The influence of experienced effort on learning and choice of solution paths in a simple navigation task

Wai-Tat Fu (wfu@cmu.edu)

Abstract

In my dissertation, I am going to study the influence of experienced effort on learning and problem-solving behavior. In a simple navigation task on a computer-simulated map, subjects have to acquire information on various levels of effort in different solution paths through experience. The experienced effort information allows subjects to improve their performance by finding faster solution paths on the map. I am planning to build ACT-R model to understand the underlying mechanisms. Specifically, I am interested in modeling the learning of the "current" effort and the "downstream" effort in ACT-R theory, and how each of them can influence problem-solving behavior in the task.

Introduction

For many problems, there are multiple solution paths that lead from the initial problem state to the goal state. Different paths may require different amount of time and effort. With experience, the problem-solver learns to choose solution paths that requires less time and effort, and as a result, performance improves. For example, if the problem is to drive to a particular destination in a city, numerous solution paths may exist. If the person is new to the city, very little knowledge is available to decide which paths to take. In this case, decisions on which paths to take may solely rely on simple heuristics, such as hill-climbing. Unless the city is extremely complex, simple heuristics are usually sufficient to lead the person to the destination. However, although simple heuristics are usually sufficient to provide a solution to the problem, there is no guarantee that the solution is good (or fast, in the current example). Fortunately, with experience, the person may be able to acquire information about the speeds for various routes in the city. With this kind of information, the person may be able to choose faster paths that lead to the destination. Although many cognitive mechanisms have been proposed to account for this kind of learning in problem solving (e.g. Anzai and Simon, 1979, Agre & Shrager, 1994, Lovett and Anderson, 1996), not many studies have directly addressed the effects of experienced effort in problem-solving, and how people learn to choose less effortful solution paths with experience. In my dissertation, I am going to design several experiments to understand how people acquire problem-specific information and how they use the information to improve their performance. Specifically, I will focus on how people learn the amount of effort involved in different solution paths, and how they improve performance by choosing the less effortful paths. I am planning to build cognitive models using ACT-R (Anderson & Lebiere, 1998) to understand the mechanisms behind this kind of learning. In this extended abstract, I will focus on describing the task and the model that I am planning to build.

The Task

I am going to use a simple navigation task in my experiments. In this task, subjects have to navigate from the starting point to the destination on a map, and multiple solution paths exist for all trials. Subjects are given maps as shown in Fig. 1, which shows the map of 3 transport systems, each represented by different colors (blue, green, and brown). Different transport systems have different speeds. Each circle in the map represents a station of one of the transport systems. To go from one transport system to the other, subjects have to use the transfers at the intersections of the transport systems. There are three different kinds of transfers (pink, orange, and black). Different transfers have different speeds.

In each trial, subjects are given a starting station (a blue circle) and a destination (a yellow circle), and the subjects are told that they have to go from the starting station to the destination. To do this, subjects have to click on the intermediate stations one by one until they reach the destination. When subjects click on an intermediate station, a little red line travels from the current station to the station just clicked on. The speeds of different transport systems are reflected by the speeds of the movement of the red line.

The subjects are instructed that different transport systems and transfers have different speeds, but are not told which one is faster and which one is slower. Since there are no numerical representations of the speeds, combining speeds of different transport systems is relatively inaccurate and difficult. This deters subjects from doing a complete mental look ahead from the starting station to the destination.

36 pairs of starting and end points are constructed so that simple hill-climbing will lead to suboptimal solutions. Since in the early trials, subjects do not have information about the speeds of different transport systems, the prediction is that simple hill-climbing

strategy will be adopted – i.e. the most straightforward paths will be chosen. However, with experience, subjects will learn the relative speeds of different transport systems and transfer, and will be able to use this information to find faster paths that go from the starting point to the destination. Action, eyetracking, and verbal data are collected to understand how subjects make their decisions across trials.

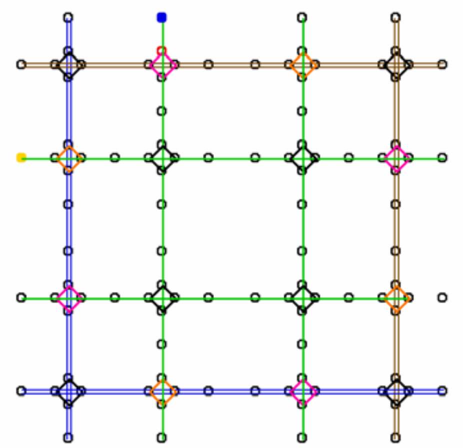

Fig. 1. The map used in the navigation task. There are three different transport systems, and three different transfers. Each of them has different speeds and represented by different colors.

The Model

I am planning to build an ACT-R model of the task. In ACT-R, procedural knowledge is represented as production rules. One of the central premises of the ACT-R theory is that the processes that act on the production rules reflect the statistical structure of the environment. For example, the process of selection among several production instantiations (conflict resolution) is based on the model's evaluation of their expected utility, and the one with the highest expected utility will be executed. The expected utility is calculated as the difference of the expected gain and the expected cost (PG-C) of executing the production. This conflict resolution mechanism allows for the influence of different levels of effort (cost) on the choice of solution paths. In my model, only effort will be taken in account. I believe this is a reasonable simplification, since eventually all paths lead to the destination, the probability of success does not play a significantly role in determining the choice of solution paths.

In ACT-R, the total effort *C* associated with a production is represented by the sum of two parameters: *a* and *b*. (C = a + b) The *a* parameter represents the current effort in executing a particular production; the *b* parameter represents the downstream effort involved between the time after the current production is executed until the time when the current goal is accomplished (popped). The higher the sum of these values, the less likely the production will be executed. Since it is the sum of these two parameters that determine the result of the production selection, it is possible that the model would choose a production rule that has a high current effort (*a*), but a low downstream effort (*b*). Or in other words, the model would initially choose a slow path if the path chosen eventually would lead to a faster overall path.

In the beginning of the experiment, the values of the effort parameters for all productions will be the same. With experience, these parameters will be updated by the formula specified in ACT-R: $a^* = (z + \Sigma effort_i)/(\alpha + \beta + m + n)$, where z is the prior effort, $\Sigma effort_i$ is the total effort taken over all past uses of the production rule, α and β are the prior number of successes and failures, m and n are the experienced number of successes and failures. The formula for b* is the same except that $\Sigma effort_i$ is the total amount of downstream effort taken over all past uses of the production rule.

Some interesting issues are whether subjects are equally sensitive to current and downstream effort, as implicitly assumed by the ACT-R theory. It is possible that people may weight the current effort more than downstream effort, especially when there is limited plan-ahead. If this is so, we may expect to see more localized improvement rather than gradual overall improvement in choosing better solution paths.

ACT-R therefore provides a theoretical basis for understanding the learning and use of experienced effort in problem-solving. By matching the data to the prediction made by the model, we should be able to have a better understanding of the underlying mechanisms for this kind of learning.

References

Agre, P. & Shrager, J. (1990). Routine evolution as the microgenetic basis of skill acquisition. *Twelfth Annual Conference of the Cognitive Science Society* (pp. 694-701), Hillsdale, NJ: Erlbaum.

Anderson, J. R., & Lebiére, C. (Eds.). (1998). *Atomic components of thought*. Hillsdale, NJ: Erlbaum.

Anzai, Y., & Simon, H. A. (1979). The theory of learning by doing. *Psychological Review*, 86 (2), 124-140.

Lovett, M. C. & Anderson, J. R. (1996). History of success and current context in problem solving. *Cognitive Psychology*, 31, 168-217.