Continual Learning Dialogue Systems - Learning during Conversation

Sahisnu Mazumder
Multimodal Dialogue & Interaction,
Intel Labs, USA
sahisnumazumder@gmail.com

Bing Liu
Department of Computer Science,
University of Illinois at Chicago, USA
liub@uic.edu

ABSTRACT

Dialogue systems, commonly known as Chatbots, have gained escalating popularity in recent years due to their wide-spread applications in carrying out chit-chat conversations with users and accomplishing various tasks as personal assistants. However, they still have some major weaknesses. One key weakness is that they are typically trained from *pre-collected* and *manually-labeled* data and/or written with handcrafted rules. Their knowledge bases (KBs) are also *fixed* and *pre-compiled* by human experts. Due to the huge amount of manual effort involved, they are difficult to scale and also tend to produce many errors ought to their limited ability to understand natural language and the limited knowledge in their KBs. Thus, when these systems are deployed, the level of user satisfactory is often low.

In this tutorial, we introduce and discuss methods to give chatbots the ability to continuously and interactively learn new knowledge during conversation, i.e. "*on-the-job*" by themselves so that as the systems chat more and more with users, they become more and more knowledgeable and improve their performance over time. The first half of the tutorial focuses on introducing the paradigm of lifelong and continual learning and discuss various related problems and challenges in conversational AI applications. In the second half, we present recent advancements on the topic, with a focus on continuous lexical and factual knowledge learning in dialogues, open-domain dialogue learning after deployment and learning of new language expressions via user interactions for language grounding applications (e.g. natural language interfaces). Finally, we conclude with a discussion on the scopes for continual conversational skill learning and present some open challenges for future research.

CCS CONCEPTS

• Computing methodologies → Discourse, dialogue and pragmatics; **Machine learning**; **• Human-centered computing**; **• Information systems → Information retrieval**;

KEYWORDS

Dialogue and interactive systems, conversational AI, lifelong and continual learning, conversational IR, learning after deployment

SIGIR '22, July 11–15, 2022, Madrid, Spain

ACM ISBN 978-1-4503-8732-3/22/07…$15.00
https://doi.org/10.1145/3477495.3532677

ACM Reference Format:
Sahisnu Mazumder and Bing Liu. 2022. Continual Learning Dialogue Systems - Learning during Conversation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '22), July 11–15, 2022, Madrid, Spain.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3477495.3532677

1 MOTIVATION

Building dialogue systems or conversational agents capable of conversing with humans in natural language (NL) and understanding human NL instructions or commands is a long-standing goal of AI. These agents, also known as chatbots, have become the front runner of AI advancement due to wide-spread applications such as assisting customers in buying products, booking flight tickets, reducing stress, and executing actions like controlling house appliances and reporting weather information. Because of the proliferation of Internet of Things (IoT) devices with NL interfaces, these agents have become ubiquitous in recent times.

Dialogue systems can be broadly categorized into two main types: (1) *Chit-chat systems* designed to engage users and provide mental support by conducting chit-chat type of conversations in a wide range of topics without having a specific goal to complete, and (2) *Task-oriented chatbots* designed to assist users to complete tasks based on users' requests. Most of the popular personal assistants such as Amazon Alexa, Apple Siri, Google Home, and Microsoft Cortana, are task-oriented bots. They are primarily designed as Natural Language Interface (NLI) systems that take human NL instructions/commands and translate them into actions to be executed by the underlying application.

Despite the wide-spread applications, existing conversation agents or chatbots still have many serious weaknesses. In this tutorial, we focus on two key weaknesses: **(1)** A great deal of manual effort is needed to label training data, write rules and compile knowledge bases (KBs) to build and improve these systems. No matter how much data is collected and used to train a chatbot, it is very hard, if not impossible, to cover all possible variations of the natural language. **(2)** The pre-compiled KBs cannot cover the rich knowledge needed in practice. Thus, when deployed, a well-trained chatbot still performs poorly.

A truly intelligent chatbot should not be limited by its offline-trained model or pre-compiled KB. It should learn continually after model deployment (i.e., during conversing or interacting with the human end-users) and thereby, improve its capability over time in a *self-supervised* and *self-motivated* manner [1, 6–8]. This is the topic of this tutorial. This topic is of critical importance for the future success of dialogue systems and conversational IR applications. We humans learn a great deal of our knowledge in our daily conversations and off conversations and the knowledge learned is used in subsequent conversations in a lifelong manner. In this tutorial, we focus on learning from end-users interactively during conversation.

In the past few years, several researchers started to address these issues to give chatbots the ability to learn from end-users in their interactions continuously after the systems have been deployed and used in practice to improve their capability over time [7, 11]. This is just like human learning on the job [6, 8]. It is well-known in learning science that about 70% of our human knowledge comes from 'on-the-job' learning, only about 10% through formal education or training, and the rest 20% through observation of others. Dialogue systems should have the same capability. In order to learn during conversation (i.e., on the job), the system has to actively communicate or interact with the users by asking them questions, for which it needs to dynamically **(1)** ***formulate an interaction strategy*** to interact with the user (e.g., deciding what to ask the user and when to ask the user), **(2)** ***execute the strategy*** to acquire the ground truth data and other relevant information, and **(3)** ***incrementally learn from the new data*** in a continual manner.

As most chatbots work in a multi-user environment, they can exploit such an environment to obtain the ground truth training data and other knowledge during actual online conversations to enable continual online learning. Below we give some example ways that the system can learn *lexical and factual knowledge* from end-users during a conversation.

(1) Extracting information from user utterances: The chatbot can extract information (e.g., real-world facts, user's preferences etc.) from user utterances (or dialogue history) directly, which can then be incorporated into the system's knowledge base or the model continually on the fly.

(2) Asking the current user: When the system (1) does not understand a user utterance, or (2) cannot answer a user query, it can initiate a new learning task. To obtain the ground truth data for learning, the chatbot can ask the current user for clarification, rephrasing, or even demonstration if it is supported in the underlying application [for situation (1)]. For situation (2), the chatbot can ask the user for some supporting facts and then infer the query answer. To obtain more knowledge, the chatbot may even ask the current user related questions in order to fully exploit the opportunity to get information that it considers missing. For example, if the user says - "*I visited London last month.*", apart from extracting the fact that *London* is a *location* (from the user utterance), it can also ask a subsequent question: "*Where is London?*" If the user answers "*London is in UK*", the system learns another piece of knowledge assuming it was not in its KB.

(3) Asking other users: When the chatbot could not answer a user query, it may also ask other users to obtain the answer. For example, if a user asks "*What is the capital city of the US?*" and the chatbot is not able to answer or infer now, it can try to find a good opportunity later to ask another user (considering its deployment and operation in a multi-user environment): "*Hey, do you happen to know what the capital city of the US is?*" If the user gives the answer "*Washington DC*", the chatbot acquires a piece of new knowledge.

Such pieces of real-world knowledge or facts learned in one conversation session can then be utilized automatically in future conversation sessions due to the continuous enrichment of its backend KBs over time and thus, can directly improve the performance of knowledge-grounded conversation modeling, conversational question answering, conversational search, interactive semantic parsing and many other conversational AI and IR applications.

In this tutorial, we will first introduce the paradigm of lifelong or continual machine learning [1], and then focus on addressing six emerging continual learning capabilities/aspects of chatbots:

(1) learning new lexical and factual knowledge in open-ended and information-seeking conversations to expand the system's knowledge base (KB).

(2) learning to ground new natural language commands or expressions via interactions with users.

(3) continual open-domain conversation modeling via learning after model deployment.

(4) learning to detect unknown intents (i.e., open-world learning or out-of-distribution detection) and incrementally learning new intents and slots for Task-oriented Dialogue systems (ToDS).

(5) learning of new conversational skills to improve response generation awareness of chatbots.

(6) dealing with wrong knowledge learned from users. The knowledge learned from end-users can be erroneous and some users may even purposely fool the system by providing wrong knowledge. We will discuss how to deal with this problem in order to ensure the credibility or trustworthiness of the learned knowledge from end-users.

2 AUDIENCE

Researchers, graduate students, and practitioners who are interested in dialogue systems and lifelong/continual learning. The tutorial will particularly benefit people who are building or intend to build dialogue systems because making such systems learn continually by themselves is becoming necessary for the next generation conversational AI. Some current techniques are ready for commercial use.

3 PREREQUISITE KNOWLEDGE

Basic knowledge of machine learning, deep learning, IR and NLP.

4 RELEVANCE TO THE IR COMMUNITY

Our tutorial focuses on building the next-generation dialogue systems that can *continually and interactively learn new knowledge (e.g., world facts, users' language usage, emotions, context etc.) from end-users during conversation* to become more and more knowledgeable and improve their performance over time. Dialogue and interactive systems is a prime topic of importance today due to its wide spread applications, considering the popularity of virtual personal assistants, customer service bots, task completion bots etc. in our daily life and their increasing use in many real-world conversational IR applications. Advancing the technology would greatly impact the state-of-the-art in IR due to its implications in improving performance of conversational search, conversational QA, conversational recommendation systems and so on. In fact, SIGIR has recognized 'Dialogue and Interactive systems' to be one of the major topics of research interests over the years and has attracted the community to contribute substantially on this topic. As more and more usage of conversational interfaces in downstream IR tasks becomes prevalent in both commercial and daily use cases, it becomes more important for the IR research community to focus on how to make these systems learn continuously and become more and more intelligent over time. As dialogue and interactive systems lie in the core of such IR systems, making them learn during conversation is a natural requirement to boost the performance of the underlying IR tasks. Thus, we believe the content of this tutorial would be interesting, beneficial and important to both IR researchers and practitioners.

5 REFERENCE TO TUTORIALS IN THE SAME AREA AT SIGIR / RELATED CONFERENCES

We are not aware of any related tutorials in SIGIR or related conferences (including NLP conferences) except a tutorial on the topic in IJCAI-2021 given by us. We are aware of several tutorials about dialogue systems in general given in various conferences, but lifelong or continual learning dialogue systems focus on a completely different and new aspect, which is more advanced and we believe should be a key component in the next generation dialogue systems (and hence, conversational IR systems).

6 FORMAT AND DETAILED SCHEDULE

• Outline of the Tutorial

(1) Introduction [25 mins]
- (a) On-the-job learning: Motivation
- (b) Dialogue & Interactive Systems: Background
- (c) Modern Dialogue systems: Weaknesses
- (d) Goal of this tutorial

(2) Lifelong and Continual Learning: An Introduction [25 mins]
- (a) Lifelong and Continual learning: Concepts and definitions
- (b) Techniques for lifelong and Continual learning
- (c) Open-World AI and learning
- (d) Lifelong Interactive Learning in Conversation (LINC): a theoretical framework

(3) Continuous Knowledge Learning in Dialogues [25 mins]
- (a) Opportunities of knowledge learning
- (b) Lexical knowledge acquisition
- (c) Interactive factual knowledge learning

(4) Continual Language Learning and Grounding [25 mins]
- (a) Learning language games through interactions
- (b) Dialogue-driven learning of self-adaptive natural language command interfaces

(5) Open-Domain Dialogue Learning after Deployment [20 mins]
- (a) Learning by extracting new training examples from conversations
- (b) Dialogue learning via role-playing games

(6) Continual Learning for Task-oriented Dialogue Systems (ToDS) [25 mins]
- (a) Open intent classification
- (b) Continual learning for slot filling
- (c) AdapterCL: End-to-end continual Learning for ToDS

(7) Continual Learning of Conversational Skills [10 mins]

(8) Challenges and future scopes [10 mins]

(9) Summary & QA [15 mins]

• Reading List

Please see the References section for the Reading list.

7 TYPE OF SUPPORT MATERIALS TO BE SUPPLIED TO ATTENDEES

We will release our slides in our tutorial website for public access and also agree to allow video recording and publication of our tutorial video on Web.

REFERENCES

[1] Zhiyuan Chen and Bing Liu. 2018. *Lifelong machine learning.* Morgan & Claypool Publishers.

[2] Braden Hancock, Antoine Bordes, Pierre-Emmanuel Mazare, and Jason Weston. 2019. Learning from Dialogue after Deployment: Feed Yourself, Chatbot!. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics.* 3667–3684.

[3] Ben Hixon, Peter Clark, and Hannaneh Hajishirzi. 2015. Learning knowledge graphs for question answering through conversational dialog. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies.* 851–861.

[4] Kazunori Komatani and Mikio Nakano. 2020. User Impressions of Questions to Acquire Lexical Knowledge. In *Proceedings of the 21th Annual Meeting of the Special Interest Group on Discourse and Dialogue.* 147–156.

[5] Han Li, Jihwan Lee, Sidharth Mudgal, Ruhi Sarikaya, and Young-Bum Kim. 2019. Continuous Learning for Large-scale Personalized Domain Classification. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers).* 3784–3794.

[6] Bing Liu. 2020. Learning on the job: Online lifelong and continual learning. In *Proceedings of the AAAI Conference on Artificial Intelligence.* 13544–13549.

[7] Bing Liu and Sahisnu Mazumder. 2021. Lifelong and Continual Learning Dialogue Systems: Learning during Conversation. In *Proceedings of the AAAI Conference on Artificial Intelligence.* 15058–15063.

[8] Bing Liu, Sahisnu Mazumder, Eric Robertson, and Scott Grigsby. 2022. AI Autonomy: Self-Initiation, Adaptation and Continual Learning. *arXiv preprint arXiv:2203.08994* (2022).

[9] Bing Liu and Chuhe Mei. 2020. Lifelong knowledge learning in rule-based dialogue systems. *arXiv preprint arXiv:2011.09811* (2020).

[10] Andrea Madotto, Zhaojiang Lin, Zhenpeng Zhou, Seungwhan Moon, Paul Crook, Bing Liu, Zhou Yu, Eunjoon Cho, and Zhiguang Wang. 2020. Continual Learning in Task-Oriented Dialogue Systems. *arXiv preprint arXiv:2012.15504* (2020).

[11] Sahisnu Mazumder. 2021. *On-the-job Continual and Interactive Learning of Factual Knowledge and Language Grounding.* Ph. D. Dissertation. University of Illinois at Chicago.

[12] Sahisnu Mazumder, Bing Liu, Nianzu Ma, Shuai Wang, and AI Amazon. 2020. Continuous and Interactive Factual Knowledge Learning in Verification Dialogues. In *NeurIPS-2020 Workshop on Human And Machine in-the-Loop Evaluation and Learning Strategies.*

[13] Sahisnu Mazumder, Bing Liu, Shuai Wang, and Sepideh Esmaeilpour. 2020. An Application-Independent Approach to Building Task-Oriented Chatbots with Interactive Continual Learning. In *NeurIPS-2020 Workshop on Human in the Loop Dialogue Systems.*

[14] Sahisnu Mazumder, Bing Liu, Shuai Wang, and Nianzu Ma. 2019. Lifelong and Interactive Learning of Factual Knowledge in Dialogues. In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue.* 21–31.

[15] Sahisnu Mazumder, Nianzu Ma, and Bing Liu. 2018. Towards a continuous knowledge learning engine for chatbots. *arXiv preprint arXiv:1802.06024* (2018).

[16] Kohei Ono, Ryu Takeda, Eric Nichols, Mikio Nakano, and Kazunori Komatani. 2016. Toward lexical acquisition during dialogues through implicit confirmation for closed-domain chatbots. In *Proc. Workshop Chatbots Conversational Agent Technologies.* 272–279.

[17] Kohei Ono, Ryu Takeda, Eric Nichols, Mikio Nakano, and Kazunori Komatani. 2017. Lexical acquisition through implicit confirmations over multiple dialogues.

In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*. 50–59.

[18] Tsugumi Otsuka, Kazunori Komatani, Satoshi Sato, and Mikio Nakano. 2013. Generating more specific questions for acquiring attributes of unknown concepts from users. In *Proceedings of the SIGDIAL 2013 Conference*. 70–77.

[19] Yilin Shen, Xiangyu Zeng, and Hongxia Jin. 2019. A progressive model to enable continual learning for semantic slot filling. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. 1279–1284.

[20] Kurt Shuster, Jack Urbanek, Emily Dinan, Arthur Szlam, and Jason Weston. 2021. Dialogue in the wild: Learning from a deployed role-playing game with humans and bots. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*. 611–624.

[21] Sida I Wang, Percy Liang, and Christopher D Manning. 2016. Learning language games through interaction. In *54th Annual Meeting of the Association for Computational Linguistics, ACL 2016*. Association for Computational Linguistics (ACL), 2368–2378.