Using Conceptual Spaces for Cognitive AI

In artificial intelligence, one can distinguish two layers of knowledge representation: On the one hand, there is the symbolic layer, where abstract knowledge is represented in a structured, logic-based format. On the other hand, there is the subsymbolic layer, where perceptual knowledge is stored in a numeric way, e.g., in the form of weights within a neural network.
The highly influential cognitive framework of conceptual spaces proposes an intermediate representational layer based on geometry: Abstract symbols from the symbolic layer are identified with regions in a similarity space whose dimensions are based on subsymbolic perceptual processing.

In my talk, I will introduce the framework of conceptual spaces as a suitable tool for cognitive AI. In order to be of practical use, three main challenges need to be solved:

1.) In order to implement conceptual spaces, one needs a thorough mathematical formalization of its key ideas. I will give a short overview of my mathematical formalization of the conceptual spaces framework which focuses on representing the correlations between domains (such as “red apples are sweet and green apples are sour”) in a geometric way.

2.) In order to connect the conceptual space to perception, one needs to devise a principled way of grounding its dimensions in subsymbolic processing. I will show first results of a hybrid approach which aims to combine artificial neural networks with psychological similarity ratings in order to learn a generalizable mapping from raw sensory input to psychologically grounded similarity spaces.

3.) In order to connect the conceptual space to abstract reasoning, one needs to identify meaningful regions inside the conceptual space and map them onto symbols in the symbolic layer. I will argue that this problem can be solved with classical machine learning techniques such as classification and clustering. I will introduce Logic Tensor Networks as an especially promising tool for this purpose.