

Repositório ISCTE-IUL

Deposited in *Repositório ISCTE-IUL*:

2019-01-10

Deposited version:

Post-print

Peer-review status of attached file:

Peer-reviewed

Citation for published item:

Ribeiro, E., Ribeiro, R. & de Matos, D. M. (2018). A study on dialog act recognition using character-level tokenization. In Gennady Agre, Josef van Genabith, Thierry Declerck (Ed.), 18th International Conference on Artificial Intelligence: Methodology, Systems, and Applications, AIMS 2018. (pp. 93-103). Varna: Springer.

Further information on publisher's website:

[10.1007/978-3-319-99344-7_9](https://doi.org/10.1007/978-3-319-99344-7_9)

Publisher's copyright statement:

This is the peer reviewed version of the following article: Ribeiro, E., Ribeiro, R. & de Matos, D. M. (2018). A study on dialog act recognition using character-level tokenization. In Gennady Agre, Josef van Genabith, Thierry Declerck (Ed.), 18th International Conference on Artificial Intelligence: Methodology, Systems, and Applications, AIMS 2018. (pp. 93-103). Varna: Springer., which has been published in final form at https://dx.doi.org/10.1007/978-3-319-99344-7_9. This article may be used for non-commercial purposes in accordance with the Publisher's Terms and Conditions for self-archiving.

Use policy

Creative Commons CC BY 4.0

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a link is made to the metadata record in the Repository
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

A Study on Dialog Act Recognition using Character-Level Tokenization*

Eugénio Ribeiro^{1,2}, Ricardo Ribeiro^{1,3}, and David Martins de Matos^{1,2}

¹ L²F – Spoken Language Systems Laboratory – INESC-ID Lisboa

² Instituto Superior Técnico, Universidade de Lisboa, Portugal

³ Instituto Universitário de Lisboa (ISCTE-IUL)

eugenio.ribeiro@l2f.inesc-id.pt

Abstract. Dialog act recognition is an important step for dialog systems since it reveals the intention behind the uttered words. Most approaches on the task use word-level tokenization. In contrast, this paper explores the use of character-level tokenization. This is relevant since there is information at the sub-word level that is related to the function of the words and, thus, their intention. We also explore the use of different context windows around each token, which are able to capture important elements, such as affixes. Furthermore, we assess the importance of punctuation and capitalization. We performed experiments on both the Switchboard Dialog Act Corpus and the DIHANA Corpus. In both cases, the experiments not only show that character-level tokenization leads to better performance than the typical word-level approaches, but also that both approaches are able to capture complementary information. Thus, the best results are achieved by combining tokenization at both levels.

Keywords: Dialog Act Recognition, Character-Level, Switchboard Dialog Act Corpus, DIHANA Corpus, Multilinguality

1 Introduction

Dialog act recognition is important in the context of a dialog system, since it reveals the intention behind the words uttered by its conversational partners [24]. Knowing that intention allows the system to apply specialized interpretation strategies, accordingly. Recently, most approaches on dialog act recognition focus on applying different Deep Neural Network (DNN) architectures to generate segment representations from word embeddings and combine them with context information from the surrounding segments [11,14,9,15]. However, all of these approaches look at the segment at the word level. That is, they consider that a segment is a sequence of words and that its intention is revealed by the combination of those words. However, there are also cues for intention at the sub-word level. These cues are mostly related to the morphology of words. For instance,

*This work was supported by national funds through Fundação para a Ciência e a Tecnologia (FCT) with reference UID/CEC/50021/2013 and by Universidade de Lisboa.

there are cases, such as adverbs of manner and negatives, in which the function, and hence the intention, of a word is related to its affixes. On the other hand, there are cases in which considering multiple forms of the same lexeme independently does not provide additional information concerning intention and the lemma suffices. Thus, it is interesting to explore dialog act recognition approaches that are able to capture this kind of information. In this paper, we explore the use of character-level tokenization with different context windows surrounding each token. Although character-level approaches are typically used for word-level classification tasks, such as Part-of-Speech (POS) tagging [23], they have also achieved interesting results on short-text classification tasks, such as language identification [8] and review rating [27]. In addition to the aspects concerning morphological information, using character-level tokenization allows us to assess the importance of aspects such as capitalization and punctuation. Additionally, we assess whether the obtained information can be combined with that obtained using word-level tokenization to improve the performance on the task. In this sense, in order to widen the scope of our conclusions, we performed experiments on two corpora, the Switchboard Dialog Act Corpus [10] and DIHANA [3], which have different characteristics, including domain, the nature of the participants, and language – English and Spanish, respectively.

In the remainder of this paper we start by providing an overview of previous approaches on dialog act recognition, in Section 2. Then, in Section 3, we discuss why using character-level tokenization is relevant for the task. Section 4 describes our experimental setup, including the used datasets, classification approach, and word-level baselines. The results of our experiments are presented and discussed in Section 5. Finally, Section 6 states the most important conclusions of this study and provides pointers for future work.

2 Related Work

Automatic dialog act recognition is a task that has been widely explored over the years, using multiple machine learning approaches, from Hidden Markov Models (HMMs) [25] to Support Vector Machines (SVMs) [6]. The article by Král and Cerisara [13] provides an interesting overview of most of those approaches on the task. However, recently, similarly to many other Natural Language Processing (NLP) tasks [16,7], most approaches on dialog act recognition take advantage of different DNN architectures.

To our knowledge, the first of those approaches was that by Kalchbrenner and Blunsom [11]. They used a Convolutional Neural Network (CNN)-based approach to generate segment representations from randomly initialized 25-dimensional word embeddings and a Recurrent Neural Network (RNN)-based discourse model to combine the sequence of segment representations with speaker information and output the corresponding sequence of dialog acts.

Lee and Dernoncourt [14] compared the performance of a Long Short-Term Memory (LSTM) unit against that of a CNN to generate segment representations from 200-dimensional Global Vectors for Word Representation (GloVe) embed-

dings [21] pre-trained on Twitter data. Those segment representations were then fed to a 2-layer feed-forward network that combined them with context information from the preceding segments. The best results were obtained using the CNN-based approach combined with information from two preceding segments in the form of their representation.

Ji et al. [9] used a Discourse Relation Language Model (DRLM) with a hybrid architecture that combined a Recurrent Neural Network Language Model (RNLM) [19] with a latent variable model over shallow discourse structure. This way, the model can learn vector representations trained discriminatively, while maintaining a probabilistic representation of the targeted linguistic element which, in this context, is the dialog act. In order to function as a classifier, the model was trained to maximize the conditional probability of a sequence of dialog acts given a sequence of segments.

The previous studies explored the use of a single recurrent or convolutional layer. However, the top performing approaches use multiple of those layers. On the one hand, Khanpour et al. [12] achieved their best results by combining the outputs of a stack of 10 LSTM units, in order to capture long distance relations between tokens. On the other hand, Liu et al. [15] combined the outputs of three parallel CNNs with different context window sizes, in order to capture different functional patterns. Both studies used Word2Vec [20] embeddings as input to the network. However, their dimensionality and training data varied.

Additionally, Liu et al. [15] explored the use of context information concerning speaker changes and from the surrounding segments. Concerning the latter, they used approaches that relied on discourse models, as well as others that combined the context information directly with the segment representation. Similarly to our previous study using SVMs [22], they concluded that providing that information in the form of the classification of the surrounding segments leads to better results than using their words. Furthermore, both studies have shown that the first preceding segment is the most important and that the influence decays with the distance.

3 Character-Level Tokenization

It is interesting to explore character-level tokenization because it allows us to capture morphological information that is at the sub-word level and, thus, cannot be directly captured using word-level tokenization. Considering the task at hand, that information is relevant since it may provide cues for identifying the intention behind the words. When someone selects a set of words to form a segment that transmits a certain intention, each of those words is typically selected because it has a function that contributes to that transmission. In this sense, affixes are tightly related to word function, especially in fusional languages. Thus, the presence of certain affixes is a cue for intention, independently of the lemma. However, there are also cases, such as when affixes are used for subject-verb agreement, in which the cue for intention is in the lemmas and, thus, considering multiple forms of the same lexeme does not provide additional information.

Information concerning lemmas and affixes cannot be captured from single independent characters. Thus, it is necessary to consider the context surrounding each token and look at groups of characters. The size of the context window plays an important part in what information can be captured. For instance, English affixes are typically short, but in other languages, such as Spanish, there are longer commonly used affixes. Furthermore, to capture the lemmas of long words, and even inter-word relations, wider context window sizes must be considered. However, using wide context windows impairs the ability to capture information from short groups of characters, as additional irrelevant characters are considered. This suggests that, in order to capture all the relevant information, multiple context windows should be used.

Using character-level tokenization also allows us to consider punctuation, which is able to provide both direct and indirect cues for dialog act recognition. For instance, an interrogation mark provides a direct cue that the intention is related to knowledge seeking. On the other hand, commas structure the segment, indirectly contributing to the transmission of an intention.

Additionally, character-level tokenization allows us to consider capitalization information. However, in the beginning of a segment, capitalization only signals that beginning and, thus, considering it only introduces entropy. In the middle of a segment, capitalization is typically only used to distinguish proper nouns, which are not related to intention. Thus, capitalization information is not expected to contribute to the task.

Finally, note that previous studies have shown that word-level information is relevant for the task. In this sense, it is interesting to assess whether that information can be captured using character-level tokenization or if there are specific aspects that require specialized approaches.

4 Experimental Setup

In order to assess the validity of the hypotheses proposed in the previous section, we performed experiments on different corpora and compared the performance of word- and character-level tokenization. The used datasets, classification approach, and word-level baselines are described below.

4.1 Datasets

In order to widen the scope of the conclusions drawn in the study, we selected two corpora with different characteristics to perform our experiments on. On the one hand, the Switchboard Dialog Act Corpus [10], henceforth referred to as Switchboard, is the most explored corpus for dialog act recognition. It features 1,155 manually transcribed human-human dialogs in English, with variable domain, containing 223,606 segments. The set is partitioned into a training set of 1,115 conversations, a test set of 19 conversations, and a future use set of 21 conversations [25]. In our experiments, we used the latter as a validation set. In

terms of dialog act annotations, we used the most used version of its tag set, which features 42 domain-independent labels.

On the other hand, the DIHANA corpus [3] consists of 900 dialogs in Spanish between human speakers and a Wizard of Oz (WoZ) telephonic train information system. The total number of annotated segments is 23,542, with 9,712 corresponding to user segments and 13,830 to system segments [2]. The set is partitioned into five folds to be used for cross-validation [17]. The dialog act annotations are hierarchically decomposed in three levels [18]. The first level represents the domain-independent intention of the segment, while the remaining are task-specific. In our experiments we focused on the first level, which has 11 different tags, out of which 5 are common to user and system segments.

4.2 Classification Approach

As a classification approach, we adapted the state-of-the-art word-level approach by Liu et al. [15] to use characters instead of words as tokens. As shown in Figure 1, the token embeddings are passed through a set of parallel temporal CNNs with different context window sizes followed by a max pooling operation. The results of those operations are then concatenated to form a representation of the segment. To achieve the state-of-the-art results, additional features concerning context information are appended to that representation before it is passed through a dimensionality reduction layer. Since our study focuses on the difference between using character- and word-level tokenization, we only included that information in a final experiment for comparison with the state-of-the-art. Finally, the reduced segment representation is passed through a dense layer with the softmax activation to obtain its classification.




Fig. 1. The generic architecture of the network used in our experiments. t_i corresponds to the embedding representation of the i -th token. w corresponds to the context window size of the CNN. The number of parallel CNNs and the respective window sizes vary between experiments. Those shown in the figure correspond to the ones used by Liu et al. [15] in their experiments.

In order to assess whether the character- and word-level approaches capture complementary information, we also performed experiments that combined both approaches. In that scenario, we used the architecture shown in Figure 2. In this case, two segment representations are generated in parallel, one based on the characters in the segment and other on its words. Those representations are then concatenated to form the final representation of the segment. The following steps do not differ from the architecture with a single branch. That is, context information can be added to the segment representation before it is passed to the two dense layers.




Fig. 2. The architecture of the network that combines the character- and -word-level approaches. c_i corresponds to the embedding representation of the i -th character while w_i corresponds to the embedding representation of the i -th word. The context window sizes of the CNNs in the character-level branch refer to those that achieved best performance in our experiments. Those on the word-level branch correspond to the ones used by Liu et al. [15] in their experiments.

We used Keras [5] with the TensorFlow [1] backend to implement the networks. The training phase stopped after 10 epochs without improvement on the

validation set. The results presented in the next section refer to the average (μ) and standard deviation (σ) accuracy values over 10 runs.

4.3 Baselines

In order to assess the performance of the character-level approach, in comparison to the word-level approach, we defined two baselines. One of them uses randomly initialized word embeddings that are adapted during the training phase, while the other uses fixed pre-trained embeddings. The latter were obtained by applying Word2Vec [20] on the English Wikipedia¹ and the Spanish Billion Word Corpus [4]. Additionally, we defined a third baseline that replicates the state-of-the-art approach by Liu et al. [15]. It consists of the baseline with pre-trained embeddings combined with context information from three preceding segments in the form of their gold standard annotations and speaker change information in the form of a flag. Similarly to Liu et al. [15], we used three parallel CNNs with context window sizes one, two, and three in all the baselines.

5 Results

Starting with the word-level baselines, in Table 1 we can see that, in comparison to using randomly initialized word embeddings, using fixed pre-trained embeddings led to an average accuracy improvement of .64 and .88 percentage points on the validation (SWBD-V) and test (SWBD-T) sets of the Switchboard corpus, respectively. However, that was not the case on the DIHANA corpus, where the improvement was negligible. This can be explained by the difference in the nature of the dialogs between corpora. Since the Switchboard dialogs have a large variability in terms of style and domain, the performance on the validation and test sets is impaired when overfitting to the training data occurs. On the other hand, since most DIHANA dialogs are similar, the cross-validation performance is not impaired and may actually benefit from it. The improvement provided by context information is in line with that reported by Liu et al. [15]. Our results on the Switchboard corpus vary from those reported in their paper mainly because they did not use the standard validation and test partitions.

Table 1. Accuracy results of the word-level baselines.

	SWBD-V		SWBD-T		DIHANA	
	μ	σ	μ	σ	μ	σ
Random	.7617	.0019	.7223	.0020	.9196	.0013
Pre-trained	.7681	.0032	.7311	.0026	.9198	.0012
Pre-trained + Context	.8129	.0030	.7835	.0036	.9826	.0004

¹<https://dumps.wikimedia.org/enwiki/>

Regarding the character-level experiments, in Table 2 we can see that, as expected, considering each character individually is not the appropriate approach to capture intention. By considering pairs of characters, the performance improved by over 5 percentage points on both corpora. Widening the window up to five characters leads to a nearly 3 percentage point improvement on the Switchboard corpus, but less than 1 percentage point on the DIHANA corpus. However, it is important to note that while the results are above 90% accuracy on the DIHANA corpus, they are below 80% on the Switchboard corpus. Thus, improvements are expected to be less noticeable on the first. Using a window of seven characters still improves the results on the Switchboard corpus, but is not relevant on the DIHANA corpus. Considering wider windows is harmful on both corpora. However, note that, similarly to what Liu et al. [15] have shown at the word level, different context windows are able to capture complementary information. Thus, it is beneficial to combine multiple windows. In our experiments, the best results on both corpora were achieved using three context windows, which considered groups of three, five, and seven characters, respectively. The sizes of these windows are relevant, since the shortest window is able to capture most affixes in English and the small affixes in Spanish, the middle window is able to capture the larger Spanish affixes and most lemmas in both languages, and the widest window is able to capture larger words and inter-word information. Finally, it is relevant to note that the results on the DIHANA corpus are already above the word-level baselines.

Table 2. Accuracy results using different token context windows.

Window Size(s)	SWBD-V		SWBD-T		DIHANA	
	μ	σ	μ	σ	μ	σ
1	.6542	.0017	.6081	.0023	.8571	.0029
2	.7221	.0047	.6752	.0054	.9154	.0014
3	.7432	.0055	.7000	.0035	.9217	.0010
4	.7456	.0019	.7064	.0049	.9222	.0014
5	.7509	.0052	.7091	.0038	.9228	.0011
7	.7535	.0023	.7086	.0034	.9224	.0013
10	.7510	.0036	.7097	.0035	.9216	.0013
(3, 5, 7)	.7608	.0033	.7208	.0042	.9244	.0012

In Section 3, we hypothesized that capitalization is not relevant for dialog act recognition. In Table 3, we can see that the hypothesis holds for the Switchboard corpus, as the results obtained when using capitalized segments do not significantly differ from those obtained using uncapitalized segments. However, on the DIHANA corpus, using capitalized segments led to an average improvement of 1.81 percentage points. Since this was not expected, we looked for the source of the improvement. By inspecting the transcriptions, we noticed that, contrarily to user segments, the system segments do not contain mid-segment

capitalization. Thus, proper nouns, such as city names which are common in the dialogs, are capitalized differently. Since only 5 of the 11 dialog acts are common to user and system segments, identifying its source reduces the set of possible dialog acts for the segment. Thus, the improvement observed when using capitalization information is justified by the cues it provides to identify whether it is a user or system segment.

Table 3. Accuracy results using different segment preprocessing approaches.

	SWBD-V		SWBD-T		DIHANA	
	μ	σ	μ	σ	μ	σ
Capitalized	.7604	.0028	.7194	.0026	.9425	.0015
Punctuated	.7685	.0021	.7317	.0032	.9371	.0007
Capitalized + Punctuated	.7673	.0025	.7314	.0040	.9548	.0004
Lemmatized	.7521	.0027	.7140	.0012	.9239	.0006

In Table 3, we can also see that, as expected, punctuation provides relevant information for the task, improving the performance around 1 percentage point on both corpora. Using this information, the character-level approach surpasses the randomly initialized word-level baseline and is in line with the one using pre-trained word embeddings on the Switchboard corpus. Also expectedly, the decrease in performance observed when using lemmatized segments proves that affixes are relevant. However, that decrease is not drastic, which suggests that most information concerning intention can be transmitted using a simplified language that does not consider variations of the same lexeme and that those variations are only relevant for transmitting some specific intentions.

Table 4. Accuracy results using the combination of word and character-level representations.

	SWBD-V		SWBD-T		DIHANA	
	μ	σ	μ	σ	μ	σ
Char + Word	.7800	.0016	.7401	.0035	.9568	.0003
Char + Word + Context	.8200	.0027	.7901	.0016	.9910	.0004

Finally, Table 4 shows the results obtained by combining the word and character-level approaches. We can see that the performance increases on both corpora, which means that both approaches are able to capture complementary information. This confirms that information at the sub-word level is relevant for the task. When using context information, the combination of both approaches leads to results that surpass the state-of-the-art word-level approach on the

Switchboard corpus by around .7 percentage points and to a nearly perfect score on the DIHANA corpus. Concerning the latter, it is not fair to compare our results with those of previous studies, since the only one that focused on Level 1 labels did not rely on textual information [26].

6 Conclusions

We have shown that there is important information for dialog act recognition at the sub-word level which cannot be captured by word-level approaches. We used character-level tokenization together with multiple context windows with different sizes in order to capture relevant morphological elements, such as affixes and lemmas, as well as long words and inter-word information. Furthermore, we have shown that, as expected, punctuation is important for the task since it is able to provide both direct and indirect cues regarding intention. On the other hand, capitalization is irrelevant under normal conditions. Finally, our experiments revealed that the character- and word-level approaches capture complementary information and, consequently, their combination leads to improved performance on the task. In this sense, by combining both approaches with context information we achieved state-of-the-art results on the Switchboard corpus and a nearly perfect score on the DIHANA corpus.

It is important to note that while one of the corpora used in our experiments features variable-domain human-human interactions in English, the other features fixed-domain interactions in Spanish between a WoZ dialog system and its users. Thus, the importance of information at the sub-word level is not domain-dependent and it is not limited to a single language.

In terms of morphological typology, although English has a more analytic structure than Spanish, both are fusional languages. Thus, as future work it would be interesting to assess whether the conclusions of this study hold for analytic languages, such as Chinese, and agglutinative languages, such as Turkish.

References

1. Abadi, M., et al.: TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. <https://www.tensorflow.org/> (2015)
2. Alcácer, N., Benedí, J.M., Blat, F., Granell, R., Martínez, C.D., Torres, F.: Acquisition and Labelling of a Spontaneous Speech Dialogue Corpus. In: SPECOM. pp. 583–586 (2005)
3. Benedí, J.M., Lleida, E., Varona, A., Castro, M.J., Galiano, I., Justo, R., de Letona, I.L., Miguel, A.: Design and Acquisition of a Telephone Spontaneous Speech Dialogue Corpus in Spanish: DIHANA. In: LREC. pp. 1636–1639 (2006)
4. Cardellino, C.: Spanish Billion Words Corpus and Embeddings. <http://crscardellino.me/SBWCE/> (2016)
5. Chollet, F., et al.: Keras: The Python Deep Learning Library. <https://keras.io/> (2015)
6. Gambäck, B., Olsson, F., Täckström, O.: Active Learning for Dialogue Act Classification. In: INTERSPEECH. pp. 1329–1332 (2011)

7. Goldberg, Y.: A primer on neural network models for natural language processing. *Journal of Artificial Intelligence Research* **57**, 345–420 (2016)
8. Jaech, A., Mulcaire, G., Hathi, S., Ostendorf, M., Smith, N.A.: Hierarchical Character-Word Models for Language Identification. In: *International Workshop on Natural Language Processing for Social Media*. pp. 84–93 (2016)
9. Ji, Y., Haffari, G., Eisenstein, J.: A Latent Variable Recurrent Neural Network for Discourse Relation Language Models. In: *NAACL-HLT*. pp. 332–342 (2016)
10. Jurafsky, D., Shriberg, E., Biasca, D.: *Switchboard SWBD-DAMSL Shallow-Discourse-Function Annotation Coders Manual*. Tech. Rep. Draft 13, University of Colorado, Institute of Cognitive Science (1997)
11. Kalchbrenner, N., Blunsom, P.: Recurrent Convolutional Neural Networks for Discourse Compositionality. In: *Workshop on Continuous Vector Space Models and their Compositionality*. pp. 119–126 (2013)
12. Khanpour, H., Guntakandla, N., Nielsen, R.: Dialogue Act Classification in Domain-Independent Conversations Using a Deep Recurrent Neural Network. In: *COLING*. pp. 2012–2021 (2016)
13. Král, P., Cerisara, C.: Dialogue Act Recognition Approaches. *Computing and Informatics* **29**(2), 227–250 (2010)
14. Lee, J.Y., Dérnoncourt, F.: Sequential Short-Text Classification with Recurrent and Convolutional Neural Networks. In: *NAACL-HLT*. pp. 515–520 (2016)
15. Liu, Y., Han, K., Tan, Z., Lei, Y.: Using Context Information for Dialog Act Classification in DNN Framework. In: *EMNLP*. pp. 2160–2168 (2017)
16. Manning, C.D.: *Computational Linguistics and Deep Learning*. *Computational Linguistics* **41**(4), 701–707 (2015)
17. Martínez-Hinarejos, C.D., Benedí, J.M., Granell, R.: Statistical Framework for a Spanish Spoken Dialogue Corpus. *Speech Communication* **50**(11–12), 992–1008 (2008)
18. Martínez-Hinarejos, C.D., Sanchis, E., García-Granada, F., Aibar, P.: A Labelling Proposal to Annotate Dialogues. In: *LREC*. vol. 5, pp. 1566–1582 (2002)
19. Mikolov, T., Karafit, M., Burget, L., Cernock, J., Khudanpur, S.: Recurrent Neural Network Based Language Model. In: *INTERSPEECH*. pp. 1045–1048 (2010)
20. Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed Representations of Words and Phrases and their Compositionality. In: *NIPS*. pp. 3111–3119 (2013)
21. Pennington, J., Socher, R., Manning, C.D.: GloVe: Global Vectors for Word Representation. In: *EMNLP*. pp. 1532–1543 (2014)
22. Ribeiro, E., Ribeiro, R., de Matos, D.M.: The Influence of Context on Dialogue Act Recognition. *CoRR* [abs/1506.00839](https://arxiv.org/abs/1506.00839) (2015), <http://arxiv.org/abs/1506.00839>
23. Santos, C.D., Zadrozny, B.: Learning Character-level Representations for Part-of-Speech Tagging. In: *ICML*. pp. 1818–1826 (2014)
24. Searle, J.R.: *Speech Acts: An Essay in the Philosophy of Language*. Cambridge University Press, Cambridge, London (1969)
25. Stolcke, A., Coccaro, N., Bates, R., Taylor, P., Van Ess-Dykema, C., Ries, K., Shriberg, E., Jurafsky, D., Martin, R., Meteer, M.: Dialogue Act Modeling for Automatic Tagging and Recognition of Conversational Speech. *Computational Linguistics* **26**(3), 339–373 (2000)
26. Tamarit, V., Martínez-Hinarejos, C.D.: Dialog Act Labeling in the DIHANA Corpus using Prosody Information. In: *V Jornadas en Tecnología del Habla*. pp. 183–186 (2008)
27. Zhang, X., Zhao, J., LeCun, Y.: Character-level Convolutional Networks for Text Classification. In: *NIPS*. vol. 1, pp. 649–657 (2015)