

GL-GIN: Fast and Accurate Non-Autoregressive Model for Joint Multiple Intent Detection and Slot Filling

Libo Qin, Fuxuan Wei, Tianbao Xie, Xiao Xu, Wanxiang Che*, Ting Liu

Research Center for Social Computing and Information Retrieval

Harbin Institute of Technology, China

{lbqin, fuxuanwei, tianbaoxie, xxu, car, tliu}@ir.hit.edu.cn

Abstract

Multi-intent SLU can handle multiple intents in an utterance, which has attracted increasing attention. However, the state-of-the-art joint models heavily rely on autoregressive approaches, resulting in two issues: *slow inference speed* and *information leakage*. In this paper, we explore a non-autoregressive model for joint multiple intent detection and slot filling, achieving more fast and accurate. Specifically, we propose a **Global-Locally Graph Interaction Network (GL-GIN)** where a local slot-aware graph interaction layer is proposed to model slot dependency for alleviating uncoordinated slots problem while a global intent-slot graph interaction layer is introduced to model the interaction between multiple intents and all slots in the utterance. Experimental results on two public datasets show that our framework achieves state-of-the-art performance while being 11.5 times faster.

1 Introduction

Spoken Language Understanding (SLU) (Young et al., 2013) is a critical component in spoken dialog systems, which aims to understand user’s queries. It typically includes two sub-tasks: intent detection and slot filling (Tur and De Mori, 2011).

Since intents and slots are closely tied, dominant single-intent SLU systems in the literature (Goo et al., 2018; Li et al., 2018; Liu et al., 2019b; E et al., 2019; Qin et al., 2019; Teng et al., 2021; Qin et al., 2021b,c) adopt joint models to consider the correlation between the two tasks, which have obtained remarkable success.

Multi-intent SLU means that the system can handle an utterance containing multiple intents, which is shown to be more practical in the real-world scenario, attracting increasing attention. To this end,

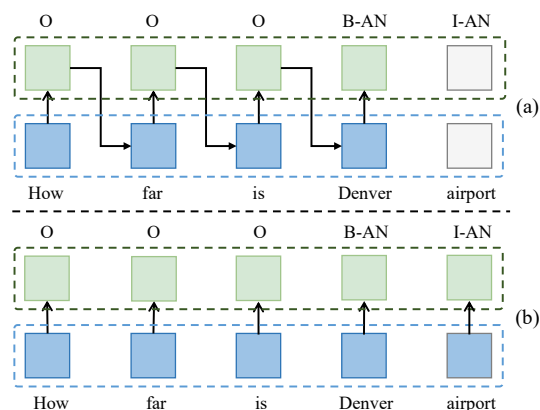


Figure 1: (a) Autoregressive model generates outputs word by word from left-to-right direction. The gray color denotes the unseen information when model decodes for the word Denver. (b) Non-autoregressive model can produce outputs in parallel. AN denotes `airport_name`.

Xu and Sarikaya (2013) and Kim et al. (2017) begin to explore the multi-intent SLU. However, their models only consider the multiple intent detection while ignoring slot filling task. Recently, Gangadharaiah and Narayanaswamy (2019) make the first attempt to propose a multi-task framework to joint model the multiple intent detection and slot filling. Qin et al. (2020b) further propose an adaptive interaction framework (AGIF) to achieve fine-grained multi-intent information integration for slot filling, obtaining state-of-the-art performance.

Though achieving the promising performance, the existing multi-intent SLU joint models heavily rely on an autoregressive fashion, as shown in Figure 1(a), leading to two issues:

- *Slow inference speed.* The autoregressive models make the generation of slot outputs must be done through the left-to-right pass, which cannot achieve parallelizable, leading to slow inference speed.

*Corresponding author.

- *Information leakage.* Autoregressive models predict each word slot conditioned on the previously generated slot information (from left-to-right), resulting in leaking the bidirectional context information.

In this paper, we explore a non-autoregressive framework for joint multiple intent detection and slot filling, with the goal of accelerating inference speed while achieving high accuracy, which is shown in Figure 1(b). To this end, we propose a **Global-Locally Graph-Interaction Network (GL-GIN)** where the core module is a proposed local slot-aware graph layer and global intent-slot interaction layer, which achieves to generate intents and slots sequence simultaneously and non-autoregressively. In GL-GIN, a local slot-aware graph interaction layer where each slot hidden states connect with each other is proposed to explicitly model slot dependency, in order to alleviate uncoordinated slot problem (e.g., *B-singer followed by I-song*) (Wu et al., 2020) due to the non-autoregressive fashion. A global intent-slot graph interaction layer is further introduced to perform sentence-level intent-slot interaction. Unlike the prior works that only consider the token-level intent-slot interaction, the global graph is constructed of all tokens with multiple intents, achieving to generate slots sequence in parallel and speed up the decoding process.

Experimental results on two public datasets MixSNIPS (Coucke et al., 2018) and MixATIS (Hemphill et al., 1990) show that our framework not only obtains state-of-the-art performance but also enables decoding in parallel. In addition, we explore the pre-trained model (i.e., Roberta (Liu et al., 2019c)) in our framework.

In summary, the contributions of this work can be concluded as follows: (1) To the best of our knowledge, we make the first attempt to explore a non-autoregressive approach for joint multiple intent detection and slot filling; (2) We propose a global-locally graph-interaction network, where the local graph is used to handle uncoordinated slots problem while a global graph is introduced to model sequence-level intent-slot interaction; (3) Experiment results on two benchmarks show that our framework not only achieves the state-of-the-art performance but also considerably speeds up the slot decoding (up to $\times 11.5$); (4) Finally, we explore the pre-trained model in our framework. With the pre-trained model, our model reaches a

new state-of-the-art level.

For reproducibility, our code for this paper is publicly available at <https://github.com/yizhen20133868/GL-GIN>.

2 Problem Definition

Multiple Intent Detection Given input sequence $x = (x_1, \dots, x_n)$, multiple intent detection can be defined as a multi-label classification task that outputs a sequence intent label $o^I = (o_1^I, \dots, o_m^I)$, where m is the number of intents in given utterance and n is the length of utterance.

Slot Filling Slot filling can be seen as a sequence labeling task that maps the input utterance x into a slot output sequence $o^S = (o_1^S, \dots, o_n^S)$.

3 Approach

As shown in Figure 2(a), we describe the proposed framework, which consists of a shared self-attentive encoder (§3.1), a token-level intent detection decoder (§3.2) and a global-local graph-interaction graph decoder for slot filling (§3.3). Both intent detection and slot filling are optimized simultaneously via a joint learning scheme.

3.1 Self-attentive Encoder

Following Qin et al. (2019), we utilize a self-attentive encoder with BiLSTM and self-attention mechanism to obtain the shared utterance representation, which can incorporate temporal features within word orders and contextual information.

BiLSTM The bidirectional LSTM (BiLSTM) (Hochreiter and Schmidhuber, 1997) have been successfully applied to sequence labeling tasks (Li et al., 2020, 2021). We adopt BiLSTM to read the input sequence $\{x_1, x_2, \dots, x_n\}$ forwardly and backwardly to produce context-sensitive hidden states $\mathbf{H} = \{\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_n\}$, by repeatedly applying the $\mathbf{h}_i = \text{BiLSTM}(\phi^{emb}(x_i), \mathbf{h}_{i-1}, \mathbf{h}_{i+1})$, where ϕ^{emb} is embedding function.

Self-Attention Following Vaswani et al. (2017), we map the matrix of input vectors $\mathbf{X} \in \mathbb{R}^{n \times d}$ (d represents the mapped dimension) to queries \mathbf{Q} , keys \mathbf{K} and values \mathbf{V} matrices by using different linear projections. Then, the self-attention output $\mathbf{C} \in \mathbb{R}^{n \times d}$ is a weighted sum of values:

$$\mathbf{C} = \text{softmax} \left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d_k}} \right) \mathbf{V}. \quad (1)$$

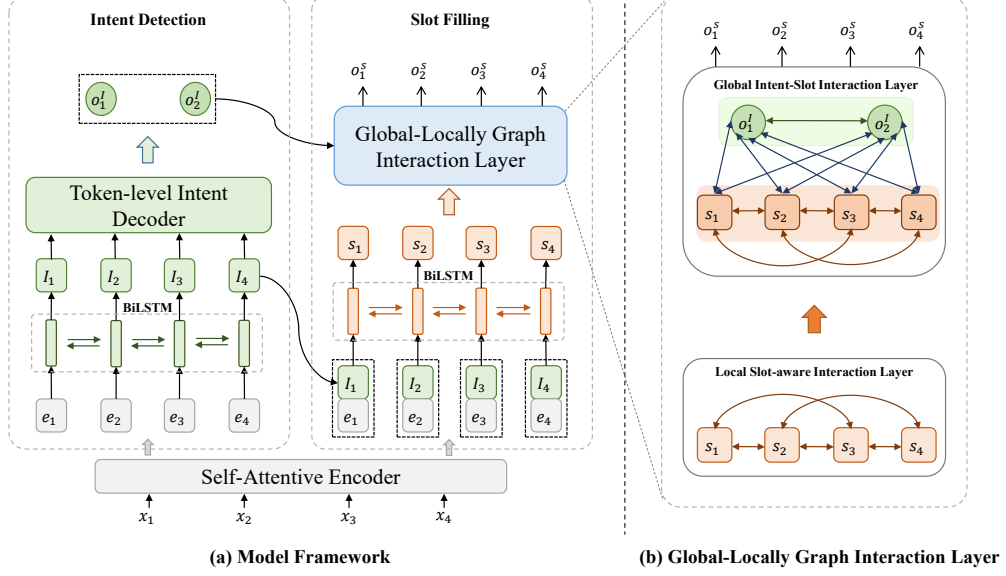


Figure 2: The overflow of model architecture (a) and global-locally graph interaction layer (b).

We concatenate the output of BiLSTM and self-attention as the final encoding representation:

$$\mathbf{E} = \mathbf{H} \parallel \mathbf{C}, \quad (2)$$

where $\mathbf{E} = \{e_1, \dots, e_n\} \in \mathbb{R}^{n \times 2d}$ and \parallel is concatenation operation.

3.2 Token-Level Intent Detection Decoder

Inspired by Qin et al. (2019), we perform a token-level multi-label multi-intent detection, where we predict multiple intents on each token and the sentence results are obtained by voting for all tokens. Specifically, we first feed the contextual encoding \mathbf{E} into an intent-aware BiLSTM to enhance its task-specific representations:

$$\mathbf{h}_t^I = \text{BiLSTM}(e_t, \mathbf{h}_{t-1}^I, \mathbf{h}_{t+1}^I). \quad (3)$$

Then, \mathbf{h}_t^I is used for intent detection, using:

$$I_t = \sigma(\mathbf{W}_I(\text{LeakyReLU}(\mathbf{W}_h \mathbf{h}_t^I + \mathbf{b}_h)) + \mathbf{b}_I), \quad (4)$$

where I_t denotes the intent results at the t -th word; σ denotes the sigmoid activation function; \mathbf{W}_h and \mathbf{W}_I are the trainable matrix parameters.

Finally, the sentence intent results o_k^I can be obtained by:

$$o^I = \{o_k^I | (\sum_{i=1}^n \mathbb{1}[I_{(i,k)} > 0.5]) > n/2\}, \quad (5)$$

where $I_{(i,k)}$ represents the classification result of token i for o_k^I .

We predict the label as the utterance intent when it gets more than half positive predictions in all n tokens. For example, if $I_1 = \{0.9, 0.8, 0.7, 0.1\}$, $I_2 = \{0.8, 0.2, 0.7, 0.4\}$, $I_3 = \{0.9, 0.3, 0.2, 0.3\}$, from three tokens, we get $\{3, 2, 1, 0\}$ positive votes (> 0.5) for four intents respectively. Thus the index where more than half of the votes ($> 3/2$) were obtained was o_1^I and o_3^I , we predict intents $o^I = \{o_1^I, o_3^I\}$.

3.3 Slot Filling Decoder

One main advantage of our framework is the proposed global-locally graph interaction network for slot filling, which is a non-autoregressive paradigm, achieving the slot filling decoding in parallel. In the following, we first describe the slot-aware LSTM (§3.3.1) to obtain the slot-aware representations, and then show how to apply the global-locally graph interaction layer (§3.3.2) for decoding.

3.3.1 Slot-aware LSTM

We utilize a BiLSTM to produce the slot-aware hidden representation $\mathbf{S} = (s_1, \dots, s_n)$. At each decoding step t , the decoder state s_t calculating by:

$$s_t = \text{BiLSTM}(I_t \parallel e_t, s_{t-1}, s_{t+1}), \quad (6)$$

where e_t denotes the aligned encoder hidden state and I_t denotes the predicted intent information.

3.3.2 Global-locally Graph Interaction Layer

The proposed global-locally graph interaction layer consists of two main components: one is a local slot-aware graph interaction network to model dependency across slots and another is the proposed global intent-slot graph interaction network to consider the interaction between intents and slots.

In this section, we first describe the vanilla graph attention network. Then, we illustrate the local slot-aware and global intent-slot graph interaction network, respectively.

Vanilla Graph Attention Network A graph attention network (GAT) (Veličković et al., 2018) is a variant of graph neural network, which fuses the graph-structured information and node features within the model. Its masked self-attention layers allow a node to attend to neighborhood features and learn different attention weights, which can automatically determine the importance and relevance between the current node with its neighborhood.

In particular, for a given graph with N nodes, one-layer GAT take the initial node features $\tilde{\mathbf{H}} = \{\tilde{\mathbf{h}}_1, \dots, \tilde{\mathbf{h}}_N\}$, $\tilde{\mathbf{h}}_n \in \mathbb{R}^F$ as input, aiming at producing more abstract representation, $\tilde{\mathbf{H}}' = \{\tilde{\mathbf{h}}'_1, \dots, \tilde{\mathbf{h}}'_N\}$, $\tilde{\mathbf{h}}'_n \in \mathbb{R}^{F'}$, as its output. The attention mechanism of a typical GAT can be summarized as below:

$$\begin{aligned} \tilde{\mathbf{h}}'_i &= \parallel_{k=1}^K \sigma \left(\sum_{j \in \mathcal{N}_i} \alpha_{ij}^k \mathbf{W}_h^k \tilde{\mathbf{h}}_j \right), \quad (7) \\ \alpha_{ij} &= \frac{\exp(\text{LeakyReLU}(\mathbf{a}^\top [\mathbf{W}_h \tilde{\mathbf{h}}_i \parallel \mathbf{W}_h \tilde{\mathbf{h}}_j]))}{\sum_{j' \in \mathcal{N}_i} \exp(\text{LeakyReLU}(\mathbf{a}^\top [\mathbf{W}_h \tilde{\mathbf{h}}_i \parallel \mathbf{W}_h \tilde{\mathbf{h}}_{j'}]))} \quad (8) \end{aligned}$$

where $\mathbf{W}_h \in \mathbb{R}^{F' \times F}$ and $\mathbf{a} \in \mathbb{R}^{2F'}$ are the trainable weight matrix; \mathcal{N}_i denotes the neighbors of node i (including i); α_{ij} is the normalized attention coefficients and σ represents the nonlinearity activation function; K is the number of heads.

Local Slot-aware Graph Interaction Layer

Given slot decode hidden representations $\mathbf{S} = (\mathbf{s}_1, \dots, \mathbf{s}_n)$, we construct a local slot-aware graph where each slot hidden node connects to other slots. This allows the model to achieve to model the dependency across slots, alleviating the uncoordinated slots problem. Specifically, we construct the graph $\mathcal{G} = (V, \mathcal{E})$ in the following way,

Vertices We define the V as the vertices set. Each word slot is represented as a vertex. Each vertex is initialized with the corresponding slot hidden representation. Thus, the first layer states vector for all nodes is $\mathbf{S}^1 = \mathbf{S} = (\mathbf{s}_1, \dots, \mathbf{s}_n)$.

Edges Since we aim to model dependency across slots, we construct a slot-aware graph interaction layer so that the dependency relationship can be propagated from neighbor nodes to the current node. Each slot can connect other slots with a window size. For node \mathbf{S}_i , only $\{\mathbf{S}_{i-m}, \dots, \mathbf{S}_{i+m}\}$ will be connected where m is a hyper-parameter denotes the size of sliding window that controls the length of utilizing utterance context.

Information Aggregation The aggregation process at l -th layer can be defined as:

$$\mathbf{s}_i^{l+1} = \sigma \left(\sum_{j \in \mathcal{N}_i} \alpha_{ij} \mathbf{W}_l \mathbf{s}_j^l \right), \quad (9)$$

where \mathcal{N}_i is a set of vertices that denotes the connected slots.

After stacking L layer, we obtain the contextual slot-aware local hidden features $\mathbf{S}^{L+1} = \{\mathbf{s}_1^{L+1}, \dots, \mathbf{s}_n^{L+1}\}$

Global Slot-Intent Graph Interaction Layer

To achieve sentence-level intent-slot interaction, we construct a global slot-intent interaction graph where all predicted multiple intents and sequence slots are connected, achieving to output slot sequences in parallel. Specifically, we construct the graph $\mathcal{G} = (V, \mathcal{E})$ in the following way,

Vertices As we model the interaction between intent and slot token, we have $n + m$ number of nodes in the graph where n is the sequence length and m is the number of intent labels predicted by the intent decoder. The input of slot token feature is $\mathbf{G}^{[S,1]} = \mathbf{S}^{L+1} = \{\mathbf{s}_1^{L+1}, \dots, \mathbf{s}_n^{L+1}\}$ which is produced by slot-aware local interaction graph network while the input intent feature is an embedding $\mathbf{G}^{[I,1]} = \{\phi^{emb}(\mathbf{o}_1^I), \dots, \phi^{emb}(\mathbf{o}_m^I)\}$ where ϕ^{emb} is a trainable embedding matrix. The first layer states vector for slot and intent nodes is $\mathbf{G}^1 = \{\mathbf{G}^{[I,1]}, \mathbf{G}^{[S,1]}\} = \{\phi^{emb}(\mathbf{o}_1^I), \dots, \phi^{emb}(\mathbf{o}_m^I), \mathbf{s}_1^{L+1}, \dots, \mathbf{s}_n^{L+1}\}$

Edges There are three types of connections in this graph network.

- *intent-slot connection*: Since slots and intents are highly tied, we construct the intent-slot connection to model the interaction between the two tasks. Specifically, each slot connects all predicted multiple intents to automatically capture relevant intent information.

- *slot-slot connection*: We construct the slot-slot connection where each slot node connects other slots with the window size to further model the slot dependency and incorporate the bidirectional contextual information.
- *intent-intent connection*: Following Qin et al. (2020b), we connect all the intent nodes to each other to model the relationship between each intent, since all of them express the same utterance’s intent.

Information Aggregation The aggregation process of the global GAT layer can be formulated as:

$$\mathbf{g}_i^{[S,I+1]} = \sigma \left(\sum_{j \in \mathcal{G}^S} \alpha_{ij} \mathbf{W}_g \mathbf{g}_j^{[S,I]} + \sum_{j \in \mathcal{G}^I} \alpha_{ij} \mathbf{W}_g \mathbf{g}_j^{[I,I]} \right), \quad (10)$$

where \mathcal{G}^S and \mathcal{G}^I are vertices sets which denotes the connected slots and intents, respectively.

3.3.3 Slot Prediction

After L layers’ propagation, we obtain the final slot representation $\mathbf{G}^{[S,L+1]}$ for slot prediction.

$$\mathbf{y}_t^S = \text{softmax} \left(\mathbf{W}_s \mathbf{g}_t^{[S,L+1]} \right), \quad (11)$$

$$\mathbf{o}_t^S = \arg \max(\mathbf{y}_t^S), \quad (12)$$

where \mathbf{W}_s is a trainable parameter and \mathbf{o}_t^S is the predicted slot if the t -th token in an utterance.

3.4 Joint Training

Following Goo et al. (2018), we adopt a joint training model to consider the two tasks and update parameters by joint optimizing. The intent detection objective is:

$$\text{CE}(\hat{y}, y) = \hat{y} \log(y) + (1 - \hat{y}) \log(1 - y), \quad (13)$$

$$\mathcal{L}_1 \triangleq - \sum_{i=1}^n \sum_{j=1}^{N_I} \text{CE}(\hat{y}_i^{(j,I)}, y_i^{(j,I)}). \quad (14)$$

Similarly, the slot filling task objective is:

$$\mathcal{L}_2 \triangleq - \sum_{i=1}^n \sum_{j=1}^{N_S} \hat{y}_i^{(j,S)} \log(y_i^{(j,S)}), \quad (15)$$

where N_I is the number of single intent labels and N_S is the number of slot labels.

The final joint objective is formulated as:

$$\mathcal{L} = \alpha \mathcal{L}_1 + \beta \mathcal{L}_2, \quad (16)$$

where α and β are hyper-parameters.

4 Experiments

4.1 Datasets

We conduct experiments on two publicly available multi-intent datasets.¹ One is the Mix-ATIS (Hemphill et al., 1990; Qin et al., 2020b), which includes 13,162 utterances for training, 756 utterances for validation and 828 utterances for testing. Another is MixSNIPS (Coucke et al., 2018; Qin et al., 2020b), with 39,776, 2,198, 2,199 utterances for training, validation and testing.

4.2 Experimental Settings

The dimensionality of the embedding is 128 and 64 on ATIS and SNIPS, respectively. The dimensionality of the LSTM hidden units is 256. The batch size is 16. The number of the multi head is 4 and 8 on MixATIS and MixSNIPS dataset, respectively. All layer number of graph attention network is set to 2. We use Adam (Kingma and Ba, 2015) to optimize the parameters in our model. For all the experiments, we select the model which works the best on the dev set and then evaluate it on the test set. All experiments are conducted at GeForce RTX 2080Ti and TITAN Xp.

4.3 Baselines

We compare our model with the following best baselines: (1) Attention BiRNN. Liu and Lane (2016) propose an alignment-based RNN for joint slot filling and intent detection; (2) Slot-Gated Atten. Goo et al. (2018) propose a slot-gated joint model, explicitly considering the correlation between slot filling and intent detection; (3) Bi-Model. Wang et al. (2018) propose the Bi-model to model the bi-directional between the intent detection and slot filling; (4) SF-ID Network. E et al. (2019) proposes the SF-ID network to establish a direct connection between the two tasks; (5) Stack-Propagation. Qin et al. (2019) adopt a stack-propagation framework to explicitly incorporate intent detection for guiding slot filling; (6) Joint Multiple ID-SF. Gangadharaiah and Narayanaswamy (2019) propose a multi-task framework with slot-gated mechanism for multiple intent detection and slot filling; (7) AGIF Qin et al. (2020b) proposes an adaptive interaction network to achieve the fine-grained multi-

¹We adopt the cleaned version that removes the repeated sentences in original dataset, which is available at <https://github.com/LooperXX/AGIF>.

Model	MixATIS			MixSNIPS		
	Overall(Acc)	Slot(F1)	Intent(Acc)	Overall(Acc)	Slot(F1)	Intent(Acc)
Attention BiRNN (Liu and Lane, 2016)	39.1	86.4	74.6	59.5	89.4	95.4
Slot-Gated (Goo et al., 2018)	35.5	87.7	63.9	55.4	87.9	94.6
Bi-Model (Wang et al., 2018)	34.4	83.9	70.3	63.4	90.7	95.6
SF-ID (E et al., 2019)	34.9	87.4	66.2	59.9	90.6	95.0
Stack-Propagation (Qin et al., 2019)	40.1	87.8	72.1	72.9	94.2	96.0
Joint Multiple ID-SF (Gangadharaiah and Narayanaswamy, 2019)	36.1	84.6	73.4	62.9	90.6	95.1
AGIF (Qin et al., 2020b)	40.8	86.7	74.4	74.2	94.2	95.1
GL-GIN	43.5*	88.3*	76.3*	75.4*	94.9*	95.6

Table 1: Main results. The numbers with * indicate that the improvement of our framework over all baselines is statistically significant with $p < 0.05$ under t-test.

Model	Decode Latency(s)	Speedup
Stack-Propagation	34.5	8.2×
Joint Multiple ID-SF	45.3	10.8×
AGIF	48.5	11.5×
GL-GIN	4.2	1.0×

Table 2: Speed comparison. Speedup is based on the ratio of the time taken by the slot decoding part of different models to run an epoch on the MixATIS dataset with batch size set to 32.

intent information integration, achieving state-of-the-art performance.

4.4 Main Results

Following Goo et al. (2018) and Qin et al. (2020b), we evaluate the performance of slot filling using F1 score, intent prediction using accuracy, the sentence-level semantic frame parsing using overall accuracy. Overall accuracy measures the ratio of sentences for which both intent and slot are predicted correctly in a sentence.

Table 1 shows the results, we have the following observations: (1) On slot filling task, our framework outperforms the best baseline AGIF in F1 scores on two datasets, which indicates the proposed local slot-aware graph successfully models the dependency across slots, so that the slot filling performance can be improved. (2) More importantly, compared with the AGIF, our framework achieves +2.7% and 1.2% improvements for MixATIS and MixSNIPS on overall accuracy, respectively. We attribute it to the fact that our proposed global intent-slot interaction graph can better capture the correlation between intents and slots, improving the SLU performance.

4.5 Analysis

4.5.1 Speedup

One of the core contributions of our framework is that the decoding process of slot filling can be significantly accelerated with the proposed

non-autoregressive mechanism. We evaluate the speed by running the model on the MixATIS test data in an epoch, fixing the batch size to 32. The comparison results are shown in Table 2. We observe that our model achieves the $\times 8.2$, $\times 10.8$ and $\times 11.5$ speedup compared with SOTA models stack-propagation, Joint Multiple ID-SF and AGIF. This is because that their model utilizes an autoregressive architecture that only performs slot filling word by word, while our non-autoregressive framework can conduct slot filling decoding in parallel. In addition, it’s worth noting that as the batch size gets larger, GL-GIN can achieve better acceleration where our model could achieve $\times 17.2$ speedup compared with AGIF when batch size is 64.

4.5.2 Effectiveness of the Local Slot-aware Graph Interaction Layer

We study the effectiveness of the local slot-aware interaction graph layer with the following ablation. We remove the local graph interaction layer and directly feed the output of the *slot LSTM* to the global intent-slot graph interaction layer. We refer it to *w/o local GAL* in Table 3. We can clearly observe that the slot F1 drops by 1.5% and 1.2% on MixATIS and MixSNIPS datasets. We attribute this to the fact that local slot-aware GAL can capture the slot dependency for each token, which helps to alleviate the slot uncoordinated problems. A qualitative analysis can be founded at Section 4.5.6.

4.5.3 Effectiveness of Global Slot-Intent Graph Interaction Layer

In order to verify the effectiveness of slot-intent global interaction graph layer, we remove the global interaction layer and utilizes the output of local slot-aware GAL module for slot filling. It is named as *w/o Global Intent-slot GAL* in Table 3. We can observe that the slot f1 drops by 0.9%, 1.3%, which demonstrates that intent-slot graph in-

Model	MixATIS			MixSNIPS		
	Overall(Acc)	Slot(F1)	Intent(Acc)	Overall(Acc)	Slot(F1)	Intent(Acc)
w/o Local Slot-Aware GAL	41.1	86.8	74.0	71.4	93.7	95.2
w/o Global Intent-Slot GAL	40.9	87.4	75.5	71.7	93.6	95.5
+ More Parameters	41.9	87.7	75.0	73.0	93.8	95.5
w/o Global-locally GAL	40.5	86.3	75.2	70.2	92.9	95.0
GL-GIN	43.5	88.3	76.3	75.4	94.9	95.6

Table 3: Ablation Experiment.

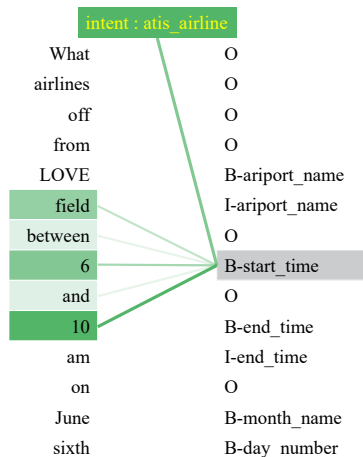


Figure 3: Visualization. We use the green color to indicate the attention value.

teraction layer can capture the correlation between multiple intents, which is beneficial for the semantic performance of SLU system.

Following Qin et al. (2020b), we replace multiple LSTM layers (2-layers) as the proposed global-locally graph layer to verify that the proposed global-locally graph interaction layer rather than the added parameters works. Table 3 (*more parameters*) shows the results. We observe that our model outperforms *more parameters* by 1.6% and 2.4% overall accuracy in two datasets, which shows that the improvements come from the proposed Global-locally graph interaction layer rather than the involved parameters.

4.5.4 Effectiveness of the Global-locally Graph Interaction Layer

Instead of using the whole global-locally graph interaction layer for slot filling, we directly leverage the output of slot-aware LSTM to predict each token slot to verify the effect of the global-locally graph interaction layer. We name the experiment as *w/o Global-locally GAL* in Tabel 3. From the results, We can observe that the absence of global GAT module leads to 3.0% and 5.2% overall accuracy drops on two datasets. This indicates that the




Figure 4: Overall accuracy Performances with Roberta.

global-locally graph interaction layer encourages our model to leverage slot dependency and intent information, which can improve SLU performance.

4.5.5 Visualization

To better understand how global-local graph interaction layer affects and contributes to the final result, we visualize the attention value of the Global intent-slot GAL. As is shown in Figure 3, we visualize the dependence of the word “6” on context and intent information. We can clearly observe that token “6” obtains information from all contextual tokens. The information from “and 10” helps to predict the slot, where the prior autoregressive models cannot be achieved due to the generation word by word from left to right.

4.5.6 Qualitative analysis

We conduct qualitative analysis by providing a case study that consists of two sequence slots which are generated from AGIF and our model. From Table 4, for the word “6”, AGIF predicts its slot label as “O” incorrectly. This is because that AGIF only models its left information, which makes it hard to predict “6” is a time slot. In contrast, our model predicts the slot label correctly. We attribute this to the fact that our proposed global intent-slot interaction layer can model bidirectional contextual information. In addition, our framework predicts the word slot “am” correctly while AGIF predicts it incorrectly (I-airport_name follows B-depart_time), indicating that the proposed local slot-

texts	What	airlines	off	from	LOVE	field	between	6	and	10	am	on	June	sixth
AGIF	O	O	O	O	B-fromloc airport_name	I-fromloc airport_name	O	O	O	B-depart_time end_time	I-toloc airport_name	O	B-depart_date month_name	B-depart_date day_number
GL-GIN	O	O	O	O	B-fromloc airport_name	I-fromloc airport_name	O	B-depart_time start_time	O	B-depart_time end_time	I-depart_time end_time	O	B-depart_date month_name	B-depart_date day_number

Table 4: Case study. Predicted slots sequence about utterance “What airlines off from LOVE field between 6 and 10 am on June sixth”

aware graph layer has successfully captured the slot dependency.

4.5.7 Effect of Pre-trained Model

Following Qin et al. (2019), we explore the pre-trained model in our framework. We replace the self-attentive encoder by *Roberta* (Liu et al., 2019c) with the fine-tuning approach. We keep other components identical to our framework and follow Qin et al. (2019) to consider the first subword label if a word is broken into multiple subwords.

Figure 4 gives the result comparison of AGIF, GL-GIN and two models with *Roberta* on two datasets. We have two interesting observations. First, the *Roberta*-based model remarkably well on two datasets. We attribute this to the fact that pre-trained models can provide rich semantic features, which can help SLU. Second, GL-GIN + *Roberta* outperforms AGIF+*Roberta* on both datasets and reaches a new state-of-the-art performance, which further verifies the effectiveness of our proposed framework.

5 Related Work

Slot Filling and Intent Detection Recently, joint models (Zhang and Wang, 2016; Hakkani-Tür et al., 2016; Goo et al., 2018; Li et al., 2018; Xia et al., 2018; E et al., 2019; Liu et al., 2019b; Qin et al., 2019; Zhang et al., 2019; Wu et al., 2020; Qin et al., 2021b; Ni et al., 2021) are proposed to consider the strong correlation between intent detection and slot filling have obtained remarkable success. Compared with their work, we focus on jointly modeling multiple intent detection and slot filling while they only consider the single-intent scenario.

More recently, multiple intent detection can handle utterances with multiple intents, which has attracted increasing attention. To the end, Xu and Sarikaya (2013) and Kim et al. (2017) begin to explore the multiple intent detection. Gangadharaiah and Narayanaswamy (2019) first apply a multi-task framework with a slot-gate mechanism to jointly model the multiple intent detection and slot fill-

ing. Qin et al. (2020b) propose an adaptive interaction network to achieve the fine-grained multiple intent information integration for token-level slot filling, achieving the state-of-the-art performance. Their models adopt the autoregressive architecture for joint multiple intent detection and slot filling. In contrast, we propose a non-autoregressive approach, achieving parallel decoding. To the best of our knowledge, we are the first to explore a non-autoregressive architecture for multiple intent detection and slot filling.

Graph Neural Network for NLP Graph neural networks that operate directly on graph structures to model the structural information, which has been applied successfully in various NLP tasks. Linmei et al. (2019) and Huang and Carley (2019) explore graph attention network (GAT) (Veličković et al., 2018) for classification task to incorporate the dependency parser information. Cetoli et al. (2017) and Liu et al. (2019a) apply graph neural network to model the non-local contextual information for sequence labeling tasks. Yasunaga et al. (2017) and Feng et al. (2020a) successfully apply a graph network to model the discourse information for the summarization generation task, which achieved promising performance. Graph structure are successfully applied for dialogue direction (Feng et al., 2020b; Fu et al., 2020; Qin et al., 2020a, 2021a). In our work, we apply a global-locally graph interaction network to model the slot dependency and interaction between the multiple intents and slots.

6 Conclusion

In this paper, we investigated a non-autoregressive model for joint multiple intent detection and slot filling. To this end, we proposed a global-locally graph interaction network where the uncoordinated-slots problem can be addressed with the proposed local slot-aware graph while the interaction between intents and slots can be modeled by the proposed global intent-slot graph. Experimental results on two datasets show that our framework achieves state-of-the-art performance with $\times 11.5$ times faster than the prior work.

Acknowledgements

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. This work was also supported by the Zhejiang Lab’s International Talent Fund for Young Professionals.

References

- Alberto Cetoli, Stefano Bragaglia, Andrew O’Harney, and Marc Sloan. 2017. [Graph convolutional networks for named entity recognition](#). In *Proceedings of the 16th International Workshop on Treebanks and Linguistic Theories*, pages 37–45, Prague, Czech Republic.
- 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*.
- Haihong E, Peiqing Niu, Zhongfu Chen, and Meina Song. 2019. [A novel bi-directional interrelated model for joint intent detection and slot filling](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5467–5471, Florence, Italy. Association for Computational Linguistics.
- Xiachong Feng, Xiaocheng Feng, Bing Qin, Xinwei Geng, and Ting Liu. 2020a. [Dialogue discourse-aware graph convolutional networks for abstractive meeting summarization](#).
- Xiachong Feng, Xiaocheng Feng, Bing Qin, and Ting Liu. 2020b. Incorporating commonsense knowledge into abstractive dialogue summarization via heterogeneous graph networks. *arXiv preprint arXiv:2010.10044*.
- Qiankun Fu, Yue Zhang, Jiangming Liu, and Meishan Zhang. 2020. [DRTS parsing with structure-aware encoding and decoding](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6818–6828, Online. Association for Computational Linguistics.
- Rashmi Gangadharaiah and Balakrishnan Narayanaswamy. 2019. [Joint multiple intent detection and slot labeling for goal-oriented dialog](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 564–569, Minneapolis, Minnesota. 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 *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 753–757, New Orleans, Louisiana. Association for Computational Linguistics.
- Dilek Hakkani-Tür, Gokhan Tur, Asli Celikyilmaz, Yun-Nung Vivian Chen, Jianfeng Gao, Li Deng, and Ye-Yi Wang. 2016. [Multi-domain joint semantic frame parsing using bi-directional rnn-1stm](#). In *Proceedings of The 17th Annual Meeting of the International Speech Communication Association (INTER-SPEECH 2016)*. ISCA.
- Charles T. Hemphill, John J. Godfrey, and George R. Doddington. 1990. [The ATIS spoken language systems pilot corpus](#). In *Speech and Natural Language: Proceedings of a Workshop Held at Hidden Valley, Pennsylvania, June 24-27, 1990*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation*, 9(8):1735–1780.
- Binxuan Huang and Kathleen Carley. 2019. [Syntax-aware aspect level sentiment classification with graph attention networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5469–5477, Hong Kong, China. Association for Computational Linguistics.
- Byeongchang Kim, Seonghan Ryu, and Gary Geunbae Lee. 2017. Two-stage multi-intent detection for spoken language understanding. *Multimedia Tools and Applications*, 76(9):11377–11390.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Changliang Li, Liang Li, and Ji Qi. 2018. [A self-attentive model with gate mechanism for spoken language understanding](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3824–3833, Brussels, Belgium. Association for Computational Linguistics.
- Yangming Li, Han Li, Kaisheng Yao, and Xiaolong Li. 2020. [Handling rare entities for neural sequence labeling](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6441–6451, Online. Association for Computational Linguistics.
- Yangming Li, lemao liu, and Shuming Shi. 2021. [Empirical analysis of unlabeled entity problem in named entity recognition](#). In *International Conference on Learning Representations*.

- Hu Linmei, Tianchi Yang, Chuan Shi, Houye Ji, and Xiaoli Li. 2019. [Heterogeneous graph attention networks for semi-supervised short text classification](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4821–4830, Hong Kong, China. Association for Computational Linguistics.
- Bing Liu and Ian Lane. 2016. [Attention-based recurrent neural network models for joint intent detection and slot filling](#). In *Interspeech 2016*, pages 685–689.
- Pengfei Liu, Shuaichen Chang, Xuanjing Huang, Jian Tang, and Jackie Chi Kit Cheung. 2019a. [Contextualized non-local neural networks for sequence learning](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):6762–6769.
- Yijin Liu, Fandong Meng, Jinchao Zhang, Jie Zhou, Yufeng Chen, and Jinan Xu. 2019b. [CM-net: A novel collaborative memory network for spoken language understanding](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1051–1060, Hong Kong, China. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019c. [Roberta: A robustly optimized BERT pretraining approach](#). *arXiv preprint arXiv:1907.11692*, abs/1907.11692.
- Jinjie Ni, Tom Young, Vlad Pandelea, Fuzhao Xue, Vinay Adiga, and Erik Cambria. 2021. [Recent advances in deep learning based dialogue systems: A systematic survey](#).
- 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 *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2078–2087, Hong Kong, China. Association for Computational Linguistics.
- Libo Qin, Wanxiang Che, Minheng Ni, Yangming Li, and Ting Liu. 2021a. [Knowing where to leverage: Context-aware graph convolutional network with an adaptive fusion layer for contextual spoken language understanding](#). *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:1280–1289.
- Libo Qin, Zhouyang Li, Wanxiang Che, Minheng Ni, and Ting Liu. 2020a. [Co-gat: A co-interactive graph attention network for joint dialog act recognition and sentiment classification](#).
- Libo Qin, Tailu Liu, Wanxiang Che, Bingbing Kang, Sendong Zhao, and Ting Liu. 2021b. [A co-interactive transformer for joint slot filling and intent detection](#).
- Libo Qin, Tianbao Xie, Wanxiang Che, and Ting Liu. 2021c. [A survey on spoken language understanding: Recent advances and new frontiers](#).
- Libo Qin, Xiao Xu, Wanxiang Che, and Ting Liu. 2020b. [AGIF: An adaptive graph-interactive framework for joint multiple intent detection and slot filling](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1807–1816, Online. Association for Computational Linguistics.
- Dechuang Teng, Libo Qin, Wanxiang Che, Sendong Zhao, and Ting Liu. 2021. [Injecting word information with multi-level word adapter for chinese spoken language understanding](#).
- Gokhan Tur and Renato De Mori. 2011. *Spoken language understanding: Systems for extracting semantic information from speech*. John Wiley & Sons.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30, pages 5998–6008. Curran Associates, Inc.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. [Graph Attention Networks](#). *International Conference on Learning Representations*. Accepted as poster.
- Yu Wang, Yilin Shen, and Hongxia Jin. 2018. [A bi-model based RNN semantic frame parsing model for intent detection and slot filling](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 309–314, New Orleans, Louisiana. Association for Computational Linguistics.
- Di Wu, Liang Ding, Fan Lu, and Jian Xie. 2020. [SlotRefine: A fast non-autoregressive model for joint intent detection and slot filling](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1932–1937, Online. Association for Computational Linguistics.
- Congying Xia, Chenwei Zhang, Xiaohui Yan, Yi Chang, and Philip Yu. 2018. [Zero-shot user intent detection via capsule neural networks](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3090–3099, Brussels, Belgium. Association for Computational Linguistics.
- Puyang Xu and Ruhi Sarikaya. 2013. [Convolutional neural network based triangular crf for joint intent detection and slot filling](#). In *2013 IEEE Workshop*

on Automatic Speech Recognition and Understanding.

Michihiro Yasunaga, Rui Zhang, Kshitijh Meelu, Ayush Pareek, Krishnan Srinivasan, and Dragomir Radev. 2017. [Graph-based neural multi-document summarization](#). In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 452–462, Vancouver, Canada. Association for Computational Linguistics.

Steve Young, Milica Gašić, Blaise Thomson, and Jason D Williams. 2013. Pomdp-based statistical spoken dialog systems: A review. *Proceedings of the IEEE*, 101(5):1160–1179.

Chenwei Zhang, Yaliang Li, Nan Du, Wei Fan, and Philip Yu. 2019. [Joint slot filling and intent detection via capsule neural networks](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5259–5267, Florence, Italy. Association for Computational Linguistics.

Xiaodong Zhang and Houfeng Wang. 2016. A joint model of intent determination and slot filling for spoken language understanding. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI'16*, page 2993–2999. AAAI Press.