YODA: The Young Observant Discovery Agent

Wei-Min Shen, Jafar Adibi, Bonghan Cho, Gal Kaminka,
Jihie Kim, Behnam Salemi, Sheila Tejada
Information Sciences Institute
University of Southern California
Email: shen@isi.edu

The YODA project at USC/ISI consists of a group of young researchers who share a passion for autonomous systems that can bootstrap their knowledge of real environments by exploration, experimentation, learning, and discovery. Our goal is to create a mobile agent that can autonomously learn from its environment based on its own actions, percepts, and missions [1].

The current YODA system has a Denning MRV-3 mobile robot and an on-board portable personal computer. The robot is a three-wheel cylindrical system with separate motors for motion and steering. It is equipped with 24 long range sonar sensors, three cameras for stereo vision, a speaker for sound emission, and a voice recognition system. The communication between the robot and the control computer is accomplished through an RS232 serial port using a remote programming interface [2]. The robot is controlled by a set of commands and the sensor readings include sonar ranges, motor status, and position vectors (vision is not used in this competition). As with any real sensing device, the sensor readings from the robot are not always reliable and this poses challenges for building a robust system.

YODA's software is implemented in MCL3.0 on a Macintosh Powerbook computer. The control architecture is designed to integrate reactive behaviors with deliberate planning. It also has facilities to accommodate learning and discovery in the future. Currently, the architecture is divided into two layers: the lower layer contains modules for navigation and reactive behaviors; and the higher layer contains modules for mission (re)planning.

The lower layer is a combination of production rules and behavior-based systems [3] with each behavior represented as a set of rules. These rules are different from traditional productions in that they have associated probabilities and make predictions. We believe the probabilities are important for robust behaviors in a real environment, and the predictions (i.e., the expected consequences of the robot's actions) are the key for self-organized learning and discovery. Given a set of goals delivered from the higher layer, the behavior rules will compete with each other based on the current sensor readings. Actions associated with the winning rules are then executed (these actions are in some sense collaborating with each other). This cycle of perception, decision, and action repeats itself until the goals are accomplished. In the case where the goals are impossible to reach, the lower layer will pass that failure to the higher layer for replanning.

To deal with imprecision in sensor readings, the system fuses information from different sensors and treat them as a whole. Decisions are never made based on any single sensor but on "patterns" of sensed information. This gives the system the ability to incorporate the idea of "signatures" for recognizing locations and features in the environment. For example, YODA has four basic built-in signatures based the 24 sonar sensors for this competition: corridor, corner, at-a-door, and at-a-foyer. These signatures, combined with the map and the (imprecise) position vector, greatly increase the reliability of determining the current location of the robot. To detect whether a room is occupied or not, the system will enter the room and ask verbally if any person in the room would like come closer to the robot. If this results some changes in the sonar readings, we conclude the room is occupied. Otherwise, after several trials of asking, we conclude the room is empty.

The movement control system of YODA is designed to bypass any obstacle during its movement towards a goal. However, if an obstacle is so large that the passage to a goal position is completely blocked, the robot will stop and wait until the obstacle moves away. Currently, we are implementing and testing the system on a real floor in the ISI office building, with furniture and people traffic present.

The higher layer in the architecture contains a planner that controls the mission-oriented and long-term behaviors. This planner uses a given map and determines a sequence of goals that must be accomplished by the robot. It can also replan based on information gathered by the lower layer (such as that a corridor is blocked or a door is closed). There are two alternate criteria for finding the best solutions: the "risky" one for finding a shortest path based on the current information, and the "conservative" one for taking into account any dynamic information that might be collected during a plan execution. We are currently investigating the trade-off between the two.

Finally, we would like to thank Dr. Ramakant Nevatia for providing us with the Denning Robot. Special thanks also to the various projects and people in ISI's Intelligent Systems Division for their moral support and their tolerance for sharing space (and occasionally "forces") with YODA.

References

[1] W.M. Shen. *Autonomous Learning from the Environment.* W.H. Freeman, Computer Science Press, 1994.

[2] Denning Robot Manual. *Denning MRV-3 Product Maunal.* Denning Mobile Robotics Inc. 1989.

[3] Ronald C. Arkin. Motor Schema-Based Mobile Robot Navigation. *International Journal of Robotics Research.* 1987.