Modeling Long-Range Context for Concurrent Dialogue Acts Recognition

Yue Yu*
yy476@georgetown.edu
Department of Computer Science
Georgetown University

Siyao Peng*
sp1184@georgetown.edu
Department of Linguistics
Georgetown University

Grace Hui Yang
huiyang@cs.georgetown.edu
Department of Computer Science
Georgetown University

ABSTRACT

In dialogues, an utterance is a chain of consecutive sentences produced by one speaker which ranges from a short sentence to a thousand-word post. When studying dialogues at the utterance level, it is not uncommon that an utterance would serve multiple functions. For instance, “Thank you. It works great.” expresses both gratitude and positive feedback in the same utterance. Multiple dialogue acts (DA) for one utterance breeds complex dependencies across dialogue turns. Therefore, DA recognition challenges a model’s predictive power over long utterances and complex DA context. We term this problem Concurrent Dialogue Acts (CDA) recognition. Previous work on DA recognition either assumes one DA per utterance or fails to realize the sequential nature of dialogues. In this paper, we present an adapted Convolutional Recurrent Neural Network (CRNN) which models the interactions between utterances of long-range context. Our model significantly outperforms existing work on CDA recognition on a tech forum dataset.

ACM Reference Format:

Yue Yu, Siyao Peng, and Grace Hui Yang. 2019. Modeling Long-Range Context for Concurrent Dialogue Acts Recognition. In *The 28th ACM International Conference on Information and Knowledge Management (CIKM '19)*, November 3–7, 2019, Beijing, China. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3357384.3358145>

1 INTRODUCTION

An utterance is a sequence of sentences that are produced by the same speaker in his or her turn in a dialogue. Dialogue Acts (DA) aim to capture the functionality of these utterances. Recognizing DAs can benefit a dialogue system in many ways. 1) A pre-trained DA model can reduce the number of utterance annotations required to train a dialogue system [11]. 2) For retrieval-based dialogue systems, DA modeling can provide meaningful patterns to reduce the search space and thus expedite response selection. Even though DA recognition is a well-studied problem, most contemporary methods concentrate on conversations of short utterances with usually a single DA per utterance, the case of the Switchboard Dialogue Act

*Both authors contributed equally to this research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CIKM '19, November 3–7, 2019, Beijing, China
© 2019 Association for Computing Machinery.
ACM ISBN 978-1-4503-6976-3/19/11...\$15.00
<https://doi.org/10.1145/3357384.3358145>

Table 1: Taxonomy of 12 dialogue acts in a tech forum.

DA	Taxonomy	%	DA	Taxonomy	%
GG	Greetings/Gratitude	40.0	FQ	Follow-up Question	8.6
PA	Potential Answer	39.7	NF	Negative Feedback	7.6
FD	Further Details	24.8	CQ	Clarifying Question	7.5
OQ	Original Question	23.3	RQ	Repeat Question	6.1
PF	Positive Feedback	10.7	JK	Junk	2.6
IR	Information Request	10.7	O	Others	1.5

U1	1 For around a month my settings would disappear ... If somebody know how to fix this , please tell me . [OQ]	
	2 Hi ... to isolate the issue ... we would like to know the following ... look forward to your response . [GG IR PA]	A1
U2	3 I want to reinstall Microsoft bingo for free I have Windows 10 I can't even open my store app ... [CQ FD]	
U2	4 I did n't have this issue before [FD]	
U2	5 There were some issues made [FD]	
	6 ... you would also have the (simpler) alternative of using a Win+Shift+Cursor move . Good luck ... [GG PA]	A2
U1	7 ...The only changes I have made prior to the issue was the installation of new RAM cards ... Windows updates ... [FD]	
U1	8 Thank you ... This works [GG PF]	
U1	9 I manage to fix the situation ... thank you ... [GG PF]	

Figure 1: Dialogue between two users and two agents.

(SwDA) corpus [5]. However, when utterances are long and complex, one DA label would not be sufficient to describe the pragmatic function of an utterance. As a result, multiple DA labels would be required and the approaches on SwDA would not apply. We term our problem Concurrent Dialogue Acts (CDA) Recognition.

Forum discussion is a type of dialogue that breeds CDAs. The sample in Figure 1 from MSDialog-Intent corpus [14] discusses a technical issue with Microsoft Bingo among two users (U1 & U2) and two agents (A1 & A2). For instance, utterance 6 is simultaneously labeled as GG for “Good Luck” and PA for “using a Win+Shift+Cursor move” (DA labels are shown in Table 1). 59.95% of utterances in MSDialog-Intent are indeed multi-labeled.

Another important characteristic of these dialogues is that a speaker can refer to any of the preceding utterances. Thus, it is unfeasible to determine a fixed window size that would incorporate all relevant context. In Figure 1, utterance 6 provides a potential answer (PA) to utterance 1; utterance 7 provides further details (FD) to the information request (IR) in utterance 2. Without the capability of capturing a wider range of utterances, the chance of recognizing DAs in these forum conversations would be reduced.

However, current approaches for DA recognition only attempt to capture either of the two characteristics but not both. 1) Sequence models, that can capture far-away context information, are often applied to DA sequence labeling. Kumar et al. [10] build a hierarchical RNN using Bi-LSTM as a base unit and CRF as the last layer; Chen et al. [2] use the CRF-Attentive Structured Network for the same hierarchical nature of DAs. Nonetheless, these models only concern single-label DA datasets (e.g., SwDA) and heavily rely on CRF, which is incapable of predicting multiple labels for each utterance. 2) Another line of research casts DA recognition as a multi-label classification problem to accommodate the CDA scenario. Qu et al. [14] apply a CNN-based text classifier proposed by Kim [8] using a fixed window to represent the context. Although capable of classifying utterances with CDAs, Qu et al. [14]’s model only concerns a strictly-local context range and thus cannot include distant information.

In this paper, we present a novel neural model that is adapted from Convolutional Recurrent Neural Network (CRNN) to both incorporate the interaction between distant utterances and generalize the DA recognition task to accommodate CDA. Our contributions can be summarized as follows: 1) In our adapted Convolutional Recurrent Neural Network (CRNN) we use the recurrent layers to gather long-range contextual information that are extracted from utterances by the convolutional layers. 2) We further optimize the model by incorporating highway connections and dynamic k -max pooling. 3) Our model significantly outperforms the state-of-the-art model by 4.68% in accuracy and 3.08% in F-score.

2 RELATED WORK

CRNN has been widely applied to classification tasks. In Natural Language Processing (NLP) research including sentence segmentation [16], sentiment analysis [17] and discourse modeling [6], CRNN achieves the state-of-the-art in single-label classification where for each input, a single output would be predicted.

Deep learning research outside of NLP proves CRNN suitable for classifications where multiple labels may be assigned to each target. Cakir et al. [1] apply CRNN to Polyphonic Sound Event Detection where multiple sound events can be detected in a timestamp. Similar to single-label CRNNs, features are extracted by the convolutional layers and further integrated into the recurrent layers to provide context information. The sigmoid activation function is used to make multi-binary decisions for each event activity. Choi et al. [4] treat music tagging as a multi-label classification task where a subset of genres, moods, etc. are assigned to each music. In Choi et al. [4]’s model, The CNN is responsible for extracting local features and the RNN presents a temporal summarization. These studies substantiate CRNN’s applicability to multi-label classification tasks, especially its ability to extract features from the current step and to integrate features from arbitrarily long distance.

In previous dialogue research, Kalchbrenner and Blunsom [6] and Kim et al. [7] incorporate CRNN into a single-label DA recognition and dialogue topic tracking system. Based on this line of research, we argue that CRNN can be applied for CDA recognition with two adaptations: binary cross-entropy loss (BCE) and sigmoid activation. Further experiments also show that adding highway connections and dynamic k -max pooling further boosts the performance.

3 TASK DEFINITION

The task is defined as a CDA recognition problem where for each utterance u_t (the t -th utterance) in a dialogue, we predict a subset of DA labels y_t that describes the functionality of the utterance from a candidate set of DA labels $\mathcal{L} = \{l_1, l_2, \dots, l_c\}$. For a dialog with s utterances, the inputs to the algorithm is $\mathcal{U} = \{u_1, u_2, \dots, u_s\}$, and the output is $\mathcal{Y} = \{y_1, y_2, \dots, y_s\}$, where y_t is the annotated DA label set for u_t , in which $y_t = \{y_t^1, y_t^2, \dots, y_t^c\}$. Here, $y_t^j = \{1, 0\}$ denotes whether the t -th utterance of the dialog is labeled with DA label l_j or not. When $\sum_{j=1}^c y_t^j > 1$, we say CDAs are recognized. Given a dialogue $\mathcal{U}$, the goal is to predict the DA sequence $\mathcal{Y}$ from the text.

4 THE PROPOSED APPROACH

The challenge of this task lies in the complexity of dialogue structures in human conversations where an utterance can express multiple DAs. In this work, we improve CDA recognition with an adapted CRNN which models the interactions between long-range context.

4.1 Convolutional Layer

The base of our architecture is a CNN module similar to Kim [8]. The module works by ‘sliding’ through the embedding matrix of an utterance with various filter sizes to capture semantic features in differently ordered n -grams. A convolution operation is denoted as

$$k_i = \tanh(\mathbf{w} \cdot \mathbf{x}_{i:i+d-1} + b_k) \quad (1)$$

where k_i is the feature generated by the i -th filter with weights $\mathbf{w}$ and bias b_k . This filter of size d is applied to an embedding matrix, which is the concatenation from the i -th to the $(i + d - 1)$ -th embedding vectors. This operation is applied to every possible window of words in an utterance of length n and generates a feature map $\mathbf{k}$.

$$\mathbf{k} = [k_1, k_2, \dots, k_{n-d+1}] \quad (2)$$

4.2 Dynamic k -Max Pooling

A max-over-time pooling operation [8] is usually applied over the feature map and takes the maximum value as the feature corresponding to this particular filter. The idea is to capture the most important features of an utterance. However, this mechanism could be problematic when the utterance is long. We experimented with Dynamic k -Max Pooling [12] to pool the most powerful features from p sub-sequences of an utterance with m words. This pooling scheme naturally deals with variable utterance length.

$$p(\mathbf{k}) = \left[\max \left\{ \mathbf{k}_{1:\lfloor \frac{m}{p} \rfloor} \right\}, \dots, \max \left\{ \mathbf{k}_{\lfloor m - \frac{m}{p} + 1 \rfloor : m} \right\} \right] \quad (3)$$

4.3 Recurrent Layer

Based on the local textual features extracted from each utterance, a bidirectional RNN is applied to gather features from a wider context for recognizing the DAs in the target utterance, u_t . We experimented with two variations of RNN: LSTM [18] and GRU [3], both of which utilize gating information to prevent the vanishing gradient problem. GRU is constructed similarly to LSTM but without using a memory cell. It exposes the full hidden state without any control, which may be computationally more efficient.

Figure 2: The proposed CRNN model architecture.

We experimented with both since it is difficult to predict which one performs better on our task.

4.4 Highway Connection

Although LSTM and GRU help capture wider context, the network training becomes more difficult with the additional recurrent layers. Inspired by the highway networks [15], we propose to add a highway connection between the convolutional layer and the last fully connected layer. With this mechanism, the information about the target utterance, u_t , can flow across the recurrent layer without attenuation. The last fully connected layer learns from the outputs of both recurrent and convolutional layers.

5 DATASET

We use the MSDialog-Intent dataset [14] to conduct experiments. In the dataset, each of the 10,020 utterances is annotated with a subset of 12 DAs. The abundance of information in a single utterance (avg. 72 tokens/utterance) breeds CDA (avg. 1.83 DAs/utterance). We observe a strong correlation between the number of DAs and utterance length, which necessitates a CDA model for forum conversations.

The dataset includes plenty of metadata for each utterance, e.g., answer vote and user affiliation. For generalizability, our model only incorporates textual content of the dialogues. Besides, unlike Qu et al. [14], we keep all the DA annotations in the dataset to preserve the meaningful DA structures within and across utterances.¹

6 EXPERIMENTS

In this section, three versions of our proposed model with incremental improvements are evaluated against a CNN baseline [8] and the state-of-the-art approach for CDA recognition [14].

- **CNN-Kim[8]:** One of the first attempts to apply CNN to text classification. The CNN model consists of three convolutional layers with the same filter size.
- **CNN-CR[14]:** The state-of-the-art approach for CDA recognition on the MSDialog-Intent dataset [14]. The CNN model incorporates context information with a window size of 3.
- **CRNN (v_1):** Our base model that adapts CRNN for CDA recognition using BCE loss and sigmoid activation function.

¹Qu et al. [14] simplified the task by removing DA labels from certain utterances and reducing rare DA combinations to top 32 frequent ones. We re-ran Qu et al. [14]’s model on the dataset preserving all DA combinations.

- **CRNN (v_2):** CRNN (v_1) with highway connections added between the convolutional layer and the fully connected layer.
- **CRNN (v_3):** CRNN (v_1) with highway connections and dynamic k -max pooling implemented.

6.1 Experimental Setup

The dataset is partitioned into training, validation, and test sets in the ratio of 8:1:1. Hyper-parameters are tuned with the validation set. For word embedding, we use the publicly available GloVe vectors that were trained on Wikipedia + Gigaword 5 [13]. We find that setting (CNN filters, CNN dropout, RNN units, RNN layers, RNN dropout and k) to (100, 0.4, 900, 2, 0.15 and 2) is the best for our model. The BCE is used as loss function and Adam [9] is used for optimization with an initial learning rate, 0.001. The models were trained on four GeForce RTX 2080 Ti GPUs.

6.2 Metrics

Following previous work [14] on multi-label classification, we adopt label-based accuracy (i.e., Hamming score) and micro- F_1 score as our main evaluation metrics. Micro-precision and micro-recall are also reported to assist the analysis. Among all, accuracy is the only metrics that is on a per utterance basis. Therefore, Student’s paired t-test is performed only on accuracy. Other metrics (P, R, F_1) provide an overall performance evaluation for all utterances.

6.3 Results

All proposed adaptations, i.e., highway connections and dynamic k -max pooling, contribute to the model, as in Table 2. The CRNN models significantly² outperform both baselines. The model with full adaptation (v_3), performs the best across all experiments: achieving the highest accuracy, recall and F_1 with LSTM and the highest precision with GRU. The divergence between LSTM and GRU may indicate that the *forget* and *output* gates in LSTM are helpful for the retrieval of related memory but increase the variance of the model.

Table 3 further exemplifies the capability of our best model in recognizing CDAs. Our model systematically outperforms the state-of-the-art approach, CNN-CR. More specifically, 1) CRNN achieves better mean accuracy for all reference sizes and significantly for

²* denotes a statistically significant difference over the best baseline (CNN-CR) with $p \leq 0.05$, ** for $p \leq 0.01$, *** for $p \leq 0.001$ and **** for $p \leq 0.0001$.

Table 2: Performance of CNN-Kim, CNN-CR, and CRNN.²

Models	Accuracy	Precision	Recall	F_1 score
CNN-Kim[8]	0.5785	0.6371	0.6745	0.6553
CNN-CR[14] ¹	0.6354	0.7108	0.6952	0.7029
CRNN (v_1) w/ LSTM	0.6668*	0.7238	0.7297	0.7267
CRNN (v_1) w/ GRU	0.6543*	0.7056	0.7065	0.7061
CRNN (v_2) w/ LSTM	0.6731****	0.7315	0.7315	0.7315
CRNN (v_2) w/ GRU	0.6734**	0.7280	0.7334	0.7307
CRNN (v_3) w/ LSTM	0.6822****	0.7254	0.7422	0.7337
CRNN (v_3) w/ GRU	0.6733***	0.7358	0.7215	0.7286

Table 3: Mean accuracy and the average number of predicted DAs grouped by the number of reference DAs.² The percentage indicates the frequency of each DA group in the test set.

# of ref DAs	%	Mean accuracy		Avg. num. of pred DAs	
		CRNN (v_3)	CNN-CR	CRNN (v_3)	CNN-CR
1	36.9	0.7704**	0.7126	1.44	1.44
2	42.8	0.6641***	0.6232	2.02**	1.89
3	16.7	0.5596*	0.5177	2.56***	2.37
≥4	3.6	0.5618	0.5339	2.68	2.74

Figure 3: Mean accuracy of CRNN (v_3) vs. CNN-CR.

sizes 1–3, which consist of 96.4% utterances. 2) The average number of DAs predicted by CRNN is the closest to the reference.

Another advantage of CRNN is the ability to handle longer dialogues. Figure 3 visualizes the mean accuracy of utterances grouped by dialogue length. Along with the increase of dialogue length, we observe an increasing advantage of CRNN (in blue) over CNN-CR (in orange), especially for dialogues longer than 6 utterances.

6.4 A Case study

Table 4 shows an example where CRNN and CNN-CR predict differently for utterance 5. Utterance 5 should respond to utterance 2, establishing a long-range dependency that cannot be captured by CNN-CR’s fixed context window. Our CRNN model succeeds in recognizing the distant context for further details (FD). Furthermore, CNN-CR wrongly recognizes utterance 5 as original question (OQ) since the OQ in utterance 1 is out of range.

7 CONCLUSION

Our proposed CRNN models for CDA recognition impose fewer restrictions on the structure of DAs and capture textual features from a wider context. The experiment results show that all of the proposed adaptations, i.e., highway connections and dynamic k -max pooling, contribute to the model optimization. Our final model

Table 4: Predictions of CRNN (v_3) and CNN-CR on a dialogue with long-range dependencies.

#	Utterance	Ref	CRNN	CNN-CR
1	How can I download Skype for Windows 8.1 ...	OQ	OQ	OQ
2	Hi ... if you are using a phone running Windows 8.1 ... no longer ... but if ... computer ... you can ...	GG PA	GG PA	GG PA
... Utterances 3 & 4 ...				
5	Hi , I am using surface tablet 8.1 Windows ... tried many times ... this app can't run on this PC	FD FQ GG	FD GG PA RQ	OQ GG

significantly outperforms the state-of-the-art approach on a tech forum dataset where the dialogues are packed with complex DA structures and information-rich utterances. Future work will consider a more fine-grained set of DAs for a deeper analysis of CDA modeling. It would also be interesting to annotate antecedent relations among utterances to structure threading dialogue flows.

8 ACKNOWLEDGEMENTS

This research is supported by U.S. National Science Foundation IIS-145374. Any opinions, findings, conclusions, or recommendations expressed in this paper are of the authors, and do not necessarily reflect those of the sponsor.

REFERENCES

- [1] E. Cakir, G. Parascandolo, T. Heittola, H. Huttunen, and T. Virtanen. 2017. Convolutional Recurrent Neural Networks for Polyphonic Sound Event Detection. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 25 (2017).
- [2] Z. Chen, R. Yang, Z. Zhao, D. Cai, and X. He. 2018. Dialogue Act Recognition via CRF-Attentive Structured Network. In *SIGIR '18*. ACM, 225–234.
- [3] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078* (2014).
- [4] K. Choi, G. Fazekas, M. Sandler, and K. Cho. 2017. Convolutional recurrent neural networks for music classification. In *ICASSP '17*. 2392–2396.
- [5] D. Jurafsky. 1997. Switchboard SWBD-DAMSL shallow-discourse-function annotation coders manual. *ICS'97* (1997).
- [6] N. Kalchbrenner and P. Blunsom. 2013. Recurrent convolutional neural networks for discourse compositionality. *ACL '13* (2013).
- [7] S. Kim, R. Banchs, and H. Li. 2016. Exploring Convolutional and Recurrent Neural Networks in Sequential Labelling for Dialogue Topic Tracking. In *ACL '16*.
- [8] Y. Kim. 2014. Convolutional Neural Networks for Sentence Classification. In *EMNLP'14*. 1746–1751.
- [9] D. Kingma and J. Ba. 2015. Adam: A method for stochastic optimization. In *ICLR'15*.
- [10] H. Kumar, A. Agarwal, R. Dasgupta, and S. Joshi. 2018. Dialogue act sequence labeling using hierarchical encoder with crf. In *AAAI'18*.
- [11] H. Kumar, A. Agarwal, and S. Joshi. 2018. Dialogue-act-driven Conversation Model: An Experimental Study. In *COLING'18*. 1246–1256.
- [12] J. Liu, W.-C. Chang, Y. Wu, and Y. Yang. 2017. Deep Learning for Extreme Multi-label Text Classification. In *SIGIR '17*. New York, NY, USA, 115–124.
- [13] J. Pennington, R. Socher, and C. D. Manning. 2014. GloVe: Global Vectors for Word Representation. In *EMNLP'14*. 1532–1543.
- [14] C. Qu, L. Yang, B. Croft, Y. Zhang, J. R. Trippas, and M. Qiu. 2019. User Intent Prediction in Information-seeking Conversations. *CHIIR '19* (2019).
- [15] P. Sanders and D. Schultes. 2005. Highway hierarchies hasten exact shortest path queries. In *European Symposium on Algorithms*. Springer, 568–579.
- [16] M. Treviso, C. Shulby, and S. Aluisio. 2017. Sentence Segmentation in Narrative Transcripts from Neuropsychological Tests using Recurrent Convolutional Neural Networks. In *EACL '17*. Association for Computational Linguistics, 315–325.
- [17] X. Wang, W. Jiang, and Z. Luo. 2016. Combination of Convolutional and Recurrent Neural Network for Sentiment Analysis of Short Texts. In *COLING '16*.
- [18] W. Zaremba and I. Sutskever. 2014. Learning to execute. (2014).