LEXICO-ACOUSTIC NEURAL-BASED MODELS FOR DIALOG ACT CLASSIFICATION

Daniel Ortega Ngoc Thang Vu

Institute for Natural Language Processing (IMS) University of Stuttgart, Germany {daniel.ortega, thang.vu}@ims.uni-stuttgart.de

ABSTRACT

Recent works have proposed neural models for dialog act classification in spoken dialogs. However, they have not explored the role and the usefulness of acoustic information. We propose a neural model that processes both lexical and acoustic features for classification. Our results on two benchmark datasets reveal that acoustic features are helpful in improving the overall accuracy. Finally, a deeper analysis shows that acoustic features are valuable in three cases: when a dialog act has sufficient data, when lexical information is limited and when strong lexical cues are not present.

Index Terms— dialog act, lexico-acoustic features

1. INTRODUCTION

Every utterance in a conversation has a level of illocutionary force [1] whose meaning induces an effect over the course of the dialog. That meaning can be categorized into dialog acts (DAs) taking into account the relationship between the words being used and the force of the utterance [2]. A DA is the expression of the speaker's attitude or intention at every utterance in a conversation. Kent Bach [2] illustrates this by pointing out that a statement expresses a belief, a request expresses a desire, and an apology expresses a regret. In this manner, dialogs can be studied and modeled by analyzing their sequence of DAs.

Automatic DA tagging is an important preprocessing step for semantic extraction in natural language understanding and dialog systems. This task has been approached using two main information sources: lexical cues from dialog transcripts and acoustic cues from speech signals. For the former, traditional statistical algorithms have been employed, such as hidden Markov models (HMMs) [3], conditional random fields (CRFs), [4] and support vector machines (SVMs) [5]. Recently, deep learning (DL) techniques, such as convolutional neural networks (CNNs) [6, 7], recurrent neural networks (RNNs) [7, 8] and long short-term memory (LSTM) models [9], have attained the state-of-the-art results in DA classification.

DAs can be ambiguous if only lexical information is considered. For example, a *Declarative Question* like "*this is your car*(?)" is hard to distinguish from a *Statement* if the question mark is not present, and can easily be misclassified due to word order. In this case, acoustic information can help disambiguate. Moreover, in real applications that involve automatic speech recognition (ASR), a DA classifier can help deal with noisy transcriptions. Hence, some researchers [3, 10] have explored acoustic and prosodic cues from the speech signal as a potential knowledge source for DA classification.

Other works [11, 12] have showed improvements exploring combinations of lexico-acoustic features. Inspired by these works, we present a neural hybrid model that takes both lexical and acoustic features as input in order to classify dialog utterances into DAs. Our model is a combination of two neuralbased models: one, which processes lexical features of the utterances and their context (based on [13]), and a second, which processes acoustic features. Our experiments show that acoustic features are helpful for improving overall accuracy and attaining state-of-the-art results on two benchmark datasets: the ICSI Meeting Recorder Dialog Act Corpus (MRDA) and the NXT-format Switchboard Corpus (SwDA). We also include an analysis of the acoustic features contribution for DA classification in three circumstances: when a DA has sufficient data, when strong lexical cues are missing and for single-word utterances.

2. MODEL

The architecture of the lexico-acoustic model (LAM) proposed in this paper is depicted in Figure 1. It contains two main parts: on the left side is the lexical model (LM) and on the right side the acoustic model (AM). Models are detailed in sections 2.1-2.3.

2.1. Lexical model

The LM, based on [13], takes the concatenation of grid-like representations of the current utterance and its n previous utterances in the dialog as input to be processed by a CNN, generating a vector representation for each of those utterances.

The CNN performs a discrete convolution using a set of different filters on an input matrix, where each column of the matrix is the word embedding of the corresponding word. We use 2D filters f (with width |f|) spanning over all embedding



Fig. 1. Architecture of the lexico-acoustic model. \oplus represents a concatenation.

dimensions d as described by the following equation:

$$(w*f)(x,y) = \sum_{i=1}^{d} \sum_{j=-|f|/2}^{|f|/2} w(i,j) \cdot f(x-i,y-j) \quad (1)$$

After convolution, an utterance-wise max pooling operation is applied. Then, the feature maps are concatenated, resulting in one vector per utterance. These are represented in Figure 1 as p_{t-2} , p_{t-1} and p_t .

The vector representations of the utterances are then processed by a context learning method, the RNN-Output-Attention (ROA) proposed in [13], in order to model the relation between each utterance and its context. ROA consists of an RNN with LSTM units followed by a weighted sum of the RNN's hidden states using an attention mechanism [14].

For each of the hidden state vector h(t - i) at time step t - i in a dialog, where t is the current time step. The attention weights α_i are computed as follows

$$\alpha_i = \frac{exp(f(h(t-i)))}{\sum_i exp(f(h(t-j)))}$$
(2)

where f is the scoring function. In our work, f is the linear function of the input h(t-i)

$$f(h(t-i)) = W^T h(t-i)$$
(3)

where W is a trainable parameter. The output l_t is the weighted sum of the hidden states sequence.

$$l_t = \sum_i \alpha_i h(t-i) \tag{4}$$

Finally, the context representation l_t is fed into a softmax layer that outputs a probability distribution over the DA set, given the current dialog utterance.

2.2. Acoustic model

We propose a CNN-based model to process acoustic features, because the speech signal of the utterances encodes important information for DA classification that is not contained in the transcripts. The acoustic features from the speech signal are not taken at word or utterance level, but at frame level, i.e. the speech signal is divided into frames of 25 ms with a shift of 10 ms, and 13 Mel-frequency-cepstral coefficients (MFCC) per frame are extracted using the openSMILE toolkit [15]. MFCC features are stacked sequentially in order to obtain a grid-like input representation of the acoustic signal.

The input is processed by a one-layer CNN using filters that span over the 13 MFCC features and 5 frames a time, with a max pooling layer in order to obtain a fixed-length vector representation. This is fed into a softmax layer for DA classification, as explained in Section 2.1.

2.3. Lexico-acoustic model

The core of this work is the LAM (depicted in Figure 1), a bi-CNN that employs lexical and acoustic cues simultaneously as input. The LAM combines a LM and a AM by concatenating the vector representations (l_t and a_t) obtained from the context processing method in the LM and the pooling layer in the AM. Both vectors represent the current utterance, and can therefore be joined at this level and passed to the softmax function to output a final probability distribution over the DA set.

3. EXPERIMENTS

3.1. Data

We test our model on two DA datasets: 1) **MRDA**: ICSI Meeting Recorder Dialog Act Corpus [16], a dialog corpus of *multiparty meetings*. The 5-tag-set used in this work was introduced by [17]. 2) **SwDA**: NXT-format Switchboard Corpus [18], a dialog corpus of 2-speaker conversations.

NXT-format Switchboard Corpus was preferred over the original Switchboard Dialog Act Corpus [19, 20] because the former provides utterance transcripts and DA annotations as well as the time stamps at word level that were useful to extract acoustic features. Nonetheless, this corpus only provides DA annotation for roughly 50% of the original dataset.

Train, validation and test splits on MRDA were taken as defined in [7]. However, on SwDA the splits were built by taking the annotated conversations from NXT-format Switchboard Corpus that appear in the split lists published in [7]. The new train, validation and test splits are roughly the half of the conversations on each original split. Summary statistics are shown in Table 1. In both datasets, the classes are highly unbalanced; the majority class is 59.1% on MRDA and 34.7% on SwDA.

Dataset	С	$ \mathbf{V} $	Train	Validation	Test
MRDA	5	12k	78k	16k	15k
SwDA	42	16k	98k	8.5k	2.5k

Table 1. Data statistics: C is the number of classes, |V| is the vocabulary size and Train/Validation/Test are #utterances.

3.2. Hyperparameters and Training

The hyperparameters of the three models for both datasets are summarized in Table 2. The LM's hyperparameters were taken from [13], while the AM's hyperparameters were obtained by varying one hyperparameter at a time while keeping the others fixed. Training was done for 25 epochs with averaged stochastic gradient descent [21] over mini-batches. The learning rate was initialized at 0.11 and reduced 10% every 2000 parameter updates. Word2vec pretrained embeddings [22] were employed and tuned during training. The context length n was taken from the original the LM, i.e. n = 3 for MRDA and n = 2 for SwDA.

Hyperparameter	LM	AM
Filter width	3, 4, 5	5
Feature maps per filter	100	100
Dropout rate	0.5	0.5
Activation function	ReLU	ReLU
Pooling size	utterance-wise	(18,1)
Word embeddings	word2vec[22] (dim. 300)	
MFCC features		13
Mini-batch size	50 (MRDA) – 150 (SwDA)	

 Table 2. Hyperparameters.

3.3. Results

Table 3 shows the results obtained from the three models on both datasets. As expected, the LM is superior to the AM, i.e. the lexical features yield more valuable information than the acoustic features for our task. On both datasets, the LM's accuracy is significantly higher than the AM's accuracy. However, for both datasets, the combined model yields improvements over both constituent models. It indicates that both cue sources complement each other.

Model	MRDA	SwDA
Lexical	84.1	73.6
Acoustic	67.8	50.9
Lexico-acoustic	84.7	75.1

Table 3. Accuracy (%) of the three models on both datasets.

4. ANALYSIS

This section's goal is to analyze the impact of joining both models, and to report and discuss which DAs benefit and which are impacted negatively, by applying a LAM versus a LM. Moreover, we also investigate the effect of the acoustic features when the question mark (?) is removed from transcripts and when utterances are very short.

On MRDA, as reported in the previous section, the LAM yielded an improvement of 0.6% over the LM. However, the improvement is not uniform over the five classes. While the prediction of the DAs *Statement*, *Disruption* and *Backchannel* obtains a benefit from the acoustic features, *Filter* is impacted negatively and *Question* stays the same. Nonetheless, in general terms, the LAM benefits the overall DA classification, specially for those DAs with a higher presence in the training set, and the degradation caused by the model does not hurt its overall performance.

On SwDA, the LAM also outperformed the LM by 1.5%. Five DAs benefited by adding acoustic features: *Statement*, *Backchannel*, *Opinion*, *Abandon* and *Agree*, *Wh_question* and *Acknowledge* were negatively affected in a minimal extent, and the remaining 35 DAs were not impacted. These results are again highly correlated to the DA distribution in the corpus – the 5 most frequent DAs obtained an improvement that is reflected in the overall accuracy. Therefore, we argue that the LAM helps when a large number of examples per DA is available. One possible reason is that we have enough training data for these particular DAs to properly train the AM part of the LAM.

Effect of removing the question mark Contrary to our initial hypothesis that acoustic features would improve the accuracy of classifying *Question*, no improvement was noted. Therefore, we analyzed how the LM and the LAM performed on this particular DA more deeply. The question mark ? in the manual transcripts plays a fundamental role for the DA *Question* in the LM; 97.7% of the utterances with question marks which are labeled as *Question* are correctly predicted (see Table 4) by the LM. For that reason, the acoustic features are not able to provide any useful information.

Consequently, we retrained and tested the LM and the LAM using transcripts from which the question mark was removed. This change also makes the transcripts more similar to transcripts from an ASR, where punctuation is not available or is not highly accurate. As expected, the overall accuracy dropped, from 84.1% to 80.8% in the LM and from 84.7% to 81.9% in the LAM. Although both models were affected by this modification, the LAM performed 1.1% better than the LM, versus the improvement of 0.6% with the original transcripts. Acoustic features slightly dampen the negative effect on the accuracy of removing the question mark.

Table 4 shows the accuracy of the LM and the LAM exclusively on utterances whose DA is *Question* and which have a question mark in the manual transcript. The second column corresponds to the models which were trained and tested on the original transcripts and the third column to the models which were trained on transcripts with question marks removed. As mentioned above, the LM has a high accuracy at correctly predicting *Question* if the utterance has the question mark. Moreover, when the acoustic features are added, the accuracy decreases by 1.6%. Nonetheless, if question marks are not present in the data, the LM's accuracy drops to 46.6%. This shows that this character is the most important cue at lexical level. The LAM's accuracy drops to 50.2%, but this time it is superior to the LM by 3.6%. This indicates that acoustic information is an important source of cues for tasks that use DA classification over data that lacks these important lexical cues, such as spoken language understanding.

Model	With '?'	'?' removed
Lexical	97.7	46.6
Lexico-acoustic	96.1	50.2

Table 4. Accuracy (%) of *Question* utterances on MRDA withquestion mark and when the question mark is removed.

Single-word utterances There exist utterances like *Right* or *Yeah* that are very common across several DAs. One of their characteristics is that they are very short and consequently they do not yield much information for classification. [13, 7, 23] have successfully explored the use of context as a way to differentiate these type of utterances. In line with these works, both the LM and the LAM (in its lexical component) encode the context.

We have shown in Section 3.3 that the LAM outperforms the LM on both datasets, however, we explored particularly the effect of using acoustic features on the utterances *Right* and *Yeah* that are frequently tagged as *Statement* and *Backchannel* on MRDA. For our analysis purposes, we extracted the predictions of the utterances that exclusively contained one word that is either *Right* or *Yeah*, from which we can artificially define four subclasses: *Statement-Right*, *Backchannel-Right*, *Statement-Yeah* and *Backchannel-Yeah*.

Table 5 shows the precision, recall and F_1 score of the LM and the LAM for the utterances *Right*. On the one hand, for the DA *Statement* the LAM achieves a higher F_1 score than the LM, while on the other hand, the F_1 score for *Backchannel* decreases slightly. This means that using acoustic features improves the classification of utterances *Right* as *Statement* without affecting those utterances tagged as *Backchannel*. A similar phenomenon is observed with utterances *Yeah*, however, in this case, the LAM improves the F_1 score for both DAs *Statement* and *Backchannel* (see Table 6).

5. COMPARISON WITH OTHER WORKS

We present a comparison between different works and our model in Table 5. On MRDA, as we used the setup proposed by [7], our results can only be compared accurately to [13] and [7], and the LAM outperforms both works. On SwDA, as we used the data available in the NXT format, and, to the best of

DA-Right	Model	Р	R	F ₁
Statement	Lexical	0.62	0.35	0.45
	Lexico-acoustic	0.60	0.45	0.52
Backchannel	Lexical	0.56	0.85	0.67
	Lexico-acoustic	0.56	0.77	0.65

Table 5. Precision, recall and F₁ score for the utterances *Right*

DA-Yeah	Model	Р	R	F ₁
Statement	Lexical	0.65	0.36	0.46
	Lexico-acoustic	0.67	0.50	0.57
Backchannel	Lexical	0.60	0.89	0.72
	Lexico-acoustic	0.64	0.87	0.74

Table 6. Precision, recall and F₁ score for the utterances Yeah

our knowledge, no other model has been trained and tested on this subset of SwDA, our results cannot be strictly compared with other works.

Model	MRDA	SwDA
LAM (Our model)	84.7	75.1
NCRL	84.3	73.8
CNN-FF	84.6	73.1
HBM	81.3	_
CNN+DAP	_	79.9
HCNN	_	73.9
HMM	_	71.0
Majority class	59.1	34.7

Table 7. Comparison of accuracy (%). *NCRL*: Neural context representation learning proposed in [13], *CNN-FF*: proposed in [7], *HBM*: hidden backoff model [24]. *CNN+DAP*:proposed by [23]. *HCNN*: hierarchical CNN [6]. *HMM* [3]. *Majority class* is the most frequent class.

6. CONCLUSION

We proposed an approach to incorporate lexical and acoustic features in a neural model for DA classification. Our experiments on two benchmark datasets reveal that adding acoustic information to the model improves the overall accuracy attaining state-of-the-art results. A deeper analysis showed that acoustic features specially help when the data for a particular DA is large enough, when lexical information is limited, as in single-word utterances, and when strong lexical cues are not present.

7. ACKNOWLEDGEMENTS

This work was funded by the National Council of Science and Technology of Mexico (CONACyT), the German Academic Exchange Service (DAAD) and the German Science Foundation (DFG), Sonderforschungsbereich 732, Project A8, at the University of Stuttgart.

8. REFERENCES

- [1] J.L. Austin, How to Do Things with Words, 1962.
- [2] Kent Bach, *Concise Routledge Encyclopedia of Philosophy*, chapter Speech Acts, 2000.
- [3] Andreas Stolcke, Noah Coccaro, Rebecca Bates, Paul Taylor, Carol Van Ess-Dykema, Klaus Ries, Elizabeth Shriberg, Daniel Jurafsky, Rachel Martin, and Marie Meteer, "Dialogue act modeling for automatic tagging and recognition of conversational speech," *Comput. Linguist.*, 2000.
- [4] Matthias Zimmermann, "Joint segmentation and classification of dialog acts using conditional random fields.," in *INTERSPEECH*, 2009.
- [5] M. Henderson, M. Gai, B. Thomson, P. Tsiakoulis, K. Yu, and S. Young, "Discriminative spoken language understanding using word confusion networks," in *IEEE SLT*, 2012.
- [6] Nal Kalchbrenner and Phil Blunsom, "Recurrent convolutional neural networks for discourse compositionality," *CoRR*, 2013.
- [7] Ji Young Lee and Franck Dernoncourt, "Sequential shorttext classification with recurrent and convolutional neural networks," *CoRR*, 2016.
- [8] Yangfeng Ji, Gholamreza Haffari, and Jacob Eisenstein, "A latent variable recurrent neural network for discourse relation language models," *CoRR*, 2016.
- [9] Sheng-syun Shen and Hung-yi Lee, "Neural attention models for sequence classification: Analysis and application to key term extraction and dialogue act detection," *CoRR*, 2016.
- [10] Elizabeth Shriberg, Rebecca A. Bates, Andreas Stolcke, Paul Taylor, Daniel Jurafsky, Klaus Ries, Noah Coccaro, Rachel Martin, Marie Meteer, and Carol Van Ess-Dykema, "Can prosody aid the automatic classification of dialog acts in conversational speech?," *CoRR*, 2000.
- [11] H. Arsikere, A. Sen, A. P. Prathosh, and V. Tyagi, "Novel acoustic features for automatic dialog-act tagging," in 2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2016.
- [12] Stanislav Ond and Jozef Juhr, "Distance-based dialog acts labeling," in *Cognitive Infocommunications (CogInfoCom)*, 2015 6th IEEE International Conference on, 2015.
- [13] Daniel Ortega and Ngoc Thang Vu, "Neural-based context representation learning for dialog act classification," in *SIGdial*, 2017.

- [14] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio, "Neural machine translation by jointly learning to align and translate," *CoRR*, 2014.
- [15] Florian Eyben, Felix Weninger, Florian Gross, and Björn Schuller, "Recent developments in opensmile, the munich open-source multimedia feature extractor," in *ACM International Conference on Multimedia*, 2013.
- [16] Elizabeth Shriberg, Raj Dhillon, Sonali Bhagat, Jeremy Ang, and Hannah Carvey, "The icsi meeting recorder dialog act (mrda) corpus," in *SIGdial*, 2004.
- [17] Jeremy Ang, Yang Liu, and Elizabeth Shriberg, "Automatic dialog act segmentation and classification in multiparty meetings," in *ICASSP*, 2005.
- [18] Sasha Calhoun, Jean Carletta, Jason M. Brenier, Neil Mayo, Dan Jurafsky, Mark Steedman, and David Beaver, "The nxt-format switchboard corpus: A rich resource for investigating the syntax, semantics, pragmatics and prosody of dialogue," *Lang. Resour. Eval.*, 2010.
- [19] John J. Godfrey, Edward C. Holliman, and Jane Mc-Daniel, "Switchboard: Telephone speech corpus for research and development," in *ICASSP*, 1992.
- [20] D. Jurafsky, E. Shriberg, and D. Biasca, "Switchboard SWBD-DAMSL shallow-discourse-function annotation coders manual," Tech. Rep., 1997.
- [21] B. T. Polyak and A. B. Juditsky, "Acceleration of stochastic approximation by averaging," *SIAM J. Control Optim.*, 1992.
- [22] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean, "Distributed representations of words and phrases and their compositionality," *CoRR*, 2013.
- [23] Yang Liu, Kun Han, Zhao Tan, and Yun Lei, "Using context information for dialog act classification in dnn framework," in *EMNLP*, 2017.
- [24] Gang Ji and Jeff Bilmes, "Backoff model training using partially observed data: Application to dialog act tagging," in *HLT-NAACL*, 2006.