SLIM: Explicit Slot-Intent Mapping with BERT for Joint Multi-Intent Detection and Slot Filling

Fengyu Cai, Wanhao Zhou, Fei Mi and Boi Faltings LIA, EPFL, Lausanne, Switzerland

{fengyu.cai,wanhao.zhou,fei.mi,boi.faltings}@epfl.ch

Abstract

Utterance-level intent detection and tokenlevel slot filling are two key tasks for natural language understanding (NLU) in taskoriented systems. Most existing approaches assume that only a single intent exists in an utterance. However, there are often multiple intents within an utterance in real-life scenarios. In this paper, we propose a multi-intent NLU framework, called SLIM, to jointly learn multiintent detection and slot filling based on BERT. To fully exploit the existing annotation data and capture the interactions between slots and intents, SLIM introduces an explicit slot-intent classifier to learn the many-to-one mapping between slots and intents. Empirical results on three public multi-intent datasets demonstrate (1) the superior performance of SLIM compared to the current state-of-the-art for NLU with multiple intents and (2) the benefits obtained from the slot-intent classifier.

1 Introduction

Natural language understanding (NLU) is an essential task for task-oriented dialog (ToD) systems. It contains two major sub-tasks, intent detection (ID) and slot filling (SF). Take "Listen to Westbam album Allergic on Google music" as an example. The task of ID is to identify the intent ("Play-Music") of the utterance. SF is a sequence labeling task to predict the slot for each token, which are [O, O, B-artist, O, B-album, O, B-service, I-service] using the "Insideoutside-begining" (IOB) tagging format (Zhang and Wang, 2016). Traditional techniques often tackle these two tasks separately (Haffner et al., 2003; Kim, 2014; Yao et al., 2014; Yang et al., 2016; Vu, 2016). Recently, models that learn these two tasks jointly achieve better performance by capturing semantic dependencies between intent and slots (Liu and Lane, 2016; Goo et al., 2018; Liu et al., 2019; Chen et al., 2019; Zhang et al., 2019; E et al., 2019; Qin et al., 2019).



(a) Gangadharaiah and Narayanaswamy (2019); Qin et al.(2020): Utterance-level intent is shared by all the slots.



(b) SLIM: Explicit slot-level intent is used by each slot.

Figure 1: Comparison between approaches of utilizing intent information for different slots.

The aforementioned methods assume only a single intent in an utterance. However, multiple intents often exist in one utterance in real-life scenarios (Kim et al., 2016; Gangadharaiah and Narayanaswamy, 2019; Qin et al., 2020; Eric et al., 2020). Regardless of single or multiple intents, the relationship between slots and intent(s) is manyto-one. In single-intent utterances, all slots share the same intent; in multi-intent scenarios, however, different slots may correspond to different intents. In Figure 1, the slot "south carolina" corresponds to GetWeather, while slots "record" and "michael jackson" should be linked to PlayMusic. Existing multi-intent NLU models (Gangadharaiah and Narayanaswamy, 2019; Qin et al., 2020) inherit the single-intent pattern, i.e., the utterance-level intent label distribution is shared by all the slots (Figure 1a). Therefore, the relationship between slots and intents are not utilized even though they are already annotated in most NLU datasets.

To this end, we propose SLIM (**SLot-Intent M**apping) that fully exploits the annotations to explicitly captures the mapping between slots and intents. We add a slot-intent classifier to predict the intent label for each slot of an utterance (Figure 1b) on top of the state-of-the-art pre-trained model (Chen et al., 2019) based on BERT (Devlin

et al., 2019) to jointly tackle ID and SF. In experiments, we compare SLIM to a wide range of NLU techniques on two simulated datasets (MixATIS and MixSNIPS Qin et al. (2020)) and a real-world dataset (DSTC4 Kim et al. (2016)). Empirical results demonstrate that SLIM achieves better performance compared to the current state-of-the-art *without* using extra annotation data. We analyze and reveal that the explicit slot-intent mapping module indeed helps the model learn faster and better. To our best knowledge, this is the first study to explicitly link slots to their intents for better NLU in task-oriented dialog systems.

2 Approach

2.1 Problem Setting

For an input utterance x with token sequence $x = (x_1, ..., x_n)$, the multi-intent NLU task is composed of (1) *utterance-level intent detection*: predict the multi-label intents $I_x \subset I$ of the utterance, where I is the set of possible intents, and (2) *token-level slot filling*: predict the slot label for each token of the input utterance from a set T of possible slots. Different from existing single-intent NLU task, the following two *assumptions* are made:

- (a) An utterance x has at least one utterancelevel intent, i.e, $|I_x| \ge 1$.
- (b) Each slot s_m = {x_{m1},...,x_{mj}} is a set of j tokens, and it is mapped to a specific slot-level intent **belonging to** the utterance-level intents, i.e, i_m ∈ I_x.

2.2 Model

The proposed model (SLIM) contains a shared encoder and three classifiers for different tasks.

Encoder We use BERT (Devlin et al., 2019) as token sequence encoder of our model. The utterance is tokenized by standard BERT tokenizer with a special token [CLS] prepended and [SEP] appended. The output $h = (h_{cls}, h_1, ..., h_n, h_{sep})$ of BERT's encoder will be utilized for three classifiers, where $h_k \in \mathbb{R}^d$.

Intent Classifier and Slot Classifier Utterancelevel intent detection is accomplished by the intent classifier. To classify the intent y^i , we use sigmoid as the activation function after feeding h_{cls} into a output network by:

$$y^{i} = Sigmoid(W^{i}h_{cls} + b^{i}), \qquad (1)$$

where $W^i \in \mathbb{R}^{|I| \times d}$, and each dimension of $y^i \in \mathbb{R}^{|I|}$ represents the probability of an intent label. In the slot filling task, to predict the slot y_k^s at position k, we apply softmax after feeding h_k into a separate slot classification network as:

$$y_k^s = Softmax(W^s h_k + b^s), \tag{2}$$

where $W^s \in \mathbb{R}^{|T| \times d}$, and $y_k^s \in \mathbb{R}^{|T|}$ is the slot probability distribution for token x_k with $k \in \{1, \ldots, n\}$. Intent classifier and slot classifier are formulated similarly as Chen et al. (2019), except that intent detection is formulated as a multi-label classification task.

Slot-Intent Classifier To explicitly capture the relation between slots and intents, we predict the slot-level intent for each slot s_m . First, we compute the representation r_m of a slot $s_m = \{x_{m_1}, ..., x_{m_j}\}$ by a *mean pooling* of the token representations in this slot by: $r_m = 1/j \sum_{i=1}^j h_{m_i}$. Afterward, we concatenate the global utterance representation h_{cls} with r_m and compute an **unconstrained** slot-intent prediction y_m^l as:

$$y_m^l = Softmax(W^l[h_{cls}|r_m] + b^l), \quad (3)$$

where $W^l \in \mathbb{R}^{|I| \times 2d}$, and | indicates vector concatenation. To better align the above slot-intent prediction with the predicted utterance intent (assumption (b)), we propose an *intent-constrained attention*. It computes a final **constrained** slot-intent prediction y_m^p by an element-wise multiplication between the utterance-level intent prediction y_m^i and unconstrained slot-intent prediction y_m^l as:

$$y_m^p = y^i \odot y_m^l \tag{4}$$

2.3 Training Objective

The global objective of SLIM is to maximize $p(y^i, y^s, y^p | x)$, and it can be decomposed for tokens inside and outside of slots as:

$$\prod_{m} \underbrace{p(y^{i}, y^{s}_{s_{m}}, y^{p}_{m} | x)}_{\text{inside each slot}} \prod_{j \notin \cup_{m} s_{m}} p(y^{i}, y^{s}_{j} | x) \quad (5)$$

$$= \prod_{m} p(y^{i}, y^{s}_{s_{m}} | x) \cdot p(y^{p}_{m} | y^{i}, y^{s}_{s_{m}}, x) \cdot \Gamma$$
(6)

$$\propto p(y^i|x) \prod_{k=1}^n p(y^s_k|x) \prod_m p(y^p_m|y^i, y^s, x)$$
(7)

$$= \underbrace{p(y^{i}|x)}_{\text{ID}} \underbrace{\prod_{k=1}^{n} p(y^{s}_{k}|x)}_{\text{SF}} \underbrace{\prod_{m} p(y^{p}_{m}|y^{i}, y^{s}_{s_{m}}, x)}_{\text{Slot-Intent Classification}}$$
(8)



Figure 2: Illustration of our model **SLIM**. The lower part is the BERT encoder. On the top, from left to right, they are the intent classifier (for ID), the slot-intent classifier, and the slot classifier (for SF) respectively.

In equation 8, the first two terms are the objectives of ID and SF, and the last term is the objective of slot-intent mapping. ID is trained with binary cross-entropy loss for multi-intent detection, while the two other terms are trained with regular cross-entropy loss. Because the slot-intent prediction y_m^p is conditional on the predicted utterance intent y^i and slot label $y_{s_m}^s$ of the slot, parameters for ID and SF will also be updated when training the slot-intent classifier. The loss of SLIM is a weighted sum of losses from these three classifiers. An overview model pipeline of SLIM is illustrated in Figure 2.

3 Experiment and Analysis

Dataset We evaluate our approach on three multiintent SLU datasets. **MixSNIPS** (Qin et al., 2020) contains 39,776/2,198/2,199 utterances for train/validation/test. It is created based on the Snips personal voice assistant (Coucke et al., 2018). **MixATIS** (Qin et al., 2020) is constructed from ATIS dataset (Hemphill et al., 1990), containing 13,161/759/828 utterances for train/validation/test. **DSTC4** (Kim et al., 2016) contains multi-intent human-human dialogues with 5,308/2,098/1,865 utterances for train/validation/test. More details on these datasets are provided in Appendix A.4.

Training Details We use English uncased BERT-Base model, containing 12 layers, 768 hidden states, and 12 heads. The max sequence length is 50, and training batch size is 32. Hyper-parameters of SLIM are tuned by a randomized search (see Appendix A.3) based on the semantic frame accuracy on validation set. Dropout rate is 0.2 for the output layers of all three classifiers. In Eq. (8), the losses of ID and slot-intent classifier are weighted by 1, and the loss of SF is weighted by 2. We train SLIM for maximum 20 epochs, and the early stop patience is 3 epochs. For the slot-intent classifier, we use slots provided by ground truth during training, while the predicted slots are used during inference.

Baselines We compare SLIM with both singleintent and multi-intent models. When evaluating single-intent models on the multi-intent task, we follow Gangadharaiah and Narayanaswamy (2019); Qin et al. (2020) to concatenate multiple intents with '#' into a single intent. On MixATIS and MixSNIPS, we compare previous single-intent models: Bi-Model (Wang et al., 2018), SF-ID (E et al., 2019), Stack-Propagation (Qin et al., 2019) and recent multi-intent models: Joint Multiple ID-SF (Gangadharaiah and Narayanaswamy, 2019), AGIF (Qin et al., 2020). On DSTC4, we compare SLIM with state-of-the-art single-intent model Stack-Propagation (Qin et al., 2019) and multiintent model AGIF (Qin et al., 2020).

Overall Results We evaluate the performance of different methods using three metrics: F1 score of slot filling (Slot F1), accuracy of intent detection (Intent Acc), and semantic frame accuracy (SeFr Acc) as in Gangadharaiah and Narayanaswamy

Model		MixATIS		MixSNIPS			
Widdel	Slot F1 Intent Acc		SeFr Acc	Slot F1	Intent Acc	SeFr Acc	
Bi-Model [†] (Wang et al., 2018)	85.5	72.3	39.1	86.8	95.3	53.9	
SF-ID [†] (E et al., 2019)	87.7	63.7	36.2	89.6	96.3	59.3	
Stack-Propagation ^{†§} (Qin et al., 2019)	86.6	76.0	42.8	93.9	96.4	75.5	
Joint Multiple ID-SF [‡] (2019)	87.5	73.1	38.1	91.0	95.7	66.6	
AGIF ^{‡§} (Qin et al., 2020)	88.1	75.8	44.5	94.5	96.5	76.4	
SLIM (w/o slot-intent classifier)	85.6	77.1	46.3	96.2	96.8	82.3	
SLIM (w/o intent-constrained attention)	87.2	75.6	46.4	96.5	96.0	83.6	
SLIM	88.5	78.3	47.6	96.5	97.2	84.0	

Table 1: Results on two multi-intent datasets. [†] or [‡] denotes single-intent or multi-intent model respectively. Results of models with [†] and [‡] are taken from Qin et al. (2020). Models with [§] indicate the previous state-of-the-art solutions. **Bold** numbers are the best results in each column.

Model	DSTC4						
WIGUEI	Slot F1	Intent Acc	SeFr Acc				
Stack-Propagation [†]	56.4	35.8	20.7				
$AGIF^{\ddagger}$	57.6	33.0	19.4				
SLIM	61.1	36.7	21.3				

Table 2: Results on DSTC4. † or ‡ denotes single-intent or multi-intent model respectively. **Bold** numbers are the best results in each column.

(2019) and Qin et al. (2020). "SeFr Acc" considers the prediction of an utterance to be correct when slots and intents are all accurate. Table 1 summarizes different models' performances on Mix-ATIS and MixSNIPS. We observe that SLIM outperforms other baselines w.r.t. all three metrics, especially on the semantic frame accuracy. On Slot F1 and Intent Acc, SLIM outperforms the previous start-of-the-art (AGIF) by 0.4%-2.5%. More importantly, SLIM improves the semantic frame accuracy compared to AGIF by 3.1% and 7.6% on two datasets respectively. This result indicates that SLIM effectively improves the joint correctness for predicting intents and slots for understanding an utterance. Table 2 summarizes the results on the realworld DSTC4 dataset, and similar performance patterns can be observed as in Table 1. Moreover, SLIM also notably improves Slot F1 on this dataset (3.5% and 4.7% gain over the two baselines), indicating that SLIM can also strengthen token-level prediction in addition to the large utterance-level improvements in Table 1.

Ablation Study We compare SLIM with two simplified versions, *w/o slot-intent classifier*¹ and *w/o intent-constrained attention*, to analyze how the slot-intent classifier and its intent-constrained attention affect the performance and the training process. From Table 1 (Bottom), we can see that



Figure 3: Intent Acc and SeFr Acc on MixATIS validation dataset after each training epoch.

dropping these two components impairs the performance of SLIM. Besides the lower semantic frame accuracy, dropping slot-intent classifier mainly degrades Slot F1, and dropping intent-constrained attention mainly degrades Intent Acc. Furthermore, we plot the learning curves of these three methods on MixATIS's validation set in Figure 3. We can observe that SLIM converges *faster and better* compared to the two simplified versions, which demonstrates the benefits of the slot-intent classifier with intent-constrained attention.

4 Conclusion

In this paper, we propose SLIM for multi-intent NLU in task-oriented dialog systems. SLIM explicitly utilizes the relation between slots and intents by mapping slots to the corresponding intent. Experimental results on three public datasets show that SLIM outperforms the previous state-of-theart NLU models. Our findings may inspire future studies to better exploit the relationship between slots and intent for complicated NLU scenarios in task-oriented dialog systems.

¹This version degenerates to Chen et al. (2019) as mentioned in Sec. 2.2

References

- Qian Chen, Zhu Zhuo, and Wen Wang. 2019. BERT for joint intent classification and slot filling. *CoRR*, abs/1902.10909.
- Alice Coucke, Alaa Saade, Adrien Ball, Théodore Bluche, Alexandre Caulier, David Leroy, Clément Doumouro, Thibault Gisselbrecht, Francesco Caltagirone, Thibaut Lavril, Maël Primet, and Joseph Dureau. 2018. Snips voice platform: an embedded spoken language understanding system for privateby-design voice interfaces. *CoRR*, abs/1805.10190.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In NAACL-HLT (1), pages 4171–4186. Association for Computational Linguistics.
- Haihong E, Peiqing Niu, Zhongfu Chen, and Meina Song. 2019. A novel bi-directional interrelated model for joint intent detection and slot filling. In *ACL* (1), pages 5467–5471. Association for Computational Linguistics.
- Mihail Eric, Rahul Goel, Shachi Paul, Abhishek Sethi, Sanchit Agarwal, Shuyang Gao, Adarsh Kumar, Anuj Kumar Goyal, Peter Ku, and Dilek Hakkani-Tür. 2020. Multiwoz 2.1: A consolidated multidomain dialogue dataset with state corrections and state tracking baselines. In *LREC*, pages 422–428. European Language Resources Association.
- Rashmi Gangadharaiah and Balakrishnan Narayanaswamy. 2019. Joint multiple intent detection and slot labeling for goal-oriented dialog. In NAACL-HLT (1), pages 564–569. Association for Computational Linguistics.
- Chih-Wen Goo, Guang Gao, Yun-Kai Hsu, Chih-Li Huo, Tsung-Chieh Chen, Keng-Wei Hsu, and Yun-Nung Chen. 2018. Slot-gated modeling for joint slot filling and intent prediction. In *NAACL-HLT (2)*, pages 753–757. Association for Computational Linguistics.
- Patrick Haffner, Gokhan Tur, and Jerry H Wright. 2003. Optimizing svms for complex call classification. In 2003 IEEE International Conference on Acoustics, Speech, and Signal Processing, 2003. Proceedings.(ICASSP'03)., volume 1, pages I–I. IEEE.
- Charles T. Hemphill, John J. Godfrey, and George R. Doddington. 1990. The ATIS spoken language systems pilot corpus. In *HLT*. Morgan Kaufmann.
- Seokhwan Kim, Luis Fernando D'Haro, Rafael E. Banchs, Jason D. Williams, and Matthew Henderson. 2016. The fourth dialog state tracking challenge. In *IWSDS*, volume 427 of *Lecture Notes in Electrical Engineering*, pages 435–449. Springer.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *EMNLP*, pages 1746– 1751. ACL.

- Bing Liu and Ian Lane. 2016. Attention-based recurrent neural network models for joint intent detection and slot filling. In *INTERSPEECH*, pages 685–689. ISCA.
- Yijin Liu, Fandong Meng, Jinchao Zhang, Jie Zhou, Yufeng Chen, and Jinan Xu. 2019. Cm-net: A novel collaborative memory network for spoken language understanding. In *EMNLP/IJCNLP* (1), pages 1051– 1060. Association for Computational Linguistics.
- Libo Qin, Wanxiang Che, Yangming Li, Haoyang Wen, and Ting Liu. 2019. A stack-propagation framework with token-level intent detection for spoken language understanding. In *EMNLP/IJCNLP* (1), pages 2078–2087. Association for Computational Linguistics.
- Libo Qin, Xiao Xu, Wanxiang Che, and Ting Liu. 2020. Towards fine-grained transfer: An adaptive graphinteractive framework for joint multiple intent detection and slot filling. In *EMNLP (Findings)*, pages 1807–1816. Association for Computational Linguistics.
- Ngoc Thang Vu. 2016. Sequential convolutional neural networks for slot filling in spoken language understanding. In *INTERSPEECH*, pages 3250–3254. ISCA.
- Yu Wang, Yilin Shen, and Hongxia Jin. 2018. A bimodel based RNN semantic frame parsing model for intent detection and slot filling. In *NAACL-HLT (2)*, pages 309–314. Association for Computational Linguistics.
- Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alexander J. Smola, and Eduard H. Hovy. 2016. Hierarchical attention networks for document classification. In *HLT-NAACL*, pages 1480–1489. The Association for Computational Linguistics.
- Kaisheng Yao, Baolin Peng, Yu Zhang, Dong Yu, Geoffrey Zweig, and Yangyang Shi. 2014. Spoken language understanding using long short-term memory neural networks. In *SLT*, pages 189–194. IEEE.
- Chenwei Zhang, Yaliang Li, Nan Du, Wei Fan, and Philip S. Yu. 2019. Joint slot filling and intent detection via capsule neural networks. In *ACL (1)*, pages 5259–5267. Association for Computational Linguistics.
- Xiaodong Zhang and Houfeng Wang. 2016. A joint model of intent determination and slot filling for spoken language understanding. In *IJCAI*, pages 2993– 2999. IJCAI/AAAI Press.

	Token Sequence x	[CLS] what	tis	the	weather	for	south	carolina	and	play	а	record	by	michael	jackson	[SEP]
	Slot Token Sequence	0 0	0	0	0	0	B-state	I-state	0	0	0	B-music	0	B-artist	I-artist	0
	Slot Representation Mask 1	0 0	0	0	0	0	1/2	1/2	0	0	0	0	0	0	0	0
	Slot Representation Mask 2	0 0	0	0	0	0	0	0	0	0	0	1	0	0	0	0
	Slot Representation Mask 3	0 0	0	0	0	0	0	0	0	0	0	0	0	1/2	1/2	0
	Slot Representation Mask 4	0 0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Padding Masks 🚽	Slot Representation Mask 5	0 0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	Slot Representation Mask 6	0 0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Figure 4: Illustration of our slot representation mechanism. In the example sentence above, we have in three actual slots and three additional padding slots. The representation r_m of *m*-th slot, where $m \in \{1, \ldots, 6\}$, is the average of hidden states of tokens inside the slot.

Hyper-parameter	Search Range
Dropout Rate	{0, 0.1, 0.2 , 0.3, 0.4}
Learning Rate	{1e-5, 5e-5 , 1e-4, 5e-4}
Loss Weight of Intent Classifier	{0.5, 1, 2}
Loss Weight of Slot classifier	{0.5, 1, 2 }
Loss Weight of Slot-Intent Classifier	{0.5, 1, 2}

Table 3: Hyper-parameter search range of our proposed **SLIM** model. **Bold** numbers indicate our choice of hyper-parameters.

Appendix

A Reproducibility Checklist

A.1 Computing Infrastructure and Computation Time

All experiments are conducted using a single GeForce GTX TITAN X GPU. When training for 10 epochs, time costs are approximately 1 hour, 45 minutes and 30 minutes on MixSNIPS, MixATIS, and DSTC4 respectively.

A.2 Number of Parameters

SLIM includes one BERT encoder and three classifiers. Compared with BERT, the parameter size of three output networks is much smaller and is dependent on the number of intent labels (|I|) and slot labels (|T|) in the target dataset. Therefore, the number of parameters in SLIM is marginally larger than BERT, which is around 110 million.

A.3 Hyper-parameter Search

In total, we have 5 hyper-parameters to configure: dropout rate, learning rate, the weight of losses for intent classifier, slot classifier and slot-intent classifier. We randomized search with 30 trials for the best setting to maximize the semantic frame accuracy on validation set. Detailed search range and our choice of hyper-parameters are given in Table 3.

Dataset	# of Intent Labels	# of Slot Labels
MixSNIPS	7	71
MixATIS	21	118
DSTC4	20	90

Table 4: Summary of the number (|I|) of intent labels and the number (|T|) of slot labels in MixSNIPS, Mix-ATIS and DSTC4.

# of Sentences # of Intent(s)	1	2	2	4
Dataset	1	2	3	-
MixSNIPS	8,277	22,499	9,000	-
MixATIS	1,118	8,444	3,599	-
DSTC4	3,963	1,240	100	5

Table 5: Summary of number of utterances with different numbers of intents in MixSNIPS, MixATIS and DSTC4.

A.4 Detailed Dataset Specifications

Table 4 summarizes the number of intent labels and slot labels in the training set of MixSNIPS, MixATIS, and DSTC4. Table 5 reports the statistics of number of intents in utterances. For DSTC4, we split the training and validation data, because the original test data are no longer accessible.

B Implementation Details on Slot Representation

Since the number of slots varies between utterances, to keep the SLIM model running in batch mode, we pad the number of slots in each utterance to 6. We observe that 6 is the 99-th percentile of the number of slots in an utterance in real-world multi-intent dataset DSTC4. We pad slots in each utterance to 6 in order to strike a balance between the memory consumption and coverage of most slots in utterances. An illustrative example is given in Figure 4. Each mask extracts a different slot in the utterance. The actual number of slots in this utterance is three, and we have three more padding masks with all zeros.