Towards Learning Through Open-Domain Dialog

Eugénio Ribeiro[1,2], **Ricardo Ribeiro**[1,3], and **David Martins de Matos**[1,2]
INESC-ID Lisboa, Portugal
Instituto Superior Técnico, Universidade de Lisboa, Portugal
Instituto Universitário de Lisboa (ISCTE-IUL), Portugal
`eugenio.ribeiro@inesc-id.pt`

Abstract

The development of artificial agents able to learn through dialog without domain restrictions has the potential to allow machines to learn how to perform tasks in a similar manner to humans and change how we relate to them. However, research in this area is practically nonexistent. In this paper, we identify the modifications required for a dialog system to be able to learn from the dialog and propose generic approaches that can be used to implement those modifications. More specifically, we discuss how knowledge can be extracted from the dialog, used to update the agent's semantic network, and grounded in action and observation. This way, we hope to raise awareness for this subject, so that it can become a focus of research in the future.

1 Introduction

The development of artificial agents able to extract knowledge from dialog without domain restrictions and use it to improve their capabilities and adapt to different situations has the potential to change how machines are built to perform certain tasks, as well as how we relate to them. This is in line with the roadmap towards machine intelligence proposed by Mikolov et al. (2016), which defines communication and learning as two of the fundamental properties that intelligent machines should have. In fact, by combining both with an appropriate body, an artificial agent would be able to naturally communicate with humans and learn in the same way they do, acquiring new knowledge and competences by connecting what is learned through dialog with what is observed in the world. Theoretically, such an agent would be able to learn virtually anything and adapt to new situations, removing the limitations that hand-programmed machines have due to the fact that programmers are not able to predict every possible situation a priori.

An agent able to learn through dialog is a dialog system at its core. However, looking into the research on dialog systems, we can see that it has mainly focused on the development of two different kinds of system. On the one hand, there are task-oriented dialog systems, which focus on the acquisition of the information required to perform a specific task required by the user (e.g. Young, 2000; Allen et al., 2001; Wen et al., 2017; Yang et al., 2020). Thus, their conversation capabilities are limited to a single or a small set of domains and restricted by a set of predefined actions that they are able to perform. On the other hand, there are conversational agents that have no restrictions in terms of domain, but only focus on keeping the user engaged in the conversation by generating appropriate responses to user utterances, even if they are not actually able to understand or extract any knowledge from them (e.g. Weizenbaum, 1966; Lowe et al., 2017). Thus, they are only developed for research or short-term entertainment purposes. Although some recent studies (e.g. Serban et al., 2017; Cuayáhuitl et al., 2019) have explored the use of reinforcement learning approaches to incrementally improve the dialog policy and generate better responses to user utterances, neither task-oriented dialog systems nor conversational agents are able to extract knowledge from the dialog and use it to improve their capabilities. Research on this subject is limited to grounding problems in simple domains (e.g. Yu et al., 2017; Thomason et al., 2017), in which the focus is typically not on the linguistic part of the dialog, but rather on the ability to identify observations of certain concepts, or to map concepts into actions.

The aim of this paper is to encourage further research towards the development of artificial agents able to learn through dialog without domain restrictions. We do that in two ways. First, by identifying the aspects of a generic dialog system that need to
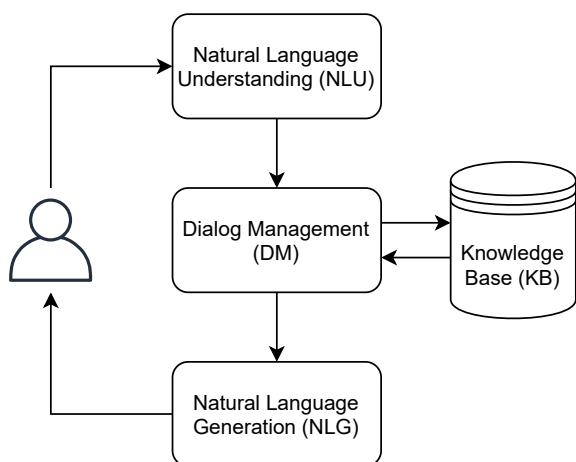
Figure 1: The generic architecture of a dialog system.

be modified in order to allow it to learn through dialog (Section 2). Second, by proposing generic approaches that can be applied to achieve the required adaptations. More specifically, we discuss how to extract conceptual knowledge from the dialog (Section 3), how to use it to update the agent's semantic network (Section 4), and how to ground it in observation and actions (Section 5).

2 Overview

Dialog systems or conversational agents are programs that interact with their users using some form of natural language. Such systems traverse most areas of Natural Language Processing (NLP), because they must be able to interpret user utterances, keep track of the dialog, select the best dialog actions in each context, and generate appropriate responses. Although recent open-domain conversational agents are based on end-to-end models (e.g. Lowe et al., 2017; Serban et al., 2017) and even task-oriented dialog systems can be developed in an end-to-end fashion that includes API calls (e.g. Byrne et al., 2020; Yang et al., 2020), most dialog systems follow the flow depicted in Figure 1, even if the division into components is implicit. Thus, in the remainder of the paper, we will refer to the components of the modular architecture, as they provide points for anchoring the modifications required to allow a dialog system to learn through dialog. Still, we believe that our considerations can be generalized to end-to-end systems as well.

Not considering the additional components required to enable spoken interaction, it has long been established that dialog systems can be split into three major components (Young, 2000; Allen et al.,

2001): Natural Language Understanding (NLU), Dialog Management (DM), and Natural Language Generation (NLG). The first refers to the interpretation of user utterances, by the identifying their domain, intent, and content. The second involves keeping track of the dialog state, communicating with information sources, such as Knowledge Bases (KBs), and selecting the best dialog actions to take according to the context of the dialog and the system's policy. The last refers to the generation of coherent sentences that transmit the information defined by the selected dialog actions and their parameterization.

The NLG ability required for an agent to learn through dialog does not differ from that required in a scenario that does not involve learning. Thus, the modifications mostly concern NLU and DM. In a task-oriented dialog system, NLU consists of the identification of the domain of the utterance, followed by or together with the identification of the domain-specific intent of the user, which includes slots to be filled, defining a frame that maps into a task that can be performed by the system. The DM component is then responsible for keeping track of the values of the slots which have already been filled and generating appropriate dialog actions that target the acquisition of the values of the remaining slots. When all the slots are filled, the execution of the corresponding task is triggered, which may involve querying a KB to obtain the information required by the user, or calling an external API to execute a command.

On the other hand, in the context of an agent that learns through open-domain dialog, we cannot define sets of slots to be filled, because there are no restrictions in terms of task nor domain. In fact, the concepts and properties that originate those slots may themselves be learned over time through dialog and the domains are fuzzy and inferred from the connections between the concepts known by the agent. Thus, in this scenario, NLU has to focus on extracting generic knowledge from the utterances, regarding the existence of concepts and the relations between them. Intent recognition is still important, but from a generic perspective which provides cues for the kinds of knowledge present in each utterance. The identification of such intents may help in the extraction process and allows the DM component to find relations between multiple utterances in the dialog.

Additionally, in order for the agent to learn, the

KB can no longer be just an information source, as it must be updatable as well. In fact, since the structure of the KB is itself learned over time based on the concepts and relations extracted from the dialog, it can be generalized as a semantic network (Sowa, 1991). The knowledge present in this semantic network can be grounded in action and observation by identifying connections between the concepts and the agent's sensors and actuators. The references to those concepts in the dialog can then be paired with the corresponding observations or actions and used to improve the agent's abilities.

Finally, if we want the agent to be able to guide the dialog towards the acquisition of further knowledge, following an active learning strategy, then the dialog policy must be adapted to also consider prompting for additional information or changing domains. However, it must also consider the context of the dialog and only apply these strategies in appropriate situations.

In the following sections, we propose some approaches for extracting conceptual knowledge from the dialog, use it to update the agent's semantic network, and ground it in action and observation.

3 Knowledge Extraction

In order to learn through open-domain dialog, an agent must be able to extract knowledge from it. Open-Domain Information Extraction (OpenIE) systems (Banko et al., 2007) are able to extract relational tuples from text without domain or task restrictions. Thus, they seem appropriate for extracting open-domain knowledge from dialog. However, they usually focus on declarative texts, such as news articles or web data. On the other hand, dialogs include non-declarative utterances, such as questions, and extensive inter-utterance dependencies in the form of answers, corrections, references, among others. Furthermore, OpenIE approaches are typically applied to large collections of documents, in which the same information appears multiple times. Thus, they focus on precision instead of recall. On the other hand, in dialogs, each piece of information is transmitted a reduced amount of times, or even only once. Thus, although OpenIE systems can serve as base for the extraction of knowledge from dialog, if an existing OpenIE system (e.g. Cui et al., 2018; Stanovsky et al., 2018; Kolluru et al., 2020) is applied directly to dialog utterances, it is bound both to extract knowledge that is not factual and to miss important extractions.

The number of missed extractions can be reduced by performing coreference resolution (Sukthanker et al., 2020) before applying the OpenIE approach. However, this does not solve inter-utterance dependencies based on function nor avoids non-factual extractions. As discussed in Section 2, the generic intention behind the utterances can be used to provide cues for these situations. The ISO 24617-2 standard for dialog act annotation (Bunt et al., 2017) defines a hierarchy of general-purpose communicative functions that can be automatically identified to a certain extent, even though the amount of annotated dialogs available is reduced (Ribeiro et al., 2022). This hierarchy includes a branch for information-transfer functions that can be used to guide the knowledge extraction process. For instance, the OpenIE approach can be applied directly to utterances with an inform function. On the other hand, although different kinds of questions are also able to provide knowledge regarding the existence of concepts or multiple alternatives, their function is to obtain some kind of information. Thus, depending on their kind, questions may include information that is incomplete, uncertain, or even incorrect. To address this problem, questions should be interpreted using approaches similar to those used in the Question Answering (QA) area (Diefenbach et al., 2018). Additionally, utterances with a responsive function, such as an answer or a correction, can be paired with the utterances they relate to by the DM component, leading to the extraction of further knowledge.

Finally, although the relational tuples extracted by OpenIE systems are easily interpretable by humans, in order to be interpreted by a learning agent, they should be extended with information regarding the temporal validity of the relations, and additional knowledge regarding the nature of arguments. In this context, it may be important to include information regarding the generic semantic roles played by the arguments (Petukhova and Bunt, 2008), or even to attempt to identify extractions that evoke similar semantic frames (QasemiZadeh et al., 2019).

4 Learning

In the previous section, we discussed means to extract knowledge from a dialog. However, in order to learn, the extracted knowledge on its own is not enough. First of all, a learning agent must be able to represent and store what it knows, that is, it must have memory. The conceptual knowledge that can

be extracted from dialog consists mainly of concepts, relations between them, and possible restrictions on the scope of their validity. By combining multiple of these extractions, a learning agent can create a semantic network (Sowa, 1991) that represents its conceptual memory and, thus, contains the concepts that the agent is aware of. Furthermore, considering that the concepts are identified by their name, this semantic network also defines the vocabulary that the agent can use to talk about its knowledge and which typically maps into a defined semantics known by its conversational partners. Thus, the semantic network is, more specifically, an ontology (Staab and Studer, 2009).

Ontologies were originally built by and shared among humans to define a formal context and avoid misunderstandings in communication. However, they can also be used as KBs queryable by automatic systems, including conversational agents, to obtain knowledge regarding the covered domains. Furthermore, an ontology can also be automatically updated in an incremental fashion, by linking (Shen et al., 2014) the concepts and relations referred to in the dialog to those present in the ontology, creating new ones if necessary.

Conflicts may arise during the learning process. Most of these conflicts are easy to identify, because the newly obtained knowledge is incompatible with that present in the KB. On the other hand, solving them is not as straightforward. Some conflicts arise due to the existence of ambiguous concepts and can be solved through context disambiguation processes. Other conflicts arise due to misunderstandings or misinformation. Trust- or confidence-based conflict solving strategies can be applied in such situations. However, considering that the agent is in an interactive context, the dialog itself can be used as a tool to solve conflicts, by prompting the conversational partners for the solution.

5   Knowledge Grounding

The processes described in the previous sections allow a learning agent to update its semantic network with conceptual knowledge acquired through dialog. However, that knowledge still has no connection to what is observed in the world. Although that might be enough for an agent whose task is purely dialog-based, it is not for an agent trying to get better at a task involving interaction with its environment through other means. In such scenarios, the knowledge present in the semantic network

has to be grounded in action and observation. For that to happen, there has to be a mapping between the agent's sensors and the primitive concepts they are able to observe, as well as between the agent's actuators and the representation of the primitive actions they are able to perform. These mappings can be added directly when a new sensor or actuator is added to the agent, or learned through the dialog. Using these connections, the agent can learn how to perform compound actions, as well as how to identify observations of derivative concepts, based on the compositional and hierarchical relations that it learns through dialog.

In order to improve its ability to recognize concept observations, the agent has to create conceptual models for the corresponding concepts, based on the features provided by the sensors related to those concepts. The models can then be improved over time using an Incremental Learning (IL) approach (Gepperth and Hammer, 2016) together with the labeled observations obtained by combining references to observable concepts in the dialog with the information provided by the sensors. This is the typical approach used in incremental natural language grounding research (e.g. Cakmak et al., 2010; Yu et al., 2017; Thomason et al., 2017). However, in those scenarios, there is a predefined set of concepts to be grounded and the whole dialog is focused on that objective. On the other hand, in the context of an agent learning through open-domain dialog, the set of observable concepts is also learned over time, not all utterances refer to observable concepts, and there may be references to past or future observations.

Still in the context of natural language grounding, the agent can also rely on the dialog to adopt active learning strategies, by prompting for the concepts which are being observed at a given time (e.g. Cakmak et al., 2010; Thomason et al., 2017), or requesting a demonstration of or feedback on a given action (e.g. Cakmak and Thomaz, 2012).

6   Conclusions

In this paper, we have raised awareness for the lack of research on artificial agents able to learn through open-domain dialog, identified the modifications required for a dialog system to be able to learn from the dialog, and proposed generic approaches that can be used to implement those modifications. This way, we hope that this subject can become a focus of research in the future.

Acknowledgments

References

James Allen, George Ferguson, and Amanda Stent. 2001. An Architecture For More Realistic Conversational Systems. In *IUI*, pages 1–8.

Michele Banko, Michael J. Cafarella, Stephen Soderland, Matt Broadhead, and Oren Etzioni. 2007. Open Information Extraction from the Web. In *IJCAI*, pages 2670–2676.

Harry Bunt, Volha Petukhova, David Traum, and Jan Alexandersson. 2017. Dialogue Act Annotation with the ISO 24617-2 Standard. In *Multimodal Interaction with W3C Standards*, pages 109–135. Springer.

Bill Byrne, Karthik Krishnamoorthi, Saravanan Ganesh, and Mihir Sanjay Kale. 2020. TicketTalk: Toward Human-Level Performance with End-to-End, Transaction-Based Dialog Systems. *Computing Research Repository*, arXiv:2012.12458.

Maya Cakmak, Crystal Chao, and Andrea Lockerd Thomaz. 2010. Designing Interactions for Robot Active Learners. *IEEE Transactions on Autonomous Mental Development*, 2(2):108–118.

Maya Cakmak and Andrea Lockerd Thomaz. 2012. Designing Robot Learners that Ask Good Questions. In *HRI*, pages 17–24.

Heriberto Cuayáhuitl, Donghyeon Lee, Seonghan Ryu, Yongjin Cho, Sungja Choi, Satish Indurthi, Seunghak Yu, Hyungtak Choi, Inchul Hwang, and Jihie Kim. 2019. Ensemble-Based Deep Reinforcement Learning for Chatbots. *Neurocomputing*, 366:118–130.

Lei Cui, Furu Wei, and Ming Zhou. 2018. Neural Open Information Extraction. In *ACL*, volume 2: Short Papers, pages 407–413.

Dennis Diefenbach, Vanessa Lopez, Kamal Singh, and Pierre Maret. 2018. Core Techniques of Question Answering Systems over Knowledge Bases: a Survey. *Knowledge and Information Systems*, 55(3):529–569.

Alexander Gepperth and Barbara Hammer. 2016. Incremental Learning Algorithms and Applications. In *ESANN*, pages 357–368.

Keshav Kolluru, Vaibhav Adlakha, Samarth Aggarwal, Mausam, and Soumen Chakrabarti. 2020. OpenIE6: Iterative Grid Labeling and Coordination Analysis for Open Information Extraction. In *EMNLP*, pages 3748–3761.

Ryan Thomas Lowe, Nissan Pow, Iulian V. Serban, Laurent Charlin, Chia-Wei Liu, and Joelle Pineau. 2017. Training End-to-End Dialogue Systems with the Ubuntu Dialogue Corpus. *Dialogue & Discourse*, 8(1):31–65.

Tomas Mikolov, Armand Joulin, and Marco Baroni. 2016. A Roadmap Towards Machine Intelligence. In *CICLing*, pages 29–61.

Volha Petukhova and Harry Bunt. 2008. LIRICS Semantic Role Annotation: Design and Evaluation of a Set of Data Categories. In *LREC*, pages 39–45.

Behrang QasemiZadeh, Miriam R. L. Petruck, Regina Stodden, Laura Kallmeyer, and Marie Candito. 2019. SemEval-2019 Task 2: Unsupervised Lexical Frame Induction. In *SemEval*, pages 16–30.

Eugénio Ribeiro, Ricardo Ribeiro, and David Martins de Matos. 2022. Automatic Recognition of the General-Purpose Communicative Functions Defined by the ISO 24617-2 Standard for Dialog Act Annotation. *Journal of Artificial Intelligence Research*, 73:397–436.

Iulian V. Serban, Chinnadhurai Sankar, Mathieu Germain, Saizheng Zhang, Zhouhan Lin, Sandeep Subramanian, Taesup Kim, Michael Pieper, Sarath Chandar, Nan Rosemary Ke, Sai Rajeshwar, Alexandre de Brebisson, Jose M. R. Sotelo, Dendi Suhubdy, Vincent Michalski, Alexandre Nguyen, Joelle Pineau, and Yoshua Bengio. 2017. A Deep Reinforcement Learning Chatbot. *Computing Research Repository*, arXiv:1709.02349.

Wei Shen, Jianyong Wang, and Jiawei Han. 2014. Entity Linking with a Knowledge Base: Issues, Techniques, and Solutions. *IEEE Transactions on Knowledge and Data Engineering*, 27(2):443–460.

John F. Sowa, editor. 1991. *Principles of Semantic Networks: Explorations in the Representation of Knowledge*. The Morgan Kaufmann Series in Representation and Reasoning. Morgan Kaufmann.

Steffen Staab and Rudi Studer. 2009. *Handbook on Ontologies*. International Handbooks on Information Systems. Springer.

Gabriel Stanovsky, Julian Michael, Luke Zettlemoyer, and Ido Dagan. 2018. Supervised Open Information Extraction. In *NAACL-HLT*, volume 1: Long Papers, pages 885–895.

Rhea Sukthanker, Soujanya Poria, Erik Cambria, and Ramkumar Thirunavukarasu. 2020. Anaphora and Coreference Resolution: A Review. *Information Fusion*, 59:139–162.

Jesse Thomason, Aishwarya Padmakumar, Jivko Sinapov, Justin Hart, Peter Stone, and Raymond J. Mooney. 2017. Opportunistic Active Learning for Grounding Natural Language Descriptions. In *CoRL*, volume 78 of *Proceedings of Machine Learning Research*, pages 67–76.

Joseph Weizenbaum. 1966. ELIZA – A Computer Program for the Study of Natural Language Communication between Man and Machine. *Communications of the ACM*, 9(1):36–45.

Tsung-Hsien Wen, David Vandyke, Nikola Mrkšić, Milica Gasic, Lina M. Rojas Barahona, Pei-Hao Su, Stefan Ultes, and Steve Young. 2017. A Network-Based End-to-End Trainable Task-Oriented Dialogue System. In *EACL*, volume 1, pages 438–449.

Yunyi Yang, Yunhao Li, and Xiaojun Quan. 2020. UBAR: Towards Fully End-to-End Task-Oriented Dialog Systems with GPT-2. *Computing Research Repository*, arXiv:2012.03539.

Steve Young. 2000. Probabilistic Methods in Spoken Dialogue Systems. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 358(1769):1389–1402.

Yanchao Yu, Arash Eshghi, and Oliver Lemon. 2017. Learning how to Learn: An Adaptive Dialogue Agent for Incrementally Learning Visually Grounded Word Meanings. In *Workshop on Language Grounding for Robotics*, pages 10–19.