

A Multi-layer LSTM-based Approach for Robot Command Interaction Modeling

Martino Mensio¹, Emanuele Bastianelli¹, Ilaria Tiddi², Giuseppe Rizzo³

Abstract—As the first robotic platforms slowly approach our everyday life, we can imagine a near future where service robots will be easily accessible by non-expert users through vocal interfaces. The capability of managing natural language would indeed speed up the process of integrating such platform in the ordinary life. Semantic parsing is a fundamental task of the Natural Language Understanding process, as it allows extracting the meaning of a user utterance to be used by a machine. In this paper, we present a preliminary study to semantically parse user vocal commands for a House Service robot, using a multi-layer Long-Short Term Memory neural network with attention mechanism. The system is trained on the Human Robot Interaction Corpus, and it is preliminarily compared with previous approaches.

I. INTRODUCTION

The area of Natural Language Understanding (NLU) has been gaining a growing consensus in recent years, also thanks to the trends dictated by new voice assistants, e.g. Amazon Echo. In parallel, the first cost-accessible commercial robots are growing in number, e.g. iRobot Roomba. As platforms evolve to expose complex services, it is legit to expect that they will be integrated with NLU capabilities. Natural language is, in fact, one of the most powerful and flexible communication tools, and vocal interface will become a mandatory feature, if we want service robots to be accessible by a wider range of users, especially non-expert ones.

Semantic parsing, the process of extracting interpretations from natural language, is a fundamental brick in NLU. In the last twenty years, a body of works have proposed solutions to this problem for virtual or real autonomous agents. Several approaches have been followed, from grammar-based ones [1], [2], to purely statistical ones [3], [4], [5], as well as hybrid approaches [6], [7], [8]. Among these, the work in [9] (henceforth BAS16) fosters the reliance on established linguistic theories to represent semantics of actions expressed in user’s commands. A SVM-based statistical semantic parser is trained over the Human-Robot Interaction Corpus [10] (HuRIC), which represents an attempt of bridging between the NLU for robots and more linguistically- and cognitively-sound theories of meaning representation, namely Frame Semantics [11] and the related FrameNet [12] resource.

Moving from BAS16, and following the successful trend in applying deep neural network in Semantic Parsing [13]

and Semantic Role Labelling tasks [14], where encoder-decoder recurrent architectures have proven particularly effective [15], [16], in this paper we propose a preliminary study of the application of a multi-layer Long-Short Term Memory (LSTM) network to parse robotic commands from the HuRIC resource. Moreover, we also aim at showing that using a multi-layer LSTM is a viable solution even in poor training condition, as HuRIC contains only 527 examples.

II. APPROACH

The objective of the system is to semantically parse transcriptions of vocal commands for a House Service robot contained in the HuRIC data set. The outputs are given in terms of *semantic frames*, which are conceptualisations of actions or more general events. Each frame is evoked in the text by a *lexical unit* and, besides its type, it also includes a set of *frame elements*, which represent the entities having a specific role in the situation described by the frame. The tasks involved aim at structuring the input sentence into actionable information, through: *i*) identifying the frame representing the user’s willingness as a label for the whole sentence (the Action Detection task, AD); *ii*) selecting the relevant text spans of frame elements in the sentence (the Argument Identification task, AI); *iii*) assigning a type to each frame element span (the Argument Classification task, AI).

Our approach builds upon the work in [17], where a two-layer LSTM is applied to the Spoken Question Answering domain, which presents some similarities with our problem. We thus propose and test two variants of a multi-layered LSTM architecture, one with two layers (2L), exactly as in [17], and one with three layers (3L). Both of them are fed with sequences of pre-trained word embeddings.

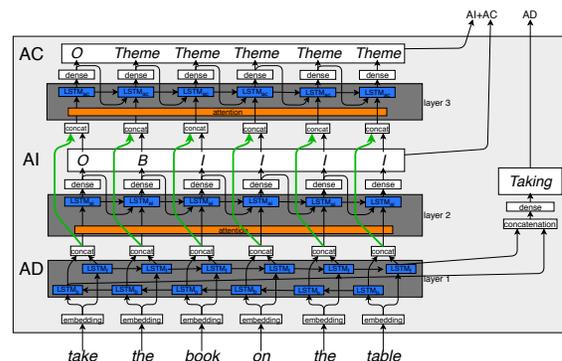


Fig. 1. The proposed 3-layer LSTM with self-attention network that addresses the three tasks of AD, AI and AC.

¹Knowledge Media Institute, The Open University, UK
{name.surname}@open.ac.uk

²Department of Computer Science, VU University Amsterdam, NL
ilaria.tiddi@vu.nl

³Istituto Superiore Mario Boella, Italy
giuseppe.rizzo@ismb.it

In the 2L setting, the first layer is a bidirectional LSTM which is used to perform the AD task, while the second, a LSTM decoder with label dependencies [18], is used to perform jointly the AI and AC tasks. The bi-directionality of the first layer should capture backward dependencies, which are crucial for frame classification. In the 3L configuration, instead, we introduce a third layer to divide the tasks of AI and AC. Every layer in the network thus solves one of the three semantic parsing tasks, as illustrated in Figure 1. The first two layers are still based on [17], but the second one only predicts the IOB labels (Inside, Outside or Beginning of an element, [19]) for the AI rather than including also information about frame element types. The third layer takes as input the outputs of the AI layer, i.e. the IOB labels, combined with the internal representation from the first layer through *highway connections* [20] (green connections in Figure 1), and outputs labels for the frame element types. For each layer we make use of a self-attention mechanism [21] that learns weights, which enable combination of word-level features. The loss function used to train the network corresponds to the sum of the cross-entropies (H_x) of each task, namely $loss = H_{AD} + H_{AI} + H_{AC}$.

III. EXPERIMENTAL EVALUATION

The Human-Robot Interaction Corpus (HuRIC) is used as data set for training, evaluation and comparison. It comprises 527 annotated sentences corresponding to vocal commands given to a robot in a house environment. FrameNet-style annotations are provided over each sentence, for a total of 16 frame types and an average of 33 examples for each frame. Hyper-parameter tuning, training and testing is performed through a 5-fold evaluation schema. Each network configuration has been tested with (ATT) and without (NO-ATT) self-attention layers. Word embeddings have been pre-trained with GloVe [22] over the Common Crawl resource [23].

TABLE I
F-MEASURE OF THE THREE SINGLE STAGES OF SEMANTIC PARSING (AD, AI, AC) AND OF THE WHOLE INTERPRETATION CHAIN.

	AD	AI	AC	Whole Chain
BAS16	94.67%	90.74%	94.93%	41.70%
3L-ATT	94.44%	94.73%	94.69%	43.67%
3L-NO-ATT	95.37%	94.90%	91.90%	41.92%
2L-ATT	96.29%	94.40%	92.30%	44.54%
2L-NO-ATT	94.44%	94.50%	92.45%	42.79%

A. Semantic Parsing

The first evaluation is performed on the semantic parsing tasks. In BAS16, the three steps (AD, AI, AC) are implemented by three independent blocks chained in a pipeline, and therefore they can be evaluated independently by using gold information. Instead, in our approach the information between the layers are implicit and gold values cannot be emulated. For this reason, the measures reported in Table I for the AI and AC are computed only on the portion of examples which are correctly classified by the preceding step. In this way, it is possible to estimate the performance

of the three different tasks independently, assuming gold information coming from the previous steps.

From Table I, we can see that the LSTM performs well in the AD task, with best results for the 2L-ATT setting. Every LSTM configuration outperforms BAS16 in the AI task. On the contrary, only the 3L-ATT configuration behaves similarly to BAS16 in the AC. The attention mechanism appears here to be crucial, as the scores drop significantly without it. This comes from the fact that the attention enables the third layer to better focus on the whole span identified by the AI task, with a softer alignment model. The AC is more complicated than the other two tasks, and the scarcity of examples seems to be a discriminant factor in these settings.

B. Whole Interpretation Chain

The second experimental setting aims at evaluating the whole interpretation chain, from the transcribed user utterance to the grounded robot command. Performances are evaluated on the fully-grounded robot commands. Each frame extracted from an utterance needs to be not only linguistically instantiated, but its arguments have also to be grounded in the environment. Each vocal command in HuRIC is paired with a semantic map representing the environment where the command has been given. A command counts as correctly grounded when all the frames in the related sentence are correctly instantiated, and all the frame entities are linked to the proper entities in the semantic map. Please refer to Section 4 of BAS16 for an in-depth definition of this task.

In Table I, under the Whole Chain column, we compare with the “Gold transcr., CoreNLP” run of BAS16, because our preliminary comparison wants to test the system over correct speech recognition transcriptions. Notice that our approach presents structural differences with BAS16, e.g. we do not need any morpho-syntactic parsing (CoreNLP), as we use word embeddings to represent words. While all the network configurations perform better than BAS16, the one that reaches highest results is 2L-ATT, which has also the best score on the AD task. This follows from the fact that the correct grounding of a command primarily depends on the correct interpretation of the frames contained in it. Misclassifying a frame but correctly recognising the frame elements, on the other hand, compromises the whole result.

IV. CONCLUSIONS

In this paper, we presented a preliminary study to semantically parse natural language robotic commands from the HuRIC resource using a multi-layer LSTM network with attention layers. For our initial tests, we compared with the work in [9], showing that a LSTM-based approach is a viable solution also in such a poor training condition (only 527 examples in HuRIC, ~ 33 examples per frame). Future works should cover the study of the attentions values to better explain the network behaviour. Such information could be used to adjust the interpretation process through some dialogue with the user. Finally, a mechanism to embed perceptual information in the LSTM framework should be investigated, as fostered in [9].

REFERENCES

- [1] J. Bos and T. Oka, "A spoken language interface with a mobile robot," *Artificial Life and Robotics*, vol. 11, no. 1, pp. 42–47, 2007.
- [2] G.-J. M. Kruijff, H. Zender, P. Jensfelt, and H. I. Christensen, "Situating dialogue and spatial organization: What, where... and why?" *International Journal of Advanced Robotic Systems*, vol. 4, no. 2, 2007.
- [3] S. Tellex, T. Kollar, S. Dickerson, M. Walter, A. Banerjee, S. Teller, and N. Roy, "Approaching the symbol grounding problem with probabilistic graphical models," *AI Magazine*, vol. 32, no. 4, pp. 64–76, 2011.
- [4] D. L. Chen and R. J. Mooney, "Learning to interpret natural language navigation instructions from observations," in *Proceedings of the 25th AAAI Conference on AI*, 2011, pp. 859–865.
- [5] J. Kim and R. J. Mooney, "Adapting discriminative reranking to grounded language learning," in *ACL (1)*. The Association for Computer Linguistics, 2013, pp. 218–227.
- [6] C. Matuszek, E. Herbst, L. S. Zettlemoyer, and D. Fox, "Learning to parse natural language commands to a robot control system," in *ISER*, ser. Springer Tracts in Advanced Robotics, J. P. Desai, G. Dudek, O. Khatib, and V. Kumar, Eds., vol. 88. Springer, 2012, pp. 403–415.
- [7] Y. Artzi and L. Zettlemoyer, "Weakly supervised learning of semantic parsers for mapping instructions to actions," *Transactions of the Association for Computational Linguistics*, vol. 1, no. 1, pp. 49–62, 2013.
- [8] J. Thomason, S. Zhang, R. Mooney, and P. Stone, "Learning to interpret natural language commands through human-robot dialog," in *Proceedings of the 24th International Conference on Artificial Intelligence (IJCAI)*, ser. IJCAI'15. AAAI Press, 2015, pp. 1923–1929.
- [9] E. Bastianelli, D. Croce, A. Vanzo, R. Basili, and D. Nardi, "A discriminative approach to grounded spoken language understanding in interactive robotics," in *Proceedings of the 2016 International Joint Conference on Artificial Intelligence (IJCAI)*, New York, USA, July 2016.
- [10] E. Bastianelli, G. Castellucci, D. Croce, L. Iocchi, R. Basili, and D. Nardi, "Huric: a human robot interaction corpus," in *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*. Reykjavik, Iceland: European Language Resources Association (ELRA), 2014.
- [11] C. J. Fillmore, "Frames and the semantics of understanding," *Quaderni di Semantica*, vol. 6, no. 2, pp. 222–254, 1985.
- [12] C. F. Baker, C. J. Fillmore, and J. B. Lowe, "The Berkeley FrameNet project," in *Proceedings of ACL and COLING*, ser. Association for Computational Linguistics, 1998, pp. 86–90.
- [13] R. Jia and P. Liang, "Data recombination for neural semantic parsing," in *ACL (1)*. The Association for Computer Linguistics, 2016.
- [14] C. N. dos Santos, B. Xiang, and B. Zhou, "Classifying relations by ranking with convolutional neural networks," in *ACL (1)*. The Association for Computer Linguistics, 2015, pp. 626–634.
- [15] J. Zhou and W. Xu, "End-to-end learning of semantic role labeling using recurrent neural networks," in *ACL (1)*. The Association for Computer Linguistics, 2015, pp. 1127–1137.
- [16] B. Yang and T. M. Mitchell, "A joint sequential and relational model for frame-semantic parsing," in *EMNLP*. Association for Computational Linguistics, 2017, pp. 1247–1256.
- [17] B. Liu and I. Lane, "Attention-based recurrent neural network models for joint intent detection and slot filling," in *INTERSPEECH*. ISCA, 2016, pp. 685–689.
- [18] Y. Dupont, M. Dinarelli, and I. Tellier, "Label-dependencies aware recurrent neural networks," *arXiv preprint arXiv:1706.01740*, 2017.
- [19] A. Ratnaparkhi, "Maximum entropy models for natural language ambiguity resolution," Ph.D. dissertation, University of Pennsylvania, 1998.
- [20] R. K. Srivastava, K. Greff, and J. Schmidhuber, "Highway networks," *arXiv preprint arXiv:1505.00387*, 2015.
- [21] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," *arXiv preprint arXiv:1409.0473*, 2014.
- [22] J. Pennington, R. Socher, and C. Manning, "Glove: Global vectors for word representation," in *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 2014, pp. 1532–1543.
- [23] (2012) Common crawl. [Online]. Available: <http://commoncrawl.org/>