SCIENCES & BIOENGINEERING SCIENCES

The Research Group
Artificial Intelligence Lab

has the honor to invite you to the public defence of the PhD thesis of

Lara Verheyen

to obtain the degree of Doctor of Sciences

Title of the PhD thesis:

Procedural Semantics for Human-like Language Understanding in Situated Environments

Promotors:
Prof. dr. Paul Van Eecke (VUB)
Prof. dr. Katrien Beuls (Université de Namur)

The defence will take place on

Tuesday, December 12, 2023 at 16h in Auditorium D.2.01

Members of the jury

Prof. dr. dr. Geraint Wiggins (VUB, chair)
Prof. dr. Beat Signer (VUB, secretary)
Prof. dr. Stefanie Keulen (VUB)
Dr. Robert Porzel (University of Bremen, Germany)
Dr. Claire Bonial (U.S. Army Research Lab, US)

Curriculum vitae

In 2019, Lara Verheyen obtained her Advanced Master's degree in Artificial Intelligence with a specialization in Speech and Language Technology from KU Leuven, one year after she obtained her master's degree in Linguistics from the same university. Then, she started her PhD at the Artificial Intelligence Laboratory, VUB, under the supervision of prof. dr. Paul Van Eecke and prof. dr. Katrien Beuls. Her research resulted in two peer-reviewed publications as first author and three more peer-reviewed papers at international conferences and workshops.

Abstract of the PhD research

In recent years, data-driven approaches have become the predominant paradigm in the field of natural language processing. These approaches mostly rely on statistical patterns inferred from large collections of textual data. Since such systems achieve impressive results on a variety of NLP tasks and because they exhibit high levels of formal and syntactic correctness, it is often assumed that these systems understand language in a human-like way. While this might appear to be the case, these systems primarily focus on the form side of language, since they are mostly learned from textual data that is not grounded in the world. It can therefore be argued that they deal with language in a way that is fundamentally different than the way in which humans do, i.e. by constructing meaning through interactions with each other and their environment.

In this thesis, I investigate how aspects of human-like language understanding can be modelled by building systems that each focus on different parts of human-like language understanding. This research resulted in three concrete contributions. A first contribution relates to the assumption that language cannot be separated from the environment in which it is used. Concretely, I present a system that is able to ground language in its environment and memory by introducing a procedural semantics that integrates these elements. This novel methodology achieves state-of-the-art results on two benchmark datasets for the task of visual dialogue. A second contribution consists in a model that starts from the idea that language is inherently connected to individual knowledge, since personal experiences shape how humans interpret language. This system integrates language with an agent's personal and dynamic knowledge system. Here, a proof-of-concept implementation demonstrates how agents can come to different interpretations of the same linguistic utterance through their individual knowledge. A third contribution starts from the assumption that language understanding becomes truly human-like when systems can reflect on their own language understanding and signal when they might fail to understand, for instance due to a lack of knowledge. It concerns the development of a system that allows an agent to monitor its own process of language understanding. This allows an agent to estimate how well it has understood a given text and to identify and signal when it has misunderstood certain aspects. This monitoring system is applied on case studies from two different tasks: a visual dialogue task and a recipe understanding task. These systems illustrate how certain aspects of human-like language understanding can be computationally modelled and thereby provide a more human-like alternative to today's data-driven NLP systems.