Introduction: Understanding Language in the Human-Machine Era

Large language models (LLMs) have revolutionized the way interactional artificial intelligence (AI) systems are developed, as they are made available to ordinary users. Significant advances have been observed in applications such as conversational AI and machine translation, and their widespread use in the so called human-machine era, where technology is integrated with our senses (Sayers et al., 2021), is undeniable; those models have produced remarkable achievements in several benchmarks (Gao et al., 2021; Hendrycks et al., 2021; Wang et al., 2019; Zhou et al., 2020), and the scientific community has discussed emergent properties (Wei et al., 2022) that result from scaling laws (Kaplan et al., 2020). Nevertheless, state-of-the-art systems are still prone to brittleness in language understanding, which raises doubts about the extent to which such systems can truly understand human language(s) (Mitchell & Krakauer, 2023).

The concept of language understanding has always been controversial (Lyons, 1990; Michael et al., 2023). As contemporary linguistic theories have shown, meaning-making relies not only on form and (immediate) meaning, but also on context. Thus, understanding natural language entails more than observing the form and the meaning withdrawn from that form; instead, harnessing meaning (Bender & Koller, 2020) requires access to grounding. Understanding language is, hence, a very complex task, even for humans (Lyons, 1990). As discourse, pragmatics, and (social) context are particularly relevant for understanding language, how to equip language models with such linguistics-grounded capabilities is yet to be fully understood.

Nevertheless, language models are seemingly capable of generalizing concepts, which could arguably be seen as some kind of meaning understanding (Piantadosi & Hill, 2022), even if modest. Understanding language in the human-machine era is, therefore, a doubly challenging task. Besides understanding the intrinsic capabilities of LLMs, it is increasingly important to investigate the requirements and impact of using such systems in real-world applications. As has been empirically demonstrated, LLMs can be used effectively in various applications, even without sophisticated language understanding skills, but the lack of theories that support these findings raises concerns about which kinds of applications, particularly those dealing directly with human interaction, pose greater risks and ethical concerns. Notable examples include the impact of language technology on teaching and language work. For instance, research is underway into using language models in educational settings, including question-generation (Leite & Lopes Cardoso,

2023); likewise, machine translation (MT) is increasingly ubiquitous, as it is used by both language professionals and general speakers at (apparently) no cost. Yet, as MT systems can take in a limited amount of context, they tend to make mistakes similar to those of human translators, who need to rely on their own knowledge to do their job more accurately.

As the way AI systems are intertwined with human expertise in language understanding is quickly changing, some have raised the question of the role played by language professionals in tasks such as translation. These professionals systematically add value to building next-generation language models that use linguistic and common-sense knowledge to provide more robust systems.

The "Language Understanding in the Human-Machine Era" (LUHME) workshop retrieves, resumes and refocuses the longstanding debate about the role of understanding in natural language use and related applications. In particular, the workshop provides insight into what language understanding is and whether it is required for computational natural language tasks, such as machine translation and natural language generation. Additionally, it furthers the discussion about the role played by language professionals (e.g., linguists, professional translators, and language teachers) in computational natural language understanding.

LUHME brings together researchers interested in the intersection between language understanding and the effective use of language technologies in human-machine interaction to discuss, among others, language understanding in/by LLMs; language grounding; psycholinguistic approaches to language understanding; discourse, pragmatics and language understanding; socio-cultural aspects in understanding language(s); effects of language misunderstanding by computational models; manifestations of language understanding; linguistic theory and language understanding by machines; linguistic, world, and common sense knowledge in language understanding; machine translation and/or interpreting and language understanding; human vs. machine language understanding; role of language professionals in the LLMs era; understanding language and explainable AI.

Rui Sousa-Silva
Henrique Lopes Cardoso
Maarit Koponen
Antonio Pareja-Lora
Márta Seresi

References

Bender, E. M., & Koller, A. (2020). Climbing towards NLU: On meaning, form, and understanding in the age of data. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 5185–5198. https://doi.org/10.18653/v1/2020.acl-main.463

Gao, L., Tow, J., Biderman, S., & et al. (2021). *A framework for few-shot language model evaluation* (Version v0.0.1). Zenodo. https://doi.org/10.5281/zenodo.5371628

Hendrycks, D., Burns, C., Basart, S., & et al. (2021). Measuring massive multitask language understanding. *International Conference on Learning Representations*. https://openreview.net/forum?id=d7KBjmI3GmQ

Kaplan, J., McCandlish, S., Henighan, T., & et al. (2020). Scaling laws for neural language models. *CoRR*, *abs/2001.08361*. https://arxiv.org/abs/2001.08361

Leite, B., & Lopes Cardoso, H. (2023). Towards enriched controllability for educational question generation. In N. Wang, G. Rebolledo-Mendez, N. Matsuda, O. C. Santos, & V. Dimitrova (Eds.), *Artificial intelligence in education* (pp. 786–791). Springer Nature Switzerland.

Lyons, J. (1990). *Language and linguistics: An introduction*. Cambridge University Press.

Michael, J., Holtzman, A., Parrish, A., Mueller, A., Wang, A., Chen, A., Madaan, D., Nangia, N., Pang, R. Y., Phang, J., & Bowman, S. R. (2023). What do NLP researchers believe? results of the NLP community metasurvey. *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 16334–16368. https://doi.org/10.18653/v1/2023.acl-long.903

Mitchell, M., & Krakauer, D. C. (2023). The debate over understanding in ai's large language models. *Proceedings of the National Academy of Sciences*, *120*(13), e2215907120. https://doi.org/10.1073/pnas.2215907120

Piantadosi, S. T., & Hill, F. (2022). Meaning without reference in large language models.

Sayers, D., Sousa-Silva, R., Höhn, S., & et al. (2021). *The dawn of the human–machine era: A forecast of new and emerging language technologies* (tech. rep.). EU COST Action CA19102 'Language In The Human–Machine Era'. https://doi.org/https://doi.org/10.17011/jyx/reports/20210518/1

Wang, A., Pruksachatkun, Y., Nangia, N., & et al. (2019). Superglue: A stickier benchmark for general-purpose language understanding systems. In *Procs. 33rd international conference on neural information processing systems*. Curran Associates Inc.

Wei, J., Tay, Y., Bommasani, R., Raffel, C., Zoph, B., Borgeaud, S., Yogatama, D., Bosma, M., Zhou, D., Metzler, D., Chi, E. H., Hashimoto, T., Vinyals, O., Liang, P., Dean, J., & Fedus, W. (2022). Emergent abilities of large language models [Survey Certification]. *Transactions on Machine Learning Research*. https://openreview.net/forum?id=yzkSU5zdwD

Zhou, X., Zhang, Y., Cui, L., & Huang, D. (2020). Evaluating commonsense in pre-trained language models. *Procs. 34th AAAI, New York, USA, February 7-12, 2020*, 9733–9740. https://doi.org/10.1609/aaai.v34i05.6523

Organising Committee

Rui Sousa-Silva
University of Porto, Portugal

Henrique Lopes Cardoso
University of Porto, Portugal

Maarit Koponen
University of Eastern Finland, Finland

Antonio Pareja-Lora
Universidad de Alcalá, Spain

Márta Seresi
Eötvös Loránd University, Hungary

Assistants

Andrés-Felipe Herrera-Ramírez
Universidad de Alcalá, Spain

Ana Sofia Meneses Silva
University of Porto, Portugal

Karen de Souza
University of Eastern Finland, Finland

Programme Committee

Aida Kostikova, *Bielefeld University, Germany*
Alex Lascarides, *University of Edinburgh, United Kingdom*
Alípio Jorge, *University of Porto, Portugal*
António Branco, *University of Lisbon, Portugal*
Belinda Maia, *University of Porto, Portugal*
Caroline Lehr, *ZHAW School of Applied Linguistics, Switzerland*
Diana Santos, *Universitetet i Oslo, Norway*
Efstathios Stamatatos, *University of the Aegean, Greece*
Ekaterina Lapshinova-Koltunski, *University of Hildesheim, Germany*
Eliot Bytyçi, *Universiteti i Prishtinës "Hasan Prishtina", Kosovo*
Hanna Risku, *University of Vienna, Austria*
Lynne Bowker, *University of Ottawa, Canada*
Nataša Pavlović, *University of Zagreb, Croatia*
Paolo Rosso, *Universitat Politècnica de València, Spain*
Ruslan Mitkov, *Lancaster University, United Kingdom*
Sule Yildirim Yayilgan, *Norwegian University of Science and Technology, Norway*

Keynote Speakers

Alexander Koller, *Saarland University*
Anders Søgaard, *University of Copenhagen*

Contents

Keynote Speakers

Keynote Speaker

Why Most Are Wrong About LLM Understanding

Anders Søgaard

University of Copenhagen

Abstract: I identify a fallacy common to many LLM no-go theorems of the form: LLMs cannot do X, because they were designed to – or trained to – do Y. I present observations that seem to challenge stochastic parrot or database views of LLMs, as well as arguments for why, contrary to popular belief, structural similarity may be sufficient for grounding.

Bio: Anders Søgaard is in a dual position as Professor of Computer Science and Professor of Philosophy at the University of Copenhagen. He is a recipient of an ERC Starting Grant, a Google Focused Research, a Carlsberg Semper Ardens Advanced, and has won eight best paper awards. He has written more than 300 articles and five academic books.

Keynote Speaker

Untrustworthy and still revolutionary: Some thoughts on how LLMs are changing NLP

Alexander Koller

Saarland University

Abstract: There is no doubt that large language models (LLMs) are revolutionizing the field of natural language processing (NLP) in many ways. There are many doubts on whether this a good thing, whether we will ever be able to overcome their inability to reliably distinguish truth from falsehood, whether there is any place left for pre-LLM models, and how to do good science any more.

I do not have definitive answers on any of these questions, and am personally torn on many of them. In this talk, I will first discuss some recent research on the limitations of LLMs for semantic parsing and on overcoming them through the use of neurosymbolic models. I will then discuss recent work on the extent to which LLMs can capture world knowledge and apply it to planning and reasoning tasks. I will conclude with some general thoughts on science and engineering in NLP in the era of LLMs.

Bio: Alexander Koller is a Professor of Computational Linguistics at Saarland University in Saarbrücken, Germany. His research interests include planning and reasoning with LLMs, syntactic and semantic processing, natural language generation, and dialogue systems. He is particularly interested in neurosymbolic models that bring together principled linguistic modeling and correctness guarantees with the coverage and robustness of neural approaches. Alexander received his PhD from Saarland University and was previously a postdoc at Columbia University and the University of Edinburgh, faculty at the University of Potsdam, and Visiting Senior Research Scientist at the Allen Institute for AI.

Workshop Presentations