KnowledgeNLP'25 2025

The 4th International Workshop on Knowledge-Augmented Methods for Natural Language Processing

Proceedings of the Workshop

May 3, 2025

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 979-8-89176-229-9

Introduction

Welcome to the 4th International Workshop on Knowledge-Augmented Methods for Natural Language Processing (KnowledgeNLP'25), held in conjunction with NAACL 2025. KnowledgeNLP will take place on May 3rd, 2025, allowing for both virtual and in-person attendance in New Mexico, USA.

Recent progress in large-scale models like ChatGPT has significantly advanced NLP capabilities. However, these models face limitations in memorizing rare information, are prone to hallucinations, and cannot access up-to-date information. Additionally, their fixed parameter size prevents them from fully encapsulating the continuously evolving world knowledge.

The field of knowledge-augmented NLP spans a wide array of techniques and applications. Acquiring relevant knowledge is challenging due to its diversity and distribution across numerous sources. Once acquired, effectively representing and utilizing this knowledge to support model predictions presents another major challenge. This workshop seeks to bring researchers together to share their insights and progress in this domain, aiming to highlight the importance of knowledge-augmented NLP.

In response to our call for papers, we received 48 submissions. Each submission was rigorously reviewed by at least three Program Committee members selected for their expertise. Based on the reviewers' feedback, we accepted 28 papers, including 6 oral presentations and 22 poster presentations. We are honored to invite five keynote speakers: Prof. Doug Downey (Allen Institute for AI and Professor Northwestern University), Prof. Graham Neubig (Carnegie Mellon University), Dr. Yunyao Li (Adobe), Prof. Yu Su (Ohio State University), Prof. Manling Li (Northwestern University).

We hope you find the workshop papers insightful and inspiring. We express our gratitude to the keynote speakers for their engaging talks, the authors for their valuable contributions, and the Program Committee members for their thorough reviews. Special thanks to the emergency reviewers for their expertise and to the NAACL 2025 workshop chairs for their support during the organization process.

Organizing Committee

Program Chairs

Weijia Shi, University of Washington
Wenhao Yu, Tencent AI Seattle Lab
Akari Asai, University of Washington
Meng Jiang, University of Notre Dame
Greg Durrett, University of Texas at Austin
Hannaneh Hajishirzi, University of Washington
Luke Zettlemoyer, University of Washington

Program Committee

Reviewers

Amit Agarwal, Oracle
Neelima Agarwal, Microsoft
Ashutosh Ahuja, Starbucks
Abhinav Balasubramanian, NVIDIA
Mithil Bangera, University Of New Haven
Yeshil Bangera
Vibha Belavadi, Adobe Systems
Zeyi Chen, University of Washington
Shristi Drolia, Novocure
Ritam Dutt, Carnegie Mellon University
Krupa Galiya
Mohammad Hasan, Indiana University Indianapolis
Yebowen Hu, University of Central Florida
Chao-Wei Huang, National Taiwan University
Zhehui Huang, University of Southern California
Shruti Jalan, Amazon
Amit Jaspal, Facebook
Bhargava Kumar, TD Securities
Yoonsang Lee, Seoul National University
Pengwei Li, Meta
Bryan Li, University of Pennsylvania
Su Liu, Amazon
Zefang Liu, J.P. Morgan Chase and Georgia Institute of Technology
Xiao Liu, University of California, Davis
Maochuan Lu, University of California, Berkeley
Elan Sopher Markowitz, University of Southern California
Bang Nguyen, University of Notre Dame
Garima Panwar, Amazon
Hitesh Laxmichand Patel, Oracle
Priyaranjan Pattnayak, Oracle
Meghana Puvvadi, NVIDIA
Yinzhu Quan, Georgia Institute of Technology
Arina Razmyslovich, DNI
Aaditya Shukla, NVIDIA
Ishneet Sukhvinder Singh, Algoverse AI Research
Dewang Sultania, Adobe Systems
Shounak Sural, Adobe
Suhas Suresha, Adobe Systems
Jiaxiang Tang, HKUST
Van-Hien Tran, NICT
Takehito Utsuro, University of Tsukuba
Ishita Verma, Adobe Systems
Qingyun Wang, University of Illinois, Urbana Champaign
Di Wu, University of California, Los Angeles
Zhenyu Wu, Xi'an Jiaotong University
Paul Youssef, Marburg University

Mengxia Yu, University of Notre Dame
Yuwei Zhang, University of California, San Diego
Andrea Zugarini, Expert.ai Srl

Table of Contents

Entity Retrieval for Answering Entity-Centric Questions
Hassan Shavarani and Anoop Sarkar 1

ELECTRA and GPT-4o: Cost-Effective Partners for Sentiment Analysis
James P. Beno 18

Retrieval of Temporal Event Sequences from Textual Descriptions
Zefang Liu and Yinzhu Quan 37

Generating Tables from the Parametric Knowledge of Language Models
Yevgeni Berkovitch, Oren Glickman, Amit Somech and Tomer Wolfson 50

Investigating Large Language Models for Text-to-SPARQL Generation
Jacopo D'Abramo, Andrea Zugarini and Paolo Torroni 66

GAVEL: Generative Attribute-Value Extraction Using LLMs on LLM-Augmented Datasets
Pollawat Hongwimol, Dong Sheng, Li Zhang, Kai Liu and Xiufei Wang 81

Leveraging Domain Knowledge at Inference Time for LLM Translation: Retrieval versus Generation
Bryan Li, Jiaming Luo, Eleftheria Briakou and Colin Cherry 91

Enhancing Cross-Language Code Translation via Task-Specific Embedding Alignment in Retrieval-Augmented Generation
Manish Bhattarai, Minh N. Vu, Javier E. Santos, Ismael Ismael and Daniel O'Malley 107

LLM Reasoning Engine: Specialized Training for Enhanced Mathematical Reasoning
Shuguang Chen and Guang Lin 118

RouteNator: A Router-Based Multi-Modal Architecture for Generating Synthetic Training Data for Function Calling LLMs
Dewang Sultania, Vibha Belavadi, Tushar Vatsa, Suhas Suresha, Ishita Verma, Tracy Holloway King, mifriedr@adobe.com mifriedr@adobe.com and Cheng Chen 129

StoC-TOT: Stochastic Tree-of-Thought with Constrained Decoding for Complex Reasoning in Multi-Hop Question Answering
Zhenyu Bi, Daniel Hajialigol, Zhongkai Sun, Jie Hao and Xuan Wang 141

EKRAG: Benchmark RAG for Enterprise Knowledge Question Answering
Tan Yu, Wenfei Zhou, leiyang@nvidia.com leiyang@nvidia.com, Aaditya Shukla, mmadugula@nvidia.com mmadugula@nvidia.com, Pritam Gundecha, Nicholas Burnett, Anbang Xu, viseth@nvidia.com viseth@nvidia.com, tbar@nvidia.com tbar@nvidia.com, Rama Akkiraju and Vivienne Zhang 152

Towards Effectively Leveraging Execution Traces for Program Repair with Code LLMs
Mirazul Haque, Petr Babkin, Farima Farmahinifarahani and Manuela Veloso 160

A Novel Multi-Document Retrieval Benchmark: Journalist Source-Selection in Newswriting
Alexander Spangher, Tenghao Huang, Yiqin Huang, Lucas Spangher, Sewon Min and Mark Dredze 180

HEAL: Hierarchical Embedding Alignment Loss for Improved Retrieval and Representation Learning
Manish Bhattarai, Ryan Barron, Maksim E. Eren, Minh N. Vu, Vesselin Grantcharov, Ismael Ismael, Valentin Stanev, Cynthia Matuszek, Vladimir I Valtchinov, Kim Rasmussen and Boian S. Alexandrov 205

Hybrid AI for Responsive Multi-Turn Online Conversations with Novel Dynamic Routing and Feedback Adaptation
Priyaranjan Pattnayak, Amit Agarwal, Hansa Meghwani, Hitesh Laxmichand Patel and Srikant Panda 215

Chain of Evidences and Evidence to Generate: Prompting for Context Grounded and Retrieval Augmented Reasoning
Md Rizwan Parvez 230

Expertly Informed, Generatively Summarized: A Hybrid RAG Approach to Informed Consent Summarization with Auxiliary Expert Knowledge
Autumn Toney, rsw66@georgetown.edu rsw66@georgetown.edu and calebs@med.umich.edu calebs@med.umich.edu 246

MSR^2: A Benchmark for Multi-Source Retrieval and Reasoning in Visual Question Answering
Kuo-Han Hung, Hung-Chieh Fang, Chao-Wei Huang and Yun-Nung Chen 259

PROPEL: Prompt Optimization with Expert Priors for Small and Medium-sized LLMs
Kawin Mayilvaghanan, Varun Nathan and Ayush Kumar 272

ClaimCheck: Automatic Fact-Checking of Textual Claims using Web Evidence
Akshith Reddy Putta, Jacob Devasier and Chengkai Li 303

Can dependency parses facilitate generalization in language models? A case study of cross-lingual relation extraction
Ritam Dutt, Shounak Sural and Carolyn Rose 317

Can dependency parses facilitate generalization in language models? A case study of cross-lingual relation extraction
Ritam Dutt, Shounak Sural and Carolyn Rose 338

DocBench: A Benchmark for Evaluating LLM-based Document Reading Systems
Anni Zou, Wenhao Yu, Hongming Zhang, Kaixin Ma, Deng Cai, Zhuosheng Zhang, Hai Zhao and Dong Yu 359