

InterNLP 2021

**First Workshop on Interactive Learning for Natural
Language Processing**

Proceedings of the Workshop

August 5, 2021
Bangkok, Thailand (online)

©2021 The Association for Computational Linguistics
and The Asian Federation of Natural Language Processing

Order copies of this and other ACL proceedings from:

Association for Computational Linguistics (ACL)
209 N. Eighth Street
Stroudsburg, PA 18360
USA
Tel: +1-570-476-8006
Fax: +1-570-476-0860
acl@aclweb.org

ISBN 978-1-954085-66-4

Message from the General Chair

Motivation: A key aspect of human learning is the ability to learn continuously from various sources of feedback. In contrast, much of the recent success of deep learning for NLP relies on large datasets and extensive compute resources to train and fine-tune models, which then remain fixed. This leaves a research gap for systems that adapt to the changing needs of individual users or allow users to continually correct errors as they emerge. Learning from user interaction is crucial for tasks that require a high grade of personalization and for rapidly changing or complex, multi-step tasks where collecting and annotating large datasets is not feasible, but an informed user can provide guidance.

What is interactive NLP?: Interactive Learning for NLP means training, fine-tuning or otherwise adapting an NLP model to inputs from a human user or teacher. Relevant approaches range from active learning with a human in the loop, to training with implicit user feedback (e.g. clicks), dialogue systems that adapt to user utterances and training with new forms of human input. Interactive learning is the converse of learning from datasets collected offline with no human input during the training process.

Goals: The goal of this workshop was to bring together researchers to:

- Develop novel methods for interactive machine learning of NLP models.
- Discuss how to evaluate interactive NLP systems, including models for realistic user simulation.
- Identify scenarios involving natural language where interactive learning is beneficial.

Previous work has been split across different tracks and task-focused workshops, making it hard to disentangle applications from broadly-applicable methodologies or establish common practices for evaluating interactive learning systems. We aimed to bring together researchers to share insights on interactive learning from a wide range of NLP-related fields, including, but not limited to, dialogue systems, question answering, summarization, and educational applications.

Concerning methodology, we encouraged submissions investigating various dimensions of interactive learning, such as (but not restricted to):

- Interactive machine learning methods: the wide range of topics discussed above, from active learning with a user to methods that extract, interpret and aggregate user feedback or preferences from complex interactions, such as natural language instructions.
- User effort: the amount of user effort required for different types of feedback; explicit labels require higher user effort than feedback deduced from user interaction (e.g., clicks, viewtime); how users cope with the system misinterpreting instructions.
- Feedback types: different types of feedback require different techniques to incorporate them into a model. E.g., explicit labels allow us to directly train while user instructions require interpretation.

A major bottleneck for interactive learning approaches is their evaluation, including a lack of suitable datasets. We, therefore, encouraged submissions that cover research into the following:

- Evaluation methods: approaches to assessing interactive methods, such as low-effort, easily reproducible approaches with real-world users and simulated user models for automated evaluation.
- Reproducibility: procedures for documenting user evaluations and ensuring they are reproducible.
- Data: Introduce novel datasets for training and evaluating interactive models.

To investigate scenarios where interactive learning is effective, we invited submissions that present empirical results for applications of interactive methods.

Organizing Committee

- Kianté Brantley (kdbrant@cs.umd.edu) is a fourth year PhD student in Computer Science at The University of Maryland College Park advised by Hal Daumé III. His research interest is in designing algorithms that efficiently integrate domain knowledge into sequential decision-making problems (e.g. reinforcement learning, imitation learning and structure prediction for natural language processing).
- Soham Dan (sohamdan@seas.upenn.edu) is a PhD student at the University of Pennsylvania working with Dan Roth on natural language understanding- specifically, in the context of grounded domains. His research involves concept learning, interactive learning and semantic parsing of instructions.
- Iryna Gurevych (gurevych@ukp.informatik.tu-darmstadt.de) is a full professor in the department of Computer Science at the Technical University of Darmstadt. She has published on interactive methods for NLP in various NLP domains such as language learning, text summarization, or entity linking.
- Ji-Ung Lee (lee@ukp.informatik.tu-darmstadt.de) is a PhD student at the Technical University of Darmstadt. His research focuses on effective model training from user feedback in low-data scenarios coupled with providing the user with instances that fit their needs.
- Filip Radlinski (filiprad@google.com) is a Research Scientist at Google, UK. His research focuses on improvements to conversational search and recommendation through better understanding and modeling user interests through natural language, improved transparency of conversational systems, as well as human-centered evaluation and personalization of information retrieval and recommendation tasks. He received his PhD from Cornell University.
- Hinrich Schütze (hinrichacl@cis.lmu.de) is a full professor at Ludwig Maximilian University, Munich and chair of computational linguistics. His research covers deep learning for NLP, semantics in NLP and linguistics, and information retrieval. He has co-organized several workshops, including two SCLeM workshops (Subword and Character level models in NLP) at EMNLP and NAACL and a Dagstuhl seminar entitled "From Characters to Understanding Natural Language".
- Edwin Simpson (edwin.simpson@bristol.ac.uk) is a lecturer at the University of Bristol working on interactive learning for NLP and machine learning for crowdsourced annotation with an interest in Bayesian methods for handling uncertainty.
- Lili Yu (liliyu@fb.com) is a research scientist in the Facebook Language Research team. Her research interest lies in summarization, conversational AI, learning from user feedback and knowledge representation and grounding.

Table of Contents

| | |
|--|----|
| <i>HILDIF: Interactive Debugging of NLI Models Using Influence Functions</i> Hugo Zylberajch, Piyawat Lertvittayakumjorn and Francesca Toni | 1 |
| <i>Apple Core-dination: Linguistic Feedback and Learning in a Speech-to-Action Shared World Game</i> Susann Boy, AriaRay Brown and Morgan Wixted | 7 |
| <i>SHAPELURN: An Interactive Language Learning Game with Logical Inference</i> Katharina Stein, Leonie Harter and Luisa Geiger | 16 |
| <i>A Proposal: Interactively Learning to Summarise Timelines by Reinforcement Learning</i> Yuxuan Ye and Edwin Simpson | 25 |
| <i>Dynamic Facet Selection by Maximizing Graded Relevance</i> Michael Glass, Md Faisal Mahbub Chowdhury, Yu Deng, Ruchi Mahindru, Nicolas Rodolfo Faucegla, Alfio Gliozzo and Nandana Mihindukulasooriya | 32 |
| <i>Active Curriculum Learning</i> Borna Jafarpour, Dawn Sepehr and Nick Pogrebnyakov | 40 |
| <i>Tackling Fake News Detection by Interactively Learning Representations using Graph Neural Networks</i> Nikhil Mehta and Dan Goldwasser | 46 |

Conference Program

August 5th, 2021

Session 1

19:00–19:05 **Opening Remarks**

19:05–19:50 *Invited talk 1*
Dan Goldwasser

19:50–20:20 *Poster short talks 1*

HILDIF: Interactive Debugging of NLI Models Using Influence Functions
Hugo Zylberajch, Piyawat Lertvittayakumjorn and Francesca Toni

Apple Core-dination: Linguistic Feedback and Learning in a Speech-to-Action Shared World Game
Susann Boy, AriaRay Brown and Morgan Wixted

SHAPELURN: An Interactive Language Learning Game with Logical Inference
Katharina Stein, Leonie Harter and Luisa Geiger

A Proposal: Interactively Learning to Summarise Timelines by Reinforcement Learning
Yuxuan Ye and Edwin Simpson

20:20–21:30 *Poster session 1*

21:30–22:15 *Invited talk 2 : Beyond Imitation Learning: New Paradigms of Querying and Integrating Different Types of Human Data for Improved Robot Learning*
Dorsa Sadigh

August 5th, 2021 (continued)

Session 2

23:00–23:05 Opening Remarks

23:05–23:50 *Invited talk 3 : Continual Language Learning through Collaborative Interaction with Users*
Yoav Artzi

August 6th, 2021

00:00–00:45 *Panel discussion 1*
Julia Kreutzer, Alison Smith-Renner, Ido Dagan

01:00–01:30 Poster short talks 2

Dynamic Facet Selection by Maximizing Graded Relevance

Michael Glass, Md Faisal Mahbub Chowdhury, Yu Deng, Ruchi Mahindru, Nicolas Rodolfo Fauceglia, Alfio Gliozzo and Nandana Mihindukulasooriya

Interactive learning from activity description

Khanh Nguyen, Dipendra Misra, Robert Schapire, Miro Dudík and Patrick Shafto

Active Curriculum Learning

Borna Jafarpour, Dawn Sepehr and Nick Pogrebnyakov

Tackling Fake News Detection by Interactively Learning Representations using Graph Neural Networks

Nikhil Mehta and Dan Goldwasser

01:30–02:45 Poster session 2

03:00–03:45 *Invited talk 4*
Percy Liang

August 6th, 2021 (continued)

Session 3

06:00–06:05 **Opening Remarks**

06:05–06:50 *Invited talk 5*
Dan Roth

07:00–07:45 *Panel discussion 2*
Seung-won Hwang, Yu Zhou, Dilek Hakkani-Tur

