

Real-time Caller Intent Detection In Human-Human Customer Support Spoken Conversations

Mrinal Rawat¹, Victor Barres¹

¹Uniphore Software Systems Inc.
{mrinal.rawat, victor}@uniphore.com

Abstract

Agent assistance during human-human customer support spoken interactions requires triggering workflows based on the caller's intent (reason for call). Timeliness of prediction is essential for a good user experience. The goal is for a system to detect the caller's intent at the time the agent would have been able to detect it (Intent Boundary). Some approaches focus on predicting the output offline, i.e. once the full spoken input (e.g. the whole conversational turn) has been processed by the ASR system. This introduces an undesirable latency in the prediction each time the intent could have been detected earlier in the turn. Recent work on voice assistants has used incremental real-time predictions at a word-by-word level to detect intent before the end of a command. Human-directed and machine-directed speech however have very different characteristics. In this work, we propose to apply a method developed in the context of voice-assistant to the problem of online real time caller's intent detection in human-human spoken interactions. We use a dual architecture in which two LSTMs are jointly trained: one predicting the Intent Boundary (IB) and then other predicting the intent class at the IB. We conduct our experiments on our private dataset comprising transcripts of human-human telephone conversations from the telecom customer support domain. We report results analyzing both the accuracy of our system as well as the impact of different architectures on the trade off between overall accuracy and prediction latency.

1 Introduction

Conversational AI finds one key application in reducing the time-to-resolution of customer support calls. It does so by assisting the support agents during a call as well as in their after call work [Adam *et al.*2021]. If after call work usually focuses on call disposition, assistance during the call consists in triggering actions at the right time. Such actions can encompass surfacing relevant information, extracting values from the ongoing call, suggest strategies, etc. In this work, we focus on

an application of conversational AI to agent assistance during the call.

A typical contact center solution provides agents with pre-defined workflows to follow depending on the caller's intent. One of the first tasks of the agent is therefore to recognize the reason for the call and select the appropriate workflow. We will focus in this work on the goal of detecting the caller's intent in real time.

Although some systems are designed to operate directly on the speech signal, we focus here on NLU solutions processing the output generated by an automated speech recognition (ASR) system. ASR systems convert audio streams into output word sequences. The frequency at which the output is generated vary. In this work however, without loss of generality, we abstract away the considerations related to ASR output frequencies by assuming that transcriptions will be available one word at a time.

Previous attempts [Chen *et al.*2019, Haihong *et al.*2019, Qin *et al.*2019] approached these problems using machine learning and deep learning based techniques in which predictions are carried out at the end of each utterance or turn. An obvious limitation is that the AI system has to wait until the entire word sequence is available before triggering actions, even when enough cues are available early on for a human to reliably trigger the same action. Consider an example, `i'm calling because i wanted to change my phone number but aah if this call gets disconnected with you call me back`, it is evident that as soon as the word `number` appears, the agent can predict the caller's intent, but a model working on entire turns would have to wait till the word `back` is available to generate its prediction.

A lag between machine prediction and human prediction can lead to increased cognitive load and overall poor user experience. This could in turn make the use of conversational AI systems counterproductive.

We propose a novel solution where we process the ASR output incrementally (one word at a time) and generate a prediction at each step. Our approach comprises of two step process 1) Intent Boundary (IB) detection: predicts whether the input word is the end of the intent statement 2) Intent detection: predicts the intent. Both these tasks are trained in a multi-task setting. Further, we propose the use of lookahead i.e. incorporating future right context to assist the intent boundary

detection. We also study the model performance without the intent boundary task. Finally, we propose two metrics Mean Turn Difference (MTD) and Mean Position Difference (MPD) that aims to understand the effect of the Intent Boundary task in the early prediction.

2 Related Work

Most of the previous research works mainly focus on intent detection using a full input sequence. [Guo *et al.*2014, Zhang and Wang2016, Kim *et al.*2017] proposes a method that aims to train the NLU tasks such as intent detection, slot filling, etc. jointly using RNN and attention based approaches. Recently, some works [Qin *et al.*2019] have also tried to use transformers based architectures such as BERT [Devlin *et al.*2018]. [Wu *et al.*2020] proposes a two-pass refine mechanism to handle uncoordinated slots problem. They claim to have better performance than CRF based architectures.

In recent years, the demand for real-time applications has increased tremendously. This includes incremental processing of the sequence outputting the prediction at each time step. [Schlangen and Skantze2009] first proposes a way that requires the input to be continually processed based on incremental input. Recently, [Shivakumar *et al.*2021] proposed a model that jointly trains two tasks 1) Intent Detection and 2) End-of-sentence(EOS) boundary using LSTM and generates the output at each time step. To imitate the real-time processing the authors randomly concatenate the sentences from the ATIS dataset and consider the last token as the EOS. [Ma *et al.*2019] also uses a similar approach where they use an incremental encoder by limiting each source word to attend to its previous input and recompute the representation. [Constantin *et al.*2019] also process the input incrementally while boosting the performance in noisy datasets using data augmentation techniques. A thorough investigation of the use of non-incremental encoders such as transformers for the incremental systems was conducted by [Madureira and Schlangen2020]. [Kahardipraja *et al.*2021] proposed an approach that uses a restart-incrementally approach along with the linear transformers [Katharopoulos *et al.*2020] that reduces the overall quadratic complexity to linear. Our approach process the input sequence incrementally using the LSTM network by incorporating the lookahead based intent detection.

3 Dataset

In this work, we use our organization’s private dataset which comprises manually annotated transcripts of human-human spoken telephone conversation from the telecom customer support domain. Transcripts were generated by our in-house Kaldi-based ASR system, fine-tuned for this particular domain. Speaker diarization and recognition was not required since agent’s and customer’s speech was recorded on separate channels. Each transcript defines a sequence of turns (single speech sequence until the other speaker starts speaking). Each turn is and each turn is further broken into utterances based on inter lexical pause duration. The time offsets for each word is provided by the the ASR engine. We use this information to analyze the impact of real-time prediction as compared to

offline prediction. Table 1 describes an example transcript where the highlighted text represents the span of the intent.

Agent	Customer
you’re welcome to XXX my name is XXX great how can i help you today	
	um yeah uh calling to i guess uh remove my grandma’s phone number from her account
⋮	⋮
⋮	⋮

Table 1: Example

We collect 1544 transcripts and asked annotators to highlight the first text span that would allow a reasonable agent to identify the caller’s intent. We also asked that the intent be classify into one of 16 different intent domain relevant classes. Although in very few transcripts the caller mentions multiple reasons for contacting customer support, the vast majority of the transcripts contain a single intent. The goal is therefore to identify the first and most often only intent stated by the customer.

For our work, we consider only the turns where the intent is present which comprises of 1917 samples and is split into (80%) train, (20%) validation and (20%) test sets. We provide the exact numbers for each split/class in Table 2 for clarity.

Split	Total
Training	1526
Validation	194
Test	197
Total	1917


Table 2: Dataset Statistics

Table 3 provides an example of training set which describes the input and output format of both tasks. We set the IB token to 1 and rest of the tokens to O.


Sentence	I	need	my	account	number	rotor	number ..
IB	O	O	O	O	I	O	O ..
Intent	O	O	O	O	AM	O	O ..

Table 3: Training example: AM refers to the ‘Account Management’ intent.

Further, we performed an analysis on the dataset to understand the advantages of employing the real-time systems. We discover that **61.84 %** of the conversations had the intent at the end of the turn which implies that there is no additional benefit of using a real-time system. In the remaining **38.16 %** transcripts, we plot the distribution of the difference in the



(a) Distribution of offset difference



(b) Cumulative Distribution of offset difference

Figure 1: Distribution of the difference in the offset of intent statement and end of the turn. Here the difference is in **seconds**

offset of the intent statement i.e. last uttered word of the highlighted intent statement and the end of the turn (See Figure 1). This helps us to measure the average time saved when our approach is used. This clearly demonstrates that in **80%** of the remaining calls the time duration between the statement of intent and the end of the turn is less than **10 seconds**.

4 Approach

In this section, we first discuss the methodology used for the offline intent detection and then describe the real-time intent detection approach (See Figure 2).

4.1 Offline Intent Detection (Turn Level)

We use the standard LSTM architecture used in supervised classification problem to detect the intent in an offline way. We use an entire input sequence i.e. turn consisting of sequence of words $w = (w_1, w_2, \dots, w_T)$ and pass it to an embedding layer. The output from the embedding layer is then passed into a BiLSTM network which finally assigns an intent class c from a pre-defined set of intents, such that :

$$\hat{c} = \arg \max_c P(c|w) \quad (1)$$

4.2 Real-time Intent Detection

We use a similar technique proposed by [Shivakumar *et al.*2021] for our dataset. The approach comprises two tasks 1) Intent Boundary Detection and 2) Intent Detection which is described below. Both the tasks use similar architecture with slight modifications. To train the network in a real-time setting, we pass the input sequence consisting of words $w = (w_1, w_2, \dots, w_T)$ incrementally, i.e. one word at a time. The input word w_t is passed to an embedding layer followed by a uni-directional LSTM. Finally, we pass it to a dense layer and generate the output at each time step. Since the processing is performed incrementally, we store the hidden states generated by LSTM and forward them to the next word w_{t+1} .

- **Intent Boundary (IB) Detection :** This task detects whether the given input word is the intent boundary or not. Since there is no notion of boundary or where to stop, thus we independently train the network with this objective. We use the sigmoid activation function at the output since the task is the binary-classification problem while keeping the rest of the architecture the same. We use the binary focal loss described in Equation 2 as there is a huge class imbalance in the IB task.

$$L_{IB} = -\frac{\alpha}{T} \sum_{t=1}^T (y_t(1-p_t)^\gamma \log(p_t) + (1-y_t)(p_t)^\gamma \log(1-p_t)) \quad (2)$$

where t is the t^{th} time step, α and γ are the hyper-parameters of the focal loss [Lin *et al.*2017], y_t is the ground truth, and p_t is the output from sigmoid function.

- **Intent Detection :** We employ a similar method described above where we predict the intent for each word. At the final level, we use the softmax function since intent detection is a multi-class classification problem. We want to highlight that the intent is the same for the entire sequence, hence we mask the loss function with an indicator function and only compute loss at the Intent boundary. The loss function used here is :

$$L_{INT} = -\sum_{t=1}^T I_{IB} \sum_{c=1}^C y_{o,c} \log \hat{y}_{o,c} \quad (3)$$

where T is the length of the sequence, I_{IB} is equal to 1 at intent boundary, $y_{o,c}$ is the actual ground truth if o belongs to class c , $\hat{y}_{o,c}$ is the output from the network.

Multi-Task Training Objective

We train both the Intent Boundary (IB) and Intent detection jointly in a multi-task setting where we share an embedding layer and create separate task-specific LSTMs. We train the network with the following objective:

$$Loss = \beta L_{IB} + (1-\beta)L_{INT} \quad (4)$$

where β is the hyperparameter used to assign weightage to the tasks. We set the β to 0.5 after performing a detailed empirical evaluation.

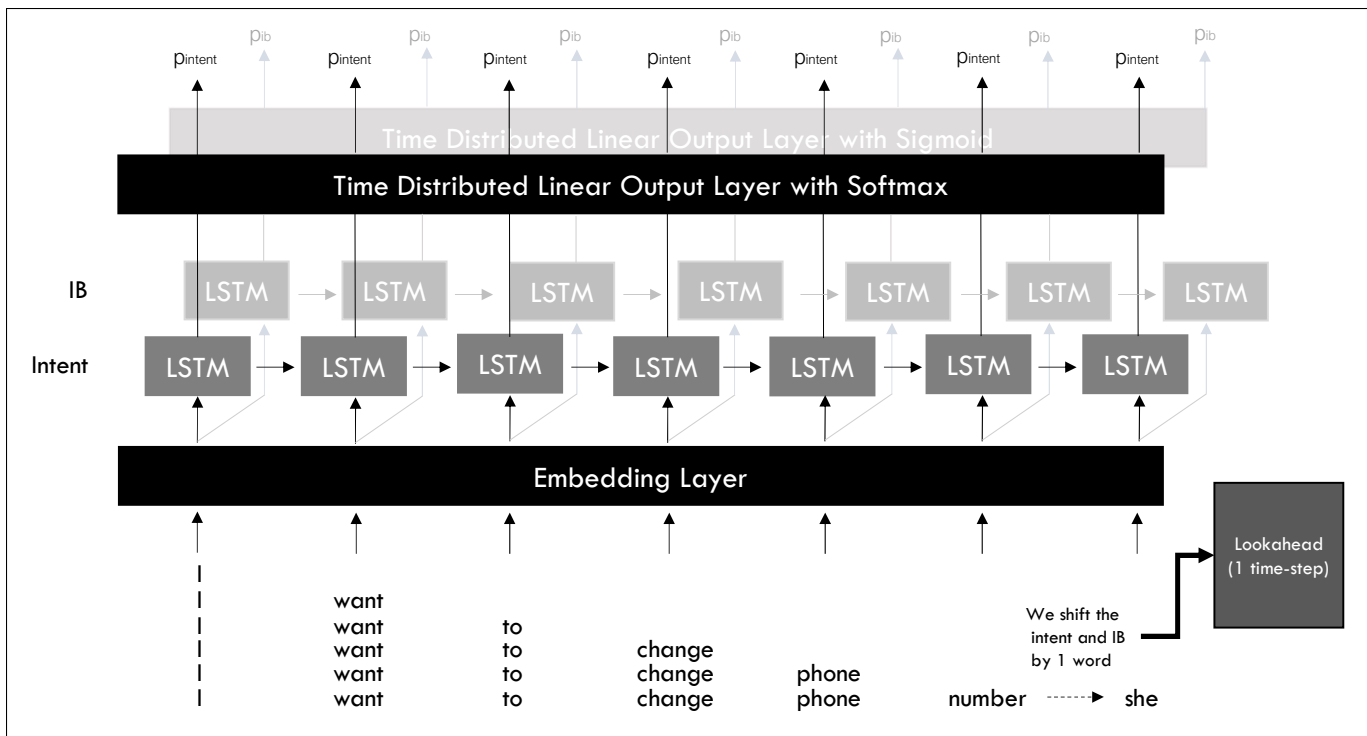


Figure 2: Real-Time Intent Detection Architecture

Lookahead based Intent Detection

Since the intent detection task is highly dependent on IB detection, improving boundary detection task leads to better performance. We analyze the impact that delaying prediction has on IB detection performance. To do so, we experiment by adding the option of a lookahead i.e. including the future right context where we shift the IB token and the intent token by k words. We evaluated lookahead of one, two and three and concluded that a lookahead of more than one word didn't correlate with better performance. Hence we set the lookahead to 1 (See Table 4). We would also like to note that in the cases where the IB is at the end of the turn, we compared appending for lookahead either random word or the first token of the next turn after the end of the turn. We found that the performance was identical, hence we use the randomly selected token for lookahead. We use the same methodology as described in the above section on the modified dataset.

Sentence	remove	my	grandma's	phone	number	she	got...
IB	0	0	0	0	1	0	0..
Intent	0	0	0	0	AM	0	0..
					→		
IB (lookahead)	0	0	0	0	0	1	0..
Intent (lookahead)	0	0	0	0	0	AM	0..

Table 4: Example sentence of lookahead based approach

Context based Detection

We further investigate the use of context in the intent detection task. We hypothesize that the previous context, i.e previous n turns, can assist the network in improving the performance of the tasks. For instance, based on our analysis we found

that intent is usually present in the turn following the agent's greeting message e.g.

AGENT: Hello my name is James and how can i help you today

CUSTOMER: Hi uh i wanted to change my phone number.

To encode the context, we take the previous 3¹ turns and concatenate them and pass them to an LSTM layer. We further initialize the network for the IB and intent detection with the hidden states generated from the context LSTM. We note that the embedding layer is shared across the tasks and the rest of the network is optimized using the method above.

5 Experimental Setup

We train both the tasks i.e 1)Intent Boundary Detection and 2) Intent Detection in a multi-task learning framework. The embedding layer is initialized randomly with the size 300 and is shared across both tasks. We use task-specific separate LSTMs with hidden dimensions of 128 and the number of layers as 2 for both tasks. We trained the network for 30 epochs using Adam optimizer with a learning rate of 0.001 and drop out as 0.25. We use the focal loss to optimize the IB task and set its hyperparameters α to 1 and γ to 8. We ran experiments multiple times to ensure model convergence. All models are implemented in Pytorch [Paszke *et al.*2019].

5.1 Evaluation Metrics

We perform evaluation at two levels:

¹Selected with empirical evaluation and manual analysis after trying 2, 3, 4

	Intent Boundary			Intent @ OB			Intent @ PB			Intent		
	P	R	F1	P	R	F1	P	R	F1	P	R	F1
Offline (Turn Level)	-	-	-	-	-	-	-	-	-	0.4271	0.3120	0.3463
With Intent Boundary												
Baseline	0.4671	0.4702	0.4686	0.8285	0.8122	0.8205	0.3521	0.4745	0.389	-	-	-
Baseline + Lookahead	0.430	0.572	0.491	0.8259	0.8183	0.8221	0.255	0.4136	0.3032	-	-	-
Baseline + Context	0.4659	0.623	0.5328	0.8551	0.8375	0.8364	0.2194	0.4152	0.2814	-	-	-
Without Intent Boundary												
Intent + Lookahead	-	-	-	-	-	-	-	-	-	0.399	0.5577	0.4532

Table 5: Evaluation of Intent Boundary (IB), Intent, Intent @ Oracle Boundary (OB), and Intent @ Predicted Boundary (PB) on the test set using precision (P), recall (R) and F1-score (F1).

Model Evaluation:

We evaluate both the tasks along with the combined performance by looking at precision, recall and f1-score computed on a test set composed only of turns containing an intent.

Real Time Evaluation: Since the goal is for the model to operate on conversations in real time, we analyze the intent detection performance on temporally unfolding conversations. Given a transcript, the customer turns are passed to the system in order, and each turn is processed one word at time. The hidden states of the network are reset before a new turn starts. We stop when an intent boundary and its associated intent class c is detected (See Algorithm 1). For the lookahead based approach, we also add a random token at the end of the turn. We evaluate the performance of the real-time processing using metrics defined below:

- **Accuracy:** It is computed by comparing the first predicted intent and the ground truth.
- **Accuracy at right turn:** It is computed by comparing the first predicted intent and the turn position with the ground truth intent and turn position.
- **Accuracy at right position:** It is computed by comparing the first predicted intent, turn position, and intent position within the turn with ground truth intent, turn position, and intent word position.
- **Mean Turn Difference (MTD):** The average difference between the ground truth intent turn position and the predicted intent turn position.

$$MTD = \frac{1}{N} \sum_{n=1}^N t_n - \hat{t}_n \quad (5)$$

where N is the number of samples in the evaluation set, t is the actual turn position and \hat{t} is the predicted turn position.

- **Mean Position Difference (MPD):** The average difference between the ground truth intent word position and the predicted intent word position.

$$MPD = \frac{1}{N} \sum_{n=1}^N p_n - \hat{p}_n \quad (6)$$

where N is the number of samples in the evaluation set, p is the actual word position and \hat{p} is the predicted word position.

Algorithm 1: Pseudo-code for the real-time evaluation

Input: transcripts t , model, true_intents

```

1  $pred\_intents \leftarrow \emptyset;$ 
2 foreach  $t_i \in t$  do
3   foreach  $turn \in t_i$  do
4      $previous\_intent\_state \leftarrow \emptyset;$ 
5      $previous\_ib\_state \leftarrow \emptyset;$ 
6     if  $turn[‘speaker’] == ‘customer’$  then
7       foreach  $token \in turn$  do
8          $o\_ib, score, previous\_ib\_state \leftarrow$ 
9            $model.predict\_ib(token, previous\_ib\_state)$ 
10        if  $o\_ib = 1 \ \& \ score > T$  then
11           $o\_int, previous\_intent\_state$ 
12             $\leftarrow model.predict\_intent(token,$ 
13               $previous\_intent\_state)$ 
14           $pred\_intents.add(o\_int);$ 
15          break {turn loop};
16  $Get\_Metrics(true\_intents, pred\_intents)$ 

```

6 Results & Discussion

We evaluate our approach using the metrics defined in Section 5.1 The overall results are described in two parts 1) Evaluation of the test dataset created using the intent statements from the transcripts 2) Evaluation of the test data by simulating the actual conversation in a real time environment.

The results for part 1 are illustrated in Table 5 using the methods described in Section 4. We break the results task wise where we describe the precision, recall, and F1 score for Intent Boundary (IB) task, Intent @ oracle boundary and Intent @ Predicted boundary. We can clearly see that incorporating lookahead yields a better F1-score than the baseline method in the IB task since we also provide the future context to the model. Further, we see that adding context yields the best performance in the Intent Boundary task and Intent @ OB task, but has the worst F1 score of intent @ PB which is the metric we desire for better performance when used in real time setting.

We also investigate the model performance by excluding the Intent Boundary objective. In contrast to the method described in Section 4, we compute the loss at each time step since we do not train with explicit intent boundary in this case. We achieve an F1-score of 0.4532 which is computed based on c

	Real-Time Evaluation				
	Acc	Acc @ RT	Acc @ RP	MTD	MPD
Offline (Turn Level)	0.7697	0.7236	-	-0.78	-
With Intent Boundary					
Baseline	0.7763	0.7631	0.3618	-0.59	-13.56
Baseline + Lookahead	0.7894	0.7236	0.4539	-0.88	-19.56
Baseline + Context	0.75	0.6578	0.4539	-0.344	-7.76
Without Intent Boundary					
Intent + Lookahead	0.8355	0.6842	0.4934	2.37	44.66

Table 6: Evaluation on the test set by simulating the conversation in real-time using accuracy (Acc), accuracy at right turn (Acc @ RT), accuracy at right position (Acc @ RP), mean turn difference (MTD), and mean position difference (MPD)

classes while excluding the O class. Similar to the previous method, we also incorporate the lookahead where we shift the intent towards the right. Finally, we report the results on the offline methodology described in Section 4.1 where we classify the entire turn into one of the c classes.


Further, we see that the addition of context did not improve the performance proving our initial hypothesis wrong. We achieve an accuracy of **75 %** which is the worst amongst all. One of the reasons could be the long sequences due to the addition of context which negatively affected the LSTM performance. However, we want to investigate this in detail by encoding the context using the more advanced methods and comparing it with other datasets as well.

Table 6 illustrates the results where we perform the online i.e. real-time evaluation. For evaluation, we use the same test dataset to compare with previous results and avoid data leakage issues. Similar to the previous results, incorporating lookahead yields better performance. We achieve an accuracy of **83.55%** without the IB task outperforming the rest of the models.


We see a consistent performance in the accuracy at the right turn and accuracy at the right position as well. We also describe the results using the Mean turn difference (MTD) and Mean Position Difference (MPD) which essentially explains the early or late prediction of the intents. Table 6 demonstrates that the inclusion of the intent boundary task helps in predicting the intents much before the actual intent boundary. Our best model can predict **0.88 turns** and **19.56 words** earlier reducing the overall latency. Further, we see that without the explicit IB task, there is a delay of **2.37 turns** in the prediction even though we achieve the best accuracy of 83 %. We also plot the distribution of the predicted turn difference w.r.t to the actual turn. Figure 3 clearly demonstrates the importance of including IB for the early prediction. This explains the trade-off between the system with the overall higher accuracy and the system that reduces the latency.

7 Conclusion and Future Work

In this work, we present a systematic method to detect intent in real-time in human-human customer support conversations. We propose a dual architecture approach to detect intent and intent boundary by incorporating the lookahead. We also investigate various methodologies such as the adding context, excluding the IB task, etc. We discuss the results on our private



(a) With IB



(b) Without IB

Figure 3: Distribution of the turn difference in the predicted IB and actual IB with and without using Intent Boundary (IB). Here the difference is in #turns. Blue color bar represents the early prediction and orange color represents the late prediction

dataset comprised of text transcripts of human-human conversations from the telecom domain. To evaluate the methodology in real-time, we also propose metrics such as Mean turn difference (MTD) and Mean Position difference (MPD). We conclude that the intent boundary task significantly helps in reducing latencies by predicting the intents earlier. Our method achieves the best overall intent accuracy of **83.55 %** when IB task is excluded.

In the future, we would like to use our method to reproduce this study with other datasets and ensure that our method is consistent. Further, we also want to explore more advanced deep learning techniques such as transformers to encode the sequences and context. We also want to extend this work to predict the entities in real-time.

References

- [Adam *et al.*, 2021] Martin Adam, Michael Wessel, and Alexander Benlian. Ai-based chatbots in customer service and their effects on user compliance. *Electronic Markets*, 31(2):427–445, June 2021. Published online:17. March 2020.
- [Chen *et al.*, 2019] Qian Chen, Zhu Zhuo, and Wen Wang. BERT for joint intent classification and slot filling. *CoRR*, abs/1902.10909, 2019.
- [Constantin *et al.*, 2019] Stefan Constantin, Jan Niehues, and Alex Waibel. Incremental processing of noisy user utterances in the spoken language understanding task. *CoRR*, abs/1909.13790, 2019.
- [Devlin *et al.*, 2018] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. *CoRR*, abs/1810.04805, 2018.
- [Guo *et al.*, 2014] Daniel Guo, Gokhan Tur, Wen-tau Yih, and Geoffrey Zweig. Joint semantic utterance classification and slot filling with recursive neural networks. In *2014 IEEE Spoken Language Technology Workshop (SLT)*, pages 554–559, 2014.
- [Haihong *et al.*, 2019] E. Haihong, Peiqing Niu, Zhongfu Chen, and Meina Song. A novel bi-directional interrelated model for joint intent detection and slot filling. In *ACL*, 2019.
- [Kahardipraja *et al.*, 2021] Patrick Kahardipraja, Brielen Madureira, and David Schlangen. Towards incremental transformers: An empirical analysis of transformer models for incremental NLU. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih, editors, *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 1178–1189. Association for Computational Linguistics, 2021.
- [Katharopoulos *et al.*, 2020] A. Katharopoulos, A. Vyas, N. Pappas, and F. Fleuret. Transformers are rnns: Fast autoregressive transformers with linear attention. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2020.
- [Kim *et al.*, 2017] Young-Bum Kim, Sungjin Lee, and Karl Stratos. Onenet: Joint domain, intent, slot prediction for spoken language understanding. pages 547–553, 12 2017.
- [Lin *et al.*, 2017] Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. *CoRR*, abs/1708.02002, 2017.
- [Ma *et al.*, 2019] Mingbo Ma, Liang Huang, Hao Xiong, Renjie Zheng, Kaibo Liu, Baigong Zheng, Chuanqiang Zhang, Zhongjun He, Hairong Liu, Xing Li, Hua Wu, and Haifeng Wang. STACL: Simultaneous translation with implicit anticipation and controllable latency using prefix-to-prefix framework. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3025–3036, Florence, Italy, July 2019. Association for Computational Linguistics.
- [Madureira and Schlangen, 2020] Brielen Madureira and David Schlangen. Incremental processing in the age of non-incremental encoders: An empirical assessment of bidirectional models for incremental NLU. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 357–374, Online, November 2020. Association for Computational Linguistics.
- [Paszke *et al.*, 2019] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- [Qin *et al.*, 2019] Libo Qin, Wanxiang Che, Yangming Li, Haoyang Wen, and Ting Liu. A stack-propagation framework with token-level intent detection for spoken language understanding. In *EMNLP*, 2019.
- [Schlangen and Skantze, 2009] David Schlangen and Gabriel Skantze. "a general, abstract model of incremental dialogue processing". volume 2, pages 710–718, 01 2009.
- [Shivakumar *et al.*, 2021] Prashanth Gurunath Shivakumar, Naveen Kumar, Panayiotis G. Georgiou, and Shrikanth Narayanan. RNN based incremental online spoken language understanding. In *IEEE Spoken Language Technology Workshop, SLT 2021, Shenzhen, China, January 19-22, 2021*, pages 989–996. IEEE, 2021.
- [Wu *et al.*, 2020] Di Wu, Liang Ding, Fan Lu, and Jian Xie. 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, November 2020. Association for Computational Linguistics.
- [Zhang and Wang, 2016] Xiaodong Zhang and Houfeng Wang. A joint model of intent determination and slot filling for spoken language understanding. In *IJCAI*, 2016.