

Data Augmentation for Spoken Language Understanding via Pretrained Models

Baolin Peng*, Chenguang Zhu*, Michael Zeng, Jianfeng Gao

Microsoft Research, Redmond

{bapeng, chezhu, nzeng, jfgao}@microsoft.com

Abstract

The training of spoken language understanding (SLU) models often faces the problem of data scarcity. In this paper, we put forward a data augmentation method with pretrained language models to boost the variability and accuracy of generated utterances. Furthermore, we investigate and propose solutions to two previously overlooked scenarios of data scarcity in SLU: i) *Rich-in-Ontology*: ontology information with numerous valid dialogue acts are given; ii) *Rich-in-Utterance*: a large number of unlabelled utterances are available. Empirical results show that our method can produce synthetic training data that boosts the performance of language understanding models in various scenarios.

1 Introduction

Spoken Language Understanding (SLU) aims at converting natural utterances into predefined semantic frames for further steps in dialogue understanding. As a concrete example, an SLU component outputs its prediction of intents and slot labels detected within a user’s utterance. SLU is widely applied in human-machine dialogue interfaces including virtual assistants and robotics (Bellegarda, 2014).

Nevertheless, as a supervised learning task, SLU suffers from the problem of data scarcity. And the problem becomes more protruding in face of new LU domains with novel definitions of intents and slot labels. Even with an existing domain, the data correlated with a certain intent or slot is often not sufficient. These problems significantly limit the applicability of SLU systems.

Recently, various successful use cases of synthetic datasets have stimulated the growth of the area of Data Augmentation (DA) (Lu et al., 2006;

Hannun et al., 2014). The usual approach is to learn a model to mimic the language style in the training data, leveraging the relationship between semantic units and their natural representations. Then, a non-generative model can modify utterances and replace slot labels from existing data (Quan and Xiong, 2019), while a generative model can produce synthetic utterances in the same distribution space of the training data (Hou et al., 2018). However, these approaches usually train the DA model on the domain’s data, which is of a small scale by itself. It is thus questionable whether the augmented data contains rich language expressibility beyond the scope of the given data.

On the other hand, the rapid development of large-scale pretrained language models has significantly improved the capacity of language understanding and generation models (Devlin et al., 2018; Liu et al., 2019; Dong et al., 2019). With a modest amount of domain-specific data, a pretrained model can quickly adapt to a new domain. For instance, SC-GPT finetunes the GPT-2 language model (Radford et al., 2018) with dialogue data (Peng et al., 2020). It can adapt to new dialogue domains very efficiently, with only a couple of labelled data samples.

In this paper, we propose to frame data augmentation as a semantically controlled generation problem. Given semantic labels, we leverage the SC-GPT model to generate corresponding utterances as synthetic training data. In the process, the general language syntax and semantics learned during the pretraining phase are fused into the generation of domain-specific utterances to increase variability and accuracy.

Furthermore, previous literature on SLU data augmentation focus on the case where only a scant number of pairs of utterance and corresponding semantic labels are given, which we denote as *Paired-Data-Only*. However, there are two other

* Equal contribution

overlooked scenarios that commonly arise in application.

- *Rich-in-Ontology*: The full ontology for the dialogue domain is given, including the definitions of intents, slot lists and possible values. Thus, the model is given a variety of valid combinations of semantic labels. What lacks is the corresponding natural utterances.
- *Rich-in-Utterance*: There are abundant historic utterances for a dialogue domain. While labellers can give semantic labels for a small number of these utterances, many utterances are without tagging information.

In this paper, we also investigate these two scenarios and propose data augmentation solutions.

For *Rich-in-Ontology*, we first finetune the pretrained model SC-GPT on the paired training data, and then apply it to the valid combination of intents and slots in the ontology information to generate additional training data.

For *Rich-in-Utterance*, following the idea of NLG model SC-GPT, we propose SC-GPT-NLU, which is pretrained on the same corpus of SC-GPT with flipped sources and targets. In detail, we feed the utterances into the model and let it generate intent and slots in a sequence. Therefore, SC-GPT-NLU can act as a language understanding module and produce semantic labels for the unlabelled utterances.

In the experiments, we evaluate the slot tagging and intent classification accuracies of a Bi-LSTM seq2seq SLU model, using various data augmentation methods. Results show that on ATIS and Snips datasets, our proposed method outperforms other baseline systems. For instance, in ATIS-Small, our method can achieve 0.5 points higher slot F1 and 3.02 points higher intent accuracy. Furthermore, when ontology information or unlabelled utterances are available, our method can produce synthetic data that further boosts the performance of SLU models.

In summary, our contribution in this paper is three-fold:

1. Employ the pretrained NLG model SC-GPT to generate synthetic utterances for SLU;
2. Investigate two previous overlooked scenarios: i) rich ontology information is given, and ii) a large number of unlabelled utterances are available;

3. Propose a pretrained NLU model for spoken language understanding: SC-GPT-NLU, and employ it to produce semantic labels for unlabelled utterances.

2 Related Work

2.1 SLU Data Augmentation

Kurata et al. (2016) proposes to add noise via perturbing the decoder states to generate variants of an utterance. Variational autoencoder (VAE) and conditional variational autoencoder (CVAE) are used to generate utterances with diversified expressions (Li et al., 2019; Yoo et al., 2019; d’Ascoli et al., 2019). Quan and Xiong (2019) uses both non-generative model like word substitution and generative model like paraphrasing and back-translation to augment training data. Hou et al. (2018) proposes a multi-stage framework to generate, filter, and rank augmented utterances. Yin et al. (2019) uses reinforcement learning to learn a generator that facilitates dialogue state tracking. Zhao et al. (2019) employs atomic templates to guide the model to generate more utterances given combination of dialogue acts. Shah et al. (2019) proposes a zero-shot framework to adapt SLU model to a new domain given descriptions and a few slot samples.

2.2 Pretraining

Pretrained models leverage the large amount of unlabelled text corpora to improve the capability of language understanding. ELMo (Peters et al., 2018) applies two unidirectional RNNs for language modeling. GPT and GPT-2 (Radford et al., 2018) utilize the transformer architecture (Vaswani et al., 2017) for the same task. BERT (Devlin et al., 2018) employs the masking technique and next-sentence-prediction task to train a bidirectional language model. UniLM (Dong et al., 2019) uses different masking patterns to unify the model structure for NLU and NLG. These pretrained models have been widely used with considerable success in various NLP applications (Nogueira and Cho, 2019; Zhu et al., 2018; Liu and Lapata, 2019).

3 Data Augmentation

In spoken language understanding (SLU), the training data consists of N sample triples. Each triple contains: i) the tokenized utterance $\mathbf{x} = (x_1, x_2, \dots, x_T)$, ii) the intent label y^i , and iii) the slot label sequence $\mathbf{y} = (y_1, y_2, \dots, y_T)$. The slot labels are in IOB format (Mesnil et al., 2014).

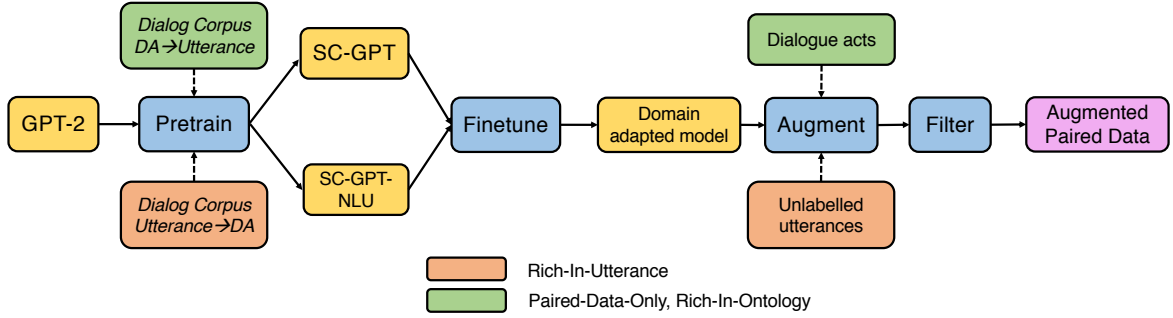


Figure 1: Illustration of data augmentation process. Yellow squares represent models, blue squares represent operations, and dashed arrow represents input data.

As the training data often contains only a few human-labelled samples, data augmentation (DA) is needed. When given a combination of intents and slot-value pairs, a DA method produces a number of corresponding utterances as synthetic data.

Formally, the input to the data augmenter is a *dialog act* \mathcal{A} , defined as the combination of intent \mathbf{I} and P slot-value pairs $\{(s_i, v_i)\}_{i=1}^P$:

$$\mathcal{A} = \left[\underbrace{\mathbf{I}}_{\text{Intent}}, \underbrace{(s_1, v_1), \dots, (s_P, v_P)}_{\text{Slot-value pairs}} \right] \quad (1)$$

The augmenter translates \mathcal{A} into a natural language utterance $\mathbf{x} = (x_1, x_2, \dots, x_T)$. For example, suppose the intent is `hotel-inform`, and the slot-value pairs are $(name = Hyatt, area = center, star = 5)$, then a generated utterance can be “I have booked the 5-star Hyatt hotel in the center area for you.”.

However, as a generation model, the data augmenter requires a considerable amount of training data by itself. Thus, previous data augmenters built upon the small-scale training data is usually limited in model capacity and expressibility.

Therefore, we leverage the pretrained model SC-GPT (Peng et al., 2020). SC-GPT initializes with GPT-2 (Radford et al., 2018), and then pretrains on a large corpus of 400K dialogue-act-utterance pairs. The benefit of SC-GPT is that it can quickly adapt to new domains with only a few domain-specific data samples. Thus, we finetune SC-GPT on the intent and slot-value labels in the training data.

Next, we augment the dialogue acts in the training data by replacing/inserting/deleting slot values to create more combinations. The finetuned model then generates multiple candidate utterances for each dialogue act. As the generated utterances may not always contain the required slot-value labels, we filter them to make sure that each utterance has

all the required input slot-values. More details are given in Section 4.2.

3.1 SLU Data Augmentation Scenarios

Previous literature on SLU data augmentation usually works for the case where only a scant number of pairs of utterance and dialogue acts, which we denote as **Paired-Data-Only**. Nevertheless, we propose that there are two more overlooked scenarios that commonly arise in application.

Rich-In-Ontology. In many cases, a detailed description of ontology of a dialogue domain is also given. Thus, the model is exposed to more valid dialogue acts. Formally, the training data consists of both labelled pairs and many dialogue acts: $\{(\mathbf{x}_1, \mathcal{A}_1), \dots, (\mathbf{x}_N, \mathcal{A}_N), \mathcal{A}_{N+1}, \dots, \mathcal{A}_M\}$.

Under this scenario, we employ the above finetuned SC-GPT model and generate utterances given these auxiliary dialogue acts.¹

Rich-In-Utterance. It is common in practice that a large number of unlabelled dialogue utterances are available, usually collected from history data. Formally, the training data consists of both labelled pairs and many unlabelled utterances: $\{(\mathbf{x}_1, \mathcal{A}_1), \dots, (\mathbf{x}_N, \mathcal{A}_N), \mathbf{x}_{N+1}, \dots, \mathbf{x}_M\}$.

To utilize these utterances, we need to produce corresponding dialogue acts. We propose to finetune pretrained language model in a reverse way: feed an utterance as input and let the model generate the dialogue act. In other words, we leverage a generative model to act as a language understanding module, denoted as SC-GPT-NLU.

SC-GPT-NLU is initialized with GPT-2 and then pretrains on the same data as SC-GPT. However,

¹In the experiment, as the training data S is constrained as a subset of the original training set T , we simulate the Rich-In-Ontology scenario by sampling dialogue acts from the whole training set T .

| Dataset | ATIS | | | | Snips | | | |
|-----------------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| Split | Small | | Medium | | Small | | Medium | |
| Model | Slot F1 | Intent Acc. | Slot | Intent | Slot | Intent | Slot | Intent |
| <i>No Data Augmentation</i> | | | | | | | | |
| No-DA | 68.91 | 84.99 | 87.30 | 90.15 | 61.30 | 93.43 | 79.83 | 97.29 |
| <i>Paired-Data-Only</i> | | | | | | | | |
| Seq2Seq | 73.71 | - | 88.72 | - | - | - | - | - |
| VAE | 74.92 | 83.65 | <u>89.27</u> | <u>90.95</u> | - | - | - | - |
| Ours | <u>75.42</u> | <u>86.67</u> | 88.61 | 90.71 | 64.96 | 93.43 | 80.62 | 97.57 |
| <i>Rich-in-Ontology</i> | | | | | | | | |
| Ours | 82.42 | 89.03 | 89.81 | 92.27 | 67.06 | 94.14 | 82.54 | 97.86 |
| <i>Rich-in-Utterance</i> | | | | | | | | |
| Ours | 78.45 | 87.46 | 88.23 | 91.94 | 63.46 | 93.43 | 80.54 | 98.14 |

Table 1: Slot F1 and intent accuracy scores on ATIS and Snips dataset. The overall highest score is in bold, and the best result in Paired-Data-Only category is underlined.

we flip the dialogue acts and utterances so that SC-GPT-NLU can produce dialogue acts given an input utterance. Then, SC-GPT-NLU augments data by generating dialogue acts for the utterances $\mathbf{x}_{N+1}, \dots, \mathbf{x}_M$.

We illustrate our data augmentation process for SLU in Figure 1.

4 Experiments

4.1 Datasets

We employ the widely used SLU benchmark dataset ATIS (Tur et al., 2010) and Snips (Coucke et al., 2018). ATIS contains around 5.8K utterances from flight reservation dialogues. It includes 120 slot labels and 21 intent types. Snips contains 14K utterances from the Snips personal voice assistant. It includes 123 slot labels and 7 intent types.

To simulate the few-shot data situations, we follow Chen et al. (2016) to use two small portions of the ATIS training set as training data: Small ($\sim 1/40$ of the original training set with 129 instances) and Medium ($\sim 1/10$ of the original training set with 515 instances). A development set of 500 instances is used. Following the same split ratio, we sampled 327 and 1308 instances in Snips for Small and Medium respectively.

4.2 Models

SLU Model. All the data augmentation methods in the experiments share the same SLU model for slot tagging and intent classification, which is a bi-directional LSTM with the same hyperparameter settings as in Hou et al. (2018).

Data augmentation. For Paired-Data-Only case, we modify the dialogue acts in the training split to construct around 300 additional combinations of DAs via dropping/inserting/replacing slots and values. For each dialogue act, we sample three utterances produced by SC-GPT. After filtering out utterances which do not contain all the slot-values, we collect around 500 synthetic utterances and add them into the original training split.

For Rich-in-Ontology case, we similarly augment 500 utterances, except that the dialogue acts are sampled from the whole training corpus.

For Rich-in-Utterance case, we sample 1,000 utterances in the training corpus and use SC-GPT-NLU to produce the most probable dialogue act. After filtering, around 500 utterance-DA pairs are added to the original training split.

We show some example generations of SC-GPT and SC-GPT-NLU in Table 2.

Baselines. The baseline systems include the seq2seq (Hou et al., 2018) and variational autoencoder (VAE) data augmentation model (Yoo et al., 2019). We also report the results for the case without data augmentation, denoted by No-DA.

Evaluation. We use F1 score to measure slot tagging quality and accuracy score to evaluate intent classification, in accordance with Chen et al. (2016).

4.3 Results

Table 1 shows the accuracy of slot tagging and intent classification for various models.

Firstly, our data augmentation method can considerably boost the model accuracy (comparing

| SC-GPT | |
|--------------------|---|
| DA | RateBook (best_rating = 6; object_select = current; object_type = textbook; rating_value = 3) |
| Utterance 1 | Give 3 out of 6 to current textbook |
| Utterance 2 | The current textbook gets a 3 out of 6 |
| Utterance 3 | I think that the current textbook should be rated 3 out of 6 |
| DA | BookRestaurant (country = Honduras; facility = indoor; restaurant_type = restaurant) |
| Utterance 1 | Book me a reservation for an indoor restaurant in Honduras |
| Utterance 2 | Book an indoor restaurant in Honduras |
| Utterance 3 | I need to book an indoor restaurant in Honduras |
| SC-GPT-NLU | |
| Utterance | 2 of us want to eat at a restaurant that serves meatballs in VT |
| DA | BookRestaurant (party_size_number = 2; restaurant_type = restaurant; served_dish = meatballs; state = VT) |
| Utterance | Add the track to the Metal Talks Metallica playlist. |
| DA | AddToPlaylist (music_item = track; playlist = metal talks Metallica) |

Table 2: Example utterances generated by SC-GPT given dialogue acts, and dialogue acts generated by SC-GPT-NLU given unlabelled utterances.

No-DA and Ours), especially when the training data size is small. For instance, in ATIS, when only paired data is available, the slot F1 increases 6.51 (Small) and 1.31 (Medium) points, while the intent accuracy increases 1.68 (Small) and 0.56 (Medium) points.

Secondly, under Rich-in-Ontology and Rich-in-Utterance scenarios, our method further boosts the slot F1 by up to 7 points and intent accuracy by up to 2.4 points. Overall, the accuracy scores are the highest when the ontology information is available. This shows that our method can take advantage of additional information and produce better synthetic training data for downstream models.

Thirdly, under the traditional Paired-Data-Only scenario, our data augmentation method outperforms all baselines in ATIS-Small, and achieves comparable results in ATIS-Medium. This shows that our method is better suited when training data is scarcer.

5 Conclusion

In this paper, we approach the problem of data scarcity in SLU with a pretrained language model SC-GPT. After finetuning with domain-specific dialogue data, it can produce abundant utterances which significantly boosts the performance of SLU model. Moreover, we provide solutions to two previously overlooked scenarios in SLU data augmentation: *Rich-in-Ontology* and *Rich-in-Utterance*. These solutions can effectively leverage the auxiliary data to produce high-quality synthetic training data. Specially, we propose a pretrained LU model SC-GPT-NLU to produce dialogue acts given utterances.

References

- Jerome R Bellegarda. 2014. Spoken language understanding for natural interaction: The siri experience. In *Natural interaction with robots, knowbots and smartphones*, pages 3–14. Springer.
- Yun-Nung Chen, Dilek Hakanni-Tür, Gokhan Tur, Asli Celikyilmaz, Jianfeng Guo, and Li Deng. 2016. Syntax or semantics? knowledge-guided joint semantic frame parsing. In *2016 IEEE Spoken Language Technology Workshop (SLT)*, pages 348–355. IEEE.
- Alice Coucke, Alaa Saade, Adrien Ball, Théodore Bluche, Alexandre Caulier, David Leroy, Clément Doumouro, Thibault Gisselbrecht, Francesco Caltagirone, Thibaut Lavril, et al. 2018. Snips voice platform: an embedded spoken language understanding system for private-by-design voice interfaces. *arXiv preprint arXiv:1805.10190*.
- Stéphane d’Ascoli, Alice Coucke, Francesco Caltagirone, Alexandre Caulier, and Marc Lelarge. 2019. Conditioned query generation for task-oriented dialogue systems. *arXiv preprint arXiv:1911.03698*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. 2019. Unified language model pre-training for natural language understanding and generation. *arXiv preprint arXiv:1905.03197*.
- Awni Hannun, Carl Case, Jared Casper, Bryan Catanzaro, Greg Diamos, Erich Elsen, Ryan Prenger, Sanjeev Satheesh, Shubho Sengupta, Adam Coates, et al. 2014. Deep speech: Scaling up end-to-end speech recognition. *arXiv preprint arXiv:1412.5567*.
- Yutai Hou, Yijia Liu, Wanxiang Che, and Ting Liu. 2018. Sequence-to-sequence data augmentation for dialogue language understanding. *arXiv preprint arXiv:1807.01554*.

- Gakuto Kurata, Bing Xiang, and Bowen Zhou. 2016. Labeled data generation with encoder-decoder lstm for semantic slot filling.
- Juntao Li, Lisong Qiu, Bo Tang, Dongmin Chen, Dongyan Zhao, and Rui Yan. 2019. Insufficient data can also rock! learning to converse using smaller data with augmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6698–6705.
- Yang Liu and Mirella Lapata. 2019. Text summarization with pretrained encoders. *EMNLP*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Xinghua Lu, Bin Zheng, Atulya Velivelli, and Chengxiang Zhai. 2006. Enhancing text categorization with semantic-enriched representation and training data augmentation. *Journal of the American Medical Informatics Association*, 13(5):526–535.
- Grégoire Mesnil, Yann Dauphin, Kaisheng Yao, Yoshua Bengio, Li Deng, Dilek Hakkani-Tur, Xiaodong He, Larry Heck, Gokhan Tur, Dong Yu, et al. 2014. Using recurrent neural networks for slot filling in spoken language understanding. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 23(3):530–539.
- Rodrigo Nogueira and Kyunghyun Cho. 2019. Passage re-ranking with bert. *arXiv preprint arXiv:1901.04085*.
- Baolin Peng, Chenguang Zhu, Chunyuan Li, Xijun Li, Jinchao Li, Michael Zeng, and Jianfeng Gao. 2020. Few-shot natural language generation for task-oriented dialog. *arXiv preprint arXiv:2002.12328*.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. *arXiv preprint arXiv:1802.05365*.
- Jun Quan and Deyi Xiong. 2019. Effective data augmentation approaches to end-to-end task-oriented dialogue. *arXiv preprint arXiv:1912.02478*.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training.
- 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, Dilek Hakkani-Tür, and Larry Heck. 2010. What is left to be understood in atis? In *2010 IEEE Spoken Language Technology Workshop*, pages 19–24. IEEE.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, pages 5998–6008.
- Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, and Qun Liu. 2019. Dialog state tracking with reinforced data augmentation. *arXiv preprint arXiv:1908.07795*.
- Kang Min Yoo, Youhyun Shin, and Sang-goo Lee. 2019. Data augmentation for spoken language understanding via joint variational generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 7402–7409.
- Zijian Zhao, Su Zhu, and Kai Yu. 2019. Data augmentation with atomic templates for spoken language understanding. *arXiv preprint arXiv:1908.10770*.
- Chenguang Zhu, Michael Zeng, and Xuedong Huang. 2018. Sdnet: Contextualized attention-based deep network for conversational question answering. *arXiv preprint arXiv:1812.03593*.