriptions are used as user utterances to eliminate the impact of speech recognition errors.

We follow the data partitioning policy as Zhu et al. (2014), which randomly selects one-half of

DSTC3 test data as the oracle training set ( $\sim 9.4k$  utterances) and leaves the other half as the evaluation set ( $\sim 9.2k$  utterances).

## 4.2 Experimental Setup

Delexicalized triples are used in case of non-enumerable slot, like “inform(food=[food])” in Table 1. There are 41 and 35 delexicalized triples for DSTC2 and DSTC3 respectively. For each triple, we prepare two short templates on average. Note that, compared to designing sentence-level templates, writing atomic templates for triples at phrase level requires much less human efforts.

For data augmentation of DSTC3, there are two ways to collect extra dialogue acts as inputs<sup>2</sup>:

- **Seed abridgement:** Abridging the delexicalized dialogue acts existing in *dstc3\_seed* data, e.g. a dialogue act with 3 triples {A,B,C} can be abridged to obtain {A,B}, {B,C}, {A,C}, {A}, {B}, {C} and itself.
- **Combination:** As all possible triples are pre-defined in the domain ontology of DSTC3 by experts. We use a general policy of triples combination, which randomly selects at most  $N_c$  triples to make up a dialogue act. ( $N_c$  is set as 3 empirically.)

Then we fill up each non-enumerable slot in the dialogue act by randomly choosing a value of this slot, which ends when each value appears at least  $N_v$  times in all the collected dialogue acts.  $N_v$  is set as 3 empirically. After that, we have 1420 and 20670 dialogue acts from the seed abridgement and combination respectively. New data samples are generated starting from these dialogue acts, through the atomic templates and the sentence generator (e.g. 1-best output is kept).

The SLU model and sentence generator use the same hyper-parameters. The dimension of word embeddings (*Glove*<sup>3</sup> word vectors are used for initialization) is 100 and the number of hidden units is 128. Dropout rate is 0.5 and batch size is 20. Maximum norm for gradient clipping is set to 5 and Adam optimizer is used with an initial learning rate of 0.001. All training consists

<sup>2</sup>The data splits, atomic templates and extra dialogue acts will be available at [https://github.com/sz128/DAAT\\_SLU](https://github.com/sz128/DAAT_SLU)

<sup>3</sup><http://nlp.stanford.edu/data/glove.6B.zip>

of 50 epochs with early stopping on the development set. We report F1-score of extracting *act-slot-value* triples by the official scoring script from <http://camdial.org/~mh521/dstc/>.

## 4.3 Systems

We first compare two SLU models to answer why the HD model is chosen for dialogue act based SLU:

- **ZS:** zero-shot learning of SLU (Yazdani and Henderson, 2015) which can adapt to unseen dialogue acts.
- **HD:** the hierarchical decoding model (Zhao et al., 2019) is adopted in our system.

We make comparisons of other data augmentation methods and the atomic templates (AT):

- **Naive:** We replace the value simultaneously existing in an utterance and its semantic labels of *dstc3\_seed* by randomly selecting a value of the corresponding slot. It ends when each value appears at least  $N_v$  times.
- **Human:** Zhu et al. (2014) proposed to design a large number of sentence-level templates for DSTC3 with lots of human efforts.
- **Oracle:** The oracle training set is used which simulates the perfect data augmentation.

Without data augmentation, the SLU models are pre-trained on DSTC2 dataset (source domain) and finetuned with *dstc3\_seed* set. In our data augmentation method, the sentence generator based on atomic concepts is also pre-trained on DSTC2 dataset and finetuned with the *dstc3\_seed*. The SLU model is first pre-trained on DSTC2 dataset, then finetuned with the augmented dataset and finally finetuned with the *dstc3\_seed*.

## 4.4 Results and Analysis

The main results are illustrated in Table 2. We can see that:

1. The hierarchical decoding (HD) model gets better performance than the zero-shot learning (ZS) method of SLU.
2. The seed data *dstc3\_seed* limits the power of the SLU model, and even the naive augmentation can enhance it.

SLU	Augmentation	F1-score
ZS	w/o	68.3
HD	w/o	78.5
	Naive	82.9
	AT (seed abridgement)	85.5
	AT (combination)	87.9
	AT (seed abr. + comb.)	<b>88.6</b>
	Human (Zhu et al., 2014)	90.4
	Oracle	96.9

Table 2: SLU performances of different systems on the DSTC3 evaluation set.

- Our data augmentation method with atomic templates (AT) improves the SLU performance dramatically. One reason may be the generated data has higher variety of semantic meaning than the naive augmentation. *Combination* can make up more dialogue acts and shows better result than *Seed abridgement*, while *Seed abridgement* provides more realistic dialogue acts. Thus, their union gives the best result.
- The best performance of our method is close to the human-designed sentence-level templates (Zhu et al., 2014), while our approach needs much less human efforts.

**Ablation Study** We conduct several ablation studies to analyze the effectiveness of different components of our method, as shown in Table 3. By removing SLU model pretraining on DSTC2 (“- *dstc2*”) and finetuning on the seed data (“- *dstc3\_seed*”), we can see a significant decrease in SLU performance. When we subsequently cast aside the sentence generator (“- *sentence generator*”, i.e. using the atomic exemplars as inputs of SLU model directly), the SLU performance decreases by 10.3%. This shows that the sentence generator can produce more natural utterances. If we replace the atomic exemplars as the corresponding *act-slot-value* triples (“- *atomic templates*”), the SLU performance drops sharply. The reason may be that the atomic templates provide a better description of corresponding semantic meanings than the surfaces of the triples. Examples of generated data samples are in Appendix A.

**Number of seed samples** Figure 2 shows the number of seed samples used versus SLU performance on DSTC3 evaluation set. For zero-shot case (no seed samples), our method is much better than the baseline. When the number of seed sam-

Method	F1-score
HD + AT(seed abr. + comb.)	88.6
- <i>dstc2</i>	86.2
- <i>dstc3_seed</i>	84.5
- <i>dstc2</i> , - <i>dstc3_seed</i>	84.3
- <i>sentence generator</i>	74.0
- <i>atomic templates</i>	70.5

Table 3: SLU performances on the DSTC3 evaluation set when removing different modules of our method.




Figure 2: SLU performance of different methods with varying number of seed samples randomly selected from the oracle DSTC3 training set.

ples increases, our method outperforms the baseline constantly.

## 5 Conclusion

In this paper, we propose a new data augmentation method with atomic templates for SLU. The atomic templates provide exemplars in natural language for *act-slot-value* triples, involving minimum human efforts. An encoder-decoder model is exploited to generate utterances based on the atomic exemplars. We believe our method can also be applied to SLU systems with other semantic representations (e.g. semantic frame). Experimental results show that our method achieves significant improvements on DSTC 2&3 dataset, and it is very effective for SLU domain adaptation with limited data.

## Acknowledgments

This work has been supported by the National Key Research and Development Program of China (Grant No. 2017YFB1002102) and the China NSFC projects (No. 61573241). We thank the anonymous reviewers for their thoughtful comments and efforts towards improving this manuscript.



## References

- Ankur Bapna, Gökhan Tür, Dilek Hakkani-Tür, and Larry P. Heck. 2017. Towards zero-shot frame semantic parsing for domain scaling. In *18th Annual Conference of the International Speech Communication Association*, pages 2476–2480.
- Paul E. Black. 2004. Ratcliff/Obershelp pattern recognition. In *Dictionary of Algorithms and Data Structures* [online], Paul E. Black, ed. Available from: <https://www.nist.gov/dads/HTML/ratcliffObershelp.html> (accessed 20 May 2019).
- Emmanuel Ferreira, Bassam Jabaian, and Fabrice Lefèvre. 2015. Zero-shot semantic parser for spoken language understanding. In *Sixteenth Annual Conference of the International Speech Communication Association*.
- Alex Graves. 2012. *Supervised Sequence Labelling with Recurrent Neural Networks*. Springer Berlin Heidelberg.
- Caglar Gulcehre, Sungjin Ahn, Ramesh Nallapati, Bowen Zhou, and Yoshua Bengio. 2016. Pointing the unknown words. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 140–149.
- Matthew Henderson, Blaise Thomson, and Jason D Williams. 2014a. 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.
- Matthew Henderson, Blaise Thomson, and Jason D Williams. 2014b. The third dialog state tracking challenge. In *2014 IEEE Spoken Language Technology Workshop (SLT)*, pages 324–329. IEEE.
- Yutai Hou, Yijia Liu, Wanxiang Che, and Ting Liu. 2018. Sequence-to-sequence data augmentation for dialogue language understanding. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1234–1245.
- Sungjin Lee and Rahul Jha. 2018. Zero-shot adaptive transfer for conversational language understanding. *arXiv preprint arXiv:1808.10059*.
- Bing Liu and Ian Lane. 2016. Attention-based recurrent neural network models for joint intent detection and slot filling. In *17th Annual Conference of the International Speech Communication Association (InterSpeech)*.
- Minh-Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Effective approaches to attention-based neural machine translation. *arXiv preprint arXiv:1508.04025*.
- Grégoire Mesnil, Yann Dauphin, Kaisheng Yao, Yoshua Bengio, Li Deng, Dilek Hakkani-Tur, Xiaodong He, Larry Heck, Gokhan Tur, Dong Yu, et al. 2015. Using recurrent neural networks for slot filling in spoken language understanding. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 23(3):530–539.
- Gokhan Tur and Renato De Mori. 2011. *Spoken language understanding: Systems for extracting semantic information from speech*. John Wiley & Sons.
- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015a. Pointer networks. In *Advances in Neural Information Processing Systems*, pages 2692–2700.
- Oriol Vinyals, Łukasz Kaiser, Terry Koo, Slav Petrov, Ilya Sutskever, and Geoffrey Hinton. 2015b. Grammar as a foreign language. In *Advances in neural information processing systems*, pages 2773–2781.
- Puyang Xu and Qi Hu. 2018. An end-to-end approach for handling unknown slot values in dialogue state tracking. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1448–1457.
- Majid Yazdani and James Henderson. 2015. A model of zero-shot learning of spoken language understanding. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 244–249.
- Kang Min Yoo, Youhyun Shin, and Sang-goo Lee. 2018. Data augmentation for spoken language understanding via joint variational generation. *arXiv preprint arXiv:1809.02305*.
- Steve Young. 2007. CUED standard dialogue acts. *Report, Cambridge University Engineering Department, 14th October*.
- Zijian Zhao, Su Zhu, and Kai Yu. 2019. A hierarchical decoding model for spoken language understanding from unaligned data. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7305–7309.
- Su Zhu, Lu Chen, Kai Sun, Da Zheng, and Kai Yu. 2014. Semantic parser enhancement for dialogue domain extension with little data. In *2014 IEEE Spoken Language Technology Workshop (SLT)*, pages 336–341. IEEE.
- Su Zhu and Kai Yu. 2018. Concept transfer learning for adaptive language understanding. In *Proceedings of the 19th Annual SIGdial Meeting on Discourse and Dialogue*, pages 391–399. Association for Computational Linguistics.

## A Examples of generated data

Dialogue act	Atomic exemplars	Generated utterance
thankyou bye	thank you good bye	Thank you good bye.
request(addr) request(phone)	the address the phone number	The address and phone number.
request(addr) request(phone)	whats the address what is the phone number	What is the address and phone number?
inform(type=coffeeshop) inform(hasinternet=true)	cafe has internet connection	I'm looking for a cafe has the internet connection.
affirm inform(type=restaurant) inform(pricerange=moderate)	yes looking for a restaurant moderately priced	Yes, I'm looking for a moderately priced restaurant.
request(childrenallowed)	does it allow children	Does it <b>have</b> children?
request(hastv) request(addr)	does it has a television the address	Does it have the television <b>address</b> and <b>ad- dress</b> ?

Table 4: Examples of generated data samples for DSTC3. It shows that the generated utterances are associated with the atomic exemplars. For some frequent triples, well-formed utterances can be produced. We can find two bad cases at the end of the table. The reason may be that both “request(childrenallowed)” and “request(hastv);request(addr)” never appear in the *dstc3\_seed*, thus the sentence generator doesn't fit well.