

KnowFM 2025

**The 3rd Workshop on Towards Knowledgeable Foundation
Models**

Proceedings of the Workshop

August 1, 2025

©2025 Association for Computational Linguistics

Order copies of this and other ACL proceedings from:

Association for Computational Linguistics (ACL)
317 Sidney Baker St. S
Suite 400 - 134
Kerrville, TX 78028
USA
Tel: +1-855-225-1962
acl@aclweb.org

ISBN 979-8-89176-283-1

Introduction

Welcome to KnowFM 2025, the 3rd workshop on knowledgeable foundation models. Co-located with ACL 2025, this workshop is scheduled for August 1, 2025, and will be held in Vienna, Austria.

Knowledge has been an important prerequisite for a variety of NLP applications and is typically sourced from either structured knowledge sources such as knowledge bases and dictionaries, or unstructured sources such as Wikipedia documents. More recently, researchers have discovered that language models already possess a significant amount of knowledge through pretraining: LLMs can be used to generate commonsense and factual knowledge for question answering. While the results are encouraging, there are still lingering questions: Where does this knowledge come from? How much do language models know? Is this knowledge reliable? If some knowledge is wrong, can we fix it?

In response to these questions, the KnowFM workshop examines the lifecycle of knowledge within foundation models: The emergence of knowledge through language model pretraining; Injection of external knowledge; Updating and modification of knowledge; Probing and generation of knowledge.

Currently, researchers focusing on different stages of this lifecycle are scattered across various sub-communities within NLP. For example, probing and editing knowledge is often associated with the interpretability track, while injecting knowledge is typically application-specific and discussed within dialog, open-domain QA, IE, or summarization tracks. This workshop seeks to bring these researchers together and facilitate collaboration to create a more holistic view of the problem.

The KnowFM workshop also addresses core challenges in LM research: reducing hallucination, improving interpretability, and enhancing model extensibility. Although these challenges remain open, knowledge clearly plays a key role: Attribution to sources or providing relevant knowledge during generation can mitigate hallucination; Locating and tracing knowledge provides insights into the LM’s inner workings; Efficiently adapting to domain knowledge or integrating updated facts improves extensibility.

This year, we received a total of 62 archival and non-archival submissions to the KnowFM workshop, of which 55 were accepted. Among these, 12 have been included in our proceedings, and 7 are included in ACL Findings.

In addition to oral and poster sessions where accepted works will be presented, the workshop will also host talks and a panel discussion with invited speakers.

Finally, we would like to express our gratitude to all the authors, committee members, invited speakers, and participants for helping make this workshop possible.

Organizing Committee

Program Chairs

Yuji Zhang, University of Illinois at Urbana-Champaign, USA
Canyu Chen, Northwestern University, USA
Sha Li, Amazon, USA
Mor Geva, Tel Aviv University, Israel
Chi Han, University of Illinois at Urbana-Champaign, USA
Xiaozhi Wang, University of Illinois at Urbana-Champaign, USA
Shangbin Feng, University of Washington, USA
Silin Gao, École Polytechnique Fédérale de Lausanne (EPFL), Switzerland
Isabelle Augenstein, University of Copenhagen, Denmark
Mohit Bansal, University of North Carolina at Chapel Hill, USA
Manling Li, Northwestern University, USA
Heng Ji, University of Illinois at Urbana-Champaign, USA

Publication Chair

Canyu Chen, Northwestern University, USA

Program Committee

Reviewers

Samir Abdaljalil
Yuzhuo Bai
Norbert Braunschweiler
Thomas Chen
Lida Chen
Zheng CHEN
Canyu Chen
Sitao Cheng
Cong-Thanh Do
Yixiong Fang
Yujie Feng
SeungYoon Han
Simeng Han
Yutong Hu
Xiaoyu Hu
Liqiang Jing
Kemal Kirtac
Dosung Lee
Zichao Li
Yanhong Li
Ming Li
Haohang Li
Ruochen Li
Laurence Liang
João Alberto de Oliveira Lima
Dong Liu
Xin Liu
Aofan Liu
Jiahong Liu
Ziming Luo
Jianfei Ma
Hudson de Martim
Kyle Montgomery
Albert Olweny Okiri
Yixin Ou
Alonso Palomino
Heramb Vivek Patil
Can Polat
Prasanth Prasanth
Swayamjit Saha
Shuzheng Si
Vincent Siu
Shane Storks
Toma Suzuki
Jianhong Tu
Sai P Vallurupalli

Mengru Wang
Zhe Yang
Wanli Yang
Yunzhi Yao
Li-Ming Zhan
Caiqi Zhang
Xiaofeng Zhang
Xinyun Zhou
Zeqi Zhou

Table of Contents

<i>Temporal Information Retrieval via Time-Specifier Model Merging</i> SeungYoon Han, Taeho Hwang, Sukmin Cho, Soyeong Jeong, Hoyun Song, Huije Lee and Jong C. Park	1
<i>EdTec-ItemGen: Enhancing Retrieval-Augmented Item Generation Through Key Point Extraction</i> Alonso Palomino, David Buschhüter, Roland Roller, Niels Pinkwart and Benjamin Paassen ...	14
<i>Teaching Large Language Models to Express Knowledge Boundary from Their Own Signals</i> Lida Chen, Zujie Liang, Xintao Wang, Jiaqing Liang, Yanghua Xiao, Feng Wei, Jinglei Chen, Zhenghong Hao, Bing Han and Wei Wang	26
<i>Knowledge-Grounded Detection of Cryptocurrency Scams with Retrieval-Augmented LMs</i> Zichao Li	40
<i>Stress-Testing Multimodal Foundation Models for Crystallographic Reasoning</i> Can Polat, Hasan Kurban, Erchin Serpedin and Mustafa Kurban	49
<i>MLAN: Language-Based Instruction Tuning Preserves and Transfers Knowledge in Multimodal Language Models</i> Jianhong Tu, Zhuohao Ni, Nicholas Crispino, Zihao Yu, Michael Bendersky, Beliz Gunel, Ruoxi Jia, Xin Liu, Lingjuan Lyu, Dawn Song and Chenguang Wang	59
<i>ToolReAGt: Tool Retrieval for LLM-based Complex Task Solution via Retrieval Augmented Generation</i> Norbert Braunschweiler, Rama Doddipatla and Tudor-catalin Zorila	75
<i>Can LLMs Recognize Their Own Analogical Hallucinations? Evaluating Uncertainty Estimation for Analogical Reasoning</i> Zheng Chen, Zhaoxin Feng, Jianfei Ma, Jiexi Xu and Bo Li	84
<i>Meetalk: Retrieval-Augmented and Adaptively Personalized Meeting Summarization with Knowledge Learning from User Corrections</i> Zheng Chen, Jiang Futian, Yue Deng, Changyang He and Bo Li	94
<i>Theorem-of-Thought: A Multi-Agent Framework for Abductive, Deductive, and Inductive Reasoning in Language Models</i> Samir Abdaljalil, Hasan Kurban, Khalid Qaraqe and Erchin Serpedin	111
<i>Reasoning or Memorization? Investigating LLMs' Capability in Restoring Chinese Internet Homophones</i> Jianfei Ma, Zhaoxin Feng, Huacheng Song, Emmanuele Chersoni and Zheng Chen	120
<i>Superfluous Instruction: Vulnerabilities Stemming from Task-Specific Superficial Expressions in Instruction Templates</i> Toma Suzuki, Yusuke Sakai, Justin Vasselli, Hidetaka Kamigaito and Taro Watanabe	140