ISSN 2499-4553

IJCoL

Italian Journal of Computational Linguistics

Rivista Italiana di Linguistica Computazionale

Volume 10, Number 2
december 2024
Special Issue

Natural Language for Artificial Intelligence in the Era of LLMs

aA Accademia university press

Registrazione presso il Tribunale di Trento n. 14/16 del 6 luglio 2016

Rivista Semestrale dell'Associazione Italiana di Linguistica Computazionale (AILC)

Associazione Italiana di Linguistica Computazionale

direttore responsabile
Michele Arnese

isbn 9791255001430

Accademia University Press
via Carlo Alberto 55
I-10123 Torino
info@aAccademia.it
www.aAccademia.it/IJCoL_10_2

IJCoL

Volume 10, Number 2
december 2024
Special Issue

Natural Language for Artificial Intelligence in the Era of LLMs

editors: *Elisa Bassignana, Dominique Brunato, Marco Polignano, Alan Ramponi*

CONTENTS

Introduction to the Special Issue on Natural Language for Artificial Intelligence in the Era of LLMs

Elisa Bassignana*
IT University of Copenhagen

Dominique Brunato**
Istituto di Linguistica Computazionale "Antonio Zampolli"

Marco Polignano†
Università di Bari Aldo Moro

Alan Ramponi‡
Fondazione Bruno Kessler

1. Introduction

The rapid advancement of Large Language Models (LLMs) has revolutionized the field of Natural Language Processing (NLP) and Artificial Intelligence (AI) in recent years. Transformer-based models such as GPT-3 and BERT have demonstrated remarkable capabilities in modeling and generating human-like text. These models have not only redefined the potential of AI systems but also revolutionized applications across a broad spectrum, including machine translation, sentiment analysis, question answering, and beyond. While the era of LLMs has significantly expanded the horizons of AI, it has also presented critical challenges in effectively and responsibly harnessing their capabilities.

This special issue was conceived as a dedicated platform for exploring the latest advancements, methodologies, and applications of LLMs. It sought to foster collaboration and knowledge exchange within the NLP and AI communities, encouraging researchers and practitioners to address real-world challenges while also delving into the theoretical foundations of language in machine learning. The issue builds on discussions initiated at the NL4AI 2023 workshop held at AIxIA 2023 (Bassignana et al. 2023), where many contributions showcased the innovative use of LLMs for natural language understanding and generation tasks.

The papers in this volume reflect a diverse and rich body of research, spanning a wide array of topics. They investigate cutting-edge applications, introduce novel methodologies, and provide theoretical insights into the use of LLMs, while also reflecting on the integration of multimodal information and cognitively inspired frameworks to tackle complex problems. More specifically, the issue covers multimodal approaches, domain-specific adaptations, and efficient model deployment. The contributions highlight both opportunities and challenges in advancing natural language technologies.

* Data Science Section. Rued Langgaards Vej 7, 2300 København. E-mail: elba@itu.dk
** Consiglio Nazionale delle Ricerche. Via Moruzzi 1, 56124 Pisa.
E-mail: dominique.brunato@ilc.cnr.it
† Dipartimento di Informatica, Università degli Studi di Bari Aldo Moro, Via E. Orabona 4, Bari.
E-mail: marco.polignano@uniba.it
‡ Digital Humanities unit, Digital Society center. Via Sommarive 18, 38123 Trento.
E-mail: alramponi@fbk.eu

2. Overview of the Issue

The issue opens with two contributions that explore multimodality and the advantages of incorporating visual information in both language understanding and generation.

Cassese *et al.* provide a detailed analysis of Transformer-based Language Models and Vision-Language Models (VLMs) in representing generalized event knowledge. By evaluating these models' ability to assess the plausibility of agent-patient interactions in minimal sentence pairs, the authors compare unimodal and multimodal models across various architectures and sizes. Their findings underscore both the potential and limitations of LLMs in leveraging real-world knowledge, demonstrating how incorporating visual information can enhance event plausibility recognition.

In his position paper, Zamparelli advocates for a paradigm shift in AI research, proposing a new class of systems for generating outputs that combine language and visual elements. He reflects on the transformative potential of such multimodal systems for advancing tasks that inherently rely on both modalities.

Ravelli and Bolognesi investigate the performance of several LLMs in approximating human judgments in a classical psycholinguistics norming study focused on estimating word specificity. Beyond evaluating the correlation with humans on the same task, they also discuss the trade-offs between performance and computational costs, offering valuable insights into the practical implications of selecting different models for such tasks.

Everitt *et al.* delve into the ability of LLMs to detect implicit bias in job descriptions, focusing on fine-grained categories such as feminine, sexuality, and racial bias. Their study evaluates several adaptation techniques, including fine-tuning, few-shot and zero-shot prompting, and self-consistency. The findings indicate that fine-tuning consistently outperforms prompting methods, with the best-performing models achieving high accuracy across all bias categories. This underscores the importance of domain adaptation in addressing nuanced bias detection tasks.

The last two contributions examine the feasibility of applying LLMs in specific domains and their adaptations to languages other than English.

Siragusa and Pirrone introduce Unipa-GPT, a virtual assistant built upon ChatGPT and fine-tuned using an Italian domain-specific corpus in a Retrieval-Augmented Generation (RAG) approach. The chatbot was designed to assist students at the University of Palermo in selecting a bachelor's degree course. Comparisons with other open-source LLMs revealed that gpt-3.5-turbo outperforms its competitors, underscoring the significant advantages of fine-tuning for domain-specific applications.

Finally, Fierens and Jodogne explore the use of the Cramming approach to train a BERT model for medical-related tasks in the French language. They present and test various pre-trained and fine-tuned model versions against comparable French LLMs. Despite computational limitations, their findings demonstrate that the Cramming approach yields competitive results, showcasing its viability as a strategy for developing effective domain-specific language models in resource-constrained environments.

References

Bassignana, Elisa, Dominique Brunato, Marco Polignano, and Alan Ramponi. 2023. Preface to the seventh workshop on natural language for artificial intelligence (NL4AI 2023). In *Proceedings of the Seventh Workshop on Natural Language for Artificial Intelligence (NL4AI 2023) co-located with 22th International Conference of the Italian Association for Artificial Intelligence (AI*IA 2023)*, Rome, Italy, November 6th-7th. CEUR-WS.org.