



Cross-Lingual Approaches for Task-Specific Dialogue Act Recognition

Jiří Martínek, Christophe Cerisara, Pavel Král, Ladislav Lenc

► To cite this version:

Jiří Martínek, Christophe Cerisara, Pavel Král, Ladislav Lenc. Cross-Lingual Approaches for Task-Specific Dialogue Act Recognition. 17th IFIP International Conference on Artificial Intelligence Applications and Innovations (AIAI), Jun 2021, Hersonissos, Crete, Greece. pp.232-242, 10.1007/978-3-030-79150-6_19 . hal-03287703

HAL Id: hal-03287703

<https://inria.hal.science/hal-03287703>

Submitted on 15 Jul 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Cross-lingual Approaches for Task-specific Dialogue Act Recognition

Jiří Martínek^{1,2}, Christophe Cerisara³, Pavel Král^{1,2}, Ladislav Lenc^{1,2}

¹ Dept. of Computer Science & Engineering
Faculty of Applied Sciences
University of West Bohemia
Plzeň, Czech Republic

² NTIS - New Technologies for the Information Society
Faculty of Applied Sciences
University of West Bohemia
Plzeň, Czech Republic

³ Université de Lorraine
CNRS, LORIA, F-54000
Nancy, France

{jimar,pkral,llenc}@kiv.zcu.cz, cerisara@loria.fr

Abstract. In this paper we exploit cross-lingual models to enable dialogue act recognition for specific tasks with a small number of annotations. We design a transfer learning approach for dialogue act recognition and validate it on two different target languages and domains. We compute dialogue turn embeddings with both a CNN and multi-head self-attention model and show that the best results are obtained by combining all sources of transferred information. We further demonstrate that the proposed methods significantly outperform related cross-lingual DA recognition approaches.

Keywords: Dialogue act recognition · Cross-lingual · Transfer learning · BERT · Multi-head self-attention

1 Introduction

Automatic dialogue act (DA) recognition has reached near-human performance on standard corpora, such as the Switchboard Dialogue Act corpus. However, various application domains lead to different types of dialogues: this variability, which is illustrated in the French corpus described in Section 5.2, severely impacts the direct application of a model trained on a standard corpus to many other types of application tasks, as shown next. Furthermore, developing DA-recognition models in other languages than English requires the costly annotation of large-enough corpora. We propose investigating cross-lingual transfer learning methods to reduce the amount of annotation required for each new domain and language.

Transfer learning aims to reuse knowledge gained from a large corpus to improve the models trained on a related task with few annotations. We investigate in this work two sources of information from which we transfer knowledge: pre-trained English

BERT sentence embeddings and an English corpus annotated with dialogue acts. Two dialogue act recognition tasks are considered: appointment scheduling in German and casual conversations in French. The amount of French annotated dialogue acts is limited to a few hundred samples that may be annotated by one application developer within a few hours. The definition of the dialogue acts are the same in English and German but different in French.

In addition to the relatively large resources available in English, we further assume the availability of an automatic translation system; we will use for this purpose Google’s translation. Section 3 describes our transfer learning strategy.

2 Related Work

New approaches in the dialogue act recognition field are mainly evaluated on English datasets. Some methods have also been tested on other languages, such as Spanish (DI-HANA corpus [3]), Czech [12], French [2] and German (Verbmobil [11] corpus).

A nice review of the state-of-the-art of the domain is summarised in [17], where the models typically reach 80% of accuracy on standard DA datasets. Nevertheless, the variety of DA labels in datasets is an obstacle for effective multi-lingual and multi-dataset research. Several interesting research efforts have thus emerged to define and exploit generic dialogue acts [14]. However, in practice, the specific requirements of most target tasks prevent a widespread usage of such standards. We rather focus next on specific types of dialogues and task-related dialogue acts.

Transfer learning has been quite a popular approach in deep learning in recent years. Such approaches have proven particularly useful in the computer vision and Natural Language Processing (NLP) domains. For instance, for automated pavement distress detection and classification [10] or for face verification [5]. A well-known transfer learning approach consists in using pre-trained word embeddings, such as Word2vec (W2V) [15], ELMO [16] and BERT [9].

In standard transfer learning, information flows from the source to the target domain (one direction only). A related approach is multi-task learning [6] where information flow across all tasks (usually more than two): information learned in each task may improve other tasks learning.

In the dialogue act recognition domain, Dai et al. [8] fine-tune BERT to classify a single utterance with quite good results, while Wu et al. [19] propose task-oriented dialogue BERT (ToD BERT).

3 Models

3.1 English DA Classifier

Our initial English dialogue act recognition model trained on the English dataset annotated with dialogue acts is a multi-layer perceptron (MLP) with BERT embeddings as inputs. We have tested various topologies for this initial model and chosen this one because of its good performances and fast training times.

Each speaker turn, composed of a variable number of words, is first encoded into a single pre-trained 1024-dimensional sentence embedding vector with BERT Large⁴.




Fig. 1. English MLP model

As shown in Figure 1, two such vectors are computed, respectively, for the previous, and the current speaker turns and concatenated as an input to the MLP. The MLP outputs 17 tags, which correspond to the dialogue act labels of the Verbmobil-EN corpus. This MLP has been trained on the training part of the Verbmobil-EN corpus.

3.2 Speaker Turn Embeddings

In our MLP model described previously, variable-length sentences are encoded into a unique speaker turn embedding vector with a pre-trained BERT model. We have further used two other models to compute the speaker turn embeddings: a convolutional neural network (CNN) and a multi-head self-attention (MH-SAtt) model.

The CNN model, shown in Figure 2, is derived from the model of Martinek et al. [13]. It takes as an input a word sequence truncated/padded to 15 words that always include the last two words of the sequence, following [7]. These 15 words are encoded into either random embeddings or W2V vectors. The CNN outputs a 256-dimensional vector for the current speaker turn.

The MH-SAtt model transforms each input random word embedding with the standard scaled dot-product multi-head self-attention module [18]⁵. A global max pooling operation is then applied to compute the speaker turn embedding. Figure 3 shows how this model is used in our experiments.

In all our models, the previous speaker turn is also encoded into a 1024-dimensional vector with BERT and injected into the classification step: this is an easy way to take into account the previous dialogue act without requiring a costly RNN, CRF and/or beam search process.

⁴ from <https://github.com/google-research/bert#pre-trained-models> – BERT-Large, Uncased (Whole Word Masking): 24-layer, 1024-hidden, 16-heads, 340M parameters

⁵ <https://github.com/CyberZHG/keras-multi-head>

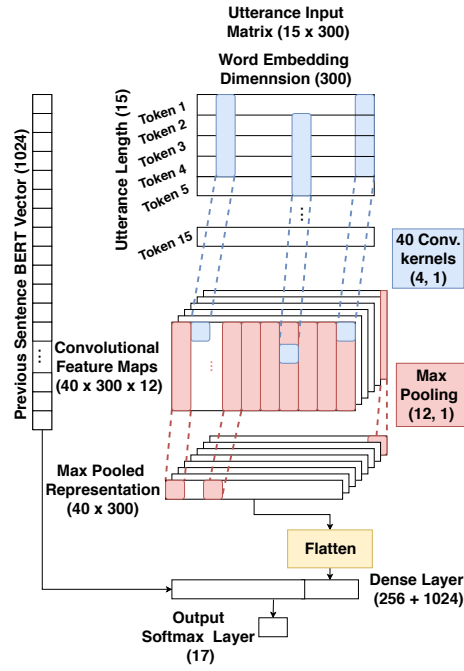


Fig. 2. CNN model for DA recognition

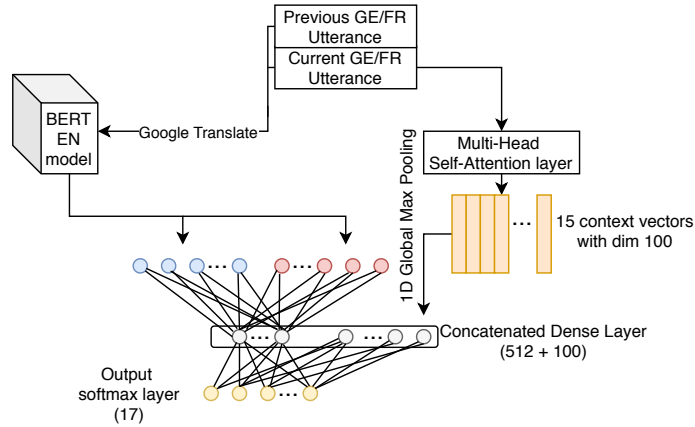


Fig. 3. Multi-head self-attention model

4 Transfer Learning Approach

Pre-trained BERT is famous for transferring general English lexical information into a large variety of downstream NLP tasks. We exploit them for cross-lingual dialogue act recognition through automatic language translation. Similarly, we evaluate pre-trained

word2vec English vectors in this context. We further investigate the option to stack the translated inputs with the original foreign representations in the MH-SAtt model to increase the robustness of the model to translation errors. Then, given a target non-English task with a few hundred annotated samples only, we evaluate the benefit from fine-tuning the English DA classifier on this target small dataset. Our global transfer learning approach thus consists of two phases:

1. Initial phase

The original English model in Figure 1 is trained on a large corpus annotated with dialogue acts.

2. Fine-tuning phase

- The foreign sentences of the task-specific corpus annotated with dialogues acts are automatically translated into English with Google Translate⁶;
- The final classification layer (in Figures 1, 2, 3) is replaced by a random layer with the same number of outputs than DAs in the target task;
- The model parameters are trained for a few epochs on the small target training corpus translated into English, as well as on the original foreign corpus for the MH-Satt model.

The diagram shown in Figure 4 summarises the two phases for both target tasks and languages that we have used in our experiments.

5 Experiments

In the following experiments, we use two target tasks, one in German and another one in French. Every experiment is run ten times and the results are averaged. The standard deviation is also computed.

5.1 German Task

Our first dialogue task consists in scheduling an appointment in German. Our target corpus is built by randomly sampling 100 utterances from the Verbmobil-GE training corpus with the same dialogue acts distribution as in the complete training corpus, see Table 1. We could have used the full Verbmobil corpus for training, but this would not have represented a realistic situation, where annotated training data for a specific task is usually scarce. However, to better evaluate the quality of our approach, the model trained on this small corpus is tested on the complete (1460 utterances), standard Verbmobil-GE test corpus.

⁶ Any other translation system may also be used

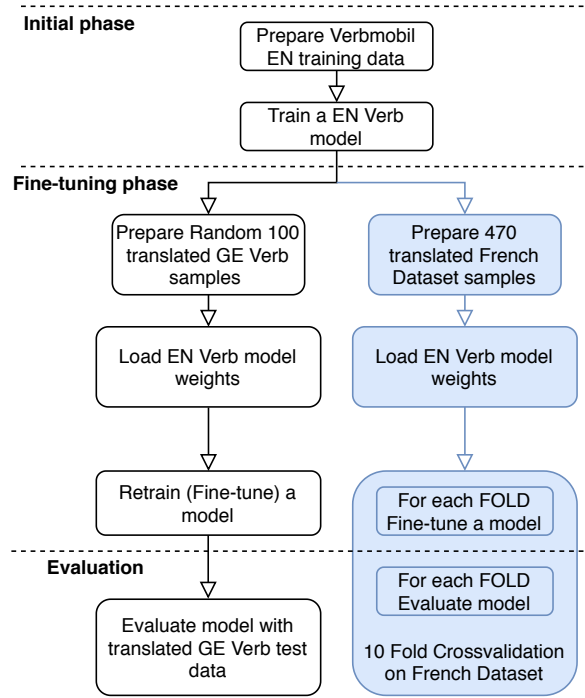


Fig. 4. Transfer learning process for both our experimental target tasks

Table 1. Distribution of DAs in Verbmobil-GE corpus

Label	Occurrence	Label	Occurrence
FEEDBACK	28%	DELIBERATE	3%
SUGGEST	19%	INTRODUCE	2%
INFORM	18%	COMMIT	1%
REQUEST	9%	CLOSE	1%
GREET	4%	POLIT. FORM.	1%
BYE	4%	THANK	1%
INIT	4%	DEFER	1%
BACKCHANNEL	3%	OFFER	1%

5.2 French Task

The French tasks consist in manual transcriptions of spoken dialogues recorded by linguists in real-life situations involving volunteer citizens. We manually annotate 470 turns from the TCOF corpus [1] with dialogue acts. The types of dialogue in our 470 turns are very different from the ones found in standard DA corpora. They involve two friends trying to find a suitable gift, three students talking about their courses while chatting with someone else on their smartphone, and an adult talking with a young girl

who is drawing. The dialogue acts distribution is shown in Table 2. This French corpus is freely distributed with a CC BY-NC-SA license⁷.

Table 2. Distribution of DAs in the French corpus

Label	Occurrence	Label	Occurrence
INFORM	26%	OPEN ANSWER	7%
AGREE	13%	DISAGREE	5%
BACKCHANNEL	10%	YES ANSWER	5%
Y/N QUESTION	10%	NO ANSWER	5%
OPEN QUESTION	9%	OTHER ANSWERS	1%
PERFORMATIVE	8%	GREETINGS	< 1%

5.3 Initial Phase: English Model

The initial English MLP model is trained on the full Verbmobil-EN corpus. The hyper-parameters of this model are tuned on the English corpus: this model is thus trained for 200 epochs with a learning rate of 0.002. Table 3 compares its accuracy when trained either from BERT embeddings or from speaker turn embeddings obtained with the CNN model. The CNN may use either initial random or pre-trained W2V word embeddings. Every experiment is run ten times and the results are averaged.

Table 3. Initial Phase – English MLP Model trained on Verbmobil-EN (9599 turns) and tested on Verbmobil-EN (1460 turns). The **Embeddings** column indicates how the word embeddings are initialised.

Model	Embeddings	Epochs	Test Acc	Std. Dev.
MLP	BERT	200	0.734	0.010
CNN	W2V	200	0.704	0.009
CNN	Random	200	0.702	0.008

5.4 Fine-Tuning Phase

Fine-tuning consists in training the last classification layers of the models presented in Section 3 on the small foreign corpus. This fine-tuning process thus involves some additional hyper-parameters, such as the learning rate and the fixed number of epochs, which are tuned on a German development corpus composed of 10,000 utterances randomly sampled from the Verbmobil-GE training corpus. The same hyperparameter values found on this German corpus are also used for the French model⁸.

⁷ <https://github.com/cerisara/TCOFDA>

⁸ All hyperparameter values along with the source code are distributed with an open-source license in <https://github.com/cerisara/TCOFDA>

5.5 Baseline Approaches

Three baseline classifiers are shown in the first part of Table 4:

1. **Majority Class Classifier (MC)**: it always predicts the most common DA in the training dataset.
2. **Training from Scratch**: Instead of transferring the model parameters from the initial English MLP model, the parameters (without the embeddings) are randomly initialised in this baseline.
3. **No-Fine-Tuning**: The model parameters are simply transferred from the initial English MLP model, and no further training is done.

5.6 Fine-Tuning Experiments

German Results Table 4 shows the results of our models on the Verbmobil-GE test set. All these models process only translated English sentences, except for the last MH-SAtt model that further includes the original German words, as shown in Figure 3.

Table 4. Accuracy on the Test Verbmobil-GE corpus

Model	Embeddings	Epochs	Test Acc	Std. Dev.
<i>Baseline MC</i>				
		–	0.279	–
<i>From scratch</i>				
CNN	W2V	15	0.410	0.012
MLP	BERT	25	0.468	0.008
<i>No-Fine-tuning</i>				
CNN	W2V	–	0.380	
MLP	BERT	–	0.479	
<i>Fine-Tuning</i>				
CNN	W2V	15	0.463	0.008
MLP	BERT	25	0.484	0.025
MH-SAtt	BERT + GE	50	0.502	0.012

With regard to transfer from pre-trained models, BERT is consistently better than W2V. Simply reusing the English models without fine-tuning gives the worst results, while fine-tuning the English models on the small target corpus systematically improves the results. The best performances are obtained with the MH-SAtt model that combines both BERT pre-trained vectors and fine-tuning transfer learning with the original and translated words sequences.

French Results We carried out 10-fold cross-validation on the 470 speaker turns of the French corpus. As we do not have enough data to create a development corpus in French, we share the same hyper-parameters as found on the German corpus. Table 5 shows the results of this experiment.

Table 5. French cross-validation experiments

Model	Embeddings	Epochs	Acc	Std. Dev.
<i>Baseline MC</i>		–	0.210	–
<i>From scratch</i>				
CNN	W2V	15	0.403	0.006
MLP	BERT	25	0.421	0.005
<i>Fine-Tuning</i>				
CNN	W2V	15	0.400	0.009
MLP	BERT	25	0.430	0.008
MH-SAtt	BERT + FR	50	0.436	0.007

The relative contributions of the various sources of information and models are similar to the German experiments. However, the differences between the results are much smaller, which is likely due to the fact that the hyper-parameters have not been tuned on French but on German.

5.7 Comparison with Related Work

We compare in Table 6 our proposal with other cross-lingual methods for dialogue act recognition that we are aware of. This related work projects German word embeddings into the English word embeddings space with the canonical correlation analysis method (CCA) [4]. The comparison is done on the German Verbmobil test corpus. Table 6 clearly shows that the proposed methods are significantly better than the previous work. To the best of our knowledge, no other study dealing with cross-lingual dialogue act recognition has been published.

Table 6. Comparison with previous cross-lingual DA recognition on the German Verbmobil test corpus

Methods	Acc
CNN + CCA [13]	0.315
BiLSTM + CCA [13]	0.340
MLP + BERT	0.484
MH-SAtt + BERT + GE	0.502

6 Conclusions

We explore two types of transfer learning for cross-lingual DA recognition: pre-trained word embeddings and classifier fine-tuning. Three types of dialogue turn embeddings are computed, based respectively on BERT, CNN and multi-head self-attention. The

objective is to leverage large available English resources annotated in dialogue acts to enable DA recognition on a specific target task and language with a limited amount of annotations.

We have validated these approaches on two tasks and languages, German and French, and released a dedicated French DA corpus with real-life dialogues recorded in quite different conditions than the existing standard DA corpora. The best results are obtained when all available sources of information are included, with both the BERT and multi-head self-attention dialogue turn embeddings.

We further demonstrated that the proposed methods significantly outperform cross-lingual DA recognition approaches developed previously.

Acknowledgements

This work has been partly supported by Cross-border Cooperation Program Czech Republic - Free State of Bavaria ETS Objective 2014-2020 (project no. 211), by Grant No. SGS-2019-018 Processing of heterogeneous data and its specialized applications and by GENCI-IDRIS (Grant 2021-AD011011668R1).

References

1. ATILF: Tcof : Traitement de corpus oraux en français (2018), <https://hdl.handle.net/11403/tcof/v2>, ORTOLANG (Open Resources and TOols for LANGUAGE) – www.ortolang.fr
2. Barahona, L.M.R., Lorenzo, A., Gardent, C.: Building and exploiting a corpus of dialog interactions between french speaking virtual and human agents. In: The eighth international conference on Language Resources and Evaluation (LREC). pp. 1428–1435 (2012)
3. Benedi, J.M., Lleida, E., Varona, A., Castro, M.J., Galiano, I., Justo, R., López, I., Miguel, A.: Design and acquisition of a telephone spontaneous speech dialogue corpus in spanish: Dihana. In: Fifth International Conference on Language Resources and Evaluation (LREC). pp. 1636–1639 (2006)
4. Brychcín, T.: Linear transformations for cross-lingual semantic textual similarity. *Knowledge-Based Systems* **187**, 104819 (2020)
5. Cao, X., Wipf, D., Wen, F., Duan, G., Sun, J.: A practical transfer learning algorithm for face verification. In: Proceedings of the IEEE international conference on computer vision. pp. 3208–3215 (2013)
6. Caruana, R.: Multitask learning. *Machine learning* **28**(1), 41–75 (1997)
7. Cerisara, C., Král, P., Lenc, L.: On the effects of using word2vec representations in neural networks for dialogue act recognition. *Computer Speech and Language* **47**, 175–193 (2018). <https://doi.org/10.1016/j.csl.2017.07.009>
8. Dai, Z., Fu, J., Zhu, Q., Cui, H., Qi, Y., et al.: Local contextual attention with hierarchical structure for dialogue act recognition. *arXiv preprint arXiv:2003.06044* (2020)
9. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018)
10. Gopalakrishnan, K., Khaitan, S.K., Choudhary, A., Agrawal, A.: Deep convolutional neural networks with transfer learning for computer vision-based data-driven pavement distress detection. *Construction and Building Materials* **157**, 322–330 (2017)

11. Jekat, S., Klein, A., Maier, E., Maleck, I., Mast, M., Quantz, J.J.: Dialogue acts in verbmobil. Tech. Rep. VM-Report 65, Universität Hamburg, DFKI GmbH, Universität Erlangen, and TU Berlin (1995)
12. Král, P., Cerisara, C., Klečková, J.: Combination of classifiers for automatic recognition of dialog acts. In: Interspeech'2005. pp. 825–828. ISCA, Lisboa, Portugal (September 2005)
13. Martínek, J., Král, P., Lenc, L., Cerisara, C.: Multi-Lingual Dialogue Act Recognition with Deep Learning Methods. In: Interspeech 2019. pp. 1463–1467. Graz, Austria (15–19 September 2019). <https://doi.org/10.21437/Interspeech.2019-1691>
14. Mezza, S., Cervone, A., Stepanov, E., Tortoreto, G., Riccardi, G.: Iso-standard domain-independent dialogue act tagging for conversational agents. In: Proceedings of the 27th International Conference on Computational Linguistics. pp. 3539–3551. Association for Computational Linguistics, Santa Fe, New Mexico, USA (Aug 2018), <https://www.aclweb.org/anthology/C18-1300>
15. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. arXiv preprint arXiv:1301.3781 (2013)
16. Peters, M.E., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., Zettlemoyer, L.: Deep contextualized word representations. arXiv preprint arXiv:1802.05365 (2018)
17. Ribeiro, E., Ribeiro, R., de Matos, D.M.: A multilingual and multidomain study on dialog act recognition using character-level tokenization. *Information* **10**(3), 94 (2019)
18. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. In: Advances in neural information processing systems. pp. 5998–6008 (2017)
19. Wu, C.S., Hoi, S., Socher, R., Xiong, C.: Tod-bert: Pre-trained natural language understanding for task-oriented dialogues. arXiv preprint arXiv:2004.06871 (2020)