

Title: Effective and Efficient Knowledge-Intensive NLP

Abstract:

Knowledge-intensive NLP tasks are the tasks that humans could not reasonably be expected to perform without access to external knowledge sources such as Google search, Wikipedia, dictionaries, and knowledge bases. They include open-domain question answering, commonsense reasoning, fact checking, etc. The state-of-the-art performance on such kinds of tasks is usually delivered by knowledge-augmented NLP solutions, named “open-book methods.” They look for useful information (referred as “knowledge”) to infer the output from input. However, the external data are heterogeneous and created independently from the task input; also, indexing and retrieval are much space and time consuming. In this talk, I will introduce three of our recent work in EMNLP 2022 and ICLR 2023 on effective and efficient knowledge augmentation. Since three conference tutorials in ACL/EMNLP and a successful workshop at AAAI 2023, this area of study has established a growing and enduring community. Please join us!

Bio:

Meng Jiang is an assistant professor in the Department of Computer Science and Engineering at the University of Notre Dame. He is interested in data mining and natural language processing. His data science research focuses on graph and text data for applications such as question answering, material science, biomedicine, and mental healthcare. He received the NSF CAREER award in 2022.

