

Searching for general models that learn compositionally

Marco Baroni
Facebook AI Research
Paris, France

Abstract

For all their impressive results, current machine learning algorithms are only able to solve one problem at a time, and they can't generalize even across closely related tasks: A system trained on chess will have to learn to play checkers from scratch. One reasonable hypothesis for the difference between algorithms and flexible intelligent beings, who do not have this limitation, is that only the latter are capable of compositional learning. That is, when faced with a new challenge, an intelligent being will 1) try to combine skills it already possesses in novel ways to solve the new task; 2) if new skills are called for, it will store these newly acquired skills in its cognitive toolkit, to recycle them in the future. In this talk, I will describe our efforts to develop a very general algorithm that displays the same compositional learning capabilities. The ideas we are exploring include: allowing the system to progressively freeze some of its weights to store previously acquired skills; separating input/output steps from computation steps, to allow the system to perform multi-step "reasoning" on its acquired knowledge; learning objectives that encourage generalization rather than rote learning.

Joint work with: Tomas Mikolov, Germn Kruszewski, Adam Liska, Rahma Chaabouni, Allan Jabri