

# Joint Learning of Dialog Act Segmentation and Recognition in Spoken Dialog Using Neural Networks

Tianyu Zhao and Tatsuya Kawahara

{zhao, kawahara}@sap.ist.i.kyoto-u.ac.jp

## Abstract

Dialog act segmentation and recognition are basic natural language understanding tasks in spoken dialog systems. This paper investigates a unified architecture for these two tasks, which aims to improve the model’s performance on both of the tasks. Compared with past joint models, the proposed architecture can (1) incorporate contextual information in dialog act recognition, and (2) integrate models for tasks of different levels as a whole, i.e. dialog act segmentation on the word level and dialog act recognition on the segment level. Experimental results show that the joint training system outperforms the simple cascading system and the joint coding system on both dialog act segmentation and recognition tasks.

## 1 Introduction

Recently the burst of interactive assistants and chatbots leads to an increasing interest of dialog systems. Natural language understanding (NLU), as an important component of dialog system, is usually responsible for dialog act (DA) or dialog intent tagging, where text classification techniques are necessary. Dialog act (also speech act) is a representation of the meaning of a sentence at the level of illocutionary force (Stolcke et al., 2000). For instance, a sentence “*How is the weather?*” belongs to the dialog act class *Question*. For another sentence, “*The weather is quite good today.*”, if it follows a previous *Question* sentence, it should be an *Answer* to the question. Otherwise it is likely to be a *Statement*. Therefore, DA recognition requires us to understand the sentence from semantic, pragmatic and syntactic aspects, and its context plays an important role as well.

<b>Words</b>	hi	my	name	is	Erica
<b>Segment</b>	E	I	I	I	E
<b>DA</b>	Greeting	Statement			

Table 1: DA segmentation and recognition.

The prerequisite for DA recognition is to split a sequence of words into segments, each of which corresponds to one DA unit. Especially for spoken dialog systems, NLU is based on Automatic Speech Recognition (ASR) hypotheses or transcripts, in which we cannot make any assumption of punctuation such as periods, commas and question marks for segmentation. Therefore DA segmentation becomes essential for spoken dialog systems. As the example given in Table 1 shows, a long utterance is firstly split into two segments “*hi*” and “*my name is Erica*”, to which DA tags “*Greeting*” and “*Statement*” are assigned afterwards. DA segmentation is a sequence labeling task and an “IE” tag coding scheme is adopted to describe segment boundaries, where “I” denotes “inside” of a segment and “E” denotes the “end”. While the DA segmentation is a pre-process of DA recognition, recognition of DA in the sequence helps segmentation. Thus, DA segmentation and recognition are two highly related tasks. We can expect that joint learning of these two tasks can improve the performance of models. In this paper we investigate architectures of joint learning of DA segmentation and recognition, and analyze their performances. DA segmentation is a sequence labeling task on the word level and DA recognition is a classification task on the DA segment level. Our model is flexible and can be applied to tasks of different levels.

The rest of this paper is structured as follows. Section 2 gives a literature review. Section 3 describes a hierarchical neural network model for

DA recognition and explains its ability to model sequential short texts. Section 4 focuses on joint learning of DA segmentation and recognition, and explores three models for joint tasks. Section 5 shows experimental results on three tasks using the proposed models. Lastly in Section 6 and 7 we discuss the results and give a conclusion.

## 2 Related Works

In the task of DA recognition, Shriberg et al. (1998) applied decision tree using rich features and emphasized the importance of prosodic features. Stolcke et al. (2000) used HMM to capture the intrinsic patterns of DA sequence. Variants of neural network have been recently used in this task. In (Lai et al., 2015), a recurrent convolutional neural network is applied to text classification, which consists of a RNN layer and a max-pooling layer over it. Ji et al. (2016) proposed a latent variable RNN for modeling discourse relations between sentences. Khanpour et al. (2016) investigated RNNs with different settings of hyperparameters. A hierarchical neural network was introduced by Lee and Derroncourt (2016). It firstly uses a RNN layer or a CNN layer to generate vector representations of short texts. Then a two-layer feedforward ANN takes a sequence of these vector representations to predict the probability distribution of output labels. Li and Wu (2016) used gated RNN for both vector representation generation and classification in their hierarchical model. We base our work on hierarchical neural network for DA recognition and propose a unified architecture for joint DA segmentation and recognition, which is discussed in Section 4.

## 3 Hierarchical Neural Network

When we regard DA recognition as a text classification task, there are two difficulties in accurately recognizing a DA. In the first place, texts in DA recognition are often limited to a small number of words while tasks such as sentiment analysis and news topic categorization aim to classify fairly long documents and can exploit mainly n-gram models. Compared with long documents, dialog utterances have much fewer words and it is difficult to extract enough information from simple word co-occurrence features. Secondly, it is of great importance to consider contexts in DA recognition. For instance, a sentence “the weather is quite good today” is regarded as an *Answer* if

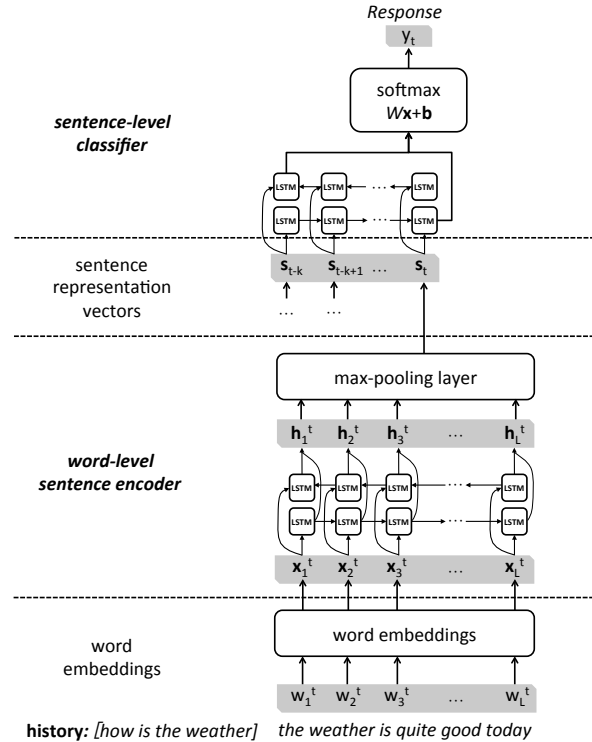


Figure 1: Hierarchical neural network: an example of input and output is given in the figure.

it appears after another speaker questioning the weather. Otherwise, it is a *Statement*.

In order to address aforementioned problems, we use a hierarchical neural network for DA recognition. The hierarchical model firstly takes distributed word representation as input, which contains richer semantic information than n-gram features do. Secondly, it exploits history information to recognize DA tags of ambiguous utterances such as “the weather is quite good today”. The general architecture of the hierarchical neural network consists of a sentence encoder and a classifier. A sentence encoder neural network encodes a sequence of words into a vector (sentence representation vector) of a fixed length, which will be explained in Section 3.1. A classifier neural network predicts the label given representation vectors of the corresponding sentence and its preceding sentences, which will be explained in Section 3.2. The architecture of hierarchical neural network used in our work is shown in Figure 1.

### 3.1 Sentence Representation

A sentence encoder encodes a sequence of words into a fixed-length vector. By training the encoder, it obtains the ability to mine useful task-related information from a word sequence. We choose Bidi-

rectional Long Short-Term Memory (BiLSTM) - a variant of RNN. LSTM (Hochreiter and Schmidhuber, 1997) can better avoid the vanishing gradient problem compared with normal RNNs, thus it is suitable for processing information through many time steps.

Input words  $w_{1:L}^t$  are firstly converted to word embeddings through a lookup table in word embedding layer. Given word embeddings  $x_{1:L}^t$ , BiLSTM outputs hidden states  $h_{1:L}^t$ , where an output hidden state  $\overrightarrow{h}_i$  is the concatenation of forward hidden state  $\overrightarrow{h}_i$  and backward hidden state  $\overleftarrow{h}_i$ . We use a max pooling layer to extract the most informative features over time, and produce a single vector  $s_t$  as the encoding of word sequence  $w_{1:L}^t$ .

### 3.2 Sequence Classification

Given sentence encoding vectors  $s_{t-k:t}$ , we use the second neural network to predict the label of the  $t$ -th sentence. We use a history of length  $k$  instead of the whole history since dialog act usually depends on the very late history and it also accelerates training. Again we use BiLSTM and it works in a similar manner as the BiLSTM sentence encoder does but without max pooling layer, since the latest sentence provides the most information for DA recognition. Given the final hidden state, a fully-connected layer with a softmax function outputs the predicted label  $y_t$ .

## 4 Joint Learning

Joint learning (multi-task learning) is an approach to learn from several related tasks in parallel so that it improves the model’s ability to generalize features, and promotes its performance on the different tasks. In natural language processing (NLP), many higher-level tasks usually depend on outputs from lower-level tasks, for example named entity recognition (NER) relies on part-of-speech (POS) tagging. Hence there are many cases where models can benefit from joint learning.

Collobert and Weston (2008) introduced a neural network-based joint architecture making minimal assumption of feature engineering, and also concluded three kinds of joint model, i.e. cascading features, shallow joint training, and deep joint training. A cascading model, however, does not include any joint learning procedure, and shallow joint training is actually to convert tags of different tasks into one tag. In the rest of this paper, we will name them cascading model, joint cod-

ing model, and joint training model for accuracy. Zheng et al. (2013) applied a joint coding method to Chinese word segmentation and POS tagging by changing POS labels using a “BIES” tag coding scheme. Peng and Dredze (2016) improved NER by word representation learnt in word segmentation task using a LSTM-CRF model. Yang et al. (2016) proposed a multi-task cross-lingual model for sequence labeling tasks using RNN-CRF structures. These approaches to joint learning mostly attempt to learn shared representation of words and characters from different tasks.

Unlike aforementioned NLP tasks, DA segmentation and recognition deal with units of different levels, i.e. word level and DA segment level. Joint learning of these two tasks does not naturally fit in the architectures above. Previous works usually use joint coding methods. Zimmermann 2006 (2006) used a hidden-event language model for sequence labeling of DA type and its boundary on word level. It also exploits prosodic features in classification. Zimmermann (2009) and Quarterono et al. (2011) applied conditional random field (CRF) and the former also investigated how different tag coding schemes affect the model’s performance. Granell et al. (2009) incorporated syntactic features and used a combination of a HMM at the lexical level and a Language Model (LM) at the DA level. Hakkani-Tür et al. (2016) used a single sequence labeling LSTM model for joint semantic frame parsing, where sentence-level intent and domain tags are predicted at the last token of the sentence. An encoder-decoder-pointer framework was used for chunking in (Zhai et al., 2017), where segmentation was done by a pointer network and labeling was done by a decoder LSTM.

To our knowledge, there is no previous work on neural network based joint model applied to joint learning of DA segmentation and recognition. In this section, we investigate cascading model, joint coding model, and joint training model using neural networks. In the joint coding model, joint tag coding is used to combine segmentation and recognition tags and leads to a word-level sequence labeling task. For cascading and joint training models, the proposed architectures can deal with tasks of different unit levels and the hierarchical model introduced in Section 3 is integrated to make use of contextual information.

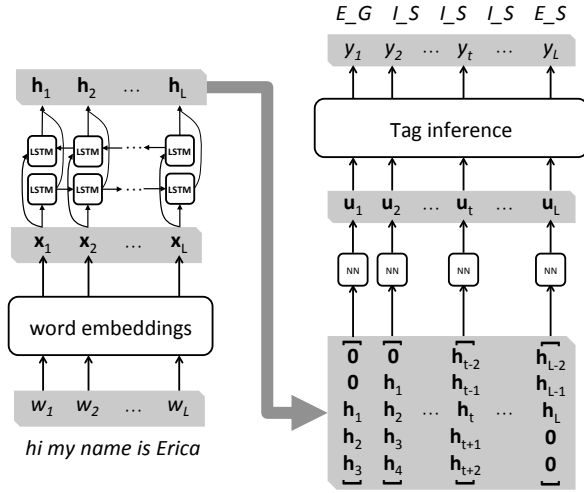


Figure 2: Joint coding model: an example of input and output is given in the figure.

## 4.1 Joint Coding Model

### 4.1.1 Joint Coding

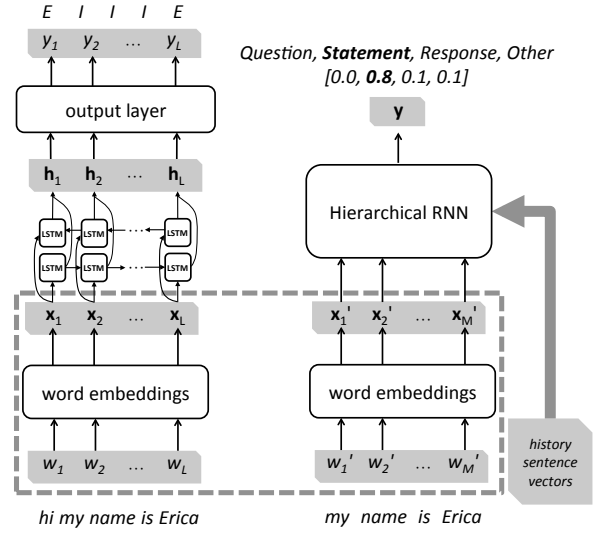
In the joint coding method, one single model predicts labels of DA segmentation and recognition at the same time, so that the units of these two tasks should keep consistency. We use a joint tag coding scheme to combine labels of segmentation and recognition and make them a word-level sequential labeling problem as shown in Table 2. This allows us to use one single sequential labeling model to solve two tasks simultaneously. The proposed joint coding model is given in Figure 2.

### 4.1.2 Tag Score

A sequence of words  $w_{1:L}$  is firstly mapped to word embeddings  $x_{1:L}$ , then we feed them to the following RNN layer which produces hidden states  $h_{1:L}$ . In order to provide contextual information explicitly, when we predict a label for the  $i$ -th step, we also make use of hidden states of preceding steps ( $h_{i-1}$ ,  $h_{i-2}$ , etc.) and succeeding steps ( $h_{i+1}$ ,  $h_{i+2}$ , etc.). The concatenated vector  $[h_{i-2}, h_{i-1}, h_i, h_{i+1}, h_{i+2}]$  is then fed to a neural network layer that computes a tag score  $s_i$ , where  $s_i \in \mathbb{R}^{|C|}$ , the  $t$ -th element in  $s_i$  indicates the score of choosing the  $t$ -th tag at the  $i$ -th step, and  $|C|$  is the number of tag classes.

### 4.1.3 Tag Inference

Since there often exists invalid tag sequences such as an “E\_S” following an “I\_Q” and we would like to penalize such invalid tag transitions, we apply a post process to compute the optimal tag sequence considering the transition probabilities by using a



(a) Segmentation component

(b) Recognition component

Figure 3: Cascading model and joint training model: in cascading model, part (a) on the left is a component for segmentation and part (b) on the right is for DA recognition. In joint training model, word embedding layers in dashed line rectangle are shared by both segmentation and recognition models. An example of input and output is given in the figure.

transition score matrix  $A_{mn}$ , which indicates the score of jumping from  $m$ -th tag to the  $n$ -th tag. Let  $s_i^t$  denotes the  $t$ -th element of  $s_i$ , the score of a tag sequence  $t_{1:L}$  is defined as:

$$score(t_{1:L}) = \sum_{i=1}^L (A_{t_{i-1}t_i} + s_i^t), \quad (1)$$

and we use the Viterbi algorithm to find the optimal tag sequence  $t_{1:L}^*$  that maximizes the sequence score:

$$t_{1:L}^* = \arg \max_{\forall t_{1:L}} score(t_{1:L}). \quad (2)$$

## 4.2 Cascading Model and Joint Training Model

DA segmentation splits a sequence of words into segments, and DA recognition assigns a dialog act type to them. Thus, these two tasks are naturally conducted in a cascading manner. The proposed cascading model is shown in Figure 3.

The left part is a sequential labeling model for segmentation. Similar to the joint coding model, a word embedding layer and a layer of RNN are used to produce a sequence of hidden states  $h_{1:L}$ .

<b>Words</b>	hi	my	name	is	Erica	nice	to	meet	you
<b>Segmentation</b>	E	I	I	I	E	I	I	I	E
<b>DA</b>	Greeting	Statement				Greeting			
<b>Joint tag coding</b>	E_G	IS	IS	IS	E_S	IG	IG	IG	E_G

Table 2: Joint “IE” tag coding scheme, where “I” and “E” refer to “inside” and “end” respectively. We concatenate segmentation tag with DA tag to produce coded tags. For example, “E.S” denotes the end of a *Statement* segment.

Then the top layer outputs predicted labels  $y_{1:L}$ . The right part of Figure 3 uses the hierarchical neural network model introduced in Section 3. In order to provide contextual information, we also maintain a history of sentence representation vectors of previous sentences using the same hierarchical neural network.

Joint training model can be seen as a cascading model with shared components. As shown in Figure 3, the proposed model uses only one set of word embeddings compared with the cascading one, while other task-specific parts are still separated. By updating the shared word embeddings using errors from both tasks, the model is expected to learn features from DA segmentation and recognition and improve its ability of generalization.

## 5 Experimental Evaluations

In order to evaluate our models, we conducted three sets of experiments:

- **Segmentation task:** evaluate models’ DA segmentation performance only.
- **Recognition task:** evaluate models’ DA recognition performance given correct segments.
- **Joint segmentation and recognition task:** evaluate models’ DA segmentation and recognition jointly, where predicted segments are given for DA recognition.

For each set of experiments, we also vary the length of history  $k$  for the cascading model and the joint training model.

### 5.1 Data Set

We use a one-to-one Japanese chatting corpus collected using a conversational android ER-ICA (Glas et al., 2016; Inoue et al., 2016). It is annotated with 4 DA tags (i.e. *Question*, *Statement*, *Response* and *Other*) following standards in (Bunt et al., 2010). Table 3 presents related statistics.

Corpus Statistics	
# of classes	4
avg. # of segments per session	165
avg. # of segments per turn	1.76
# of training sessions	30
# of test sessions	8

Table 3: Corpus statistics.

## 5.2 Implementation

We implemented the proposed models with *TensorFlow*<sup>1</sup>. Neural networks are trained using the Adam optimizer (Kingma and Ba, 2014). We use an initial learning rate of 0.0001 which decays in half when the objective loss does not decrease. A dropout layer of 0.25 dropout probability is added before every RNN layer for regularization. We choose 128 as the word embedding dimension since it works well in most NLP tasks. To find out how history length  $k$  affects the models, we experimented on  $k$  of 1, 3, 5, 10, 20. We also test CRFs implemented by *CRF++*<sup>2</sup> for comparison in our experiments.

## 5.3 Evaluation Metrics

For the segmentation task, we use the DA Segmentation Error Rate (DSER) in (Zimmermann et al., 2005). The DSER is the percentage of segments that are incorrectly segmented, i.e. its left or right boundary differs from the reference boundaries. Accuracy, Macro F1 measure, and Weighted F1 measure are used for evaluation of the recognition task since it is a text classification problem. For the joint task, we use the DA Error Rate (DER) which is the same as the DSER but also considers the DA type for correctness. Table 4 demonstrates the calculation of DSER and DER.

<sup>1</sup><https://www.tensorflow.org/>

<sup>2</sup><https://taku910.github.io/crfpp/>



Reference Prediction	E_G	LS	LS	E_S	I_Q	I_Q	E_Q	LS	E_S
	E_G	LS	LS	LS	E_S	I_Q	E_Q	LR	ER
<b>DSER</b>	✓		×			×			✓
<b>DER</b>	✓		×			×			×

Table 4: An example of DSER and DER metrics. The DSER equals 0.5 (2/4) and the DER equals 0.75 (3/4).

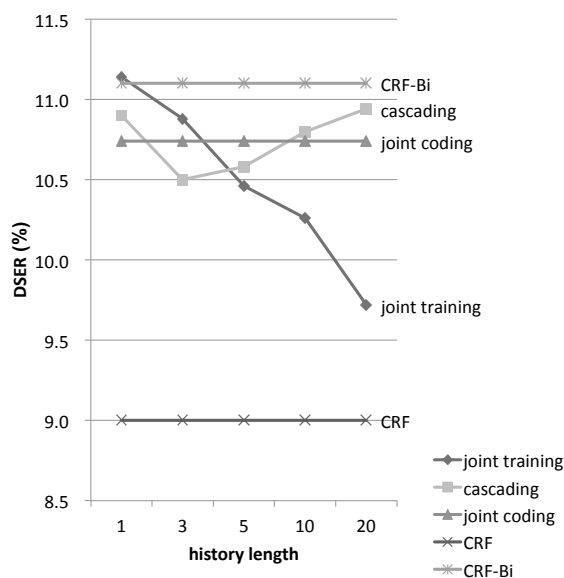


Figure 4: Results of the segmentation task.

## 5.4 Segmentation Result

In this task, a set of experiments are conducted to evaluate the model performances on DA segmentation specifically. We apply our models to segment every turn in dialogs, where one turn refers to a sequence of consecutive words uttered by one speaker without being interrupted by another speaker. We compare the cascading model, the joint coding model and the joint training model. Two CRFs are used for comparison. A simple CRF uses unigrams and bigrams of  $w_{t-2}$ ,  $w_{t-1}$ ,  $w_t$ ,  $w_{t+1}$ ,  $w_{t+2}$  as features, where  $w_t$  is the current word. Another CRF additionally uses the last tag as a feature and is named CRF-Bi in the following experiments.

Figure 4 shows the results of the segmentation task. The joint training model achieves 9.7% DSER at a history length of 20, which is comparable to CRF’s 9.0% (without a significant difference), and it outperforms the cascading model’s 10.5% at a history length of 3. The joint coding model lags behind slightly with the DSER of 10.7% and CRF-Bi gets a DSER of 11.1%.

## 5.5 Dialog Act Recognition Result

In the task of DA recognition, we evaluate the model performances of DA recognition. Therefore, we directly use ground-truth segments as inputs and predict a DA label for these segments. We only compare the cascading model and the joint training model in the recognition task because the joint coding model and CRFs are sequence labeling models and they do not naturally fit the text classification task.

As shown in Figure 5, the joint training model gets better results than the cascading model according to all three metrics. The joint training model achieves the best results of 78.0% accuracy, 77.4% Macro F1 measure and 78.1% Weighted F1 measure at the history length of 10, while the cascading model reaches 77.2%, 76.5% and 77.4% respectively. A significant improvement is obtained when we increase the history length from 1 to 3.

## 5.6 Joint Segmentation and Recognition Result

In the joint task, segmentation and recognition performances are evaluated jointly. We firstly segment each turn in dialogs into segments, then use the predicted segments as inputs of DA recognition. As in the segmentation task, we compare our proposed models and two CRFs.

Figure 6 shows the results of the joint task. The joint training model has the lowest error rate of 27.3% at the history length of 10, gaining an absolute improvement of 1.6% compared with the cascading model’s 28.9%. CRF-Bi, CRF and the joint coding model got the DSER of 33.5%, 36.8% and 37.1% respectively.

## 6 Discussion

In the segmentation task, CRF using only n-gram features obtains a result of 9.0%, though there is not a statistically significant difference in performance between CRF and the joint training model, whose DSER is 9.7%. We conjecture that the re-




Figure 5: Recognition results evaluated using Accuracy, Macro F1 measure and Weighted F1 measure.

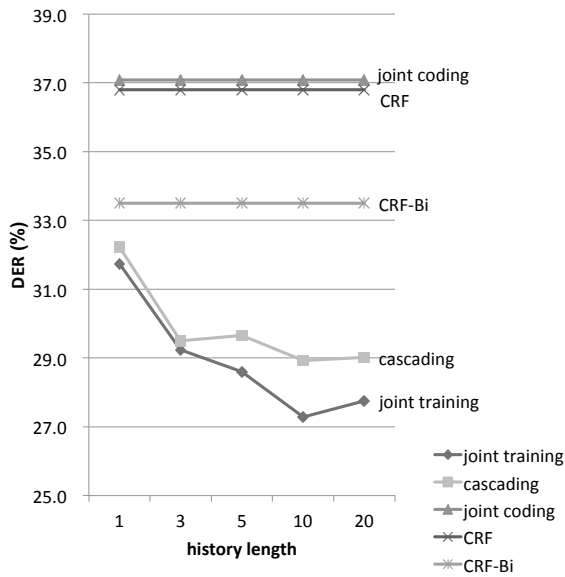


Figure 6: Results of the joint task.

result is due to that the boundaries of DA segments in Japanese are usually marked with words such as “*masu*”, “*desu*”, “*kedo*”, and considering simple n-gram features can cover most of the cases in DA segmentation. CRF-Bi, which uses a previous tag as additional feature, has a higher error rate of 11.1% compared with CRF using n-grams. It implies that extra information (previous segmentation tag and dialog act tag) may hurt the model’s performance in the sense of segmentation. However, by comparing the results of the cascading joint model and the joint training model (10.5% DSER of the cascading model and 9.7% DSER of

the joint training model), we can see that our joint training method is able to extract useful information from DA recognition task for segmentation and improve our models’ ability of segmentation.

In the DA recognition task, we compare the best results of the cascading model and the joint training model at a history length of 10. Similar to aforementioned conclusion, joint training helps the joint training model outperform the cascading model by 0.8% in accuracy, 0.9% in Macro F1 measure and 0.7% in Weighted F1 measure. Thus we can conclude that joint training can also learn features from segmentation task to help recognize DA tags.

In the joint evaluation, we observe that even though CRFs obtain fairly good results in the segmentation task, they significantly lag behind the proposed cascading model and the joint training model. The hierarchical neural network introduced in Section 3 makes use of contexts and contributes to the improvement. Similarly, CRF-Bi gets a better result of 33.5% DER than CRF’s 36.8%, which implies that contextual information (previous tag) plays an important role in DA recognition.

Finally the joint coding model has an acceptable result of segmentation but a very low performance in the joint task. There are three possible reasons: (1) The joint coding model only uses surrounding words as context instead of previous sentences, thus it fails to capture DA relations while the proposed hierarchical neural network is able to. (2) Failure in learning from the recognition

task can degrade the model’s performance on the segmentation task. (3) Lastly but the most important, DA recognition requires understanding of the whole segment. In the architecture of joint coding model, however, it has to predict the type of DA tag at the very beginning of a segment. Although a short context (i.e. 5 words in our experiments) is available, the model still cannot make use of complete information of the corresponding segment.

## 7 Conclusion

In this work, we explored joint learning of dialog act (DA) segmentation and recognition. To exploit contextual information in DA recognition, we introduced a hierarchical neural network architecture that incorporates history utterances. Based on the hierarchical neural network, we investigated three models for joint DA segmentation and recognition, i.e. cascading model, joint coding model, and joint training model. Our proposed models can (1) integrate the hierarchical neural network and (2) combine tasks of different levels (word level and DA segment level) in a unified architecture.

Three sets of experiments were carried out to evaluate the proposed models’ performances on the segmentation task, the DA recognition task and the joint task. Experimental results showed that (1) contextual information plays an important role in DA recognition; (2) the cascading model and the joint training model outperform CRF baselines significantly (4.6% and 6.2% in DER respectively) in the joint task while having comparable performances in the segmentation task; (3) the joint training model outperforms the cascading model in all three tasks. The result demonstrates that joint training can learn useful generalized features.

## Acknowledgments

This work was supported by JST ERATO Ishiguro Symbiotic Human-Robot Interaction program (Grant Number JPMJER1401), Japan.

## References

Harry Bunt, Jan Alexandersson, Jean Carletta, Jae-Woong Choe, Alex Chengyu Fang, Kôiti Hasida, Kiyong Lee, Volha Petukhova, Andrei Popescu-Belis, Laurent Romary, Claudia Soria, and David R. Traum. 2010. Towards an ISO standard for dialogue act annotation. In Nicoletta Calzolari, Khalid

Choukri, Bente Maegaard, Joseph Mariani, Jan Odijk, Stelios Piperidis, Mike Rosner, and Daniel Tapias, editors, *Proceedings of the International Conference on Language Resources and Evaluation, LREC 2010, 17-23 May 2010, Valletta, Malta*. European Language Resources Association.

Ronan Collobert and Jason Weston. 2008. A unified architecture for natural language processing: deep neural networks with multitask learning. In William W. Cohen, Andrew McCallum, and Sam T. Roweis, editors, *Machine Learning, Proceedings of the Twenty-Fifth International Conference (ICML 2008), Helsinki, Finland, June 5-9, 2008*. ACM, volume 307 of *ACM International Conference Proceeding Series*, pages 160–167.

Dylan F. Glas, Takashi Minato, Carlos Toshinori Ishi, Tatsuya Kawahara, and Hiroshi Ishiguro. 2016. ERICA: the ERATO intelligent conversational android. In *25th IEEE International Symposium on Robot and Human Interactive Communication, RO-MAN 2016, New York, NY, USA, August 26-31, 2016*. IEEE, pages 22–29.

Ramón Granell, Stephen G. Pulman, and Carlos D. Martínez-Hinarejos. 2009. Simultaneous dialogue act segmentation and labelling using lexical and syntactic features. In Patrick G. T. Healey, Roberto Pieraccini, Donna K. Byron, Steve J. Young, and Matthew Purver, editors, *Proceedings of the SIGDIAL 2009 Conference, The 10th Annual Meeting of the Special Interest Group on Discourse and Dialogue, 11-12 September 2009, London, UK*. The Association for Computer Linguistics, pages 333–336.

Dilek Hakkani-Tür, Gökhan Tür, Asli Çelikyilmaz, Yun-Nung Chen, Jianfeng Gao, Li Deng, and Ye-Yi Wang. 2016. Multi-domain joint semantic frame parsing using bi-directional RNN-LSTM. In Nelson Morgan, editor, *Interspeech 2016, 17th Annual Conference of the International Speech Communication Association, San Francisco, CA, USA, September 8-12, 2016*. ISCA, pages 715–719.

Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9(8):1735–1780.

Koji Inoue, Pierrick Milhorat, Divesh Lala, Tianyu Zhao, and Tatsuya Kawahara. 2016. Talking with erica, an autonomous android. In *Proceedings of the SIGDIAL 2016 Conference, The 17th Annual Meeting of the Special Interest Group on Discourse and Dialogue, 13-15 September 2016, Los Angeles, CA, USA*. The Association for Computer Linguistics, pages 212–215.

Yangfeng Ji, Gholamreza Haffari, and Jacob Eisenstein. 2016. A latent variable recurrent neural network for discourse relation language models. In *NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17,*



2016. The Association for Computational Linguistics, pages 332–342.
- Hamed Khanpour, Nishitha Guntakandla, and Rodney D. Nielsen. 2016. Dialogue act classification in domain-independent conversations using a deep recurrent neural network. In Nicoletta Calzolari, Yuji Matsumoto, and Rashmi Prasad, editors, *COLING 2016, 26th International Conference on Computational Linguistics, Proceedings of the Conference: Technical Papers, December 11-16, 2016, Osaka, Japan*. ACL, pages 2012–2021.
- Diederik P. Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *CoRR* abs/1412.6980.
- Siwei Lai, Liheng Xu, Kang Liu, and Jun Zhao. 2015. Recurrent convolutional neural networks for text classification. In Blai Bonet and Sven Koenig, editors, *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA*. AAAI Press, pages 2267–2273.
- Ji Young Lee and Franck Dernoncourt. 2016. Sequential short-text classification with recurrent and convolutional neural networks. In Kevin Knight, Ani Nenkova, and Owen Rambow, editors, *NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016*. The Association for Computational Linguistics, pages 515–520.
- Wei Li and Yunfang Wu. 2016. Multi-level gated recurrent neural network for dialog act classification. In Nicoletta Calzolari, Yuji Matsumoto, and Rashmi Prasad, editors, *COLING 2016, 26th International Conference on Computational Linguistics, Proceedings of the Conference: Technical Papers, December 11-16, 2016, Osaka, Japan*. ACL, pages 1970–1979.
- Nanyun Peng and Mark Dredze. 2016. Improving named entity recognition for chinese social media with word segmentation representation learning. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 2: Short Papers*. The Association for Computer Linguistics.
- Silvia Quarteroni, Alexei V. Ivanov, and Giuseppe Riccardi. 2011. Simultaneous dialog act segmentation and classification from human-human spoken conversations. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing, ICASSP 2011, May 22-27, 2011, Prague Congress Center, Prague, Czech Republic*. IEEE, pages 5596–5599.
- Elizabeth Shriberg, Andreas Stolcke, Daniel Jurafsky, Noah Coccaro, Marie Meteer, Rebecca Bates, Paul Taylor, Klaus Ries, Rachel Martin, and Carol Van Ess-Dykema. 1998. Can prosody aid the automatic classification of dialog acts in conversational speech? *Language and speech* 41(3-4):443–492.
- Andreas Stolcke, Klaus Ries, Noah Coccaro, Elizabeth Shriberg, Rebecca Bates, Daniel Jurafsky, Paul Taylor, Rachel Martin, Carol Van Ess-Dykema, and Marie Meteer. 2000. Dialogue act modeling for automatic tagging and recognition of conversational speech. *Computational Linguistics* 26(3):339–373.
- Zhilin Yang, Ruslan Salakhutdinov, and William W. Cohen. 2016. Multi-task cross-lingual sequence tagging from scratch. *CoRR* abs/1603.06270.
- Feifei Zhai, Saloni Potdar, Bing Xiang, and Bowen Zhou. 2017. Neural models for sequence chunking. In Satinder P. Singh and Shaul Markovitch, editors, *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*. AAAI Press, pages 3365–3371.
- Xiaoqing Zheng, Hanyang Chen, and Tianyu Xu. 2013. Deep learning for chinese word segmentation and POS tagging. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, EMNLP 2013, 18-21 October 2013, Grand Hyatt Seattle, Seattle, Washington, USA, A meeting of SIGDAT, a Special Interest Group of the ACL*. ACL, pages 647–657.
- Matthias Zimmermann. 2009. Joint segmentation and classification of dialog acts using conditional random fields. In *INTERSPEECH 2009, 10th Annual Conference of the International Speech Communication Association, Brighton, United Kingdom, September 6-10, 2009*. ISCA, pages 864–867.
- Matthias Zimmermann, Yang Liu, Elizabeth Shriberg, and Andreas Stolcke. 2005. Toward joint segmentation and classification of dialog acts in multiparty meetings. In Steve Renals and Samy Bengio, editors, *Machine Learning for Multimodal Interaction, Second International Workshop, MLMI 2005, Edinburgh, UK, July 11-13, 2005, Revised Selected Papers*. Springer, volume 3869 of *Lecture Notes in Computer Science*, pages 187–193.
- Matthias Zimmermann, Andreas Stolcke, and Elizabeth Shriberg. 2006. Joint segmentation and classification of dialog acts in multiparty meetings. In *2006 IEEE International Conference on Acoustics Speech and Signal Processing, ICASSP 2006, Toulouse, France, May 14-19, 2006*. IEEE, pages 581–584.