On the Way to Grounding Referential Behavior

Matthias Scheutz (mscheutz@indiana.edu)
Cognitive Science/Computer Science, Indiana University
Bloomington, IN 47404

Introduction

Reference is undoubtedly still one of the most important, yet bothersome notions that the philosophy of cognitive science is confronted with. Although significant progress has been made in the logical analysis of this concept, the cognitive aspects of how humans manage to refer to (arbitrary) things is not as of yet well understood. How do humans acquire the skills to refer to arbitrary things? What are the mechanisms at work? I believe that a purely logico-philosophical analysis is neither sufficient to explain nor does justice to a discussion of human referential behavior.[1] Therefore, I propose an empirical model which takes the *dynamics* and *constraints of the real world* seriously (Scheutz and Tillotson, 1996). This embodied model, "situated" in its environment, should *learn* to establish "links" between two different modalities and to use those links to "produce and resolve reference" (i.e., referential behavior of the robot should "emerge" from more primitive behaviors). Among the many reasons to prefer a physical agent over a mere simulation of the agent--as described in Brooks (1991)--I would like to mention the potential of this approach to "ground" reference by coupling sensor and motor actions with the environment (see Harnad, 1990).

The Project in a Nutshell

The (autonomous) robot "Eeyore" used for the project is currently under development by a group of graduate students at IU.[2] It consists of a three-wheeled body, carrying a PC as well as the batteries that supply two motors for locomotion. A head with a camera and two microphones is mounted on the body allowing the robot to look around and localize sounds. In addition, speakers will be added later to permit speech production. For this project Eeyore will first have to learn to *survive* and *orient itself* in the given environment (described below). Later, abstracting over individual instances, it will learn to associate certain types of visual and auditory stimuli. These associations can be used to give commands to the robot, and will also serve as a touchstone to check whether the robot has made the right connection between a "name" and its "denotation", i.e., whether it has learned to *resolve* the reference (with the additional speech production system, it could also *produce* reference). Although learning to associate names with the objects (and/or actions) they refer to, and (re)producing those names to refer to such objects (and/or actions), is obviously still very far from grounding human language, I believe that learning to use "names" in a *dynamic* environment where the robot has to function under real-time conditions will offer valuable insights into the nature of referential behaviors in humans.

The Pathfollowing Task

So far, only the first part of the project has been implemented; it concerns the mere "survival" of the robot in its environment, a network of red brick paths leading through a small forest in front of the computer science department of IU. The robot's task is to navigate on these paths through the forest without leaving them. In order to do this, the robot's visual field is divided into various parts of different size, so-called "sensors", which determine whether or not they are "sensing" (parts of) the path. The actual configuration of all sensors, in turn, defines the current state of the system in which certain actions have to be undertaken depending on whether the robot is secure in the middle of the path or is in danger of leaving it. A global "goodness" function is computed expressing "how much of the path" the robot actually sees (the more path, the better) which is then used for reinforcement learning: whenever the goodness value is reduced because the robot is close to getting off the path, either a stored maneuver is retrieved or a new one is created non-deterministically to improve goodness. That way the robot adapts to the current situation and stores ("reinforces") successful maneuvers, while forgetting unsuccessful ones as it discovers the world (in other words, the state space of the finite state machine is learned and altered dynamically!).

Acknowledgments

I would like to thank all members of the "Eeyore" robot group for their suggestions, their help, and their support, as well as the Cognitive Science Program at IU for funding the robot project.

References

Brooks, A. R. (1990). Intelligence Without Reason. *Computers and Thought*, IJCAI-91.

Harnad, S. (1990). The Symbol Grounding Problem. *Physica D 42*.

Scheutz, M., & Tillotson, J. (1996), A Dynamic View of Reference. *Proceedings of MAICS96*.

[1] I will use the term "human referential behavior" to emphasize my conviction that *reference* cannot be fully understood without taking the "human factor" seriously!

[2] Eeyore is intended as a multi-purpose platform for and sponsered by the IU cognitive science program to test cognitive models that need a "physical agent"