

Guesswhat?! - Learning strategies for visually grounded dialogue

Olivier Pietquin

Google Brain

Paris, France

Abstract

In this talk we will present a methodology for end-to-end learning of dialogue systems. Especially, the emergence of a grounded language in goal oriented dialogues through a fully data-driven approach will be addressed. To do so, we will present the Guesswhat?! game and the associated database. Guesswhat?! is a language-based game supported by an image. A database of 150k dialogues has been collected and is freely available for research. Code for supervised learning baselines is also available. In addition, we will present recent work on Reinforcement Learning applied to that environment and some improvements brought to the supervised learning approach based on conditioning on language a feature-wise modulation of convnets.