CS/NEUR125 Brains, Minds, and Machines

Assignment 3: Reasoning about Actions in a Mobile Robot

Due: Tuesday, March 14

This Assignment is a guided reading of the paper, Navigation with Learned Spatial Affordances, by Dr. Susan Epstein and her colleagues, which was presented at the 2015 Annual Meeting of the Cognitive Science Society. Reading this article will prepare us to discuss this work during our third Journal Club in class on Friday, March 17.

To begin, create a copy of this Google document and modify the title of the copy to include your name. Questions that you should submit answers to are shown in blue. As with labs, you'll turn in this Assignment by sharing your copy of this Google document with Ellen and Mike.

This paper by Epstein et al. is a primary research article. Unlike our first two Journal Club papers that addressed questions in empirical neuroscience, this article addresses questions about the computations that enable a biological or robotic system to perform a certain task. Similar to other primary research literature, the authors are likely to assume knowledge in the reader, or leave out details that are familiar to experts. We'll try to fill in some of the assumed knowledge with this document, but we again have to accept that we won't be able to digest and understand every line in this technical paper. Our goal is not to understand every line, but to explore the general idea about how a robot (or biological agent) might reason about its actions as it navigates around a complex environment.

Assignments such as this will help to develop your ability to express ideas about computational approaches more clearly and precisely. For this assignment our emphasis is on helping you to understand the paper rather than testing your ability to interpret it on your own. As in previous assignments, part of what we're practicing, is choosing where to put our reading effort in order to get what we want to get out of the paper. In this case, we want to get a basic understanding of the authors' approach to the design of their system, and to learn enough about their methods to understand what they're trying to show in each figure, and how they are evaluating the performance of their system. Many of the questions below are trying to get you to look up information in specific paragraphs or sentences of the paper. If you use phrases from the paper in answering the questions, you must put them in quotation marks, and you should try to reformulate the idea in your own words.

Because it's easy to get bogged down in technical details in a paper, we first want to understand what is the question or hypothesis the authors are trying to address with their study. That way you can try to relate everything else you read to answering that question--and if it doesn't help address the main question, you might be able to safely ignore it.

Q1. What navigation skill are the authors trying to simulate in this study? What is the authors' basic *thesis* about how this skill may be learned? Note that a *hypothesis* is a statement that can be proved or disproved, while a *thesis* is a short, direct statement that summarizes the main point or claim that is supported by the research.

Q2. How do the authors characterize the *world* in which they want their system to be able to navigate? What are examples of such worlds noted by the authors? Give one example of a different world that you think has these same characteristics.

In the context of the work described in this paper, one can think of a *cognitive architecture* as a formal system for representing knowledge and using this knowledge to reason and make decisions about actions to perform in the pursuit of particular goals. The authors use a cognitive architecture known as FORR (FOr the Right Reasons), which was proposed by Epstein (1994) as a model for developing expertise at problem solving in a specific domain. This cognitive architecture is applied to the domain of spatial navigation by a mobile robot, in a system the authors call *SemaFORR*.

Q3. SemaFORR is designed in a way that enables a robot to *explain the reasons for its actions*. Why is this important? In other words, what is an intended application of the robots for which this capability would be useful?

Q4. What is the *sensory information* available to the robot as it moves around its environment?

In addition to the above sensory information, it is assumed that the robot has access to its current (x,y) location in a 2D coordinate frame, provided by overhead cameras. A *task* for the robot is to visit a *target* in its environment, whose (x,y) location is specified by the user.

Q5. Imagine that you are blindfolded and placed at a location outside the Leaky Beaker Café in the Science Center. Suppose you are given your current coordinates (e.g. (0,0)) and heading direction (e.g. 0º, toward the Science Library), and are asked to navigate to a new location (e.g. (10,20), with units in meters). Assume that you have no memory of the layout of the space around you, and that you can only sense the environment by reaching out with your arms. What are some challenges that you may face, while trying to reach your goal?

Q6. The robot senses its environment at *decision points*, which are locations at which the robot needs to make a decision about which action to perform next. What information is available to the robot for making a decision about which action to perform? What is the set of fixed actions that the robot can select from, at each decision point?

In the Introduction, the authors describe a *spatial affordance* as "a spatial abstraction that supports navigation." It is an abstraction in the sense that it removes "perceived but irrelevant details from spatial knowledge." This term is especially relevant in a spatial environment with lots of obstacles to movement, like walls. ***Intuitively, "spatial affordances" refer to areas you can move through, like rooms, pathways, or exit doors; because they "afford" one the opportunity to move through that spot.*** SemaFORR's spatial affordances are intended to facilitate the movement of the robot to reach a target. There are three types of spatial affordance described in this paper: (1) corrected paths (referred to as *trails* in later work), (2) conveyors, and (3) regions with exits. Over time, the system learns the particular spatial affordances (trails, conveyors, and regions) that exist in its environment. The collection of conveyors and regions learned by the robot for a particular world constitute its *spatial model* of the world.

Q7. What is a *corrected path* (see example in Figure 2(a))? What irrelevant details about the robot's *true path* are removed from this abstraction?

Q8. What are *conveyors* (e.g. shown in Figure 2(b))? How might they be helpful for future navigation tasks?

Q9. What is a *region*, and why do the regions depicted in Figure 1(b) have different sizes? Note that an *exit* from a region is a place on its perimeter where the robot has previously traveled, so it represents a place where the robot can safely move out of a region.

In SemaFORR, a decision about which action to perform next is based on the assessment of a set of *Advisors* that each embody a distinct rationale for either promoting or discouraging certain actions. The robot can only consider actions that are advocated for by a particular Advisor. Three Advisors have a special status and are referred to as Tier 1 Advisors: `VICTORY`, `AVOIDWALLS`, and `NOTOPPOSITE`. These Advisors either force a particular action to take place, or remove certain actions from consideration.

Q10. For each of the three Tier 1 Advisors, what is their rationale and which actions do they support or veto?

When there is no action mandated by one of the Tier 1 Advisors, the remaining *Tier 3 Advisors* contribute collectively to the decision of which action to perform. More specifically, each Tier 3 Advisor considers each of the potential actions (5 moves, 10 turns, or pause), and associates each action with a numerical *strength* that captures how much the Advisor supports or opposes the action. This strength is based on the Advisor's rationale, the current state of the robot, and the current spatial model of the world. Consider, for example, the `CLOSEIN` Advisor. When the current state of the robot indicates that it is close to the target, this Advisor associates high strength with "move" actions that would bring the robot closer to the target, and also associates high strength with "turn" actions that would rotate the robot so that it is facing in a direction closer to the target. Three of the Tier 3 Advisors make use of the "spatial affordances" described earlier: `CONVEY` uses information about conveyors--specifically which ones have been used a lot in the past, while `EXIT` and `UNLIKELY` use information about current regions--specifically memories of where the known exits from each region are located, and regions where the target was <u>not</u> seen previously.

Q11. Select three Tier 3 Advisors (other than `CLOSEIN`), and for each one, describe its rationale and the actions that this Advisor might support or oppose.

To decide on an action based on the assessment of the Tier 3 Advisors, SemaFORR tallies the strengths provided by each of the Tier 3 Advisors for each action, and then selects the action that receives the most support. Decisions about moves and turns of the robot are made on separate *decision cycles*, that is, the system alternates between deciding on the next move (which could also be a pause) and deciding on the next turn.

Q12. Suppose two Tier 3 Advisors disagree on which action to perform next, for example, one Advisor supports a certain action while another Advisor opposes it. Can such disagreements be resolved? Explain why or why not.

The authors tested the SemaFORR architecture through computer simulations in which they created three artificial worlds and simulated many trials of the robot navigating to 40 different targets in each of these worlds. They compared performance of SemaFORR with that of a standard method used in robotics known as A*, which computes the optimal (i.e. shortest) path between any two locations, given a detailed map of the environment (SemaFORR, in contrast, does not use an explicit map of its environment!)

Q13. To test the efficacy of various Advisors used to guide the actions of the robot, the authors tested versions of the system (called "Navigators") that included different subsets of the Tier 3 Advisors. Drawing an analogy to Mike's distinction between *mimic*, *block*, and *measure* experiments in the study of biological systems, which type of experiment are the authors performing here? Which Advisors (listed in Table 1) were tested in the versions of SemaFORR labeled as SemaFORR-B, SemaFORR-E, SemaFORR-C, and SemaFORR-R, in Table 2?

Q14. For each of the three worlds (A, B, C), Table 2 lists the **Success** for each Navigator. How is Success measured, i.e. what do the percentages mean? Considering only this measure of Success, what can you conclude about the usefulness of the various Advisors? In particular, how does the use of Advisors *that are able to learn* about conveyors and regions impact the performance of the system?

Q15. Considering the work on human spatial cognition and navigation behavior described briefly in the section on Related Work, list three examples of aspects of SemaFORR that the authors believe are supported by human studies.

Q16. Please submit two questions you have about terms, figures, concepts or anything in this paper that confused you or that you'd like to pursue further during our Journal Club discussion. For example, one question might be related to a technical detail, and another might be broader (e.g. related to assumptions, methods, interpretation, or open questions for future research).

Reference

Epstein, S. L. (1994) For the Right Reasons: The FORR Architecture for Learning in a Skill Domain, *Cognitive Science 18*, 479-511.