

Future of Work in the Age of LLMs

Zora Wang, Yijia Shao, David Nguyen, Diyi Yang

<https://future-of-work-llm-tutorial.github.io/>

The recent development of large language models (LLMs) has revolutionized the landscape of human work. These models possess the ability to follow complex human instructions and operate versatile computer software, enabling them to participate in, augment, or even automate realistic occupational tasks that once thought to be exclusive to humans. As LLMs are increasingly integrated into workplaces, they are already reshaping labor dynamics (Hoffmann et al., 2024; Demirci et al., 2025) and raising urgent concerns about job displacement, diminished human agency, and overreliance on automation (Hazra et al., 2025). As a result, the future of work is undergoing a profound transformation: How will human occupations and task requirements evolve over time? And what roles will LLM-based systems play, as they become increasingly capable collaborators and autonomous workers? And how can we build technological and data infrastructures to support human-AI collaboration? This tutorial will provide an overview of the future of work shaped by the interplay of LLMs and humans, examining the emerging challenges, opportunities, and ethical considerations in this dynamic landscape. We begin by examining the economic landscape of work and how NLP technologies drive automation, followed by methods for developing LLMs that augment human labor and recent advances in LLM-based agents. We then cover evaluation approaches for workplace contexts, including datasets, benchmarks, and metrics, and conclude with open questions on technical, human, and societal implications.

Zora Wang

email: zhiruow@cs.cmu.edu

website: <https://zorazrw.github.io/>

Zora Zhiruo Wang is a PhD student at Carnegie Mellon University, Language Technologies Institute, advised by Professors Graham Neubig and Daniel Fried. Her research focuses on using programmatic approaches to solve real-world problems, especially those related to human occupations. Zora has organized Deep Learning for Code workshops at ICLR, NeurIPS, and ICML, given a tutorial about LLMs for tabular data at SIGIR 2024, and presented her work at and served as reviewers for top-tier NLP/ML conferences including NeurIPS, ICLR, ACL, EMNLP, and KDD. She has been recognized with Google PhD Fellowship and CMU Presidential Fellowship.

Yijia Shao

email: shaoyj@stanford.edu

website: <https://cs.stanford.edu/~shaoyj/>

Yijia Shao is a PhD student at Stanford University, advised by Professor Diyi Yang. Her research focuses on human-agent collaboration, specifically positioning NLP models into broader systems (e.g. LLM agents) and optimizing them to collaborate with humans. Her work appears in top ML and NLP conferences, including NeurIPS, ICLR, ACL, EMNLP, and NAACL. She is supported by Stanford School of Engineering Fellowship and Thinking Machines Fellowship.

David Nguyen

email: davidngu@stanford.edu

website: <https://digitaleconomy.stanford.edu/person/david-nguyen/>

David Nguyen is a Research Scientist at the Stanford Digital Economy Lab within the Institute for

Human-Centered Artificial Intelligence (HAI), and a Research Associate at the Economic Statistics Centre of Excellence (ESCoE). His research explores new ways to analyze and measure changing, modern, and digital economies. He is particularly interested in advancing economic indicators and statistics on economic output and welfare.

Diyi Yang

email: diyiy@cs.stanford.edu

website: <https://cs.stanford.edu/~diyiy/>

Diyi Yang is an assistant professor in the Computer Science Department at Stanford University. Her research focuses on human-centered natural language processing and computational social science. Diyi has organized four workshops and seven tutorials at *CL conferences, with recent ones including “Language Agents: Foundations, Prospects, and Risks” at EMNLP 2024, “Social Intelligence in the Age of LLMs” at NAACL 2025, and “Human-AI Collaboration: How AIs Augment Human Teammates” at ACL 2025.