

IJCAI - 2021  
**Tutorial Proposal**

**Continual Learning Dialogue Systems - Learning on the Job  
after Model Deployment**

**Sahisnu Mazumder , Bing Liu**

Department of Computer Science, University of Illinois at Chicago, USA  
sahisnumazumder@gmail.com, liub@uic.edu

**Type of Tutorial:** Topic of emerging importance for AI.

**Duration:** 1:45 hours

## 1 Short Two-sentence Description

The tutorial focuses on the research topic of building the next-generation dialogue systems that can continuously and interactively learn from end-users during conversation (on-the-job learning after deployment) to become more and more powerful over time. We will provide a background of the topic, and discuss existing techniques and open challenges, which we believe, will shape the future of dialogue systems.

## 2 Abstract for Webpage Overview

Dialogue systems, commonly known as Chatbots, have gained escalating popularity in recent times 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 owing 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

command interfaces). Finally, we conclude with a discussion on the scopes for continual conversational skill learning and present some of the open challenges for future research.

## 3 Brief Description

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 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 to 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. Thus, when deployed in practice, a well-trained chatbot still performs poorly. (2) The pre-compiled KBs cannot cover the rich knowledge needed in practice. A truly intelligent chatbot should not be limited by its offline-trained model or pre-compiled KB. It should learn continually after model deployment and during conversing or interacting with the (human) end-users and thereby, improve its capability over time in a

self-supervised manner (Liu, 2020). This is the topic of this tutorial. This topic is of critical importance for the future success of dialogue systems. 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 this tutorial, we focus on learning from end-users 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 (Liu, 2020). This is just like human learning on the job. 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 from end-users during a conversation.

1. *Extracting information from user utterances*: The chatbot can extract information from user utterances (or dialogue history) directly which can be real-world facts, user’s preferences, etc., 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 initiates a new learning task. To obtain the ground truth data for learning, for (1), the chatbot can ask the current user for clarification, rephrasing, or even demonstration if it is supported in the underlying application. For (2), the chatbot may 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, the user said “I visited London last month.” Apart from extracting London as a location, it can also ask a subsequent question: “Where is London?” If the user answers “London is in UK,” the system learned another piece of knowledge if it is not already in its knowledge base.
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: “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.

In this tutorial, we will first introduce the paradigm of life-long or continual machine learning (Chen and Liu, 2018), and then focus on five emerging continual learning capabilities of chatbots as follows:

- (1) learning to classify intents and detect unknown intents (i.e., open-world learning or out-of-distribution detection) and incrementally learning new intents if sufficient training data is available.
- (2) learning new lexical and factual knowledge in open-ended and information-seeking conversations to expand the system’s knowledge base (KB).
- (3) continual open-domain conversation modeling via learning after model deployment.
- (4) learning to ground new natural language commands or expressions via interactions with end-users.
- (5) learning of new conversational skills to improve response generation awareness of the chatbot over time.

Learning from end-users comes with a risk. That is, 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 to ensure the credibility or trustworthiness of the learned knowledge from end-users.

## 4 Target 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 dialogue systems. Some current techniques are ready for commercial use.

## 5 Estimate of Audience Size

We estimate that there will be more than 100 attendees because dialogue systems are not only popular in academia but also very popular in industry. Numerous dialogue systems are used in practice, which makes learning during dialogues interesting and important to both researchers and practitioners.

This tutorial has not been given anywhere before.

## 6 Prerequisite Knowledge

Basic knowledge of machine learning and NLP.

## 7 Outline of the Tutorial

1. Introduction [15 minutes]
  - (a) Goal of this tutorial
  - (b) Dialogue & Interactive Systems: Background
  - (c) Opportunities for self-motivated learning during conversation

2. Lifelong and Continual Learning: an Introduction [15 minutes]
  - (a) Diverse tasks of lifelong learning and their definitions
  - (b) Introduction to existing techniques
  - (c) Lifelong Interactive Learning in Conversation
3. Intent Classification and Out-of-Distribution Detection [15 minutes]
  - (a) Open-world AI and learning
  - (b) Detecting out-of-distribution user utterances
4. Continuous Knowledge Learning during Conversation [15 minutes]
  - (a) Different aspects of knowledge learning
  - (b) Interactive factual knowledge learning in dialogues
  - (c) Lexical knowledge acquisition via dialogue
5. Open-Domain Dialogue Learning After Deployment [15 minutes]
  - (a) Learning by extracting new training examples from conversations
  - (b) Dialogue learning via role-playing games
6. Continual Language Learning and Grounding [15 minutes]
  - (a) Learning language games through interactions
  - (b) Dialogue-driven learning of application-independent natural language command interfaces
7. Continual Learning of Conversational Skills [10 minutes]
8. Summary & QA [5 minutes]

## 8 Tutorial-Level Publications

1. Bing Liu and Sahisnu Mazumder. Lifelong and Continual Learning Dialogue Systems: Learning during Conversation. *Proceedings of AAAI Conference on Artificial Intelligence (AAAI-2021, senior member track)*, 2021.
2. Zhiyuan Chen and Bing Liu. *Lifelong Machine Learning*. Morgan & Claypool, 2018 (2nd edition), 2016 (1st edition).

Bing Liu has given and will be giving several invited talks on the topic.

## 9 Presenter # 1 Information

### 9.1 Contact Information

**Sahisnu Mazumder**  
 Department of Computer Science  
 University of Illinois at Chicago  
 851 S. Morgan (M/C 152)  
 Chicago, IL 60607-7053  
 Email: sahisnumazumder@gmail.com  
 Website: <https://sahisnu.github.io/>

## 9.2 Short Biography

**Sahisnu Mazumder** is a final year PhD student of Computer Science at the University of Illinois at Chicago (UIC), USA. He received his Masters in Computer Science from Indian Institute of Technology (IIT) - Roorkee, India. His research interests include Lifelong and Continual Learning, Dialogue and Interactive Systems, Open-world Learning, Knowledge base Reasoning, and Sentiment Analysis. He has published several research papers in leading AI, NLP and Dialogue conferences like AAAI, IJCAI, ACL, EMNLP, SIGDIAL and served as PC Member / Reviewer of premier conferences like AAAI, IJCAI, ACL, EMNLP, NAACL, EACL and journals like ACM TALLIP and IEEE TNNLS. He has also worked as a Research Intern at Huawei Research USA on user activity & interest mining and at Microsoft Research - Redmond on Natural Language Interaction (NLI) system design. Prior to joining UIC, he worked as a Junior Research Fellow at IIT Roorkee on News Data Analytics and Location-aware Review Mining and delivered a tutorial on the topic at the International Conference on Big Data Analytics (BDA - 2014).

### 9.3 Background in the Tutorial Area: publications/presentations

**Sahisnu Mazumder** has been working on the topic of continual and interactive learning in dialogues (as his main PhD research topic) for past five years and together with his advisor Bing Liu (presenter # 2) published several research articles on the topic [see publications # 6-10 in Section 14] and presented works in premier dialogue and AI conferences and workshops like IJCAI, AAAI, SIGDIAL and NeurIPS Workshops on Dialogue and Interactive Systems. He and Bing Liu were one of the earliest to introduce the idea of a continual knowledge learning engine for chatbots, which was featured in Data Science Central in March, 2018 [<https://www.datasciencecentral.com/profiles/blogs/make-your-chatbot-smarter-by-talking-to-it>]. Besides, he has also co-authored several lifelong learning papers in sentiment analysis domain in recent years.

Please visit <https://sahisnu.github.io/> to get a more detailed information about his list of publications, talks and tutorials.

### 9.4 Teaching Experience

**Sahisnu Mazumder** has served as a Graduate Teaching Assistant in the Department of Computer Science at UIC from Aug, 2015 - May, 2019 and has sufficient experience in classroom teaching and student interactions. Besides, he has delivered a tutorial in international conference BDA-2014 and various talks in Huawei Research USA and Microsoft Research Redmond on topics related to machine learning and natural language interface (NLI) design during his internships.

## 10 Presenter#2 Information

### 10.1 Contact Information

**Bing Liu**  
 Department of Computer Science  
 University of Illinois at Chicago

851 S. Morgan (M/C 152)  
Chicago, IL 60607-7053  
Email: liub@uic.edu  
Website: <https://www.cs.uic.edu/~liub/>

## 10.2 Short Biography

**Bing Liu** is a distinguished professor of Computer Science at the University of Illinois at Chicago (UIC). He received his PhD in Artificial Intelligence from the University of Edinburgh. His research interests include lifelong and continual learning, continual learning dialogue systems, sentiment analysis, open-world AI/learning, natural language processing (NLP), and machine learning. He has published extensively in top conferences and journals. He also authored four books: one about lifelong learning, two about sentiment analysis and one about Web mining. Three of his papers received Test-of-Time awards: two from KDD and one from WSDM. Another of his papers received Test-of-Time award - honorable mention also from WSDM. He is a Fellow of the ACM, AAAI, and IEEE. He has given tutorials in numerous NLP, AI and data mining conferences, including ACL-2007, EACL-2012, EMNLP-2016, AAAI-2011, IJCAI-2015, KDD-2016, and WWW-2008 and. He has also given numerous keynotes and invited talks on related topics.

## 10.3 Background in the Tutorial Area: publications/presentation

See the corresponding section of the first presenter as Bing Liu is the PhD advisor of the first presenter. His group has published extensively on the topic of lifelong and continual learning. More details can be found at <https://www.cs.uic.edu/~liub/lifelong-learning.html>.

## 10.4 Teaching Experience

Bing Liu is a distinguished professor and has taught at the university level for more than 25 years.

## 10.5 Evidence of Scholarship in AI or Computer Science

He is a Fellow of AAAI, ACM and IEEE.

## 11 Special Requirements

No special requirements.

## 12 Reading List

1. Zhiyuan Chen and Bing Liu. *Lifelong Machine Learning*. Morgan & Claypool, 2018 (2nd edition), 2016 (1st edition).
2. Braden Hancock, Antoine Bordes, Pierre-Emmanuel Mazare, and Jason Weston. Learning from Dialogue after Deployment: Feed Yourself, Chatbot!. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL-2019)*, July 28 - Aug 2, Florence, Italy.
3. Ben Hixon, Peter Clark and Hannaneh Hajishirzi. Learning knowledge graphs for question answering

through conversational dialog. *Proceedings of Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT 2015)*, Denver, Colorado, May 31 – June 5, 2015.

4. Kazunori Komatani and Mikio Nakano. User Impressions of Questions to Acquire Lexical Knowledge. *Proceedings of Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL-2020)*. 2020.
5. Bing Liu. Learning on the Job: Online Lifelong and Continual Learning. *Proceedings of AAAI Conference on Artificial Intelligence (AAAI-2020)*, Feb 7-12, 2020, New York City, USA.
6. Bing Liu and Sahisnu Mazumder. Lifelong and Continual Learning Dialogue Systems: Learning during Conversation. In *Proceedings of AAAI Conference on Artificial Intelligence (AAAI-2021)*, 2021.
7. Bing Liu and Chuhe Mei. Lifelong Knowledge Learning in Rule-based Dialogue Systems. arXiv:2011.09811 [cs.AI], 2020.
8. Sahisnu Mazumder, Bing Liu, Shuai Wang, and Sepideh Esmaeilpour. An Application-Independent Approach to Building Task-Oriented Chatbots with Interactive Continual Learning. In *NeurIPS-2020 Workshop on Human in the Loop Dialogue Systems*.
9. Sahisnu Mazumder, Bing Liu, Nianzu Ma, Shuai Wang. Continuous and Interactive Factual Knowledge Learning in Verification Dialogues. In *NeurIPS-2020 workshop on Human and Model in the Loop Evaluation and Training Strategies*. 2020.
10. Sahisnu Mazumder, Bing Liu, Shuai Wang, Nianzu Ma. Lifelong and Interactive Learning of Factual Knowledge in Dialogues. *Proceedings of Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL-2019)*, 11-13 September, Stockholm, Sweden.
11. Sahisnu Mazumder, Nianzu Ma, and Bing Liu. Towards a continuous knowledge learning engine for chatbots. *arXiv preprint arXiv:1802.06024* (2018).
12. Kohei Ono, Ryu Takeda, Eric Nichols, Mikio Nakano and Kazunori Komatani. Lexical acquisition through implicit confirmations over multiple dialogues. *Proceedings of the Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL 2017)*, 15-17 August, Saarbrücken, Germany.
13. Kohei Ono, Ryu Takeda, Eric Nichols, Mikio Nakano, and Kazunori Komatani. Toward lexical acquisition during dialogues through implicit confirmation for closed-domain chatbots. *Proceedings of Second Workshop on Chatbots and Conversational Agent Technologies (WOCHAT)*, 2016.
14. Tsugumi Otsuka, Kazunori Komatani, Satoshi Sato and Mikio Nakano. Generating more specific questions for acquiring attributes of unknown concepts from users.

*Proceedings of the Annual Meeting of the Special Interest Group on Discourse and Dialogue Conference (SIGDIAL-2013)*, 22-24 August, Metz, France.

15. Kurt Shuster, Jack Urbanek, Emily Dinan, Arthur Szlam, and Jason Weston. Deploying Lifelong Open-Domain Dialogue Learning. arXiv preprint arXiv:2008.08076 (2020).
16. Han Li, Jihwan Lee, Sidharth Mudgal, Ruhi Sarikaya, and Young-Bum Kim. Continuous learning for large-scale personalized domain classification. arXiv preprint arXiv:1905.00921 (2019).
17. Yilin Shen, Xiangyu Zeng and Hongxia Jin. A Progressive Model to Enable Continual Learning for Semantic Slot Filling. *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP-2019)*, pp. 1279-1284. 2019.
18. Andrea Madotto, Zhaojiang Lin, Zhenpeng Zhou, Seungwhan Moon, Paul Crook, Bing Liu, Zhou Yu, Eunjoon Cho, and Zhiguang Wang. Continual Learning in Task-Oriented Dialogue Systems. arXiv preprint arXiv:2012.15504 (2020).
19. Sida I. Wang, Percy Liang and Christopher D. Manning. Learning Language Games through Interaction. *Proceedings of Annual Meeting of the Association for Computational Linguistics (ACL-2016)*, Berlin, Germany, August 7-12, 2016.