

4th Workshop on Knowledge-Augmented Methods for Natural Language Processing (KnowledgeNLP 2025)

Albuquerque, New Mexico, USA
3 May 2025

ISBN: 979-8-3313-1920-5

Printed from e-media with permission by:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571



Some format issues inherent in the e-media version may also appear in this print version.

Copyright© (2025) by the Association for Computational Linguistics
All rights reserved.

Printed with permission by Curran Associates, Inc. (2025)

For permission requests, please contact the Association for Computational Linguistics
at the address below.

Association for Computational Linguistics
209 N. Eighth Street
Stroudsburg, Pennsylvania 18360

Phone: 1-570-476-8006
Fax: 1-570-476-0860

acl@aclweb.org

Additional copies of this publication are available from:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: 845-758-0400
Fax: 845-758-2633
Email: curran@proceedings.com
Web: www.proceedings.com

Table of Contents

<i>Entity Retrieval for Answering Entity-Centric Questions</i> Hassan Shavarani and Anoop Sarkar	1
<i>ELECTRA and GPT-4o: Cost-Effective Partners for Sentiment Analysis</i> James P. Beno	18
<i>Retrieval of Temporal Event Sequences from Textual Descriptions</i> Zefang Liu and Yinzhu Quan	37
<i>Generating Tables from the Parametric Knowledge of Language Models</i> Yevgeni Berkovitch, Oren Glickman, Amit Somech and Tomer Wolfson	50
<i>Investigating Large Language Models for Text-to-SPARQL Generation</i> Jacopo D’Abramo, Andrea Zugarini and Paolo Torroni	66
<i>GAVEL: Generative Attribute-Value Extraction Using LLMs on LLM-Augmented Datasets</i> Pollawat Hongwimol, Dong Sheng, Li Zhang, Kai Liu and Xiufei Wang	81
<i>Leveraging Domain Knowledge at Inference Time for LLM Translation: Retrieval versus Generation</i> Bryan Li, Jiaming Luo, Eleftheria Briakou and Colin Cherry	91
<i>Enhancing Cross-Language Code Translation via Task-Specific Embedding Alignment in Retrieval-Augmented Generation</i> Manish Bhattarai, Minh N. Vu, Javier E. Santos, Ismael Ismael and Daniel O’Malley	107
<i>LLM Reasoning Engine: Specialized Training for Enhanced Mathematical Reasoning</i> Shuguang Chen and Guang Lin	118
<i>RouteNator: A Router-Based Multi-Modal Architecture for Generating Synthetic Training Data for Function Calling LLMs</i> Dewang Sultania, Vibha Belavadi, Tushar Vatsa, Suhas Suresha, Ishita Verma, Tracy Holloway King, mifriedr@adobe.com mifriedr@adobe.com and Cheng Chen	129
<i>StoC-TOT: Stochastic Tree-of-Thought with Constrained Decoding for Complex Reasoning in Multi-Hop Question Answering</i> Zhenyu Bi, Daniel Hajialigol, Zhongkai Sun, Jie Hao and Xuan Wang	141
<i>EKRAG: Benchmark RAG for Enterprise Knowledge Question Answering</i> 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
<i>Towards Effectively Leveraging Execution Traces for Program Repair with Code LLMs</i> Mirazul Haque, Petr Babkin, Farima Farmahinifarahani and Manuela Veloso	160
<i>A Novel Multi-Document Retrieval Benchmark: Journalist Source-Selection in Newswriting</i> Alexander Spangher, Tenghao Huang, Yiqin Huang, Lucas Spangher, Sewon Min and Mark Dredze	180
<i>HEAL: Hierarchical Embedding Alignment Loss for Improved Retrieval and Representation Learning</i> 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

<i>Hybrid AI for Responsive Multi-Turn Online Conversations with Novel Dynamic Routing and Feedback Adaptation</i>	
Priyaranjan Pattanayak, Amit Agarwal, Hansa Meghwani, Hitesh Laxmichand Patel and Srikant Panda	215
<i>Chain of Evidences and Evidence to Generate: Prompting for Context Grounded and Retrieval Augmented Reasoning</i>	
Md Rizwan Parvez	230
<i>Expertly Informed, Generatively Summarized: A Hybrid RAG Approach to Informed Consent Summarization with Auxiliary Expert Knowledge</i>	
Autumn Toney, rsw66@georgetown.edu rsw66@georgetown.edu and calebs@med.umich.edu calebs@med.umich.edu	246
<i>MSR²: A Benchmark for Multi-Source Retrieval and Reasoning in Visual Question Answering</i>	
Kuo-Han Hung, Hung-Chieh Fang, Chao-Wei Huang and Yun-Nung Chen	259
<i>PROPEL: Prompt Optimization with Expert Priors for Small and Medium-sized LLMs</i>	
Kawin Mayilvaghanan, Varun Nathan and Ayush Kumar	272
<i>ClaimCheck: Automatic Fact-Checking of Textual Claims using Web Evidence</i>	
Akshith Reddy Putta, Jacob Devasier and Chengkai Li	303
<i>Can dependency parses facilitate generalization in language models? A case study of cross-lingual relation extraction</i>	
Ritam Dutt, Shounak Sural and Carolyn Rose	317
<i>Can dependency parses facilitate generalization in language models? A case study of cross-lingual relation extraction</i>	
Ritam Dutt, Shounak Sural and Carolyn Rose	338
<i>DocBench: A Benchmark for Evaluating LLM-based Document Reading Systems</i>	
Anni Zou, Wenhao Yu, Hongming Zhang, Kaixin Ma, Deng Cai, Zhuosheng Zhang, Hai Zhao and Dong Yu	359