

Hands-On Tutorial: Labeling with LLM and Human-in-the-Loop

Ekaterina Artemova¹ Akim Tsvigun² Dominik Schlechtweg³
Natalia Fedorova¹ Sergei Tilga¹ Boris Obmoroshev¹

¹ Toloka AI, ² University of Amsterdam, ³University of Stuttgart

Correspondence: katya-art@toloka.ai

Abstract

Training and deploying machine learning models relies on a large amount of human-annotated data. As human labeling becomes increasingly expensive and time-consuming, recent research has developed multiple strategies to speed up annotation and reduce costs and human workload: generating synthetic training data, active learning, and hybrid labeling. This tutorial is oriented toward practical applications: we will present the basics of each strategy, highlight their benefits and limitations, and discuss in detail real-life case studies. Additionally, we will walk through best practices for managing human annotators and controlling the quality of the final dataset. The tutorial includes a hands-on workshop, where attendees will be guided in implementing a hybrid annotation setup. This tutorial is designed for NLP practitioners from both research and industry backgrounds who are involved in or interested in optimizing data labeling projects.

Keywords: language models, data labeling, active learning, hybrid labeling, dataset quality, synthetic data

1. Introduction

Labeled data is the primary tool for providing supervision to learnable systems. In the last decade, labeling data on crowdsourcing platforms or by hiring expert annotators has become the de facto standard approach to creating training data. However, human labeling can be time- and resource-consuming, hindering the ability to scale data labeling projects. Recent research has developed multiple strategies to use machine learning methods with human-in-the-loop to obtain larger amounts of high-quality data in shorter periods.

This tutorial explores recent strategies to speed up annotation and reduce costs and human workload. Emerging automatic approaches use language models (LMs) to *generate synthetic training data* (§3.2). By fine-tuning or prompting LMs, we can create data that closely mirrors the distribution of the target dataset. This synthetic dataset can be further used to train models from scratch or to augment existing datasets, particularly in low-resource settings. However, this approach may result in biased or low-quality data, highlighting the need for human efforts in labeling.

Automatic approaches for data labeling cannot replace entirely human annotation. Limitations of automatic approaches (§3.6) lie in the nature of the language models used and are associated with bias in generated data, the difficulty in accurately mirroring the nuances of real-world data, and the challenges in processing subjective, culture-specific, or knowledge-intensive domains.

Maintaining a human-in-the-loop for data labeling is essential to ensure high-quality and accurate annotations. By involving humans in the labeling

process, we can continually adapt to new data, and address edge cases. *Active learning* (§3.3) selects the most informative instances for human labeling to maximize model performance. *Hybrid labeling* (§3.5) combines human and model efforts, with the model handling simple instances and humans handling complex ones. These approaches reduce the amount of labeled data needed, lowering the annotation budget and improving data accuracy.

The quality of the resulting dataset depends on automatic methods and effective management of human workers. We will cover best practices in writing guidelines and using quality control methods (§3.4). Additionally, we emphasize the importance of maintaining ethical standards when working with human annotators, ensuring fair payment, continuous communication, and avoiding conditions harmful to their well-being.

The tutorial includes a 30-minute hands-on session (§3.7), where the instructor will cover the setup of a hybrid data annotation project step-by-step. This session will feature a code walk-through and a real-life case study demonstrating how combining human and LM-based labeling yields high-quality data affordably and quickly.

Aimed at NLP practitioners involved in data labeling projects, this tutorial focuses on practical applications. Each section of the tutorial will feature a real-life case study to demonstrate labeling approaches. These case studies, drawn from the instructor's research or industrial experience, provide practical insights and share effective strategies.

By the end, participants will understand how to optimize their annotation workflows for natural language processing tasks, such as text or token classification.

2. Outline

Part 1: Introduction (20 min)

- **Introduction:** discussing the need for large amounts of labeled data
- **Key concepts:** annotation schema and guidelines, quality control, inter-annotator agreement.
- **Annotation lifecycle:** the typical setup of an annotation project; transition from human workers to LLMs.
- **Costs and duration of an annotation project:** comparison of projects with and without using LLMs.
- **Benefits and limitations of LLMs:** reduced time and labor but lower quality.
- **Conclusion:** motivating the need for LLMs and human-in-the-loop approaches, overview of next sections.

Part 2: LM workflows (30 min)

Overview This section will demonstrate best practices for common workflows involving language models (LMs) and large language models (LLMs). These workflows aim to (i) create efficient LMs with acceptable performance optimized for labeling data, and (ii) generate synthetic data for data augmentation.

- **Models:** large-scale LMs, such as FLAN (Longpre et al., 2023), Llama 3.1 (team, 2024) or XLM-R (Conneau et al., 2020)
- **Performance optimization:** prompt engineering; parameter-efficient fine-tuning (Lialin et al., 2023).
- **Efficiency optimization:** quantization (Tao et al., 2022); pruning (Wang et al., 2020); knowledge distillation (Gu et al.).
- **Synthetic train data generation:** attribute-aware generation (Yu et al., 2024), zero-shot generation (Ye et al., 2022), model collapse caused by training on synthetic data only (Shumailov et al., 2023).
- **Case study:** Sentiment analysis in media monitoring context.

Part 3: Active learning with LMs (40 min)

Overview This section presents Active learning (AL) in data annotation. We discuss key strategies for both generative and non-generative AL, their applications, advantages, and limitations.

- **Introduction to AL:** definition and core principles of AL; benefits of AL in reducing annotation costs and improving model performance; AL workflow.
- **AL strategies for text classification:** basic strategies (Least Confidence (Lewis and Gale,

1994), Breaking Ties (Luo et al., 2005), etc.); gradient-based strategy BADGE (Ash et al., 2020), contrastive AL (Margatina et al., 2021).

- **Generative AL with LLMs:** leveraging LLMs to generate informative samples; controlling LLM-generated samples' diversity and relevance.
- **Human-based vs. LLM-based annotation in AL:** strengths and weaknesses; strategies for combining human and LLM annotations; quality control measures for AL pipelines.
- **Efficient AL:** methods to accelerate the acquisition stage; batched-mode AL.
- **Case study:** AL for text classification in the law domain.

Part 4: Quality control and managing human workers (30 min)

Overview This section focuses on quality control and best practices in working with human annotators.

- **Pipeline setup:** sourcing & selection of supply of annotators (including smart matching for the general crowd based on skills and demographics); onboarding, exams & training of annotators; control tasks (Sabou et al., 2014; Gordon et al., 2022).
- **Data production:** dynamic overlaps; platform tools for experts (time keeping, copy-paste & word count limitations, spell check, etc.); anti-fraud rules (Drutsa et al., 2020), regular audits.
- **Data acceptance and working with data:** automatic metrics; audit by annotation or domain experts; aggregation tools for general crowd tasks; ML assessment of dataset quality; LLM-checks (Klie et al., 2024).
- **Managing human annotation work:** writing annotation guidelines; psychological characteristics of doing annotation tasks; inter-annotator agreement; communication with the workers. (Saal et al., 1980; Tseng et al., 2020).

Part 5: Hybrid pipelines (40 min)

Overview This section presents developing hybrid pipelines, e.g. effectively combining human and model labeling to achieve the best balance of quality, cost, and speed.

- **Model confidence estimation:** prediction quality depending on various confidence levels; confidence estimation in open-source and proprietary models (Geng et al., 2023; Mukherjee and Awadallah, 2020).
- **Aggregation of model and human responses:** setting confidence thresholds to route tasks between human and model labeling; combining model and human responses

using overlaps, human response confidence, weighting techniques (Wei et al., 2023).

- **Balancing between quality and automation:** strategies for balancing quality cost, and speed of labeling through confidence thresholding; optimal combinations to maximize cost-efficiency and quality for various setups.
- **Hybrid pipelines:** quality assessment; continuous improvement of model performance (Wang et al., 2024b).
- **Case study:** hybrid labeling applications for product recommendations and ad search relevance.

Part 6: Limitations (20 min)

Overview This section addresses the challenges of labeling tasks with LLMs, the various reasons behind these difficulties, and future research directions to lift these limitations.

- **Tasks challenging to label with LLMs:** tasks which require subjective evaluation (Dillion et al., 2023), deep domain or culture-specific knowledge (Thirunavukarasu et al., 2023; Wang et al., 2024a).
- **Limitations of LLMs in labeling:** biases towards various social groups (Gallegos et al., 2024); compute requirements (Kaplan et al., 2020); inevitable tendency to hallucinate inaccurate labels (Xu et al., 2024); model collapse caused by training on synthetic data (Shumailov et al., 2024).

Part 7: Hands-on session: Hybrid data annotation (30 min)

Overview In this hands-on session, we will implement a hybrid approach on a real-world dataset and demonstrate improvements in annotation quality.

- **Task:** implementing a hybrid approach on a real-world dataset to show how combining human labeling with LLM-generated labels can produce high-quality annotation.
- **Data:** an open dataset for multiclass classification with high-quality expert annotations as ground truth.
- **Human labeling:** obtain noisy labels via Toloka's crowd platform.
- **Labeling with LLMs:** (i) prompt LLMs in zero-shot and few-shot setups, and (ii) fine-tune a custom LLM on a held-out dataset portion for label generation.
- **Hybrid labeling:** demonstrating strategies from Part 5 to combine LLM and human labels, evaluating their performance wrt ground truth.

Setup The demonstration will be conducted using Google Colab Notebooks. All materials will be

published in advance for the audience to download and reproduce. To save time, the crowdsourcing annotation and the model training will be performed beforehand. The annotated dataset and annotation guidelines will be published for further uptake. Attendees can follow along without needing to code independently.

3. Tutorial Information

Tutorial type and duration This is a *cutting-edge, half-day, in-person* tutorial. The increasing demand for labeled data justifies discussing automatic and hybrid approaches for data annotation. The widespread adoption of LLM-based technologies introduces novel data labeling approaches that accelerate annotation, reduce human workload in redundant tasks, and allow human workers to focus on complex instances. This tutorial differs from prior tutorials on data labeling (Suhr et al., 2021; Drutsa et al., 2021), which deal with labeling by humans only, and from prior tutorials on human-AI interaction (Yang et al., 2024; Bondi-Kelly et al., 2023), which focus on a broad range of tasks in which AI and humans can interact.

Tutorial audience and pre-requisites This tutorial is aimed at NLP practitioners who develop novel datasets and data collection protocols. While it provides an overview of the current best practices and state-of-the-art approaches in hybrid labeling with LLMs and human-in-the-loop, it is designed to fit an entry-level audience. The tutorial assumes a basic understanding of supervised classification, such as learning from labeled data and the concept of a classifier's confidence. Understanding of current large language models is recommended but not necessary. We will introduce the necessary terms and notations in each section of the tutorial. The final part of the tutorial includes a high-level code walkthrough for those interested in implementing hybrid approaches on their own.

Software Throughout the tutorial, we will use the Python language and Google Colab with GPU access as an environment.

Approximate count We expect this tutorial to attract over 100 attendees, reflecting the increasing demand for labeled data in both industrial and research settings.

Tutorial materials Tutorial resources are made available online for further uptake¹.

4. Diversity Considerations

The topic of the tutorial is linked to the ultimate goal of achieving a representation of diversity in NLP. We place special emphasis on the limitations of current

¹toloka.ai/coling-2025-human-w-llm-tutorial

automatic approaches to data labeling, particularly in tasks where the diversity of annotator representation is important. We discuss strategies to ensure ethical and fair use of manual labor in data annotation that guarantees better work conditions for annotators. Finally, we emphasize the importance of efficient approaches that reduce compute load.

The presenters are diverse in terms of gender, language, career stages, background, occupation, affiliation, and location.

5. Reading List

1. **Labeling with humans:** handbook (Ide and Pustejovsky, 2017), quality control with CrowdKit (Ustalov et al., 2024)
2. **Synthetic data and labeling with LLMs:** surveys (Honovich et al., 2023; Tan et al., 2024),
3. **Active learning:** surveys (Schröder and Niekler, 2020; Zhang et al., 2022), active learning with LLMs (Dor et al., 2020; Zhang et al., 2023)
4. **Hybrid labeling:** surveys (Wang et al., 2021; Wu et al., 2022), blog entry (Tilga, 2024)

6. Presenters

Ekaterina Artemova is a Senior Research Scientist at Toloka AI, holding a PhD from FRC CSC RAS. Her research focuses on data-centric NLP, specifically benchmarking strategies, low-resource settings, and LLM evaluation. Ekaterina has published in leading NLP and AI conferences and journals, including NLP, LRE, EMNLP, and *ACL. She has co-organized the 1st NLP Power! workshop at ACL '22, a tutorial on artificial text detection at INLG '22, and multiple shared tasks at the Dialogue conference.

Akim Tsvigun is a Senior ML Architect at Nebius and an MSc student at the University of Amsterdam. He has a strong publication record in top-tier NLP conferences, including ACL, EMNLP, and NAACL. His research primarily focuses on uncertainty estimation and active learning in natural language processing tasks, with an emphasis on improving model robustness and efficiency.

Sergei Tilga is the Head of R&D at Toloka AI, with over 7 years of industrial experience applying ML models and AI systems in consumer products and production processes. Sergei obtained a Master's degree from MSU. His research interests encompass AI model alignment, LLM evaluation, and data quality.

Natalia Fedorova is the Head of Academic Programs at Toloka. She is in charge of establishing long-term collaboration with universities, research labs, and educational organizations. Natalia co-organized a series of workshops and tutorials about crowdsourcing at NeurIPS, RecSys, ECIR, and WSDM conferences. She has a Master's degree in social psychology. Her research interests include the efficiency of education, the role of personality in education, role of humans in AI development.

Dominik Schlechtweg is an independent research group leader at the Institute for Natural Language Processing. He leads the Stuttgart group of the 6-year research program 'Change is Key!' where his group is responsible for the automation of semantic annotation processes and the application of such processes to the task of detecting semantic changes. His group develops several online text annotation systems. He organized several shared tasks on semantic change detection including SemEval-2020 Task 1 and created various Word-in-Context datasets.

Boris Obmoroshev is the Data and R&D Director at Toloka, with over 8 years of industrial experience in practical Data Science applications. He leads a team of Data Science and Machine Learning engineers, focusing primarily on adopting modern natural language processing technologies for large-scale data labeling projects. His main research interests are data quality and AI model evaluation. Boris obtained a Master's degree from MSU.

7. Ethics Statement

During the tutorial, we will address the limitations of using LLMs for data annotation, such as their predisposition to social biases and significant compute requirements. We will highlight the importance of respectful and responsible use of human annotators, ensuring fair compensation, informed consent, clear instructions, and reasonable working hours to avoid exploitation in any form.

While we acknowledge the risks of over-relying on automated approaches for data labeling, we believe the benefits of hybrid strategies outweigh these risks. This tutorial emphasizes leveraging a balanced mix of human and automated labeling techniques to ensure higher quality and fairer data production.

Use of AI-assistants We improve and proofread the text of this paper using Grammarly² to correct grammatical, spelling, and style errors and paraphrasing sentences. Therefore, specific segments

²[grammarly.com](https://www.grammarly.com)

of our proposal can be detected as AI-generated, AI-edited, or human-AI-generated.

8. References

- Jordan T. Ash, Chicheng Zhang, Akshay Krishnamurthy, John Langford, and Alekh Agarwal. 2020. [Deep batch active learning by diverse, uncertain gradient lower bounds](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Elizabeth Bondi-Kelly, Krishnamurthy Dvijotham, and Matthew Taylor. 2023. [How to work with real humans in human-ai systems](#). NeurIPS 2023.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Édouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451.
- Danica Dillion, Niket Tandon, Yuling Gu, and Kurt Gray. 2023. Can ai language models replace human participants? *Trends in Cognitive Sciences*, 27(7):597–600.
- Liat Ein Dor, Alon Halfon, Ariel Gera, Eyal Shnarch, Lena Dankin, Leshem Choshen, Marina Danilevsky, Ranit Aharonov, Yoav Katz, and Noam Slonim. 2020. Active learning for bert: an empirical study. In *Proceedings of the 2020 conference on empirical methods in natural language processing (EMNLP)*, pages 7949–7962.
- Alexey Drutsa, Valentina Fedorova, Dmitry Ustalov, Olga Megorskaya, Evfrosiniya Zermnova, and Daria Baidakova. 2020. Crowdsourcing practice for efficient data labeling: Aggregation, incremental relabeling, and pricing. In *Proceedings of the 2020 ACM SIGMOD international conference on management of data*, pages 2623–2627.
- Alexey Drutsa, Dmitry Ustalov, Valentina Fedorova, Olga Megorskaya, and Daria Baidakova. 2021. [Crowdsourcing natural language data at scale: A hands-on tutorial](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Tutorials*, pages 25–30, Online. Association for Computational Linguistics.
- Isabel O Gallegos, Ryan A Rossi, Joe Barrow, Md Mehrab Tanjim, Sungchul Kim, Franck Deroncourt, Tong Yu, Ruiyi Zhang, and Nesreen K Ahmed. 2024. Bias and fairness in large language models: A survey. *Computational Linguistics*, pages 1–79.
- Jiahui Geng, Fengyu Cai, Yuxia Wang, Heinz Koepl, Preslav Nakov, and Iryna Gurevych. 2023. A survey of language model confidence estimation and calibration. *arXiv preprint arXiv:2311.08298*.
- Mitchell L Gordon, Michelle S Lam, Joon Sung Park, Kayur Patel, Jeff Hancock, Tatsunori Hashimoto, and Michael S Bernstein. 2022. Jury learning: Integrating dissenting voices into machine learning models. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pages 1–19.
- Yuxian Gu, Li Dong, Furu Wei, and Minlie Huang. Minillm: Knowledge distillation of large language models. In *The Twelfth International Conference on Learning Representations*.
- Or Honovich, Thomas Scialom, Omer Levy, and Timo Schick. 2023. [Unnatural instructions: Tuning language models with \(almost\) no human labor](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14409–14428, Toronto, Canada. Association for Computational Linguistics.
- Nancy Ide and James Pustejovsky. 2017. *Introduction: The Handbook of Linguistic Annotation*. Springer.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*.
- Jan-Christoph Klie, Richard Eckart de Castilho, and Iryna Gurevych. 2024. Analyzing dataset annotation quality management in the wild. *Computational Linguistics*, pages 1–48.
- David D. Lewis and William A. Gale. 1994. [A sequential algorithm for training text classifiers](#). In *Proceedings of the 17th Annual International ACM-SIGIR Conference on Research and Development in Information Retrieval. Dublin, Ireland, 3-6 July 1994 (Special Issue of the SIGIR Forum)*, pages 3–12. ACM/Springer.
- Vladislav Lialin, Vijeta Deshpande, and Anna Rumshisky. 2023. Scaling down to scale up: A guide to parameter-efficient fine-tuning. *arXiv preprint arXiv:2303.15647*.

- Shayne Longpre, Le Hou, Tu Vu, Albert Webson, Hyung Won Chung, Yi Tay, Denny Zhou, Quoc V Le, Barret Zoph, Jason Wei, et al. 2023. The FLAN collection: Designing data and methods for effective instruction tuning. In *International Conference on Machine Learning*, pages 22631–22648. PMLR.
- Tong Luo, Kurt Kramer, Dmitry B. Goldgof, Lawrence O. Hall, Scott Samson, Andrew Remsen, and Thomas Hopkins. 2005. [Active learning to recognize multiple types of plankton](#). *J. Mach. Learn. Res.*, 6:589–613.
- Katerina Margatina, Giorgos Vernikos, Loïc Barraud, and Nikolaos Aletras. 2021. [Active learning by acquiring contrastive examples](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 650–663. Association for Computational Linguistics.
- Subhabrata Mukherjee and Ahmed Hassan Awadallah. 2020. Uncertainty-aware self-training for text classification with few labels. *arXiv preprint arXiv:2006.15315*.
- Frank E Saal, Ronald G Downey, and Mary A Lacey. 1980. Rating the ratings: Assessing the psychometric quality of rating data. *Psychological bulletin*, 88(2):413.
- Marta Sabou, Kalina Bontcheva, Leon Derczynski, and Arno Scharl. 2014. Corpus annotation through crowdsourcing: Towards best practice guidelines. In *LREC*, pages 859–866. Citeseer.
- Christopher Schröder and Andreas Niekler. 2020. A survey of active learning for text classification using deep neural networks. *arXiv preprint arXiv:2008.07267*.
- Iliia Shumailov, Zakhar Shumaylov, Yiren Zhao, Yarin Gal, Nicolas Papernot, and Ross Anderson. 2023. The curse of recursion: Training on generated data makes models forget. *arXiv preprint arXiv:2305.17493*.
- Iliia Shumailov, Zakhar Shumaylov, Yiren Zhao, Nicolas Papernot, Ross Anderson, and Yarin Gal. 2024. Ai models collapse when trained on recursively generated data. *Nature*, 631(8022):755–759.
- Alane Suhr, Clara Vania, Nikita Nangia, Maarten Sap, Mark Yatskar, Samuel R. Bowman, and Yoav Artzi. 2021. [Crowdsourcing beyond annotation: Case studies in benchmark data collection](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: Tutorial Abstracts*, pages 1–6, Punta Cana, Dominican Republic & Online. Association for Computational Linguistics.
- Zhen Tan, Alimohammad Beigi, Song Wang, Ruocheng Guo, Amrita Bhattacharjee, Bohan Jiang, Mansooreh Karami, Jundong Li, Lu Cheng, and Huan Liu. 2024. Large language models for data annotation: A survey. *arXiv preprint arXiv:2402.13446*.
- Chaofan Tao, Lu Hou, Wei Zhang, Lifeng Shang, Xin Jiang, Qun Liu, Ping Luo, and Ngai Wong. 2022. Compression of generative pre-trained language models via quantization. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4821–4836.
- Llama team. 2024. [The llama 3 herd of models](#). *arXiv*.
- Arun James Thirunavukarasu, Darren Shu Jeng Ting, Kabilan Elangovan, Laura Gutierrez, Ting Fang Tan, and Daniel Shu Wei Ting. 2023. Large language models in medicine. *Nature medicine*, 29(8):1930–1940.
- Sergei Tilga. 2024. [LLMs & humans: The perfect duo for data labeling](#). Accessed: 2024-07-04.
- Tina Tseng, Amanda Stent, and Domenic Maida. 2020. Best practices for managing data annotation projects. *arXiv preprint arXiv:2009.11654*.
- Dmitry Ustalov, Nikita Pavlichenko, and Boris Tseitlin. 2024. Learning from crowds with crowd-kit. *Journal of Open Source Software*, 9(96):6227.
- Bin Wang, Zhengyuan Liu, Xin Huang, Fangkai Jiao, Yang Ding, Aiti Aw, and Nancy Chen. 2024a. Seaeval for multilingual foundation models: From cross-lingual alignment to cultural reasoning. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 370–390.
- Xinru Wang, Hannah Kim, Sajjadur Rahman, Kushan Mitra, and Zhengjie Miao. 2024b. Human-llm collaborative annotation through effective verification of llm labels. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pages 1–21.
- Ziheng Wang, Jeremy Wohlwend, and Tao Lei. 2020. Structured pruning of large language models. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6151–6162.

- Zijie J Wang, Dongjin Choi, Shenyu Xu, and Diyi Yang. 2021. Putting humans in the natural language processing loop: A survey. In *Proceedings of the First Workshop on Bridging Human-Computer Interaction and Natural Language Processing*, pages 47–52.
- Jiaheng Wei, Zhaowei Zhu, Tianyi Luo, Ehsan Amid, Abhishek Kumar, and Yang Liu. 2023. To aggregate or not? learning with separate noisy labels. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 2523–2535.
- 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*, 135:364–381.
- Ziwei Xu, Sanjay Jain, and Mohan Kankanhalli. 2024. Hallucination is inevitable: An innate limitation of large language models. *arXiv preprint arXiv:2401.11817*.
- Diyi Yang, Sherry Tongshuang Wu, and Marti A. Hearst. 2024. [Human-AI interaction in the age of LLMs](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 5: Tutorial Abstracts)*, pages 34–38, Mexico City, Mexico. Association for Computational Linguistics.
- Jiacheng Ye, Jiahui Gao, Qintong Li, Hang Xu, Jiangtao Feng, Zhiyong Wu, Tao Yu, and Lingpeng Kong. 2022. Zerogen: Efficient zero-shot learning via dataset generation. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11653–11669.
- Yue Yu, Yuchen Zhuang, Jieyu Zhang, Yu Meng, Alexander J Ratner, Ranjay Krishna, Jiaming Shen, and Chao Zhang. 2024. Large language model as attributed training data generator: A tale of diversity and bias. *Advances in Neural Information Processing Systems*, 36.
- Ruoyu Zhang, Yanzeng Li, Yongliang Ma, Ming Zhou, and Lei Zou. 2023. LLMaAA: Making Large Language Models as Active Annotators. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 13088–13103.
- Zhisong Zhang, Emma Strubell, and Eduard Hovy. 2022. A survey of active learning for natural language processing. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 6166–6190.