Modeling Individual Differences in Learning a Navigation Task

Diana Gordon (gordon@aic.nrl.navy.mil)*
Devika Subramanian (devika@cs.rice.edu)**
Mary Haught (airth@itd.nrl.navy.mil)*
Robynn Kobayashi (rkobayas@aic.nrl.navy.mil)*
Sandra Marshall (smarshall@sciences.sdsu.edu)***
*Naval Research Laboratory, Code 5510, Washington, D.C. 20375
**Department of Computer Science, Rice University, Houston TX 77005
***Department of Psychology, San Diego State University, San Diego, CA 92120

Abstract

Our goal is to develop a cognitive model of how humans acquire skills on complex, sensorimotor tasks. To achieve this goal, we collected data from subjects learning the NRL Navigation task, then used the data to construct a model that reflects the basic, cognitive elements required to learn and thereby succeed at this task (Gordon & Subramanian, 1997). This paper describes a new experiment with human subjects on the task. Data from this experiment not only confirms the key cognitive element of our model, but also helps us better understand individual differences in learning this task. Four evaluation metrics indicate that we are able to model important trends in the evolution of action choice.

Introduction

Our goal is to model how humans acquire skills on complex, cognitive tasks. We are pursuing this goal by designing computational architectures for the NRL Navigation task, which requires competent sensorimotor coordination. To achieve this goal, we first constructed a model reflecting the basic, cognitive elements required to learn and thereby succeed at the task. The model was engineered from human subjects' data. This model is reported in Gordon and Subramanian (1997), and is briefly summarized here. The metric for evaluating the degree of fit between the subjects and the model is a learning curve, which captures improvement in success rate over time. With respect to this metric, the model is a good match for learning behavior in our subjects.

Two questions from our previous research motivated a new experiment with human subjects on the task. First, a key cognitive element is a focus of attention heuristic, which is used to switch between two subtasks of the Navigation task. Can we confirm this heuristic objectively? Second, although nearly all subjects use the basic cognitive elements, a deeper analysis of the data suggests subjects acquire additional cognitive elements that vary between individuals. Can we better understand these individual differences in what is learned? Here, we describe the new experiment, which uses an eyetracker to monitor subjects' visual focus of attention. The results of this experiment not only confirm the focus heuristic, but they also provide sufficient data for understanding and modeling individual differences. For our subjects, acquisition of a new, cognitive element is accompanied by a shift in perception and action strategy. We do not model the cognitive elements per se. However, using a popular machine learning tool, we model significant shifts in action strategy that are correlated with changes in eyetracker pattern. The verbal utterances that occur during these shifts indicate that they are associated with deep, conceptual shifts which radically alter the subject's view of the task. At least one of the conceptual shifts appears to be motivated by an individual's dislike of a certain type of failure. Therefore, to better understand this and other other subjects' shifts, the success rate evaluation metric is supplemented with two additional metrics which reveal the nature of subjects' failures before and after learning. We also add a fourth evaluation metric consisting of perception and action probability distributions. This is a much stricter performance metric than learning curves, but because we are now modeling individuals' action strategies, it is plausible that we can obtain a reasonable fit, even using this stricter criterion.

The main novelty of this work is the modeling of shifts in action strategy that coincide with conceptual and perceptual shifts, using a suite of revealing evaluation metrics. The evaluation results suggest that important learning trends are captured.

This paper begins with a description of the task and a brief review of the cognitive model in Gordon and Subramanian (1997). The new, human subjects experiment is then described, followed by our current modeling of individuals. The paper concludes with related work and directions for future research.

The NRL Navigation Task

The NRL navigation and mine avoidance domain, developed by Alan Schultz at the Naval Research Laboratory (NRL) and hereafter abbreviated the "Navigation task," is a 2-D computer simulation that can be run either by humans through a graphical interface, or by an automated agent (Gordon, et al., 1994). The task involves a single agent who controls an autonomous, underwater vehicle (AUV) that has to avoid mines and rendezvous with a stationary target (goal) before exhausting its fuel. Time is divided into episodes. An episode begins with the agent on one side of the mine field, and random target and mine locations; it ends with one of three possible outcomes: the agent reaches the goal (success), hits a mine and explodes (failure), or the simulation times out because fuel is exhausted (failure). The outcome is received at the end of each episode.

When human subjects run this task, the sensory input is through visual gauges, and the motor output is controlled by a joystick. A sonar gauge for detecting mines consists of seven squares in a row that provide a 90 degree forward field of view for a short distance. Mines appear as circles in the squares; the mapping between mines and circles is often not one-to-one. Circle size in a square is proportional to mine proximity in that direction. A range gauge provides the target distance, a bearing gauge in clock notation indicates the target direction (12 o'clock means target ahead, 6 o'clock behind), and a time gauge indicates the remaining fuel. The AUV's turn and speed are controlled by joystick motions.

A Cognitive Model

Our goal is to build the simplest model that accounts for human subject data in learning performance. Initial experiments were run in 1994 with five human subjects, using a task configuration of no sensor noise and 25 mines. A cognitive model was constructed from the verbal protocol data alone, and is reported in (Gordon & Subramanian, 1996; 1997).

The verbal protocol data from the 1994 experiments reveals that the most salient aspect of learning and reasoning on this task is a decomposition of the task into two subtasks: avoid mines and navigate to the target. The subtask on which the subject is focused determines his/her action choice.

Our cognitive model M_{focus}, which inputs numeric sensor values and outputs numeric actions, reflects the basic, cognitive elements in its architectural structure. The key element is the model's focus of attention heuristic for selecting actions, which is the following. When the sonar values are below an empirically determined threshold (indicating mines nearby), use sonar predictions to select the best action to take; otherwise use bearing predictions to select the best action. The learning curves generated by M_{focus} match those gathered from our human subjects (Gordon & Subramanian, 1997). It is quite interesting that although verbal protocols can in general be quite unreliable, in this case they provided useful guidance for model engineering.

Experiment

Experimental and Task Configuration

Five subjects ran the task with a task configuration of 60 mines and no sensor noise.[1] An Applied Systems Laboratories (ASL) Model 4000 eyetracker was placed on the head of each subject. The gauge sizes and the visual distances between gauges were sufficiently large to enable the eyetracker to distinguish subjects' focus in almost all cases.[2]

The joystick, custom-made by Thrustmaster, Incorporated, was used to input the turn and speed of the AUV. Joystick conversion routines, written by James Ballas at NRL, convert the joystick position to one of 17 discrete turn values and one of 9 discrete speed values, which are forwarded to the simulation. Turn ranges from -32 (sharp right) to 32 (sharp left), and speed from 0 to 40.

Data Collection Procedure

Subjects ran consecutive episodes during the hour. The number of episodes per hour varied from around 60 to 160. Each episode varied from a few to 200 time steps (action decisions).

Data was collected on three different media: (1) *execution traces* of sequential snapshots of every set of gauge readings and actions taken, along with success/failure feedback at the end of each episode, (2) fixation files of every visual fixation, and (3) videotapes recording the pictorial gauges seen by the subjects on the computer screen, along with a white square denoting the eyetracker's recording of the subject's visual focus of attention, and all verbal utterances of the subject.

All subjects ran for five one-hour daily sessions. At the beginning of the first session, they were told they had to navigate through a minefield to get to a target location and were instructed on how to operate the joystick. Subjects only saw the gauges view of the task. Between episodes, the experimenter occasionally asked them to verbalize what they were thinking and learning.

Data Analysis and Results

One of the most striking results from the eyetracker data is confirmation of the focus heuristic. Novice subjects distribute their focus of attention rather randomly among the gauges. The three subjects who developed expertise at the task eventually converged upon an eyetracker pattern restricted to only the sonar and bearing gauges. When the sonar squares are empty, focus is on the bearing; otherwise, focus is on the sonar. This is

[1] Five undergraduates at San Diego State University participated in this experiment and received $10 per hour as compensation.

[2] To control the brightness in the room, a photographic light meter was used, and the readings were consistently between 9.4 and 9.8 exposure values. Sometimes the eyetracker had to be recalibrated once or twice mid-session, incurring a loss of about 5-10 minutes for each recalibration.

precisely the focus of attention heuristic for switching between subtasks embedded in M_{focus}.

A more detailed analysis reveals striking individual differences. The next section summarizes our analysis and modeling of individuals.

Modeling Individuals' Conceptual Shifts

Recall that we have four evaluation metrics: success rate, two failure rate metrics, and the probability distributions over actions and perceptions. In particular, the two failure rate metrics are the explosion rate and the timeout (fuel exhaustion) rate.

With this suite of metrics, we turn to the modeling. M_{focus} is not able to capture individuals well enough to satisfy our more exacting evaluation metrics. For these metrics, function fitters (e.g., decision trees and neural networks) seem most appropriate. To select one function fitting method, we ran a comparison of the fit to the execution trace (sensorimotor data) of one of our subjects. C4.5 (Quinlan, 1986) performs well and provides the most understandable strategies of all the systems tested; therefore it is selected for further modeling of individuals. From the subject's execution traces, C4.5 learns a decision tree model of the subject's action strategy, which can be summarized by two functions: $sensors \rightarrow turn$ and $sensors \rightarrow speed$. It does not model internal, cognitive elements such as the focus of attention heuristic; future work will address adding the cognitive elements. But first we need to understand what conceptual shifts may motivate the development of these elements, and what are their associated sensorimotor shifts.

We begin by examining the trends in the data to be modeled. When using the timeout, explosion, and success rate measures, we notice that individual subjects go through periods of relatively stable performance, punctuated by substantial improvements in performance along at least one of these three dimensions. Further examination of the data reveals that the performance leaps are associated with radical shifts in conceptualization of the task coupled with shifts in perception then action strategies. The remainder of this section focuses on a study and initial modeling of two of these conceptual shifts, one for Subject 4 and another for Subject 5 (two of the subjects who became experts). Both of these subjects show suggestive evidence for their shifts well before they verbalize them conclusively.

Shifts in both subjects occur gradually and unevenly, but once cemented they correspond to a leap in performance. Let us examine Subject 5's shift first. During session 2, around episode 45, Subject 5 first verbalizes the shift as a hypothesis by stating "only the middle sonar can kill me." By this, the subject means that she can safely ignore all sonar squares other than the middle one, i.e., only a circle in the middle square (which senses mines straight ahead) determines whether the AUV will hit a mine. At this point, the eyetracker pattern shifts from attention on all gauges to attention on only the bearing and sonar gauges. When looking at the sonar, attention is more closely clustered near the middle square, as seen in 50-episode fixation and transition summaries in Figures 1 and 2.[3] In these figures, the row of adjacent squares near the middle of the figure is the sonar gauge. The bearing gauge is in the square just above the middle sonar. Other gauges and regions of interest are denoted with rectangles.

By episode 67, the subject states that her hypothesis is confirmed, and a change in action strategy occurs. In particular, Subject 5's pre-shift strategy is forward motion and more random turn decisions. The post-shift strategy consists of slowing down when the circles get larger, "sweeping" the AUV left and right in an attempt to see the direction with least obstruction, then proceeding in that direction. She keeps the bearing straighter toward the target (12 o'clock) post-shift. Figures 3 and 5 show how Subject 5's action probability distributions changed. All figures are obtained by subtracting the post-shift minus the pre-shift distribution. Positive numbers imply an increase in frequency from pre- to post-shift. Figure 7 shows the change in bearing distribution resulting from her change in action strategy, and Figure 9 shows her accompanying substantial performance improvement.[4] It is very interesting to note that the performance improvement is *exclusively* along the dimension of reduced explosions. This is consistent with Subject 5's stated philosophy that "Timeouts are less bad than explosions."

C4.5 learns a separate pair of functions to model Subject 5 before and after the shift. The results are in Figures 4, 6, 8, and 10. Note that although the magnitudes produced by the model only coarsely approximate those produced by the subject, most trends are captured. For example, both model and subject increase the number of full stops, go straight slightly less often, increase the number of timeouts, reduce the number of explosions, and keep the success rate nearly constant after the conceptual shift. Only the bearing trend is not correctly modeled. When running the simulation with the model's action strategy, one can see that post-shift the model mimics the subject's strategy: slow down when seeing a mine, sweep, then move toward a "hole."

Next, consider Subject 4's conceptual shift. During session 3, Subject 4 shows the seeds of the shift as early as episode 85, with a seemingly purposeful scan across the sonar. At episode 122, Subject 4 shows signs that

[3] Since each sonar square is only 0.6 inch, it is hard to stay focused in exactly the middle square. Also, we hypothesize the eyetracker calibration was off slightly, causing the focus to be slightly left shifted.

[4] Based on empirical data, episodes 48-66 are selected for pre-shift, and episodes 67-82 for post-shift.

a shift has taken place by a change in eyetracker pattern. Before the shift, he checks other gauges periodically. Post-shift, his attention seems tightly focused between sonar and bearing only (see Figures 11 and 12). At this time, the subject tentatively admits to a conceptual shift. Subject 4 states that he can perform the task more easily by visualizing the mines as "stones." By this, he means that he clusters multiple circles in consecutive sonar squares together into a single "stone," with smaller circles on the ends represent the receding sides of the stone. This perceptual clustering correlates with a change in action strategy. After the conceptual shift, Subject 4 goes faster, goes straight more often, and keeps the bearing to target straighter. The result is an improvement in the success rate, accompanied by a substantial reduction in timeouts. Figures 13, 15, 17, and 19 show his post- minus pre-shift performance differences.[5] Subject 4 does not verbalize the conceptual shift conclusively until episode 136.

C4.5's model of Subject 4 is shown in Figures 14, 16, 18, and 20. Again, the trends, but not the magnitudes, are closely modeled. Like Subject 4, post-shift the model tends to go faster and straighter, and increases the frequency that the bearing is at 12 o'clock. Success rate is increased, timeouts decreased, and explosions increased. Only with timeouts is the magnitude close.

Although most trends are nicely captured, the action distributions generated by C4.5 are statistically (using a chi-squared test) significantly different at the 99% level from those of the subject. It appears that modeling of internal state is crucial to meet this most stringent fit criterion. Therefore, we are currently exploring stochastic finite state automata (SFSAs, also known as hidden Markov models), which so far seem to provide a much better fit to action probability distributions. Since SFSAs model internal states, they facilitate integrating the internal, cognitive elements of M_{focus}. Additionally, we can include the cognitive elements learned during the shifts just described. The cognitive element acquired by Subject 5 is a refinement of the focus of attention heuristic, with attention more toward the middle sonar square. Subject 4's cognitive element will be more challenging to characterize and model. We surmise it may be a mental model of mines.

Discussion and Related Work

Sun and Peterson (1997) use their CLARION architecture to model learning on the NRL Navigation task. Comparisons between our model M_{focus} and theirs on this task are in progress. Gray and Kirschenbaum (1997) also study strategy selection on a complex task. The research of John and Lallemont (1997) is closely related to ours because they also (1) model learning on a complex task, (2) study individual differences, and (3) study shifts in strategy choice. Nevertheless, the novelty of *our* work is that we study the action strategy shifts that coincide with conceptual and perceptual shifts, and we identify a suite of performance measures for which individuals reach differing levels of expertise in unique ways. Finally, literature on insights is related – these are what appear to prompt the conceptual shifts. Our findings confirm Metcalfe's (1986) experimental results, for example, which show that on insight types of problems her subjects demonstrate lack of confidence at problem solving until the moment of insight, at which point confidence jumps to a high level. Both Subjects 4 and 5 expressed frustration and low confidence prior to their conceptual shifts.

Future work will focus on developing and fusing the cognitive elements of M_{focus} with the function fitting approach, using SFSAs.

[5] Based on empirical data, episodes 106-122 are selected for pre-shift, and episodes 123-142 for post-shift.

Acknowledgements

This research was sponsored by the Office of Naval Research N00014-95-WX30360, N00014-95-1-0846 and N00014-96-1-0538. Thanks to Bill Spears for useful comments on this paper.

References

Gordon, D., Schultz, A., Grefenstette, J., Ballas, J., & Perez, M. (1994). *User's guide to the navigation and collision avoidance task* (AIC-94-013). Washington, D.C.: Naval Research Laboratory.

Gordon, D., & Subramanian, D. (1997). A cognitive model of learning to navigate. *Proceedings of the 18th Annual Conference of the Cognitive Science Society* (pp. 271-276). Mahwah, NJ: Lawrence Elrbaum Associates.

Gray, W., & Kirschenbaum, S. (1997). The precis of project Nemo, Phase 1: Subgoaling and schemas for submariners. *Proceedings of the 18th Annual Conference of the Cognitive Science Society* (pp. 283-288). Mahwah, NJ: Lawrence Elrbaum Associates.

John, B. & Lallement, Y. (1997). Strategy use while learning to perform the Kanfer-Ackerman air traffic controller task. *Proceedings of the 18th Annual Conference of the Cognitive Science Society* (pp. 337-342). Mahwah, NJ: Lawrence Elrbaum Associates.

Metcalfe, J. (1986). Premonitions of insight predict impending error. *Journal of Experimental Psychology, 12(4)*, 623-634.

Quinlan, J. R. (1986). Induction of decision trees. *Machine Learning, 1*, 81-107.

Sun, R., & Peterson, T. (1997). A hybrid model for learning sequential navigation. *Proceedings of the Conference on Computational Intelligence in Robotics and Automation* (pp. 234-239).