

Towards Understanding Lifelong Learning for Dialogue Systems

Mark Cieliebak and Olivier Galibert and Jan Deriu

Abstract Lifelong learning is the ability of a software system to adapt to new situations during its lifetime. We explore how this paradigm can be applied to dialogue systems, how it might be implemented, and how we can evaluate the lifelong learning progress.

1 Introduction

Chatbots, dialogue systems, conversational user interfaces - the names may differ, but the basic idea is the same: "intelligent" computer systems that can interact with humans in natural language. These systems have become more and more popular in the past years, and there is an increasing interest in spoken and written dialogue systems in research and industry. Prominent examples include automatic customer support agents, smart home devices such as Amazon Alexa or Apple's Siri, and in-car operating systems. While implementing a successful and reliable dialogue systems is already a challenge, "lifelong learning" even adds an additional twist: the system should be able to adapt to new situations during its lifetime. More precisely, the dialogue system learns to handle new situations by interacting with its environment (e.g. asking a domain expert, scraping the web), instead of being retrained by a machine learning expert. For instance, a chatbot for travel advice might be confronted with a new location that is not yet in its knowledge base. One strategy to deal with this situation could be to ask the user to give additional information (e.g.

Mark Cieliebak
Zurich University of Applied Sciences (ZHAW), Winterthur, Switzerland, e-mail: ciel@zhaw.ch

Olivier Galibert
Laboratoire national de mtrologie et d'essais, Trappes, France, e-mail: Olivier.Galibert@lne.fr

Jan Deriu
Zurich University of Applied Sciences (ZHAW), Winterthur, Switzerland, e-mail: deri@zhaw.ch

in which country, GPS coordinates etc.), then explore the web to find information about the location (e.g. databases, Wikipedia or travel reviews), and finally analyze, structure, and integrate the information into the chatbots' knowledge base.

In this paper, we attempt to make a step forward towards understanding what lifelong learning in the context of dialogue systems means. In order to achieve this, we first briefly introduce both concepts independently and discuss typical settings and applications. Then we describe the the impact of applying lifelong learning to dialogue system (in Section 4). Finally, we turn to the important question how we could measure the success (or failure) of lifelong learning in the context of dialogue systems (Section 5).

2 What is a Dialogue System?

In the following, we introduce the concept of a dialogue system. *A dialogue system allows the user to converse with a computer system using natural language.* Such systems can be applied to a variety of tasks, e.g.:

- *Virtual Assistants*, which are developed to aid its users in every-day tasks, such as scheduling appointments. They usually operate on predefined actions, which can be triggered by voice command.
- *Interaction with Information Systems*, by asking questions or finding a piece of information (e.g. the most suitable hotel in town).
- *Training environments*, where the dialogue systems are developed to train students in the interaction with medical patients or train military personnel in questioning a witness.
- *Answering Questions*, where the dialogue system can answer specific questions of a user. These might be factoid questions or more complex questions.

Dialogue systems usually structure dialogues in *turns*, each turn is defined by one or more *utterances* from one speaker. Two consecutive turns between two different speakers is called an *exchange*. Multiple exchanges constitute a *dialogue*. Another correlated view, is to interpret each turn or each utterance as an action (more on this later). The main component of a dialogue system is the dialogue strategy, which defines the content of the next utterance and thus the behaviour of the dialogue system. There are many different approaches to design a dialogue strategy, which are partly dictated by the application of the dialogue system. However, there are three broad classes of dialogue systems, which we encounter in the literature: task-oriented systems, conversational agents and interactive question answering systems¹. We identified the following characteristic features, which help differentiate between the three different classes: is the system developed to solve a task, does the dialogue follow a structure, is the domain restricted or is it open domain, does the dialogue span over

¹ In recent literature, the distinction is made only between the first two classes of dialogue systems [1, 2, 3]. However, interactive question answering systems cannot be completely placed in either of the two categories.

multiple turns, are the dialogues rather long or efficient, who takes the initiative, and what is the interface used (text, speech, multi-modal).

- *Task-oriented systems* are developed to help the user solve a specific task as efficiently as possible. The dialogues are characterized by following a clearly defined structure, which is derived from the domain. The dialogues follow mixed initiative: both, the user and the system can take the lead. Usually, the systems found in the literature are built for speech input and output. However, task-oriented systems in the domain of assistance are built on multi-modal input and output.
- *Conversational agents* display a more unstructured conversation, as their purpose is to have open-domain dialogues with no specific task to solve. Most of these systems are built to emulate social interactions and thus longer dialogues are desired.
- *Question Answering (QA) systems* are built for the specific task of answering questions. The dialogues are not defined by a structure as with task-oriented systems, however they mostly follow the question and answer style pattern. QA systems may be built for a specific domain, but also be tilted towards more open domain questions. Usually, the domain is dictated by the underlying data, e.g. knowledge bases or text snippets from forums. Traditional QA systems work on a single-turn interaction, however, there exist systems that allow multiple turns to cover follow-up questions. The initiative is mostly done by the user who asks questions.

3 What is Lifelong Learning?

In the most abstract way, *Lifelong Learning (LL)* is the ability of a system to use past experiences to adapt to future challenge. There exist various definitions of LL in the literature, for instance in ???. For the purpose of this paper, we exploit the definition of LL from Chen and Liu [4], which we summarize in the following:

Lifelong learning is a continuous learning process. Given that the learner has learned N tasks. When faced with the $(N + 1)$ th task the learner leverages past knowledge to help learn the new task. The goal is to optimize on both the new task and the previous tasks. The three components are: continuous learning, knowledge accumulation and maintenance and leverage past knowledge to learn new tasks. There are some additional considerations to be made considering the above definition.

- The learner learns new tasks continuously, however, in contrast to transfer-learning, the learner improves or at least does not deteriorate its performance on the old tasks. Ideally, by learning new tasks, the performance on the previous tasks improves.
- The learner is not restricted to a certain task or domain. On the contrary, the learner is encouraged to learn different types of tasks (e.g. sentiment analysis, named entity recognition, etc) and domains.

- Ideally, the learner is *self-motivated* and able to find its own learning tasks and data by interacting with the environment.

Note that this definition is strongly focused on "knowledge improvement", whereas in the setting of LL for dialogue systems, there are also other goals, as we discuss in the following section. In addition, we would like to mention that the concept of a "task" may cover situations of varying complexity, ranging from single new instances (e.g. a new person in a face recognition system) up to new domains (e.g. switching from cooking to car tuning for a questions answering systems). Finally, note that several other terms have been coined and used for very similar learning paradigms of systems that improve over time, such as continuous learning [5, 6], meta-learning [7], active learning [8], or transfer learning [9]. For a more elaborate introduction of LL, see the recent book by Chen and Liu [4], which gives a good overview of LL in general and describes applications in various fields.

4 Lifelong Learning for Dialogue Systems

While the definition of lifelong learning given by [4] is very general, we attempt not to apply the definition to dialogue systems. These systems allow its users to converse with a computer via natural language. This implies a high level of interactivity. Thus, the focus of applying LL to dialogue systems should lie in the interactive nature of the dialogue. Furthermore, LL describes the capability of the dialogue system to learn to handle new situations throughout its deployment, i.e. without being re-trained by a machine-learning expert. Ideally, the learning takes place in a self-driven and autonomous manner. This does not exclude (it rather encourages) the assistance of a "domain expert", i.e. a type of user who takes the role of a teacher.

We assume that the dialogue system is an agent that interacts with its environment. The environment includes humans as well as having access to structured and unstructured knowledge sources (e.g. knowledge bases, Wikipedia, Twitter). When faced with a new situation, the dialogue system has to learn how to handle this new situation. This does not necessarily means that the dialogue system directly adapts to the new situations. Rather, through interaction with its environment, the dialogue system learns to handle the situations over time.

There are various aspects to a dialogue system which can be improved over time:

- *Language Understanding*: Here, the dialogue systems' capability of parsing the user input is the focal point. This is the case, for instance, when new request types occur over time, for instance if a system was only faced with simple factual questions until now, and the system is suddenly confronted with complex questions.
- *Dialogue Behaviour* is concerned with the "soft" quality factors of a dialogue, such as human-likeness, appropriateness of responses, efficiency of reaching a goal, engaging utterances etc. Typically, the DS asks after a user interaction for feedback, which is then used to improve the behaviour over time. Thus, the DS

leverages past experiences to improve its future behaviour. Note that in this case, there are no external catalysts that trigger LL, but there is an "intrinsic motivation" of the system itself.

- *Knowledge Induction* is concerned with accumulating more information. This means adapting the knowledge base (KB) with new or updated knowledge, which can be factual knowledge or unstructured. Here, new situations are in the context of handling new entities and relations which are not in the current knowledge base.
- *Capability Improvement* is concerned with extending the functionality of the DS. This can range from domain adaption (e.g. moving from asian recipes to the pasta domain) up to integrating new skills (e.g. reporting weather forecasts for a personal assistant)

In each case the dialogue system needs to improve its aspects over time. Each time it is faced with a new situation one or more of the aforementioned aspects need to be adapted. In the context of dialogue systems, this adoption can be done by means of interacting with a "domain" expert. More precisely, the goal is to remove the need to rely on a dialogue system expert who would retrain the different components of the dialogue system. Rather the domain-expert teaches the dialogue system how to handle a new situation through interaction. Note that in some cases the system may be able to learn how to handle the situation autonomously, especially in the case where it can aggregate data from some external sources. Thus, a LL enhanced dialogue system is able to learn to adapt to new situations by interacting with its environment and not by means of retraining components.

5 Evaluation of Lifelong Learning for Dialogue Systems

The above definition of LL for dialogue systems sets a strong focus on learning to handle new situations by interacting with its environment. Thus, the LL component of the dialogue system needs to be trained and evaluated with this in mind. More precisely, the interaction with an environment lies at the centre of the training and evaluation. The environment should enable the interaction the dialogue system will encounter during deployment. This includes the interaction with a domain expert.

In general, LL evaluation methods need to be reproducible in order to measure improvements and changes over time. One straightforward way of doing this is to deploy a dialogue system and let humans interact with it. However, this is very time consuming and expensive, and alternatives with less or no humans in the loop are desired. One major issue that is particular for evaluating dialogue systems is that they produce their "result" - the dialogue - during the interaction with their environment. Thus, any automated environment environment has to provide artificial users, and building them can be as complex as building the dialogue system itself.

When it comes to LL evaluation, additional complexity arises due to the fact that the interaction with the expert needs to be simulated as well. For instance, the dialogue system may ask an expert for advice about a new entity or topic. In general,

the evaluation system cannot know in advance *which* questions the dialogue system will ask - hence, it is hard to simulate.

Experimental Evaluation Environment

We are currently working on an experimental setting to automate the evaluation of knowledge acquisition and capability improvement. We work in the cooking domain, where the dialogue system is developed to assist the user by answering questions about cooking, e.g. providing recipes, giving advice or providing tutorials. Typical question might be "How do i prepare linguini, which is answered with a corresponding recipe from the database.

In order to evaluate the LL capabilities of the dialogue system, we deploy it in a simulated environment, which consists of:

- Evaluation agent: provides the questions and evaluates the answers given by the dialogue system. The agent institutes new situations by asking about entities which were not present in any training set of the dialogue system (e.g. enchilada), by asking types of questions which the dialogue system did not encounter yet (e.g. "How do i clean my oven?"), or by asking questions about unseen domains (e.g. Chinese food).
- Expert: provides advice to the dialogue system when stuck. The dialogue system can ask clarifying questions to the expert before it tries to answer the question of the evaluation agent. However, this comes at a cost, i.e. each interaction with the expert has its fee, and thus, the system should learn to efficiently interact with the expert.

We envision that the dialogue system asks questions from a list of predefined templates, which the (automated) expert can easily parse and answer. These are, for instance, "What is <X>?" or "Is <X> a relevant entity for this question?".

The domain expert has at its hand a large collection of pre-recorded dialogues on the domain, and returns extracts of these dialogue that match to the clarification question.

The evaluation measures the capability of the LL component to adapt to the new situations. This capability is measured by the number of interactions needed with the expert system before answering the initial question correctly. A system with a strong LL component should adapt to new situations quickly.

6 Conclusion

Implementing lifelong learning for a dialogue system may aim at 1. extending the underlying knowledge base (Knowledge Induction); 2. handling more complex user interactions (Language Understanding); 3. improving the perceived quality of the

resulting dialogues (Dialogue Behaviour); or 4. extending the functionality of the system (Capability Improvement) over time.

While stating these goals is simple, implementing a system that achieves any of these four goal is far from trivial. To the best of our knowledge, most approaches in research currently tackle the first dimension (Knowledge Induction), while there is almost no solution (yet) for the other three.

One important challenge is to evaluation the progress of LL in such systems. In order to avoid time-consuming and costly human evaluations, automated environments are required. We are currently working on such a system, which shall be presented as shared task in 2020.

7 Acknowledgements

This paper has been partially funded by the LIHLITH project, supported by ERA-NET CHIST-ERA and Swiss National Science Foundation (20CH21_174237).

References

1. Iulian Vlad Serban, Ryan Lowe, Peter Henderson, Laurent Charlin, and Joelle Pineau. A survey of available corpora for building data-driven dialogue systems. *arXiv preprint arXiv:1512.05742*, 2017.
2. Hongshen Chen, Xiaorui Liu, Dawei Yin, and Jiliang Tang. A Survey on Dialogue Systems: Recent Advances and New Frontiers. *SIGKDD Explor. Newsl.*, 19(2):25–35, November 2017.
3. Daniel Jurafsky and James H. Martin. *Speech and Language Processing*, chapter Dialog Systems and Chatbots. Draft of 3rd edition edition, 2017.
4. Zhiyuan Chen and Bing Liu. *Lifelong Machine Learning*. Morgan & Claypool Publishers, 2016.
5. Ian J Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville, and Yoshua Bengio. An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv preprint arXiv:1312.6211*, 2013.
6. German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *arXiv preprint arXiv:1802.07569*, 2018.
7. Ricardo Vilalta and Youssef Drissi. A perspective view and survey of meta-learning. *Artificial Intelligence Review*, 18(2):77–95, 2002.
8. Burr Settles. Active learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 6(1):1–114, 2012.
9. Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Trans. on Knowl. and Data Eng.*, 22(10):1345–1359, October 2010.