

Multi-level Gated Recurrent Neural Network for Dialog Act Classification

Wei Li

Key Laboratory of Computational
Linguistics, Peking University
Beijing, China
liweitj47@pku.edu.cn

Yunfang Wu

Key Laboratory of Computational
Linguistics, Peking University
Beijing, China
wuyf@pku.edu.cn

Abstract

In this paper we focus on the problem of dialog act (DA) labelling. This problem has recently attracted a lot of attention as it is an important sub-part of an automatic dialog model, which is currently in great demand. Traditional methods tend to see this problem as a sequence labelling task and deal with it by applying classifiers with rich features. Most of the current neural network models still omit the sequential information in the conversation. Henceforth, we apply a novel multi-level gated recurrent neural network (GRNN) with non-textual information to predict the DA tag. Our model not only utilizes textual information, but also makes use of non-textual and contextual information. In comparison, our model has shown significant improvement over previous works on the Switchboard Dialog Act (SWDA) data by over 6%.

1 Introduction

Dialog act labelling is one of the ways to find the shallow discourse structures of natural language conversations. It represents the meaning or intention of each short sentence within a conversation by giving a tag to each sentence (Austin and Urmson, 1962; Searle, 1969). DA can be of help to many tasks, for example, the DA of the current sentence provides very important information for answer generation in an automatic question answering system. This converts a complex system into a classification problem, enabling many existing systems to fit in the problem.

Traditional methods apply classifiers with rich human-crafted features to tag the sentences. One can view each sentence in the dialog as a separate one and label it accordingly, such as the work of (Silva et al., 2011), but this results in the loss of sequential information in the conversation context. Stolcke et al. (2000) used a segmented version of switchboard dialog act (SWDA) (Godfrey et al., 1992) with 43 tags based on the DAMSL labelling system, and proposed to use a hidden Markov model with rich features to predict the DA of each sentence. Although their model produces relatively good results, the feature construction and tuning consume too much human effort, and also make the adaptation between tasks difficult.

Using the deep learning framework, researchers have developed various systems to deal with DA and related problems like sentiment analysis and sentence classification. One can build a simple CNN architecture like Kim (2014) to do the labelling work. However, the sentences in a conversation are highly variant in length, some of which can be as short as one to two words or may even include nothing but some telephone script symbols. For example, a lot of sentences consist of nothing but "*<laughter>*." and "*Okay*". To be specific, in the SWDA data, 3,253 sentences consist of a single word and the length of 41.4% sentences are under 5 words. Figure 1 shows the distribution of sentence lengths in detail. As is shown in the figure, most of the sentences (61%) are under 10 words, which implies that a significant portion of the overall accuracy can be attributed to short sentences.

Most of previous models tend to do poorly on these extremely short sentences because of the lack of information. To deal with short texts, one must uncover more information, such as context sentences,

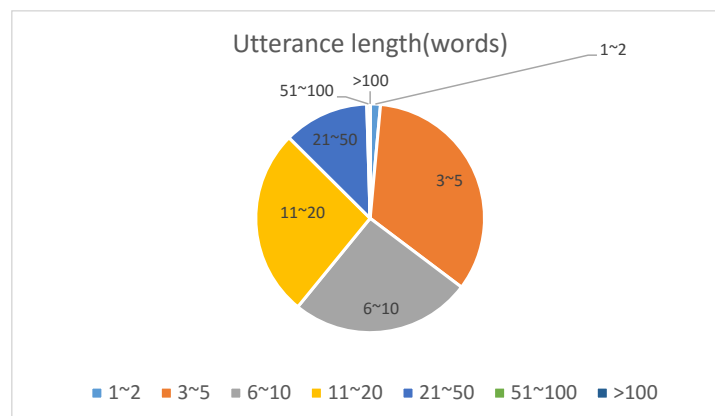


Figure 1: Sentence length distribution in the SWDA corpus

to facilitate the labelling process. In fact, the most important character of DA labelling that is different from simple sentence classification is that utterances appear sequentially in a conversation. Lee and Derroncourt (2016) tried to make use of historical information by feeding previous sentences in a fixed window together with the current one to a feed forward neural network. This makes a good attempt in applying contextual information. However, this approach loses long distance dependency, thus giving very little improvement when compared with the CNN baseline. Zhou et al. (2015) tried to capture sequential information with the conditional random field (CRF) on the basis of a heterogeneous neural network. While their model works very well, we must also be keen to note that the RNN family models surpass CRF in sequence prediction tasks, as pointed out by Irsoy and Cardie (2014) and Yao et al. (2014).

Apart from textual and contextual information, non-textual information can also be considered. Hu et al. (2013) applied a restricted Boltzmann machine to combine textual and non-textual features in a community question answering problem. Their work makes good use of the non-textual features by combining them with textual features in an unsupervised manner.

To deal with the limitations of previous works, we propose a multi-level GRNN with non-textual features to predict the DAs. Our contributions can be highlighted in the following aspects:

- We apply a two-level GRNN to predict the DA. The low level GRNN is designed for modelling textual information of each sentence, and the top level GRNN is designed to make use of historical information in a conversation. This method produces an obvious improvement over the previous works as it automatically selects what information in the context to remember and forget.
- We use a feed forward neural network to capture the non-textual information. Then we feed the hidden layer as sentence level non-textual information to the top level GRNN.
- We conduct extensive experiments for DA labelling on the open SWDA corpus by exploiting different neural network models. With the new framework applied, our model achieves a significant improvement over previous works in SWDA task by over 6% from 73.1 to 79.37.

2 Related Work

2.1 Traditional methods on dialog act labelling

Dialog acts are to represent the intention of each sentence within a conversation. Allen and Core (1997) proposed the Dialog Act Markup in Several Layers (DAMSL) scheme to provide a top level structure for

annotating dialogs, which was applied by many dialog annotation systems (Jurafsky et al., 1997; Dhillon et al., 2004). Bunt et al. (2012) gave a detailed summary over the standard of dialog acts annotation in semantic annotation framework.

Dialog act labelling was traditionally viewed as a sequence labelling or sentence modelling problem. Most of the previous works try to predict the DA by calculating the probability of each label. Reithinger and Klesen (1997) used a language model to predict the probability of a certain DA. However, the effort to predict probability using a language model results in a severe loss of information, thereby leading to a poor result. Louwse and Crossley (2006) introduced n-gram features to predict the DA, which is widely used in NLP tasks. This model uncovers more information from the text, but it fails to capture long-distance dependency. Surendran and Levow (2006) used SVM on individual sentences then viterbi decoding to make use of contextual information in a HMM style. This model builds a rather good framework for sequential labelling, as it not only feeds each sentence to a strong classifier SVM, but also makes use of context information in a probability graph. (Kim et al., 2010) further proposed to use CRF to deal with the problem, using both traditional bag of words features and new features such as dialog structures and dependencies between utterances. The common weakness of these methods is that they depend heavily on the features selected, and the feature construction process consumes much human effort.

2.2 Deep Learning models

As deep learning becomes increasingly popular, researchers have been trying to apply deep learning frameworks to deal with natural language processing and understanding tasks, including sentence modelling, DA labelling and many other tasks. Collobert and Weston (2007), Collobert and Weston (2008) and Collobert et al. (2011) constructed deep neural network structures for natural language processing tasks, which project one-hot word representations into distributed representations with a look-up table (or a projection layer) and build either feed forward or convolutional neural network upon them. This type of models seek to free researchers from laborious feature engineering, and allow the systems to easily adapt to different tasks.

Kalchbrenner et al. (2014) proposed a dynamic convolution neural network with multiple layers of convolution and k-max pooling to model a sentence. As imagined, this model is computationally expensive due to the many layers. Conversely, the CNN model proposed by Kim (2014) takes just one convolution and pooling layer with multi-channel word embeddings, followed by a softmax classifier. This model succeeded in many NLP tasks, such as sentence classification, sentiment analysis and so on.

Apart from CNN like architectures, researchers also applied recurrent neural network (RNN) and its variants to model sentences. Originally proposed by Elman (1990), RNN is expected to propagate information through time, which means one can make use of past information as latent variables. Mikolov et al. (2010) applied RNN to language modelling and got some very interesting results for word embedding. However, this vanilla RNN suffers from the same problem as other deep neural networks, the problem of vanishing gradient. More specifically, gradients can either explode or vanish through time (Bengio et al., 1994). To tackle this problem, Hochreiter and Schmidhuber (1997) proposed long short term memory (LSTM), which uses a cell with input, forget and output gates to prevent the vanishing gradient problem. This makes RNN family networks much more powerful by memorizing information from long distance.

Recently, inspired by the gating idea, Cho et al. (2014) proposed another variant of RNN named gated recurrent neural network, which only uses a reset gate and a update gate to encode and decode sentences in a translation system. As reported in Chung et al. (2014), GRNN can achieve better results than LSTM in most tasks.

Palangi et al. (2015) proposed to sequentially take each word in a sentence, extract its information, and embed it into a semantic vector. This way, one can access the sentence level vector and use it to deal with other tasks such as information retrieval. Shen and Lee (2016) introduced one type of attention mechanism to sentence modelling based on LSTM, they also tested their model on SWDA task, which we will reference as a baseline. Their model performed better on longer sentences by highlighting the important parts of the sentence. But, as aforementioned, the most important part of this problem is




Figure 2: Gated recurrent neural network for sentence representation based on textual information

not about long sentences, but the short ones, which take the majority share of the corpus. Lee and Dernoncourt (2016) regarded this problem as a sequential short text classification problem, which is a good direction. However, although they tried to capture the historical information, they failed to seize long distant information in a conversation, because they only feed a fixed window to the neural network and the capability of the feed forward neural network is very limited.

3 Our Approach

In this paper, we propose to utilize a multi-level GRNN architecture to mine the information from both within the sentence and between the sentences. Gated recurrent neural network is a variant of the recurrent neural network. The GRNN allows information to flow over time without the problem of vanishing gradient, and is expected to memorize long distance dependency.

Equations 1 to 2 show the method to calculate the output h_t at time stamp t , with the input x_t and history information h_{t-1} , which is the output at time stamp $t - 1$. In each gated recurrent unit, the *reset gate* (Equation 1) and the *update gate* (Equation 2) are designed to decide which latent information is to be discarded and which is to be held. Equation 3 calculates the candidate unit similar to vanilla RNN unit, except that it uses a *reset gate* to filter history information, and Equation 4 uses the *update gate* and the candidate unit to get the final output unit.

In our model, we first use the low level GRNN on the scale of words to learn sentence level vector, then we use GRNN to propagate the information between sentences over time within the same conversation. To discover more information on the sentence level, we also apply a feed forward neural network to capture the non-textual information such as the length of the sentence, the index of the utterance and so on.

$$z_t = \sigma(W_z x_t + U_z h_{t-1}) \quad (1)$$

$$r_t = \sigma(W_r x_t + U_r h_{t-1}) \quad (2)$$

$$\tilde{h}_t = \tanh(W x_t + U(r_t \odot h_{t-1})) \quad (3)$$

$$h_t = (1 - z_t)h_{t-1} + z_t \tilde{h}_t \quad (4)$$

3.1 Textual information

Textual information is the basis of our end-to-end labelling system. We use a GRNN with max-pooling to encode the sentence into a vector.

As is shown in Figure 2, we treat each word as a separate unit. We first look up the corresponding embedding in a lookup table, which gives a matrix of $D * L$, D is the dimension of word embedding and L is the sentence length. Then we feed each word in the sentence into the low level GRNN, one word per time step, and then perform max pooling on the output of the GRU cells over the whole sentence.

caller	utterance index	sub-utterance index	act tag	text
A	5	2	qy	{F Um, } {F uh, } do you live right in the city itself? /
B	6	1	nn	No, /
B	6	2	sd	I'm more out in the suburbs, /
B	6	3	sd	{C but } I certainly work near a city. /
A	7	1	bk	Okay, /
A	7	2	qy	{C so } [ca-, +

Table 1: Utterance examples in SWDA corpus

3.2 Non-textual information

Although the aforementioned low level GRNN can capture the textual information within the sentence itself, it fails to make use of information from a higher level. For instance, in our DA labelling, the length of sentence plays an important role in identifying the tag of sentence, because the distribution of sentence length varies between different DAs. For sentences under the label of *acknowledge*, most sentences are below 10 words; whereas for sentences under the label of *statement non-opinion*, the sentences have more varied length distribution. As a matter of fact, it is shown in our experiment that this sentence length feature alone gives a much better prediction than random guesses.

Feed forward neural network (FFNN) is one of the simplest form of deep neural networks, and does a good job in many tasks. In this part of the neural network, we feed four shallow non-textual features to a FFNN. We use the hidden layer as the vector representing the non-textual information of the sentence. The four features we used are listed below. To better understand the features, Table 1 shows some examples from the original scripts.

- **Utterance index:** A conversation consists of multiple natural utterances, which are further split into lines of sentences for the convenience of tagging. Utterance index is the index of utterances, which can span multiple sentences. For example in Table 1, caller B says three sentences, and these three sentences share the same utterance index (6), but have different sub-utterance index. This feature may help when different acts take place in different parts of the conversation, for instance, conversations tend to begin with greetings.
- **Sub-utterance index:** Utterances can be broken across lines, sub-utterance index gives the internal position of the current sentence in the utterance. For example, in Table 1, the 6th utterance has three sentences or sub-utterances indexing from 1 to 3. This feature helps when different acts appear in different parts of an utterance. For example, questions tend to appear at the end of each utterance.
- **Same speaker:** This feature is a boolean feature of 0 or 1, indicating whether the identity of the speaker changes. Unlike the features above, this feature is deduced from the sub-utterance index. If the sub-utterance index is 1, then this feature is set to 1, otherwise 0.
- **Sentence length:** As explained earlier, the length of sentence plays an important role in predicting the label. As sentence lengths vary a lot, we normalize the lengths using Equation 5, where l is the word-wise sentence length.

$$l_{norm} = \frac{l - range(l)/2}{std(l)} \quad (5)$$

After we have the vector for textual and non-textual information aforementioned, we concatenate them together to get a combined vector for the sentence, as shown in the lower part of Figure 3.




Figure 3: Gated recurrent neural network on sentence feature

3.3 Context information

GRNN is designed to remember valuable information while discarding useless information. In the DA labelling problem, the segmentation of sentences is not very strict. Many sentences are very short, which makes it very difficult to classify a sentence based on only little textual information and sentence level non-textual information. Therefore, GRNN can fit this problem very well.

In our model, we try to use GRNN to capture the structure between sentences, as shown in Figure 3. This enables our model to utilize information from longer distances, unlike the structure proposed by (Lee and Derroncourt, 2016), which uses a fixed window to capture history information. Learning distant information is crucial for the fact that the dialog turn changes with no pattern, whereas some utterances consist of a single sentence while others consist of multiple sentences, which makes it impossible to learn the words from both speakers within a fixed window, as words of one speaker in the current sentence can be distant from the last words from the other speaker.

4 Experiment

4.1 Settings

We conducted experiments on the switchboard dialog act corpus, which extends the Switchboard-1 Telephone Speech Corpus with turn/utterance-level dialog-act tags. The tags summarize syntactic, semantic, and pragmatic information about the associated turn. There are over 200 tags in the corpus. Jurafsky et al. (1997) defines a system for collapsing them down to 44 tags.

In our experiments, we use the same data version as Stolcke et al. (2000), where there are 1,115 conversations (1.4M words, 198K utterances) in the training set, and 19 conversations (29K words, 4K utterances) in the test set. We use the same valid set as Lee and Derroncourt (2016), which consists of 19 randomly chosen conversations.¹

In our experiment, we build our model upon tensorflow by Abadi et al. (2015)², which is a popular package developed by Google for deep learning.

We use all the tokens of the utterances including texts and other telephone related symbols to train word embeddings with word2vec³ (Mikolov et al., 2013), and set the dimension of word embeddings to 300. We use the Adam stochastic optimization method (Kingma and Ba, 2014) to minimize the negative log-likelihood cost with fine-tuning on the word embeddings. To try to avoid the over-fitting problem, we run each experiment for 10 epochs, and use the hyper-parameters from the epoch with the highest validation accuracy. We use rectified linear unit (relu) as the activation function.

4.2 Baselines

We conduct extensive experiments on the SWDA corpus by utilizing various neural network models.

¹The train/validation/test splits were found at <https://github.com/Franck-Derroncourt/naacl2016>

²available in <https://www.tensorflow.org>

³available in <https://code.google.com/archive/p/word2vec/>

Method	Accuracy
Sequential short-text classification (Lee and Derroncourt, 2016)	73.1
Neural attention (Shen and Lee, 2016)	72.6
Our model	79.37

Table 2: Experimental results compared with previous state-of-the-art methods

- **CNN:** We implemented a convolutional neural network following the framework of (Kim, 2014). We use filters of length 2,3 and 4, and for each window length there are 100 feature maps. So each sentence has a vector of 300 real numbers. After the convolution and max-pooling layer, there is a softmax layer to predict the DA of each sentence.
- **non-textual:** We feed the four non-textual features to a typical three-layer feed forward neural network as described in Section 3.2. We set the unit number of the hidden layer to 300 and use the output of the softmax layer to predict the label.
- **CNN+non-textual:** This model is a combination of CNN and non-textual. We concatenate the pooled feature maps of CNN and the hidden layer of non-textual FFNN, and feed this new combined vector to a softmax layer to predict the label.
- **single-level GRNN:** This model follows the description in section 3.1. We feed the word embedding to the GRU cells, each word per cell. After we get the output of the GRU cells from each time step, we perform a max-pooling over them and get the sentence vector. Then we feed the sentence vector to a softmax layer to predict the tag.
- **single-level GRNN + non-textual:** This model combines the max-pooled sentence vector from single-level GRNN and the hidden layer of non-textual FFNN in the same way as CNN+non-textual. Then the concatenated vector is fed to a softmax layer to predict the tag.
- **non-textual+GRNN:** We feed the hidden layer of the non-textual FFNN to a GRNN. Then we feed the output of each GRU cell to the softmax layer to predict the labels.
- **CNN+GRNN:** We feed the sentence vector from CNN to GRNN. Then we feed the output of each GRU cell to the softmax layer to predict the labels.
- **multi-level GRNN:** We feed the sentence vector from lower level GRNN to the upper level GRNN. Then we feed the output of each GRU cell to the softmax layer to predict the labels.
- **CNN+non-textual+GRNN** We feed the combination of sentence vector from CNN and hidden layer from non-textual FFNN to a GRNN. Then we feed the output of each GRU cell to the softmax layer to predict the labels.
- **multi-level GRNN+non-textual:** This is our model in this paper. In this model, we feed the combination of sentence vector from lower level GRNN and hidden layer from non-textual FFNN to the upper level GRNN. Then we feed the output of each GRU cell to the softmax layer to predict the labels.

4.3 Comparison with previous models

Table 2 shows our result compared with other state-of-the-art results. By utilizing information from previous time stamp with GRNN, we achieve significant improvement over the previous works. As seen in Table 2, we improve the performance by 6.27% over Lee and Derroncourt (2016) and 6.77% over Shen and Lee (2016), as we better capture both the sentence level knowledge and contextual information in a conversation.

Method	Accuracy
CNN	68.25
single-level GRNN	69.75
non-textual	43.60
CNN+non-textual	70.86
single-level GRNN + non-textual	71.90
non-textual+GRNN	48.09
CNN+GRNN	77.14
multi-level GRNN	77.65
CNN+non-textual+GRNN	78.40
multi-level GRNN+non-textual	79.37

Table 3: Results of different neural networks in our experiment

Text	standard	single-level GRNN	final model
{F Um, } {F uh, } do you live right in the city itself? /	qy	qy	qy
No, /	nn	nn	nn
I'm more out in the suburbs, /	sd	sd	sd
{C but } I certainly work near a city. /	sd	sd	sd
Okay, /	bk	fo_o_fw_by_bc	bk
{C so } [ca-, +	qy	sd	qy

Table 4: The tagging DA results using two different neural network models, where "standard" means the golden standard tag in the data.

4.4 Comparison with baseline models

The experimental results in Table 3 show that both CNN and single-level GRNN with textual information can give relatively good results (68.25 & 69.75) for the DA labelling task.

Non-textual information can further improve the accuracy as they provide information about the whole sentence, instead of just individual words. This is verified by the fact that CNN+non-textual improves 2.61% over CNN and single-level GRNN+non-textual improves 2.15% over single-level GRNN. In fact, non-textual itself gives a surprisingly good result compared with random guess.

It is the GRNN which captures long distance dependency from context that produces the most significant improvement to the problem. As a matter of fact, the role of GRNN is so important that GRNN based on the weak classifier non-textual FFNN improves the result by almost 5% over the non-textual FFNN alone, and GRNN on the basis of CNN improves the result by almost 10% over the raw CNN. Altogether, our model of "multi-level GRNN+non-textual" surpasses the CNN baseline significantly by over 11%.

4.5 Analysis

In Table 4 we show the tagging results of the examples listed in Table 1. These results are from the single-level GRNN (one of our baselines) and our final model, respectively. The sentences are selected from the first conversation in the test set.

From the examples, we can observe that sentences with obvious features can be easily recognized by both models, for instance the first sentence "do you" is correctly tagged as "qy" (*Yes-No-Question*). However, when the sentence is short and ambiguous or can appear in multiple circumstances, such as "Okay", the simpler model mistakes the "bk" (*Response Acknowledgement*) for "fo_o_fw_by_bc" (*other*), while our final model which utilizes contextual information succeeds in predicting the right tag.

5 Conclusion

In this paper, we propose a multi-level GRNN model combined with non-textual features to deal with the dialog act labelling problem. We manage to mine multi-level information out of the conversation. Our model does a very good job on predicting short sentences in the SWDA corpus. Our results surpass previous state-of-the-art results significantly without much feature engineering, which makes our model easier to adapt to similar tasks. In the future, we hope to introduce the attention mechanism into our model and make better use of contextual information.

Acknowledgement. This work is supported by National Natural Science Foundation of China (61371129), National High Technology Research and Development Program of China (2015AA015403) and Key Program of Social Science foundation of China (12&ZD227). The corresponding author of this paper is Yunfang Wu.

References

- Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2015. TensorFlow: Large-scale machine learning on heterogeneous systems. Software available from tensorflow.org.
- James Allen and Mark Core. 1997. Draft of damsl: Dialog act markup in several layers. *Unpublished manuscript*, 2.
- John Langshaw Austin and JO Urmson. 1962. *How to Do Things with Words. The William James Lectures Delivered at Harvard University in 1955.*[Edited by James O. Urmson.]. Clarendon Press.
- Yoshua Bengio, Patrice Simard, and Paolo Frasconi. 1994. Learning long-term dependencies with gradient descent is difficult. *Neural Networks, IEEE Transactions on*, 5(2):157–166.
- Harry Bunt, Jan Alexandersson, Jae-Woong Choe, Alex Chengyu Fang, Koiti Hasida, Volha Petukhova, Andrei Popescu-Belis, and David R Traum. 2012. Iso 24617-2: A semantically-based standard for dialogue annotation. In *LREC*, pages 430–437. Citeseer.
- Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.
- Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*.
- Ronan Collobert and Jason Weston. 2007. Fast semantic extraction using a novel neural network architecture. In *Annual meeting-association for computational linguistics*, volume 45, page 560.
- Ronan Collobert and Jason Weston. 2008. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *Proceedings of the 25th international conference on Machine learning*, pages 160–167. ACM.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *The Journal of Machine Learning Research*, 12:2493–2537.
- Rajdip Dhillon, Sonali Bhagat, Hannah Carvey, and Elizabeth Shriberg. 2004. Meeting recorder project: Dialog act labeling guide. Technical report, DTIC Document.
- Jeffrey L Elman. 1990. Finding structure in time. *Cognitive science*, 14(2):179–211.
- John J Godfrey, Edward C Holliman, and Jane McDaniel. 1992. Switchboard: Telephone speech corpus for research and development. In *Acoustics, Speech, and Signal Processing, 1992. ICASSP-92., 1992 IEEE International Conference on*, volume 1, pages 517–520. IEEE.

- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Haifeng Hu, Bingquan Liu, Baoxun Wang, Ming Liu, and Xiaolong Wang. 2013. Multimodal dbn for predicting high-quality answers in cqa portals.
- Ozan Irsoy and Claire Cardie. 2014. Opinion mining with deep recurrent neural networks. In *EMNLP*, pages 720–728.
- Dan Jurafsky, Elizabeth Shriberg, and Debra Biasca. 1997. Switchboard swbd-damsl shallow-discourse-function annotation coders manual. *Institute of Cognitive Science Technical Report*, pages 97–102.
- Nal Kalchbrenner, Edward Grefenstette, and Phil Blunsom. 2014. A convolutional neural network for modelling sentences. *arXiv preprint arXiv:1404.2188*.
- Su Nam Kim, Lawrence Cavedon, and Timothy Baldwin. 2010. Classifying dialogue acts in one-on-one live chats. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 862–871. Association for Computational Linguistics.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. *arXiv preprint arXiv:1408.5882*.
- Diederik Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Ji Young Lee and Franck Dernoncourt. 2016. Sequential short-text classification with recurrent and convolutional neural networks. *arXiv preprint arXiv:1603.03827*.
- Max M Louwerse and Scott A Crossley. 2006. Dialog act classification using n-gram algorithms. In *FLAIRS Conference*, pages 758–763.
- Tomas Mikolov, Martin Karafiát, Lukas Burget, Jan Cernocký, and Sanjeev Khudanpur. 2010. Recurrent neural network based language model. In *INTERSPEECH*, volume 2, page 3.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Hamid Palangi, Li Deng, Yelong Shen, and Jianfeng Gao. 2015. Deep sentence embedding using long short-term memory networks: Analysis and application to information retrieval. *IEEE/ACM Transactions on Audio Speech & Language Processing*, 24(4):694–707.
- Norbert Reithinger and Martin Klesen. 1997. Dialogue act classification using language models. In *EuroSpeech*. Citeseer.
- John R Searle. 1969. *Speech acts: An essay in the philosophy of language*, volume 626. Cambridge university press.
- Sheng-syun Shen and Hung-yi Lee. 2016. Neural attention models for sequence classification: Analysis and application to key term extraction and dialogue act detection. *arXiv preprint arXiv:1604.00077*.
- Joao Silva, Luísa Coheur, Ana Cristina Mendes, and Andreas Wichert. 2011. From symbolic to sub-symbolic information in question classification. *Artificial Intelligence Review*, 35(2):137–154.
- Andreas Stolcke, Noah Coccaro, Rebecca Bates, Paul Taylor, Carol Van Ess-Dykema, Klaus Ries, Elizabeth Shriberg, Daniel Jurafsky, Rachel Martin, and Marie Meteer. 2000. Dialogue act modeling for automatic tagging and recognition of conversational speech. *Computational linguistics*, 26(3):339–373.
- Dinoj Surendran and Gina-Anne Levow. 2006. Dialog act tagging with support vector machines and hidden markov models. In *INTERSPEECH*.
- Kaisheng Yao, Baolin Peng, Yu Zhang, Dong Yu, Geoffrey Zweig, and Yangyang Shi. 2014. Spoken language understanding using long short-term memory neural networks. In *Spoken Language Technology Workshop (SLT), 2014 IEEE*, pages 189–194. IEEE.
- Yucan Zhou, Qinghua Hu, Jie Liu, and Yuan Jia. 2015. Combining heterogeneous deep neural networks with conditional random fields for chinese dialogue act recognition. *Neurocomputing*, 168:408–417.