

User or Labor: An Interaction Framework for Human-Machine Relationships in NLP

Ruyuan Wan

University of Notre Dame
rwan@nd.edu

Naome Etori

University of Minnesota
etori001@umn.edu

Karla Badillo-Urquiola

University of Notre Dame
kbadill13@nd.edu

Dongyeop Kang

University of Minnesota
dongyeop@umn.edu

Abstract

The bridging research between Human-Computer Interaction and Natural Language Processing is developing quickly these years. However, there is still a lack of formative guidelines to understand the human-machine interaction in the NLP loop. When researchers crossing the two fields talk about humans, they may imply a user or labor. Regarding a human as a user, the human is in control, and the machine is used as a tool to achieve the human's goals. Considering a human as a laborer, the machine is in control, and the human is used as a resource to achieve the machine's goals. Through a systematic literature review and thematic analysis, we present an interaction framework for understanding human-machine relationships in NLP. In the framework, we propose four types of human-machine interactions: Human-Teacher and Machine-Learner, Machine-Leading, Human-Leading, and Human-Machine Collaborators. Our analysis shows that the type of interaction is not fixed but can change across tasks as the relationship between the human and the machine develops. We also discuss the implications of this framework for the future of NLP and human-machine relationships.

1 Introduction

Research at the intersection of Natural Language Processing (NLP) and Human-Computer Interaction (HCI) is developing rapidly. Humans and machines are now both engaged in each step of the end-to-end NLP pipeline (Wang et al., 2021; Wu et al., 2022). NLP systems are trained on data created and annotated by humans, such as news articles (Da San Martino et al., 2019), Wikipedia pages (Faruqui et al., 2018), product reviews (Badlani et al., 2019), and social media posts (Joseph et al., 2021). Yet, humans are also empowered

by NLP systems, such as writing assistants (Chen et al., 2012), collaborative text revision (Du et al., 2022), and translators (Gu et al., 2016). Nonetheless, humans and machines play distinct roles in these scenarios. From the HCI perspective, researchers usually consider humans as the users of certain technology within the interaction, while many NLP researchers emphasize the labor responsibility of humans to improve the performance of NLP models in tasks.

Humans and machines naturally have different strengths, such as trustworthiness, automation, and assessment ability (Shneiderman, 2020; Maes, 1995). Humans stand out for trustworthiness, while machines are known for automation. Both humans and machines have the competence to evaluate each other, given their own specialties.

While more and more interdisciplinary research and systems are being built to bridge the HCI and NLP fields, we still lack a normative understanding of how human and machine interaction works within the NLP context. More importantly, does the interaction work well? To fill this gap, we address two main research questions:

RQ1: *How does human-machine interaction happen in NLP?*

RQ2: *How do humans and machines interact with each other in NLP?*

We address our research questions by conducting a systematic literature review on the interactive NLP research from leading HCI and NLP venues. Our goal was to define a generalizable human-machine interaction framework in NLP to explain current implementation, guide the design of human-machine interaction, and inspire future research in interactive NLP systems. Based on our synthesis, we defined three properties of interaction: *continuity*, *variety* of interaction options, and *medium* (RQ1). We also conceptualized four re-

relationships summarizing the roles of human and machine interaction patterns in different scenarios: *Human-Teacher and Machine-Learner*, *Machine-Leading*, *Human-Leading*, and *Human-Machine Collaborators* (RQ2). We used these properties and conceptualizations to contribute a theoretical interaction framework for human-machine relationships in NLP. Future research from HCI and NLP can use this framework to design and evaluate human-machine interaction for human and machine needs within their role and responsibility.

2 Related Work

Traditionally, NLP models are trained, fine-tuned, and tested on existing datasets before being deployed to solve real-world problems. While HCI research usually involves users designing, developing, implementing, and evaluating systems. To encourage collaboration between HCI and NLP, combining both approaches is required for successful 'HCI+NLP' applications (Heuer and Buschek, 2021). Further, Heuer and Buschek (2021) proposed five methods for designing and evaluating HCI+NLP Systems: user-centered NLP, co-creating NLP, crowdsourcing, and user models.

The human-machine interaction in NLP focuses on how people interact with machines and how NLP systems can be designed to support humans. In recent years, many researchers and practitioners have made significant advances in interactive NLP systems, such as text classification (Godbole et al., 2004), text summarization (Passali et al., 2021), semantic parsing and entity linking (Liang et al., 2020; He et al., 2016; Klie et al., 2020; Zhong and Chen, 2020), dialogue systems (Sanguinetti et al., 2020; Liu et al., 2018), topic modeling (Smith et al., 2018).

Wang et al. (2021) summarized recent human-in-the-loop NLP work based on their tasks, goals, human interactions, and feedback learning methods. According to Wang et al. (2021), a good human-in-the-loop NLP system must clearly communicate to humans what the model requires, provide user-friendly interfaces for collecting feedback, and effectively learn from it. For example, humans can provide various types of feedback, such as training data providers, annotations, and evaluators of the system's output to improve the model's performance, interpretability, or usability.

Also, it is vital to understand the constraints of collecting and interpreting human inputs cor-

rectly. Many design considerations influence the efficiency and effectiveness of interactive learning from human feedback (Cui et al., 2021). For example, noise caused by human error (such as when a human teacher fails to provide the conventional ground truth) could be challenging in designing human-machine interaction systems. In addition, data collected from humans may be poor quality Hsueh et al. (2009). Social bias Fiske (2019) can also be introduced automatically during data collection Garrido-Muñoz et al. (2021) and language model training, which emerging the need of developing fair and responsible models Hutchinson et al. (2020); Mehrabi et al. (2021). Therefore, it is critical to appropriately interpret collected human data and analyze the effects of various interaction types on learning outcomes (Cui et al., 2021).

3 Survey and Analysis Methods

We conducted a systematic literature review to understand the types of interactions humans and machines have within the NLP and HCI context. To select the targeted papers, we searched from ACL anthology (NLP database) and ACM Digital Library (HCI database) for articles that have been published over the last two years and included papers with keywords such as 'human-in-the-loop,' 'interactive,' 'collaborative,' 'active learning,' 'HCI + NLP,' 'human-machine,' and 'human-AI.' We also searched for workshops held in ACL or ACM conferences, such as HCI+NLP in EACL 2021 (Blodgett et al., 2021). Further, we conducted a backward reference search on the articles we found to identify any additional missing articles. This resulted in a total of 73 papers at the intersection of NLP and HCI. We narrowed down our search to only include articles that made algorithmic contributions, system contributions, or empirical contributions. We excluded papers that only contributed opinions, theories, surveys, or datasets. This resulted in a total of 54 papers.

Next, we included only papers that had simultaneous interactions between humans and machines. We discarded 21 articles that had no interaction, or the interaction between the human and machine was asynchronous. This resulted in a final set of 33 articles included in our analysis.

Further, We conducted a thematic analysis (Braun and Clarke, 2012) of our dataset. We began by reading through each article and taking notes on observations and insights regarding the inter-

Relationships	Papers
Human-Teacher, Machine-Learner	OUG Wiechmann et al. (2021), IUG Stiennon et al. (2020), IVN Jandot et al. (2016), IVM Wallace et al. (2019), IVM Liu et al. (2018), IVM Settles (2011), IVM Godbole et al. (2004)
Machine-Leading	OUG Khashabi et al. (2020), OUG Lawrence and Riezler (2018), IUG Lertvittayakumjorn et al. (2020), IUG He et al. (2016), IUN Liang et al. (2020), IVG Simard et al. (2014), IVM Lo and Lim (2020), IVM Smith et al. (2018), IVM Ross et al. (2021)
Human-Leading	SUN Bhat et al. (2021), SUN Rao and Daumé III (2018), SVG Kim et al. (2021), SVM Coenen et al. (2021), IVM Chung et al. (2022), IVM Passali et al. (2021)
Human-Machine Collaborators	OUG Kreutzer et al. (2018), OUN Khashabi et al. (2021), OVG Head et al. (2021), SUN Ashktorab et al. (2021), IVG Karmakharm et al. (2019), IVN Hancock et al. (2019), IVN Van Heerden and Bas (2021), IVN Klie et al. (2020), IVM Clark and Smith (2021), IVM Trivedi et al. (2019), IVM Kim et al. (2019),

Table 1: Human-Machine Relationships Mapping Interaction Properties: The properties of the interaction in each paper are coded by the first letter of their three interaction properties: **O/S/I** represents **One-time/Sequential/Iterative**. **U/V** represents **Unitary/Various**. **G/N/M** represents **GUI/NUI/MUI**.

actions between humans and machines. We used these notes to develop a set of guiding questions that helped us generate our initial codes:

- *What is the frequency of the interaction?*
- *What are the different ways humans can interact with machines?*
- *In what form does the interaction take place?*
- *Who starts the interaction?*
- *Who ends the interaction?*
- *Who benefits from the interaction?*

4 Findings

After iterating on our codes informed by the guiding questions, we conceptualized the codes into two major dimensions: 1) the properties of human-machine interaction based on the first three questions, and 2) the types of human-machine relationships based on the last three questions. Table 1 summarizes the mapping between our dimensions, codes, and dataset.

4.1 Properties of Interaction

Within our first dimension, we identified three major properties of how interactions happen, which address the first three guiding questions respectively: 1) continuity, 2) variety of interaction options, and 3) medium of interactions.

4.1.1 Continuity

Continuity measures the frequency of interaction which can be one-time, sequential, or iterative, to perform a single task.

One-time interaction is when the human-machine interaction is designed to happen just once, usually in active learning. For example, ActiveAnno (Wiechmann et al., 2021) offers annotation generator functionality: human annotators will manually label documents at the beginning; if it reaches a threshold of the annotation generator, it will trigger the machine annotation generator to label the remaining documents in the project.

Sequential interaction means that one agent acts first, and the other responds. Furthermore, the latter interactions are based on the previous exchanges. Over multiple rounds to complete the project, the interaction and the model will not update. For example, in the word guessing game to study the effects of communication directionality (Ashktorab et al., 2021), the machine and human play as giver and guesser by giving word hints and word guessing together to let the guess win the game in 10 rounds. The human player is playing with the same machine player in each game, but previous rounds influence the words given in later rounds.

Iterative interaction: the model performance improves through iterating over several interactions. For example, the Interactive Classification and Extraction (ICE) system enables humans to define the appropriate features through interactive features while humans can monitor their classifier’s progress (Simard et al., 2014). Researchers usually set up a limitation of the rounds of the interaction

for experiments, but they show that theoretically, the continuous interaction (either sequential or iterative interactions) can continue with no end.

4.1.2 Variety of Interaction Actions

Variety represents the number of ways humans can interact with machines. Some interactions are limited to a specific task (unitary), while some are flexible to multiple options (various).

Unitary interaction is a single option of the interaction action, like labeling data (Wiechmann et al., 2021), having humans perform one action to evaluate submissions (Khashabi et al., 2021), ranking candidate questions (Rao and Daumé III, 2018), giving natural language responses to query (Liang et al., 2020), or selecting an answer from the multiple-choice questions (Lertvittayakumjorn et al., 2020; He et al., 2016).

Various interaction is that there are multiple options for an interaction action. Humans can select what to do from multiple options, such as choosing continuation writing, rewriting, or filling in an interactive editor (Coenen et al., 2021). In CYOA (Clark and Smith, 2021), human has many options, such as deleting suggestions, submitting suggestions as-is, or writing a new suggestion.

4.1.3 Medium of Interactions

Medium in interactive NLP systems can be Graphical User Interfaces (GUI), Natural Language User Interfaces (NUI), or Mixed User Interfaces (MUI):

Graphical User Interfaces (GUI) allow humans to select given options or highlight text. Such as, in FIND (Lertvittayakumjorn et al., 2020), humans answer multiple-choice questions about whether a given word cloud is relevant to a class to disable irrelevant hidden features. Journalist-in-the-loop (Karmakharm et al., 2019) uses a web-based interface rumour analysis that takes user feedback. (Kreutzer et al., 2018) collects reinforcement signals from humans using a 5-star rating interface.

Natural Language User Interfaces (NUI) let the agent respond with natural language. For instance, GENIE (Khashabi et al., 2021) uses text generation tasks. Liang et al. (2020)'s ALICE utilizes contrastive natural language explanations to improve data efficiency in learning.

Mixed User Interfaces (MUI) contain both graphical and natural language interfaces. For example, in the Interactive NLP in Clinical Care

(Trivedi et al., 2019), physicians can add or remove highlighted predicted sentences in a report and write natural language feedback in a separate box. Also, Wordcraft allows humans to start writing with a prompt and change text and selection options for machine editing in the side GUI (Coenen et al., 2021). In CYOA (Clark and Smith, 2021), a human can write a storyline alone and delete suggestions, while models provide suggestions of the story, and then the human can submit the suggestion. Participants are asked to score on a Likert-scale and open-ended questions about the systems and suggestions they received after submitting their stories. In DUALIST (Settles, 2011) humans can label documents by clicking the appropriate class from the drop-down menus below each text, In addition, each column has a text box where users can "inject" domain knowledge by typing in random words. Users must click a large submit button at the top of the screen to retrain the classifier and get a new set of queries. Smoothed dictionary features (Jandot et al., 2016) use a methodology to solicit features from humans or teachers. In Passali et al. (2021) human has the option to view visualization and choose different decoding strategies.

4.2 Relationships of Human and Machine

For our second dimension, we conceptualized four relationships (each based on the last three guiding questions): 1) *Human-Teacher and Machine-Learner*, 2) *Machine-Leading*, 3) *Human-Leading*, and 4) *Human-Machine Collaborators*.

Human-Teacher and Machine-Learner When humans initiate the interaction, then machines learn from humans to mimic the task, after that, humans evaluate and give feedback on the machines' learning results in the end. Through this interaction, humans benefit from machines' automation in finishing tasks, and machines also benefit from learning to improve their performance.

For instance, humans annotate and label data at the beginning in ActiveAnno (Wiechmann et al. (2021)), ICE (Simard et al. (2014)), DUALIST (Settles, 2011) and HIClass (Godbole et al., 2004). Once the machine has learned enough from the human annotation, it can predict labels for the remaining documents. In Trick-Me-If-You-Can (Wallace et al., 2019), human authors are guided to break the model by writing adversarial questions, and the machine exposes the predictions and interpretations of the answers to hu-




Figure 1: Interaction Framework for Human-Machine relationships in NLP

mans. In smoothed dictionary features (Jandot et al., 2016), the machine elicits dictionary features from a teacher or human.

Machine-Leading Machines initiate interactions with their optimal competence, then humans respond with suggestions, and later, machines iterate on the task based on the intermediate suggestions. From the interaction, machines benefit from humans’ knowledge in improving their own capacity, while humans don’t earn anything from the process.

For instance, ALICE (Coenen et al., 2021), FIND (Lertvittayakumjorn et al., 2020), Interactive NLP in clinical care (Trivedi et al., 2019), and Human-in-The-Loop Parsing (He et al., 2016), have a baseline model to perform the corresponding NLP tasks. Then human steps into interactively selecting useful features, promoting machine learning efficiency. In Interactive Entity Linking (Lo and Lim, 2020), an instance chooses mentions for human annotation using an active learning approach. The goal is to improve entity linking accuracy while updating the embedding model. The machine triggers the interaction, the human assists the machine in the annotation process, and human annotation feedback is used to update the model. Finally, high machine accuracy determines the end of the interaction process.

Human-Leading Humans initiate the task, then machines give suggestions based on their pre-trained expertise, later, humans select the way they want to interact with machines. Via the interaction, humans gain help from machines, but machines don’t take any benefit from these interactions.

Like Wordcraft (Coenen et al. (2021)), human plays the lead role because human triggers the interaction by writing a prompt, selecting collaborative writing options, and eventually deciding what to write and when to end. In Towards Human-Centered Summarization (Passali et al., 2021), human actively participates and takes a leading role in the summarization process, such as deciding on the decoding strategies during the inference stage, choosing visualization color and viewing visualization of the attention weights, combining sentences from the various summaries to create a new one that can be used in fine-tuning the model.

Human-Machine Collaborators Either a human or machine initiates the task, then the other one gives the response. There is no explicit benefit for humans or machines during the interaction.

For example, in Ashktorab et al. (2021), humans and machines can play as giver and guesser and cooperate for a common goal: to help the guesser answer the right word. In GENIE (Khashabi et al., 2021), human annotators work with a machine to assess leaderboard submissions on various axes such as correctness, conciseness, and fluency and compare their responses to a variety of automated metrics. Journalist-in-the-loop (Karmakharm et al., 2019) is a rumor annotation service that continuously allows journalists to provide feedback on social media posts. The feedback is then used to improve the neural network-based rumor classification mode. Also, human-machine collaboration improves the journals’ productivity. In AfriKI (Van Heerden and Bas, 2021), human authors collaborate with machines to gener-




Figure 2: Human-Machine Relationships Mapping Interaction Properties

ate Afrikaans poetry through phrase selection and vertically arranging poetry lines. The collaboration promotes human creativity and also improves the robustness of the model. From Zero to Hero (Klie et al., 2020), an interactive machine learning annotation supports that guides the user in locating text entities and deciding on the correct knowledge base entries. As a result, the annotation speed and quality improve tremendously and reduce human cognitive load. In CYOA (Clark and Smith, 2021), humans and machines take turns writing a story, and two models generate suggestions for the following line after the human writes the story’s first line. The human suggestions provide interpretability of the model performance and improve human creativity.

5 Discussion

Based on our findings, we developed an interaction framework for the four types of human-machine relationships in one interaction cycle (see Figure 1). The "human-in-the-loop" concept has become popular in NLP. But it is too broad to describe the nuance in similar but different human-machine interactions. Our interaction framework depicts human and machine actions under different roles.

It is important for NLP and HCI practitioners to define human-machine interactions when designing interactive NLP systems because it builds up a clearer understanding of the human and machine’s strengths and weaknesses. We can use the identi-

fied strengths to complement the weaknesses. For example, in *Human-Teacher and Machine-Learner interactions*, the strengths of automation in annotation tasks can be leveraged to complete a task more quickly, while the human’s strength in assessment can validate the machine’s output. For *Human-Leading* or *Machine-Leading*, the leader may have limited knowledge and capacity that could be complemented with external expertise.

Next, we share how NLP and HCI practitioners can use our conceptual framework in practice. While *Human-Leading* is similar to *Human-Teacher, Machine Learner* (since both are driven by humans who take initial and final action with machines’ intermediate response), the framework captures the different statuses of machines and the corresponding actions of humans and machines. As a learner, the machine is a novice at the beginning and is launched by mimicking humans finishing their remaining tasks. However, the machines in the *Human-Leading* relationship have their own expertise to offer suggestions. Accordingly, this also leads to different humans’ final actions, i.e., humans as teachers will validate machines’ learning results, while humans would choose machines’ suggestions based on personal interests during human-leading interaction.

Additionally, the framework can be used as a guideline to check whether the interaction design is appropriate. For instance, the machine directly overwrites a human’s control will be problematic when *Human-Leading* is required. For instance,

someone is writing in colloquial English, using 'ain't' as a general preverbal negator (Rickford, 1999). A machine trained in standard English will detect this as an error and would like to correct it, but it violates the writer's intention.

Figure 2 shows some trends of interactions: *Human-Leading* relationship usually happens with sequential interactions; *Machine-Leading* relationship uses more GUI and has iterative interaction; *Human-Machine Collaborators* mainly use NUI; *Human-Teacher*, *Machine-Learner* relationship usually happens iteratively as well.

6 Limitations

From the papers we reviewed, we didn't find any interaction that machines validate humans' actions. This might be because humans overall are more trustworthy than machines and may be influenced by the range of literature we reviewed. Additionally, 'collaboration' can sometimes be interchangeable with 'interaction' based on their semantic meanings. But we strive to name our precise definitions of each relationship with clear and straightforward enough abstract names.

In addition, we initially coded 'what are the outcomes of the research paper' to understand how interaction can influence the research outcome, such as efficiency, creativity, etc. But it didn't synthesize sufficiently with other codes. More importantly, the research outcome is usually the goal of those research papers so we couldn't derive the causal relationship from interaction design to the research outcome. However, this guides us in the future to study how we can manipulate desired research improvement by designing human-machine interaction.

7 Conclusion

Humans and machines interact with each other in a variety of ways. For example, humans may be involved in providing input to a machine learning algorithm, labeling data for training purposes, or evaluating the output of a machine learning system. Additionally, humans may interact with machine learning systems through natural language interfaces, such as chatbots or virtual assistants. We contribute an interaction framework for human-machine interactions through a systematic literature review and thematic analysis, which conceptualizes four human-machine relationships based on three different interaction properties, to help re-

searchers and practitioners better understand and manage human-machine interactions in NLP.

References

- Zahra Ashktorab, Casey Dugan, James Johnson, Qian Pan, Wei Zhang, Sadhana Kumaravel, and Murray Campbell. 2021. Effects of communication directionality and ai agent differences in human-ai interaction. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–15.
- Rohan Badlani, Nishit Asnani, and Manan Rai. 2019. Disambiguating sentiment: An ensemble of humour, sarcasm, and hate speech features for sentiment classification. *W-NUT*, 2019:337–345.
- Advait Bhat, Saaket Agashe, and Anirudha Joshi. 2021. How do people interact with biased text prediction models while writing? In *Proceedings of the First Workshop on Bridging Human-Computer Interaction and Natural Language Processing*, pages 116–121.
- Su Lin Blodgett, Michael Madaio, Brendan O'Connor, Hanna Wallach, and Qian Yang. 2021. Proceedings of the first workshop on bridging human-computer interaction and natural language processing. In *Proceedings of the First Workshop on Bridging Human-Computer Interaction and Natural Language Processing*.
- Virginia Braun and Victoria Clarke. 2012. *Thematic analysis*. American Psychological Association.
- Mei-Hua Chen, Shih-Ting Huang, Hung-Ting Hsieh, Ting-Hui Kao, and Jason S Chang. 2012. Flow: a first-language-oriented writing assistant system. In *Proceedings of the ACL 2012 System Demonstrations*, pages 157–162.
- John Joon Young Chung, Wooseok Kim, Kang Min Yoo, Hwaran Lee, Eytan Adar, and Minsuk Chang. 2022. *TaleBrush: Sketching Stories with Generative Pretrained Language Models*. Association for Computing Machinery, New York, NY, USA.
- Elizabeth Clark and Noah A Smith. 2021. Choose your own adventure: Paired suggestions in collaborative writing for evaluating story generation models. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3566–3575.
- Andy Coenen, Luke Davis, Daphne Ippolito, Emily Reif, and Ann Yuan. 2021. Wordcraft: a human-ai collaborative editor for story writing. *arXiv preprint arXiv:2107.07430*.
- Yuchen Cui, Pallavi Koppol, Henny Admoni, Scott Niekum, Reid G Simmons, Aaron Steinfeld, and Tesca Fitzgerald. 2021. Understanding the relationship between interactions and outcomes in human-in-the-loop machine learning. In *IJCAI*, pages 4382–4391.

- Giovanni Da San Martino, Seunghak Yu, Alberto Barrón-Cedeno, Rostislav Petrov, and Preslav Nakov. 2019. Fine-grained analysis of propaganda in news article. 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)*, pages 5636–5646.
- Wanyu Du, Zae Myung Kim, Vipul Raheja, Dhruv Kumar, and Dongyeop Kang. 2022. Read, revise, repeat: A system demonstration for human-in-the-loop iterative text revision. *arXiv preprint arXiv:2204.03685*.
- Manaal Faruqui, Ellie Pavlick, Ian Tenney, and Dipanjan Das. 2018. Wikiatomicedits: A multilingual corpus of wikipedia edits for modeling language and discourse. *arXiv preprint arXiv:1808.09422*.
- Susan T Fiske. 2019. 13.4 prejudice, discrimination, and stereotyping. *Introduction to Psychology*.
- Ismael Garrido-Muñoz, Arturo Montejo-Ráez, Fernando Martínez-Santiago, and L Alfonso Ureña-López. 2021. A survey on bias in deep nlp. *Applied Sciences*, 11(7):3184.
- Shantanu Godbole, Abhay Harpale, Sunita Sarawagi, and Soumen Chakrabarti. 2004. Document classification through interactive supervision of document and term labels. In *European Conference on Principles of Data Mining and Knowledge Discovery*, pages 185–196. Springer.
- Jiatao Gu, Graham Neubig, Kyunghyun Cho, and Victor OK Li. 2016. Learning to translate in real-time with neural machine translation. *arXiv preprint arXiv:1610.00388*.
- 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*, pages 3667–3684.
- Luheng He, Julian Michael, Mike Lewis, and Luke Zettlemoyer. 2016. Human-in-the-loop parsing. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2337–2342.
- Andrew Head, Kyle Lo, Dongyeop Kang, Raymond Fok, Sam Skjonsberg, Daniel S. Weld, and Marti A. Hearst. 2021. [Augmenting scientific papers with just-in-time, position-sensitive definitions of terms and symbols](#).
- Hendrik Heuer and Daniel Buschek. 2021. Methods for the design and evaluation of hci+ nlp systems. In *Proceedings of the First Workshop on Bridging Human-Computer Interaction and Natural Language Processing*, pages 28–33.
- Pei-Yun Hsueh, Prem Melville, and Vikas Sindhwani. 2009. Data quality from crowdsourcing: a study of annotation selection criteria. In *Proceedings of the NAACL HLT 2009 workshop on active learning for natural language processing*, pages 27–35.
- Ben Hutchinson, Vinodkumar Prabhakaran, Emily Denton, Kellie Webster, Yu Zhong, and Stephen DeNuyt. 2020. Social biases in nlp models as barriers for persons with disabilities. *arXiv preprint arXiv:2005.00813*.
- Camille Jandot, Patrice Simard, Max Chickering, David Grangier, and Jina Suh. 2016. Interactive semantic featurer for text classification. *arXiv preprint arXiv:1606.07545*.
- Kenneth Joseph, Sarah Shugars, Ryan Gallagher, Jon Green, Alexi Quintana Mathé, Zijian An, and David Lazer. 2021. (mis) alignment between stance expressed in social media data and public opinion surveys. *arXiv preprint arXiv:2109.01762*.
- Twin Karmakharm, Nikolaos Aletras, and Kalina Bontcheva. 2019. Journalist-in-the-loop: Continuous learning as a service for rumour analysis. 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): System Demonstrations*, pages 115–120.
- Daniel Khashabi, Tushar Khot, and Ashish Sabharwal. 2020. More bang for your buck: Natural perturbation for robust question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 163–170.
- Daniel Khashabi, Gabriel Stanovsky, Jonathan Bragg, Nicholas Lourie, Jungo Kasai, Yejin Choi, Noah A Smith, and Daniel S Weld. 2021. Genie: A leaderboard for human-in-the-loop evaluation of text generation. *arXiv preprint arXiv:2101.06561*.
- Hannah Kim, Dongjin Choi, Barry Drake, Alex Endert, and Haesun Park. 2019. Topicsifter: Interactive search space reduction through targeted topic modeling. In *2019 IEEE Conference on Visual Analytics Science and Technology (VAST)*, pages 35–45. IEEE.
- Jeonghwan Kim, Junmo Kang, Suwon Shin, and Sung-Hyon Myaeng. 2021. Can you distinguish truthful from fake reviews? user analysis and assistance tool for fake review detection. In *Proceedings of the First Workshop on Bridging Human-Computer Interaction and Natural Language Processing*, pages 53–59.
- Jan-Christoph Klie, Richard Eckart de Castilho, and Iryna Gurevych. 2020. From zero to hero: Human-in-the-loop entity linking in low resource domains. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6982–6993.

- Julia Kreutzer, Shahram Khadivi, Evgeny Matusov, and Stefan Riezler. 2018. Can neural machine translation be improved with user feedback? *NAACL HLT 2018*, pages 92–105.
- Carolin Lawrence and Stefan Riezler. 2018. Counterfactual learning from human proofreading feedback for semantic parsing. *arXiv preprint arXiv:1811.12239*.
- Piyawat Lertvittayakumjorn, Lucia Specia, and Francesca Toni. 2020. Find: human-in-the-loop debugging deep text classifiers. *arXiv preprint arXiv:2010.04987*.
- Weixin Liang, James Zou, and Zhou Yu. 2020. Alice: Active learning with contrastive natural language explanations. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4380–4391.
- Bing Liu, Gokhan Tur, Dilek Hakkani-Tur, Pararth Shah, and Larry Heck. 2018. Dialogue learning with human teaching and feedback in end-to-end trainable task-oriented dialogue systems. *arXiv preprint arXiv:1804.06512*.
- Pei-Chi Lo and Ee-Peng Lim. 2020. Interactive entity linking using entity-word representations. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1801–1804.
- Pattie Maes. 1995. Agents that reduce work and information overload. In *Readings in human-computer interaction*, pages 811–821. Elsevier.
- Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2021. A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)*, 54(6):1–35.
- Tatiana Passali, Alexios Gidiotis, Efstathios Chatzikiriakidis, and Grigorios Tsoumakas. 2021. Towards human-centered summarization: A case study on financial news. In *Proceedings of the First Workshop on Bridging Human-Computer Interaction and Natural Language Processing*, pages 21–27.
- Sudha Rao and Hal Daumé III. 2018. Learning to ask good questions: Ranking clarification questions using neural expected value of perfect information. *arXiv preprint arXiv:1805.04655*.
- John Russell Rickford. 1999. *African American vernacular English: Features, evolution, educational implications*. Wiley-Blackwell.
- Andrew Ross, Nina Chen, Elisa Zhao Hang, Elena L Glassman, and Finale Doshi-Velez. 2021. Evaluating the interpretability of generative models by interactive reconstruction. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–15.
- Manuela Sanguinetti, Alessandro Mazzei, Viviana Patti, Scalerandi Marco, Dario Mana, and Rossana Simone. 2020. Annotating errors and emotions in human-chatbot interactions in italian. In *The 14th Linguistic Annotation Workshop*, pages 1–12. Association for Computational Linguistics.
- Burr Settles. 2011. Closing the loop: Fast, interactive semi-supervised annotation with queries on features and instances. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 1467–1478.
- Ben Shneiderman. 2020. Human-centered artificial intelligence: Reliable, safe & trustworthy. *International Journal of Human-Computer Interaction*, 36(6):495–504.
- Patrice Simard, David Chickering, Aparna Lakshmiratan, Denis Charles, Léon Bottou, Carlos Garcia Jurado Suarez, David Grangier, Saleema Amershi, Johan Verwey, and Jina Suh. 2014. Ice: enabling non-experts to build models interactively for large-scale lopsided problems. *arXiv preprint arXiv:1409.4814*.
- Alison Smith, Varun Kumar, Jordan Boyd-Graber, Kevin Seppi, and Leah Findlater. 2018. Closing the loop: User-centered design and evaluation of a human-in-the-loop topic modeling system. In *23rd International Conference on Intelligent User Interfaces*, pages 293–304.
- Nisan Stiennon, Long Ouyang, Jeff Wu, Daniel M Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul Christiano. 2020. Learning to summarize from human feedback. *arXiv preprint arXiv:2009.01325*.
- Gaurav Trivedi, Esmael R Dadashzadeh, Robert M Handzel, Wendy W Chapman, Shyam Visweswaran, and Harry Hochheiser. 2019. Interactive nlp in clinical care: identifying incidental findings in radiology reports. *Applied clinical informatics*, 10(04):655–669.
- Imke Van Heerden and Anil Bas. 2021. Afriki: Machine-in-the-loop afrikaans poetry generation. *arXiv preprint arXiv:2103.16190*.
- Eric Wallace, Pedro Rodriguez, Shi Feng, Ikuya Yamada, and Jordan Boyd-Graber. 2019. Trick me if you can: Human-in-the-loop generation of adversarial examples for question answering. *Transactions of the Association for Computational Linguistics*, 7:387–401.
- Zijie J Wang, Dongjin Choi, Shenyu Xu, and Diyi Yang. 2021. Putting humans in the natural language processing loop: A survey. *arXiv preprint arXiv:2103.04044*.
- Max Wiechmann, Seid Muhie Yimam, and Chris Biemann. 2021. Activeanno: General-purpose document-level annotation tool with active learning integration. In *Proceedings of the 2021 Conference of the North American Chapter of the Association*

for Computational Linguistics: Human Language Technologies: Demonstrations, pages 99–105.

Xingjiao Wu, Luwei Xiao, Yixuan Sun, Junhang Zhang, Tianlong Ma, and Liang He. 2022. A survey of human-in-the-loop for machine learning. *Future Generation Computer Systems*.

Zexuan Zhong and Danqi Chen. 2020. A frustratingly easy approach for entity and relation extraction. *arXiv preprint arXiv:2010.12812*.