Automatic Dialogue Flow Extraction and Guidance

Patrícia Ferreira

December 19, 2022

Automatic Dialogue Flow Extraction and Guidance

Patrícia Ferreira*

CISUC, Universidade de Coimbra
patriciaf@dei.uc.pt

Abstract. Nowadays, human agents are often replaced by conversational software agents, designed to communicate with humans through natural language, often based on Artificial Intelligence, namely Natural Language Processing (NLP) and Machine Learning (ML). This work will have as one of the main goals the improvement of communication between customers and human agents in a call-center to make this work more efficient or, not only in call-centers but in a common conversation between actors, suggesting appropriate responses, thus anticipating their interventions. You will start by identifying and annotating sets of dialogues, written in Portuguese. The guidance given can be supported by a history of interactions, where information is extracted and frequent dialog flows are discovered, allowing a representation of them to guide humans. The approach will be divided into three components: the Extraction component to process dialogues and use the information to describe interactions; the Representation component to discover the most frequent dialogue flows, represented by interaction graphs; and the Guidance component to guide the agent during a new dialogue.

Keywords: Natural Language Processing · Dialog Analysis · Dialog Information Extraction · Representation of dialog flows · Assisted guidance

1 Background and Related Work

A dialog is composed of utterances, which instantiate dialog acts (DAs), that is, abstract representations of intentions. There are several dialogue datasets, mainly for English [2], however, this work will focus on Portuguese, where public dialogue datasets are scarce.

There are several approaches for automatic classification of DAs (DAC) [1]. Most are based on supervised learning, with models trained in dialog datasets where the DAs are manually annotated [3]. Others use traditional classification[3][4]. However, since there may be a dependency between the current interaction and previous ones, DAs can be tackled as a sequence classification problem, with methods such as Hidden Markov Models (HMM)[5] or Conditional Random Fields (CRF) [6]. DAs and transition graphs between them are

* Supervised by Hugo Gonçalo Oliveira (hroliv@dei.uc.pt), Catarina Silva (catarina@dei.uc.pt) e Ana Alves (ana@dei.uc.pt)

useful representations of dialogs. When applied to large sets of dialogs, models trained will allow the discovery of different types of interactions and the most common dialog flows. These flows can even be discovered without annotating DAs, through unsupervised approaches (clustering) [8]. Furthermore, as real-time call monitoring, these flows will be useful to help interpret the dialog and support the participants. [9] [10].

2 Methodology

Overall, this work consists of researching, implementing and testing a solution that aims to improve communication between participants in a dialog by guiding their actions, which can be supported by previous interactions, where information should be automatically extracted and where frequent dialog flows. Besides task-oriented dialogues experiments will be extended to conversations between common users, such as, in a social network, where it will be possible to analyze communication trends.

The experiments will be conducted with data in Portuguese, which will be a differentiating factor from the state of the art. They will also be limited to written text, i.e., written conversations or transcripts of oral communication.

The data can be created following the Wizard-of-Oz (WOZ) [11], where a conversation takes place between two participants with different roles. One plays the role of an ordinary user who is assigned a certain task and must interact, using natural language, with another that will have access to information about the domain (e.g., a database) and will be able to provide appropriate answers.

Data can also result from available transcriptions of existing dialogues into Portuguese, such as CORAA [12]; from customer support services, such as conversations with telecom operators on Twitter; or from movie subtitles [13]. One last possibility will be the translation of English datasets (e.g., DailyDialogue [14], MultiWOz [15]) into Portuguese, from where existing annotation can be imported.

The data will be used in the development of a framework consisting of three components. The first will process real dialog transcripts and extract useful information from them to represent interactions, such as keywords, entities or actions. The extraction of some of these items may resort to an NLP pipeline [16]. The extracted information can be used to better describe the utterances, by classifying intentions and filling slots. However, the performance of these tasks is usually based on supervised learning, which implies data annotation. The extracted information can also be used to group similar utterances, using clustering. This process can also resort to Sentence Embedding techniques [17].

The second component will aim at discovering the most frequent dialog flows, represented by graphs, where the vertices represent speech classes or clusters, and the arcs represent transitions between them, with associated probabilities. In this component one can apply the classification of interactions into more generic classes or, if there is a lack of data to make the system less domain-dependent, perform a clustering that approximates these acts.

Finally, the guidance component will take advantage of past dialogs, represented according to previous components. In each interaction, previous interactions will be considered, while anticipating the next interaction. This component will analyze the ongoing dialog, using an NLP pipeline, and, whenever possible, map an expression and its context, represented by previous interactions, to expressions represented in the dialog flow graph. This mapping can be done through Semantic Textual Similarity [19] mechanisms, using techniques that consider only the words used and their relevance (e.g. TF-IDF based [20]). When this mapping is successful, the possible transitions from that expression will be collected and presented to the agent. Techniques used by Recommender Systems [21] to make recommendations given a context will be explored.

A final evaluation of the results of integrating the three components into the framework should be done, as the approaches explored will be evaluated on the data gathered and created using metrics for classification when annotations are produced, or metrics for clustering when they are not.

3 Objectives and Expected Results

The main objective of this work is to investigate and develop approaches to improve communication in a dialog, in Portuguese, supporting guidance of human agents, such as for example supporting a human in a call-center.

NLP techniques will be explored, focusing on dialog modeling [18], in order to, based on the history of interactions, identify the most common ones in each application domain, discover and interpret flows, and take advantage of the latter to guide interlocutors, who may thus anticipate their interventions. The dialog modeling will focus on dialog and intent classification acts as well as flow discovery in a scenario where there is no annotated data.

We aim at approaches, applicable to written dialogues, e.g. between users of a social network, or to transcriptions of task-oriented dialogues (e.g., call-center) to assist call-center operators in providing more efficient service.

We believe that the work will result in innovative approaches, and highlight the fact that, regardless of the possible adaptation to other languages, it will be focused on Portuguese. The work will be divided into some specific objectives, namely: identify and create sets of dialogues, in Portuguese; study, develop and experiment approaches for extracting structured dialogue information from the various interactions; for representing interactions and dialog flows extracted from those interactions; for guiding the human by exploiting the knowledge extracted from dialogues, dialogue type and interactions, and common flows.

To achieve the defined goals, the following tasks were established: 1. To deepen the study of the state of the art; 2. Definition of the data to be used; 3. Exploring approaches for the three components; 6. Proposed framework encompassing the approaches explored; 7. Testing and final evaluation; 8. Writing of the thesis and dissemination papers. The approaches resulting from task 3 will be evaluated independently during their respective tasks, but a final evaluation of the results of their integration into the framework will be required. The experiences are regularly described in writing in the doctoral thesis. We further believe that from tasks 2 to 7 will result contributions relevant to write papers.

Acknowledgements

This work was financially supported by the project FLOWANCE (POCI-01-0247-FEDER-047022), cofinanced by FEDER , through PT2020 , and by COMPETE 2020 ; and by national funds through FCT, within the scope of the project CISUC (UID/CEC/00326/2020) and by European Social Fund, through the Regional Operational Program Centro 2020.

References

1. Ruizhe Li, Chenghua Lin, Matthew Collinson, Xiao Li, and Guanyi Chen. A dual attention hierarchical recurrent neural network for dialogue act classification. arXiv preprint arXiv:1810.09154, 2018
2. Hugo Gonçalo Oliveira, Patrícia Ferreira, Daniel Martins, Catarina Silva, and Ana Alves. A brief survey on textual dialogue corpora. In Proceedings of the 13th International Conference on Language Resources and Evaluation (LREC 2022), page In press, 2022.
3. Srinivas Bangalore, Giuseppe Di Fabbrizio, and Amanda Stent. Learning the structure of task driven human–human dialogs. IEEE Transactions on Audio, Speech, and Language Processing, 16(7):1249–1259, 2008.
4. Konstantina Christakopoulou, Filip Radlinski, and Katja Hofmann. Towards conversational recommender systems. In Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining, pages 815–824, 2016.
5. [5] R Eddy Sean. What is a hidden markov model. Nat Biotechnol, 22(10):1315–1316, 2004.
6.] Harshit Kumar, Arvind Agarwal, Riddhiman Dasgupta, and Sachindra Joshi. Dialogue act sequence labeling using hierarchical encoder with crf. In Proceedings of the aaai conference on artificial intelligence, volume 32, 2018.
7. Fei Sun, Jun Liu, Jian Wu, Changhua Pei, Xiao Lin, Wenwu Ou, and Peng Jiang. Bert4rec: Sequential recommendation with bidirectional encoder representations from transformer. In Proceedings of the 28th ACM international conference on information and knowledge management, pages 1441–1
8. Ronald M Kaplan, Jill Burstein, Mary Harper, and Gerald Penn. Human language technologies: the 2010 annual conference of the north american chapter of the association for computational linguistics. In Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics, 2010.
9. Su Nam Kim, Lawrence Cavedon, and Timothy Baldwin. 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, 2010.
10. Michael McTear. The role of spoken dialogue in user–environment interaction. Human-Centric Interfaces for Ambient Intelligence, pages 225–254, 2010.
11. Anders Green, Helge Huttenrauch, and K Severinson Eklundh. Applying the wizard-of-oz framework to cooperative service discovery and configuration. In RO-MAN 2004. 13th IEEE International Workshop on Robot and Human Interactive Communication (IEEE Catalog No. 04TH8759), pages 575–580. IEEE, 2004.

12. Arnaldo Candido Junior, Edresson Casanova, Anderson Soares, Frederico Santos de Oliveira, Lucas Oliveira, Ricardo Corso Fernandes Junior, Daniel Peixoto Pinto da Silva, Fernando Gorgulho Fayet, Bruno Baldissera Carlotto, Lucas Rafael Stefanel Gris, et al. Coraa: a large corpus of spontaneous and prepared speech manually validated for speech recognition in brazilian portuguese. arXiv preprint arXiv:2110.15731, 2021.
13. Pierre Lison and J¨org Tiedemann. Opensubtitles2016: Extracting large parallel corpora from movie and tv subtitles. European Language Resources Association, 2016.
14. Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. Dailydialog: A manually labelled multi-turn dialogue dataset. arXiv preprint arXiv:1710.03957, 2017.
15. Pawe l Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Inigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gasic. Multiwoz–a large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling. arXiv preprint arXiv:1810.00278, 2018.
16. Xavier Schmitt, Sylvain Kubler, J´er´emy Robert, Mike Papadakis, and Yves Le-Traon. A replicable comparison study of ner software: Stanfordnlp, nltk, opennlp, spacy, gate. In 2019 Sixth International Conference on Social Networks Analysis, Management and Security (SNAMS), pages 338–343. IEEE, 2019.
17. Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert networks. arXiv preprint arXiv:1908.10
18. Andreas Stolcke, Klaus Ries, Noah Coccaro, Elizabeth Shriberg, Rebecca Bates, Daniel Jurafsky, Paul Taylor, Rachel Martin, Carol Van Ess-Dykema, and Marie Meteer. Dialogue act modeling for automatic tagging and recognition of conversational speech. Computational linguistics, 26(3):339–373, 2000.
19. Daniel Cer, and Diab M. Agirre Semeval-2017 task 1: Semantic textual similarity-multilingual and cross-lingual focused evaluation, arXiv preprint arXiv:1708.00055, 2017.
20. Akiko Aizawa. An information-theoretic perspective of tf–idf measures. Information Processing Management 39.1 (2003): 45-65.
21. Paul Resnick, and Varian Hal R. Recommender systems. Communications of the ACM 40.3 (1997): 56-58.