A Rational Theory of Cognitive Strategy Selection and Change

Quanfeng Wu and John R. Anderson
Department of Psychology
Carnegie Mellon University
Pittsburgh, PA 15213
wu@psy.cmu.edu, anderson@psy.cmu.edu

Abstract

This paper presents a rational theory of cognitive strategy selection and change in which the cognitive agent in consideration is proposed to be adaptive in choosing the "best" or optimal strategy from a set of strategies available to be employed. The optimal strategy is assumed to maximize the difference between the expected utility of the goal which the selected strategy would lead to and the computational cost associated with achieving this goal. We considered an example of strategy selection and change in computer programming and interpreted the results from a set of experimental studies we had conducted in this domain in the light of this rational framework. We also substantiated our theoretical claims by developing a computer simulation of this example. The simulation was implemented in ACT-R, a cognitive model constrained by rational analysis as well as by experimental data.

Introduction

Cognitive research on strategy selection and change has flourished in the last three decades. Wason (1960) studied the hypothesis-testing behavior of human subjects on a concept attainment task and found that the dominant strategy used by subjects on the task was a kind of "biased" confirmation strategy which had been supposed to be "irrational" according to Popper's (1959) philosophy of science. Wason's research has since promoted a large body of controversy and triggered a series of ensuring studies (Tukey, 1986). Recently, from the perspective of information theory, Hoenkamp (1989) has suggested that "the biased strategy is not necessarily a bad one; moreover, it reflects a healthy propensity of subjects to optimize the expected information on each trial" (p.651). Early in 1960s, Bruner and his colleagues were also among the first to investigate how subjects chose among various strategies in concept attainment (Bruner, Goodnow, & Austin, 1962). Contrary to what Wason's research seemed to reveal, however, they arrived at the conclusion that "in a formal sense it may be said that the subjects in this experiment were seeking to maximize the expected utility of their decisions and in this way to regulate the risk involved in their problem-solving behavior" (p.124). Consequently, according to Hoenkamp and Bruner et al., subjects' performance on concept attainment tasks can in fact be characterized as "rational" or "adaptive" in terms of information or utility maximization.

Recently, there has also been research on strategy selection and change in other cognitive domains—e.g., in memory retrieval (Reder, 1987), in arithmetic (Siegler & Jenkins, 1989), and in programming (Wu & Anderson, 1991). Though conducted in different cognitive domains, these related studies have so far yielded convergent evidence suggesting that subjects are quite adaptive in their strategy selection in that they are highly sensitive to problem types and that they choose appropriate strategies accordingly. Reder & Ritter (1992) have further demonstrated that there is a general tendency of shifting from a computing-on-site strategy to a retrieving-from-memory strategy as subjects practice more and more on problems of a certain type. Thus, taken together, this line of research shows that subjects' strategy selection behavior is not only very sensitive to environmental cues (i.e., problem types) but also strongly influenced by their experience of learning and practicing.

The present paper consists of three parts: First, we attempt to outline a rational theory of strategy selection and change. As mentioned above, the notion of rational strategy choices has existed in the cognitive psychology literature for some time; our objective in this regard is simply to articulate this notion in a more formal manner and to put it in a more general framework of studying human cognition—namely, the Rational Analysis (RA) perspective. Second, we illustrate this rational theoretical framework with an example of strategy selection and change in computer programming. Third, we put forward a simulation model for these results to substantiate some of the theoretical claims we make.

A Rational Framework of Strategy Selection and Change

RA is a new theoretical framework for understanding human cognitive behavior (Anderson, 1990). A fundamental assumption underlying this approach is the *Principle of Rationality* which basically claims that the human cognitive system is adaptive to its informational environment. It is worth noting that the "rationality" here does not mean optimization without any limitation; it merely implies that an individual can achieve an optimal solution only within his or her cognitive or computational constraints. Since the difference between optimization -without-limitation and optimization-with-limitation is somewhat subtle, this issue has in fact caused considerable controversy in evolutionary theory, in economics, as well as in psychology (Simon, 1983).

The idea of human rationality in decision making rooted in economics; in fact, the expected utility model of strategy selection proposed by Bruner et al. (1962) came directly from economics. The traditional notion of rationality in economics, nevertheless, was that there were not cognitive or computational limits on people for decision making. This was an unrealistic assumption of human rationality. In opposing this traditional view, Simon has long been propounding the notion of bounded rationality and arguing that people only adopt satisficing solutions but never optimal ones. Considered that the RA approach only assumes optimization within computational constraints, RA and Simon's notion of bounded rationality can in fact be reconciled (for more detailed discussion, see Anderson, 1990; pp. 246-250).

In Anderson (1990), much attention has been paid to developing RA theories of memory, categorization, causal inference, and problem solving. Recently, Anderson and Kushmerick have also endeavored to develop an RA theory of strategy selection (Ch. 5; Anderson, in preparation). Basically, strategy selection can be conceptualized as choosing among a set of branches in a conceptual tree with each branch leading to a certain goal. For each branch, there would be a probability of success P_i, an utility value for its goal U_i, and a computational cost C_i associated with. Under this characterization, a rational choice among all available strategies would amount to choosing the branch which satisfies the following:

$$\text{Max}\{P_i U_i - C_i \mid i \text{ goes through all available strategies}\}. \quad (1)$$

Note that the set consisting of all available strategies may be dynamic; that is, some new strategies may only become available during the course of problem solving. Thus, (1) is the basic tenet of our rational theory of strategy selection and change presented in a formal fashion; it directly corresponds to the underlying concept of the cost-benefit analysis approach in economics (Varian, 1987). In cognitive psychology, Russo & Dosher (1983) and Payne, Bettman, & Johnson (1988) have similarly argued that human strategy selection involves not only maximizing expected utility but minimizing cognitive effort as well.

Relating to the memory structure of human cognition (e.g., see Anderson, 1983), there are basically two factors involved in the computational cost term C_i in (1): the cost of retrieving from long-term memory (LTM) and the cost of calculation. In other words, to choose and apply a certain strategy, there would be a cost associated with retrieving that strategy from LTM as well as a cost of calculating the details and actions of that strategy. As these two factors are contributing to the same term, we would expect that a tradeoff between these two factors may sometimes be involved in strategy selection and change. We would further expect that the more one practices with a certain strategy, the lower the retrieving cost for that strategy would be, and consequently the more often that strategy would be selected. In fact, this is the shift from calculating-on-site to retrieving-from-memory which Reder & Ritter (1992) observed in their experiments. On the other hand, if an individual is only naive with a certain strategy, we can conceive that in such a case the calculation cost would be smaller than the retrieval cost. In other words, what is already in the individual's current consciousness may influence his or her strategy selection behavior in subsequent situations; in consequence, it can be expected that there be some lateral transfer effect occurring from solving earlier problems to solving later problems for such an individual.

Iterative Strategy Selection and Change in Programming

We had conducted a set of experimental studies on how PASCAL programmers would choose and change iterative strategies in their programming. Since the quantitative results had been published elsewhere (Wu & Anderson, 1991), here we only show some major points of these results. However, we will analyze an episode of strategy selection and an episode of strategy change in some detail since these episodes are very illustrative of our theoretical claims of rational strategy selection and change and that they had not been reported previously.

There are two indefinite looping constructs in PASCAL—namely, the **while...do** and the **repeat...until** constructs. For convenience, these two constructs or strategies will hereafter be referred to as **W**- and **R**-constructs or strategies. To implement any kind of looping program, either construct alone would suffice; nonetheless, in certain cases using the **W**-construct would produce a more concise and well-structured program than using the **R**-construct, and in other cases it is just the opposite. The general principle for choosing between the **W**- and the **R**-constructs is to use the **W**-construct for looping programs where the looping body may not be executed at all and to use the **R**-construct for cases where the looping body must be executed at least once. According to this principle, we can classify looping problems either as **W**-problems, for which it is easier and more natural to use the **W**-construct, and **R**-problems for which it is better to use the **R**-construct. An example of **W**-problem and an example of **R**-problem, which we actually used in our experiments, along with their modal PASCAL solutions are illustrated in Figure 1.

The first of our experiments was to investigate how programmers would choose between the two looping strategies on different types of problems. The subjects involved in the experiment were recruited from CMU; their programming experience ranged from having just finished an introductory programming course in PASCAL to highly skillful (e.g., having more than ten years of programming experience). The results from the experiment turned out to be that there was a minority of the subjects (about 20%) who idiosyncratically used only one type of looping construct over all the problems tested in the experiment while the majority (the rest 80%) did vary their choices of looping strategies on different problems. For these 80% subjects, Figure 2 (a) shows the pattern of their choices on the two types of looping problems. As the figure shows, these subjects were in fact very sensitive to problem types and quite adaptive in choosing appropriate looping strategies; a one-way ANOVA performed on the data for the **W**-strategy in programming revealed that the effects due to problem types were significant. In another experiment, we tried to see whether subjects' performance would deteriorate if they were forced to use an unnatural strategy, i.e., to use the **W**-construct on **R**-problems or *vice versa*. Figure 2 (b) shows the major results from this experiment; statistical analyses revealed that the effect due to experimental manipulations and its interaction with problem types were both significant. Therefore,these results clearly indicated that when the subjects were forced to use a non-perferred strategy their performance in terms of programming time did suffer. As to be shown, this performance deterioration can be accounted for by the higher computational cost associated with the non-perferred iterative strategy.

A W-problem:

Copy a file into another file until a record with its name field as 'END' is reached, but excluding this record.

Its modal solution:

(only main portion shown):

```
. . . . . . . . .
while File1^.Name <> 'END' do
   begin
      Read(File1, TempRecord);
      Write(File2, TempRecord)
   end;
. . . . . . . . .
```

A R-problem:

Copy a file into another file until a record with its name field as 'END' is reached, including this record as well.

Its modal solution:

(only main portion shown):

```
. . . . . . . . .
repeat
      Read(File1, TempRecord);
      Write(File2, TempRecord)
until TempRecord.Name = 'END';
. . . . . . . . .
```

Figure 1. Examples of W-problem and R-problem and their modal PASCAL solutions.

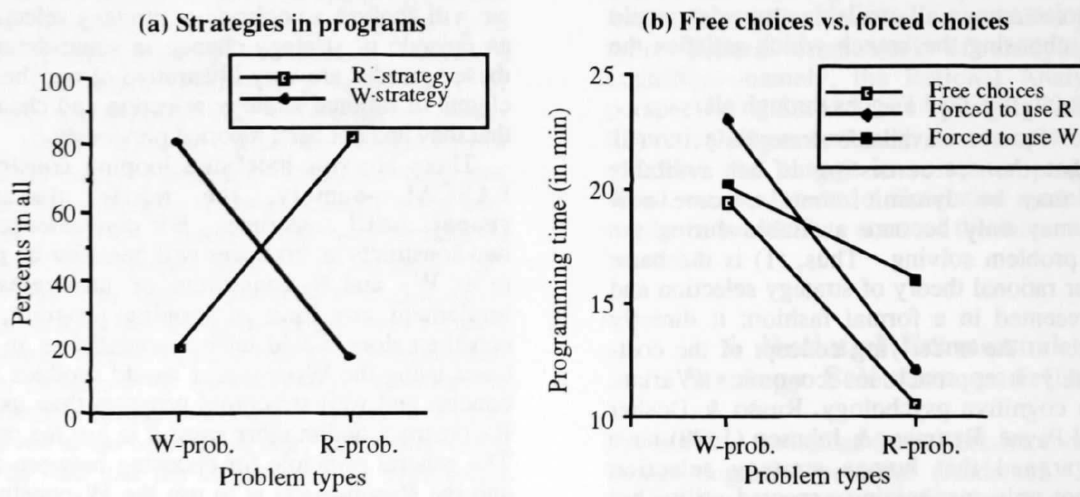

Figure 2. Major results from experiments on iterative strategy selection and change.

Among the 80% adaptive subjects, certainly some were highly adaptive and others were only moderately adaptive (for more details, see Wu & Anderson, 1991). As discussed before, we can expect transfer from solving earlier problems to solving later problems if the subject was only moderately adaptive; furthermore, if the effect was of a negative type, this would amount to the well-known Einstellung effect in problem solving (Luchins & Luchins, 1959). We collected verbal protocols in our experiments, and some programming episodes showed that this type of negative lateral transfer effect did occur in the experiments. Figure 3 presents an example of such transfer effect. The subject involved in this example solved the **W**-problem shown in Figure 1 (a) first; his solution to the problem, however, was not well-structured since he used the **R**-construct instead and an inside-loop **if...then** statement which would be unnecessary otherwise. Simply judging from this solution, we could infer that the subject was a novice programmer. The problem subsequent to the **W**-problem was the **R**-problem shown in Figure 1 (b). When the subject came to this **R**-problem, he (mentally) retrieved his solution to the preceding **W**-problem and only made a minor change of it to fit the new problem. Although on this **R**-problem the subject did use the **R**-construct, his solution was obviously not concise compared to the modal solution shown in Figure 1 (b). This example clearly showed that the subject was negatively transferring what he had done before to what he had to do subsequently, and this negative transfer can easily be interpreted as a result of the subject's attempt to reduce his computational cost, or in other words, to spare his cognitive effort of devising a new but better solution.

For our current example of iterative strategy selection and change, since both the **W**- and the **R**-strategies would lead to working programs, we could suppose that the success probabilities for both strategies are the same, i.e., both equal to one. On the other hand, as the two strategies produce solutions of different styles and of different execution efficiencies, we would assume that they had different utilities. Specifically, we would assume that the strategy leading to more concise, more well-structured, and more efficient programs has a higher utility than the other strategy on a particular type of problems. Do subjects indeed evaluate the utility of their selected strategy? and how would they do this? For those highly-adaptive subjects, they seemed to adopt the preferred strategy at the very beginning of their programming in most cases, and their evaluation of utility seemed to be realized subconsciously without any manifestation in their verbal protocols. On the other hand, for those moderately-adaptive subjects, they usually performed this evaluation very explicitly, and their verbal protocols would have corresponding episodes for the evaluation process. Moreover, in some cases there were strategy changes occurring in such subjects' courses of programming. Figure 4 shows an example of strategy change which reveals the utility evaluation process. The problem was again the **W**-problem shown in Figure 1 (a). The subject first tried to use the **R**-construct; however, having almost completed her first solution, she paused for a while and found that the **if...then** statement would be redundant if using the **W**-construct instead. Upon the reflection, the subject modified her solution to be a **W**-constructed one, and her verbal protocols in fact revealed that she was evaluating the utility of her first choice during the pause and before making the change.

Modeling Strategy Selection and Change

As pointed out in Anderson (1990), the RA approach should not be pursued independently from other existing practices in cognitive science; rather, there should be an intimate interplay among building RA theories, computer modeling, and experimentation. The recent research by Anderson and his group has vigorously employed multiple practices and involved strong interactions among them (see Anderson, in preparation). We have so far proposed a rational theory of strategy selection and change and interpreted some of the results from our experiments on iterative strategy selection and change in PASCAL programming in terms of this rational

The subject's solution to the W-problem:

```
. . . . . . . . .
  repeat
    Read(File1, T);
    if T.Name <> 'END' then
      Write(File2, T);
  until T.Name = 'END';
. . . . . . . . .
```

The subject's solution to the R-problem:

```
. . . . . . . . .
  repeat
    Read(File1, T);
    if T.Name <> 'END' then
      Write(File2, T);
  until T.Name = 'END';
  Write(File2, T);
. . . . . . . . .
```

Figure 3. An example of lateral transfer effect in strategy selection.

The subject's 1st solution to the W-problem:	The subject's 2nd solution to it:

```
The subject's 1st solution to the W-problem:
   .........
   Read(File1, Predata);
   if Predata.Name <> 'END' then
      repeat
         Write(File2, T);
         Read(File1, T)
      until Predata.Name = 'END';
   .........

The subject's 2nd solution to it:
   .........
   Read(File1, Predata);
   while Predata.Name <> 'END' do
      begin
         Write(File2, T);
         Read(File1, T)
      end;
   .........
```

Figure 4. An example of strategy change involving utility evaluation.

theory. In what follows we will present a computer simulation we developed for our example of iterative strategy selection and change; we hope that the simulation will exemplify some of our theoretical claims in a more concrete way.

Our simulation was accomplished in the ACT-R (Adaptive Character of Thought—Rationality) cognitive model. The ACT-R model maintains the same basic architecture as its predecessor—the ACT* model (Anderson, 1983), having a declarative LTM with data chunks as its knowledge elements, a procedural LTM with productions as its knowledge elements, and a dynamic WM. Compared to ACT*, however, the ACT-R model has some new features derived from RA. One such feature which is particularly relevant to the present rational theory of strategy selection and change is that each production in ACT-R has an associated probability of succeeding its goal actions and a cost of retrieving and applying it. Another particularly relevant feature is related to the conflict resolution mechanism embedded in ACT-R. Basically, each time when a production is matched or instantiated, the conflict resolution mechanism estimates the utility of its goal and the computational cost for executing the production, evaluates the formula expressed in (1), and chooses the optimal one among all the productions instantiated at that point of time. Based on these new features of ACT-R, Anderson & Kushmerick have successfully simulated subjects' route-choosing behavior on a computer-based navigation task (Anderson, in preparation).

For simulating iterative strategy selection and change, we considered the two looping constructs as two high-level productions; these productions can also be conceived as plans with details to be filled in (for the notion of plan in programming, see Soloway et al., 1988). Under this consideration, for those idiosyncratic subjects who used only one strategy on all problems, it was reasonable to suppose that they had only developed one such production—i.e., either P1 or P2 shown in Figure 5—or that they had developed both productions, but one was highly practiced while the other seldom practiced. In the latter case, it was conceivable that the retrieval cost associated with the frequently-used production was overwhelmingly lower than the cost for the other production; as a consequence, these subjects also displayed the type of idiosyncratic behavior observed in our experiments.

As to the adaptive subjects, since there was a difference manifested in their protocols between those who were highly-adaptive and those who were only moderately-adaptive, we simulated their strategy selection behavior differently with different productions. In ACT-R, it is a premise that only declarative knowledge coming into WM can be verbalized; the mere execution of a production is procedural and not verbalizable without much reflection. Thus, to simulate the strategy selection behavior of those highly adaptive subjects we used two adaptive productions, i.e., P1' and P2' shown in Figure 5, for which the evaluation of their utilities is an integrated part of themselves and is automatically performed by the conflict resolution mechanism. On the other hand, for the moderately-adaptive subjects, since their verbal protocols revealed an explicit process of utility evaluation and of conflict resolution, it seemed most suitable to use the non-adaptive productions P1 and P2 together with the assistance of the WM mechanism to model their performance. Specifically, we assumed that these subjects had developed both P1 and P2 associated with comparable costs of retrieval. Consequently, since either production could fire, the simulation would indiscriminately choose any one looping construct and then proceeded along with it. During the next several steps of simulation while the elements in WM are accumulating, some further productions would evaluate whether there was any redundant statement (i.e., the same statement in more than one place) in the constructed program and then, depending on the evaluation, either proceed to finish the selected course or to make a change. This was basically the type of behavior we observed in the example presented in Figure 4.

Conclusion

In this paper we have presented a rational perspective for understanding human cognitive behavior of strategy selection and change; we have

Non-adaptive looping-construct selection productions

PRODUCTION P1:
IF the program involves an indefinite loop;
THEN choose the **while...do** construct.

PRODUCTION P2:
IF the program involves an indefinite loop;
THEN choose the **repeat...until** construct.

Adaptive looping-construct selection productions

PRODUCTION P1':
IF the program involves an indefinite loop;
AND the loop may be executed zero times;
THEN choose the **while...do** construct.

PRODUCTION P2':
IF the program involves an indefinite loop;
AND the loop will at least be executed once;
THEN choose the **repeat...until** construct.

Figure 5. High-level productions for simulating looping-construct selection.

also illustrated this theoretical perspective with a concrete example of strategy selection and change in computer programming and substantiated our theoretical claims in terms of a computer simulation for this example. To summarize, we would conclude from this work the following points:

1. Several related studies conducted within different cognitive domains, including our own one on iterative programming, have provided convergent evidence indicating that subjects are highly adaptive in their strategy selection; that is, they choose their strategies appropriately in response to environmental information.

2. Subjects' behavior of strategy selection and change can be better understood within a rational framework which proposes that in making strategical choices subjects usually attempt to optimize the difference between the utility (e.g., information gain, execution efficiency, or design elegance) of the expected goal and the cognitive or computational cost associated with attaining the goal; this framework is integrative in that it incorporates within it such earlier models as proposed by Bruner et al. and by Hoenkamp.

3. In the course of choosing a strategy, subjects' evaluation of the expected utility of the chosen strategy may be performed implicitly (subconsciously) without any manifestation in their verbal protocols or explicitly (consciously) with corresponding episodes in their verbal protocols.

4. In the ACT-R cognitive model, cognitive strategies can be modeled as high-level productions. Specifically, strategy selection with implicit evaluation of expected utility can be modeled as instantiation of specialized productions, whereas strategy selection with explicit evaluation relies on some evaluation of working memory elements.

References

Anderson, J. R. 1983. *The Architecture of Cognition.* Cambridge, MA: Harvard University Press.

Anderson, J. R. 1990. *The Adaptive Character of Thought.* Hillsdale, NJ: Erlbaum.

Anderson, J. R. In Preparation. *Rules of the Mind.* Hillsdale, NJ: Erlbaum.

Bruner, J. S.; Goodnow, J. J.; & Austin, G. A. 1962. *A Study of Thinking.* New York: Science Editions, Inc.

Hoenkamp, E. 1989. "Confirmation bias" in rule discovery and the principle of maximum entropy. *The Proceedings of the Eleventh Annual Conference of the Cognitive Science Society.* 651-658. Hillsdale, NJ: Erlbaum.

Payne, J. W.; Bettman, J. R.; and Johnson, E. J.1988. Adaptive strategy selection in decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition,* 14: 534-552.

Popper, K. R. 1959. *The Logic of Scientific Discovery* (2nd ed.). New York: Basic Books.

Reder, L. R. 1987. Strategy selection in question answering. *Cognitive Psychology,* 19: 90-138.

Reder, L. R.; and Ritter, F. 1992. The effect of feature frequency on feeling of knowing and strategy selection for arithmetic problems. *Memory and Cognition,* Forthcoming.

Russo, J. E.; and Dosher, B. A. 1983. Strategies for multi-attribute binary choice. *Journal of Experimental Psychology: Learning, Memory, and Cognition,* 9: 676-696.

Siegler, R. S.; and Jenkins, E. A. 1989. *How Children Discover New Strategies?* Hillsdale, NJ: Erlbaum.

Simon, H. A. 1983. *Reasons in Human Affairs.* Stanford, CA: Standford University Press.

Soloway, E.; Adelson, B.; and Ehrlich, K. 1988. Knowledge and processes in the comprehension of computer programs. In M. T. H. Chi; R. Glaser; and M. J. Farr (Eds.), *The Nature of Expertise.* Hillsdale, NJ: Erlbaum.

Tukey, D. D. 1986. A philosophical and empirical analysis of subjects' modes of inquiry in Wason's 2-4-6 task. *The Quarterly Journal of Experimental Psychology,* 38A: 5-33.

Varian, H. R. 1987. *Intermediate Microeconomics: A Modern Approach.* New York: Norton.

Wason, P. C. 1960. On the failure to eliminate hypotheses in a conceptual task. *The Quarterly Journal of Experimental Psychology,* 12: 129-140.

Wu, Q.; and Anderson, J. R. 1991. Strategy selections and changes in PASCAL programming. In J. Koenemann-Belliveau, T. G. Moher; and S. P. Robertson (Eds.), *Empirical Studies of Programmers: Fourth Workshop.* pp. 227-238. Norword, NJ: Ablex.