Shute, V. J. (1993). A comparison of learning environments: All that glitters... In S. P. Lajoie & S. J. Derry (Eds.), Computers as Cognitive Tools (pp. 47-74), Hillsdale, NJ: Lawrence Erlbaum Associates.

2 A Comparison of Learning Environments: All That Glitters . . .

Valerie J. Shute
Armstrong Laboratory, Brooks Air Force Base

Computerized learning environments can be characterized by the amount of learner control supported during the learning process. This dimension can be viewed as a continuum ranging from minimal (e.g., rote or didactic environments) to almost complete learner control (e.g., discovery environments). Two differing perspectives, representing the ends of this continuum, have arisen in response to the issue of the most optimal learning environment to build in intelligent tutoring systems (ITS). One approach is to develop an environment containing assorted tools and allow the learner freedom to explore and learn, unfettered (e.g., Collins & Brown, 1988; Shute, Glaser, & Raghavan, 1989; White & Horowitz, 1987). Advocates of an opposing perspective argue that it is more efficacious to develop straightforward learning environments that do not permit "garden path" digressions (e.g., Anderson, Boyle & Reiser, 1985; Corbett & Anderson, 1989; Sleeman, Kelly, Martinak, Ward, & Moore, 1989). This disparity between the positions becomes more complicated because the issue is not just which is the better learning environment; but rather, which is the better environment for what type(s) of persons, a classic aptitude-treatment interaction question (Cronbach & Snow, 1977).

Many kinds of learner characteristics (e.g., incoming knowledge and skills) affect what is learned in an instructional setting. This chapter focuses on another individual differences measure, *learning styles.* Baron (1985) defines styles as, ". . . general behavioral dispositions that characterize performance in mental tasks; they are intellectual personality traits" (p. 366). Thus, learning styles may be seen as reflecting different *approaches* to learning and may include such traits as being holistic versus analytic, verbal versus spatial, reflective versus impulsive, or exploratory versus passive (e.g., Baron, 1985; Glushko & Cooper, 1978; Hunt & MacLeod, 1979; Kyllonen & Shute, 1989; Pask & Scott, 1972; Pellegrino, Mumaw, & Shute, 1985).

To illustrate individual differences on one of these style dimensions, exploratory versus passive disposition, compare the following hypothetical persons. After receiving a new word-processing program, Ann immediately loads the program onto her computer, tosses aside the manual, and learns the new knowledge and procedures by trial-and-error. In contrast, Bob studies the accompanying manual, reads it cover-to-cover, and only then loads the software onto his computer. After 2 weeks, both are using the new word-processing program with comparable efficiency. Which method is better? Which should be supported by a tutor's learning environment? Is there a trade-off between learning time and quality of learning? These questions become very important when developing computerized instructional systems.

This chapter systematically explores the possible interaction between learning environment and learner style on various learning outcome measures. This experimental method has, in the past, been referred to as aptitude-treatment interaction (ATI) research (see Cronbach & Snow, 1977) where aptitudes are defined in the broadest sense of a person's incoming knowledge, skills, personality traits, and so on. The point of ATI research is to provide information about initial learner states that can be used to select the best learning environment for a particular student. To justify such an approach, evidence is needed that individuals do perform better or worse under different learning conditions (or environments).

ATI research was very popular in the 1960s and 1970s, then popularity declined. The main reason contributing to the decline was that the older ATI research typically involved studies conducted in classroom environments. Data were confounded by many extraneous variables (e.g., personality of the teacher, instructional materials, classroom dynamics) making ATIs hard to find and difficult to interpret. The current study circumvents this problem of "noisy data" by using a rigorously controlled learning environment.

An ITS instructing basic principles of electricity was used as the learning task, manipulated to yield two learning environments. These environments differed only in the type of information provided by the tutor to the student. I posited that active, exploratory learning behaviors would facilitate knowledge and skill acquisition, especially in conjunction with the environment supporting inductive learning behaviors. Less exploratory behaviors were hypothesized to be better suited to the structured learning environment.

Learning Task. The intelligent tutoring system used in this study taught basic principles of electricity: Ohm's and Kirchhoff's laws. It was originally developed at the Learning Research and Development Center, University of Pittsburgh (Lesgold, Bonar, Ivill, & Bowen, 1989) and then modified extensively at the Armstrong Laboratory, Human Resources Directorate. In particular, I created the two learning environments, developed learning indicators, rewrote the feedback, established the mastery criterion, and modified the interface. Subjects learned by solving problems presented by the computer. They also could read definitions about concepts that were written in a hypertext structure (i.e., nested concepts within concepts), but that was optional. Another optional activity included using a meter with positive and negative leads to obtain readings from different parts of a circuit (e.g., measuring voltage drop, current). Additionally, learners were free to change component values (e.g., increase a resistor's value) to see the effects on the circuit.

The two environments differed solely in terms of the feedback provided to learners. In both environments, following the solution of each circuit problem, learners were informed of the correctness of their solution. I called the first environment "rule application" because after the "right" or "wrong" feedback was given, the computer presented the relevant rule. Subjects then applied this rule or principle in solving subsequent problems. To illustrate, the computer would comment, *"Great! (or Sorry!) You are correct (or incorrect). The principle involved in this kind of problem is that current is the same before and after a voltage source."* Thus the principle was explicated after each problem solution (for both correct and incorrect responses) until learners reached the mastery criterion, which I set as three consecutively correct answers for a given problem type.

I labeled the second learning environment "rule-induction." Here, the correctness of the problem solution was again provided to learners, in conjunction with the relevant variables in the problem, but not their relationship(s). For learners in this environment, the computer might respond. *Great! (or Sorry!) You are correct (or incorrect). What you need to know to solve this problem is how current behaves—both before and after a voltage source.* The inductive environment thus required subjects to generate for themselves the relationships among variables during the solution of problems.

The curriculum consisted of a set of basic principles. Some of these principles were: (1) The current at one point in an uninterrupted wire is equal to the current at another point in an uninterrupted piece of wire, (2) The current before a resistor is equal to the current after a resistor in a parallel net, (3) Voltage is equal to the current multiplied by the resistance ($V = I*R$). There were a total of 26 different principles or "problem types" to be learned.

Problems were generated by the computer based on each of these principles. Each problem was unique to each individual, not preprogrammed, based on the particular subject's response history. For example, if a student needed more work on current flow across resistors, the system would generate a problem satisfying specific constraints such as it must be a "current problem" involving at least one resistor, perhaps requiring a more quantitative solution, and so forth.

Figure 2.1 shows an example of the main screen. On the screen's left, a parallel circuit is depicted with various component values. The upper right of the

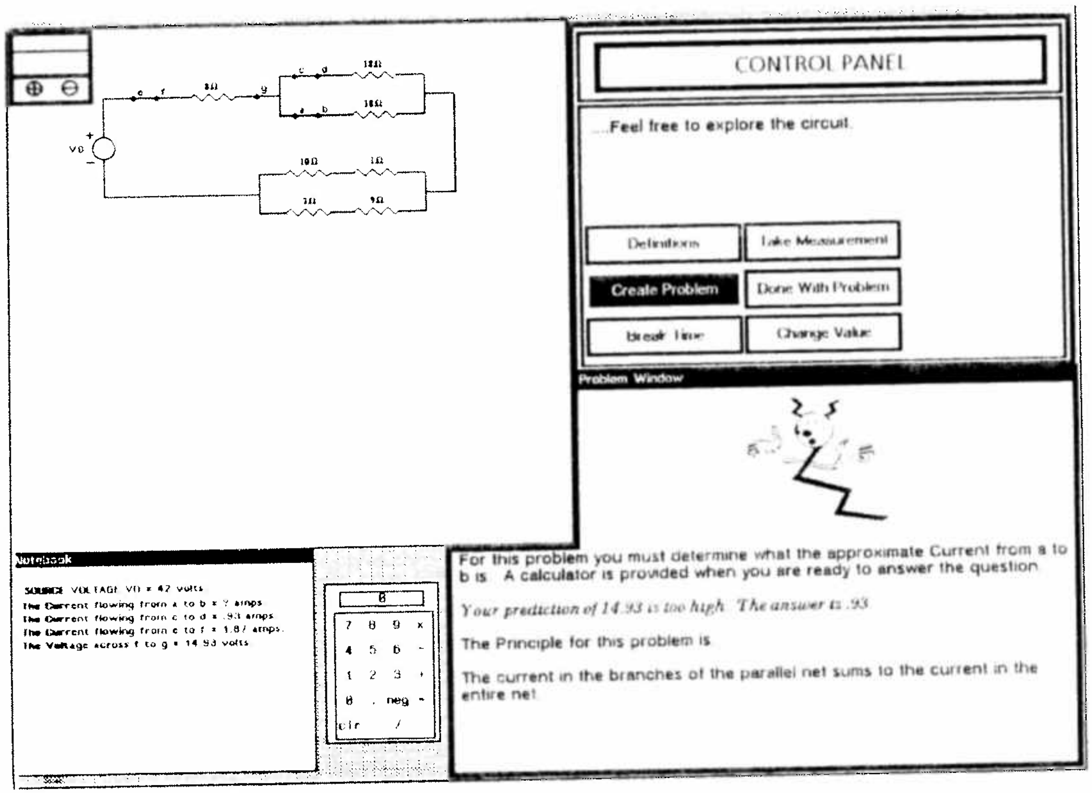

FIG. 2.1. The main screen of the electricity tutor (rule-application environment).

screen presents the learner with his or her main options (e.g., look at definitions, take a measurement on the circuit). Problems are presented in the lower right quadrant of the screen with feedback given in the same window. In this problem, the student was supposed to determine the current (in amps) from points a to b. The values from c to d and from e to f were given. A notebook present at the lower left of the screen allowed for the storage of information. If a student chose to explore the circuit, he or she could store new information in the notebook and compare it to the old data. Finally, an on-line calculator was always available for the solution of more complex, quantitative problems.

Optional (Exploratory) Behaviors. Some individuals like to control what they do and when they do it during the learning process. In this tutor, in addition to solving a problem, there were three different elective activities: viewing definitions, taking measurements on a circuit, and changing circuit component values. The first exploratory behavior was declarative (i.e., looking up terms and definitions) while the second and third activities were procedural. That is, taking measurements and changing components actually required the learner to *do* something to the circuit rather than more passively reading definitions.

If the subject chose to see *definitions*, the screen would clear and a menu of items would appear: ammeter, ampere, charge, circuit, current, ohm, resistance, resistor, parallel circuit, series circuit, voltage, voltage source, and voltmeter. Selecting any of these terms would cause a large window to open that would contain three parts: a relevant diagram, a definition (formal), and an explanation (informal). Bold-faced words would appear within the definitions and explanations. Selecting a bold-faced word would move the learner to the related concept (see Fig. 2.2). In come cases, simulations were available for the learner to run: (a) Comparison between current flow in a series versus a parallel circuit (see Fig. 2.3), and (b) Comparison between voltage drop in a series versus a parallel circuit. A dynamic display would appear on the computer screen illustrating how current (or voltage) operated differentially in the two circuit types, presented side-by-side.

Following problem solution, a subject could elect to *take measurements* on the circuit. For instance, Fig. 2.4 shows what happens when someone chose to measure the voltage drop across a resistor, from point g to point h. Positive and negative leads allow the learner to meter on two parts of the circuit and obtain a reading. Subjects could employ either the voltmeter (giving readings in volts) or the ammeter (giving readings in amps). This option was available at all times. However, if the subject had not yet answered the immediate problem, he or she was not allowed to take measurements that would yield the answer. For example, if the problem to be solved involved "current across a resistor in a series circuit," the learner could only take a voltage reading. After the problem was solved, then it was possible to obtain readings for both voltage and current.

Another optional activity involved changing component values. Again, after

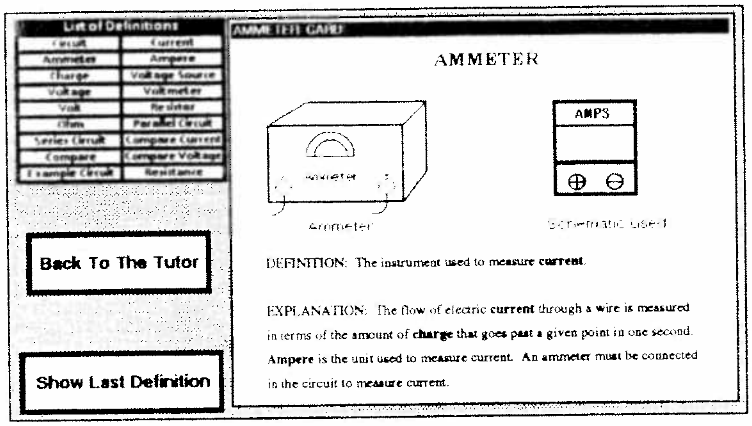

FIG. 2.2. Example of a definition (ammeter).

solving a problem, a component value (i.e., voltage source or resistor) on the circuit could be modified. Subjects could then observe how that particular change effected other parts of the circuit. To make a change, the subject would button on, for instance, the voltage source. He or she would then type in the new value (e.g., from 76 volts to 55 volts). Results from the changes appear automatically on the screen and in the notebook (see Fig. 2.5).

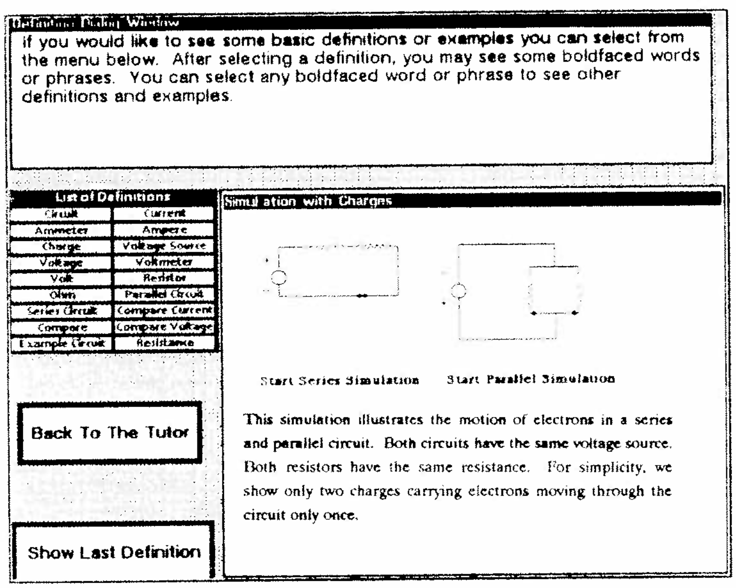

FIG. 2.3. Example of a simulation (comparing movement of charges in series and parallel circuits).

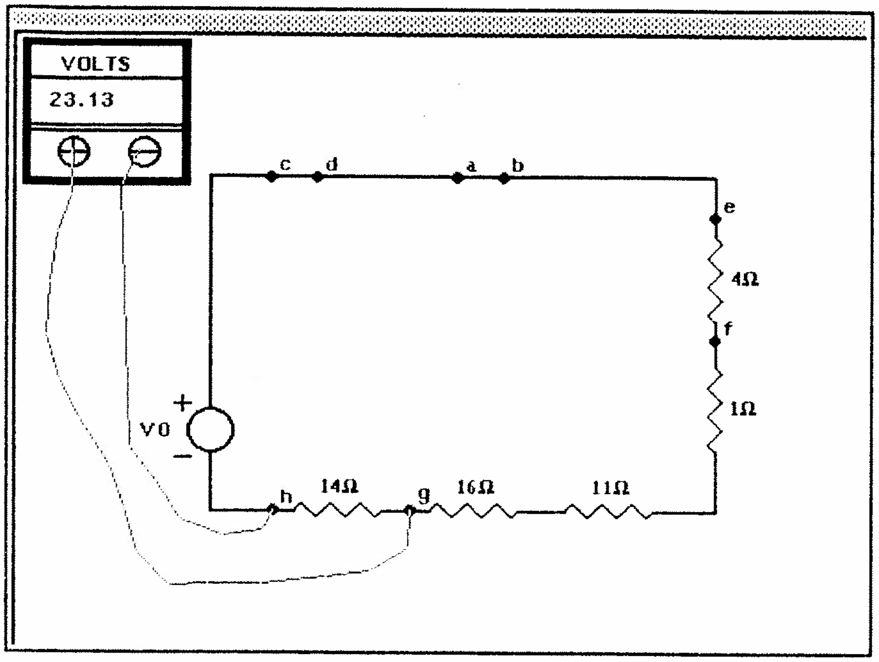

FIG. 2.4. Using the voltmeter on a circuit to obtain a reading between two points.

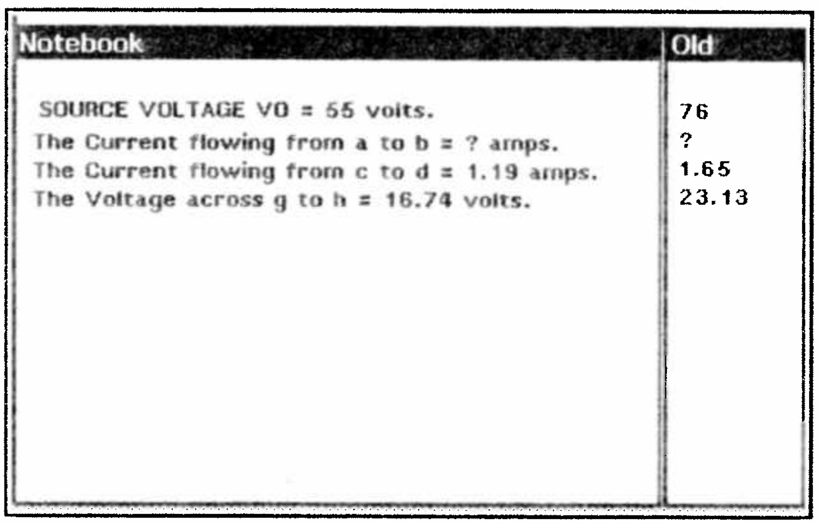

FIG. 2.5. The on-line notebook after the voltage source was changed (from 76 to 55 volts).

The computer tallied data on subjects' use of these elective activities. At a global level, I distinguished two kinds of exploratory indices. I computed the first one, "declarative exploration," as the time spent looking at definitions divided by the total time spent on the tutor (because time-on-tutor differed for everyone). The second index was "procedural exploration." This too was a proportion involving the time spent using a meter plus the time spent changing a component's value divided by the total time on the tutor.

Subjects. There were 309 subjects (84% males, 16% females) who completed this study on the acquisition of basic principles of electricity. Each subject participated for 7 days (45 hours). All subjects were high school graduates (or equivalent) with a mean age of 22 and an age range from 18 to 28. Subjects were obtained from two local temporary employment agencies and none had any prior electronics instruction or training. All subjects were paid for their participation.

Subjects were tested in groups of 15–20 at Lackland Air Force Base, Texas. They occupied individual testing stations and all instructions, testing, and feedback were computer administered with proctors available to answer questions. The ITS was administered on Xerox 1186 computers with standard keyboards and high resolution monochromatic displays on 19″ monitors. On the morning of Day 1, subjects were given a brief orientation to the electricity study and then randomly assigned to one of two learning conditions.

Pretests. Two pretests were included in the study to assess individuals' incoming domain-related knowledge. The first pretest measured **declarative knowledge** of different electrical components and devices involved in electronics. The concepts that were covered included: ammeter, ampere, charge, circuit, current, ohm, parallel circuit, resistance, resistor, series circuit, volt, voltage, voltage source, and voltmeter. This test included multiple choice and true/false questions. An example multiple choice item from the test asked: Which statement is most true about a *voltage source*? (a) It supplies electricity to a circuit, (b) It cannot store electricity for later use, (c) It does not have to be a physical device, (d) It is necessary to measure the current flowing through a circuit, or (e) It restricts the amount of current going through a circuit. Some example true/false items were: A parallel circuit requires two voltage sources, Unlike charges are attracted to one another, and Resistance is measured in amps.

The second pretest measured **conceptual understanding** of Ohm's and Kirchhoff's laws. These questions did not require any computations. Half of the items in this test contained pictures of circuits along with the questions, and the other half did not have pictures. To illustrate a question without a picture: If current was measured before and after a resistor in a series circuit, would the measurement before the resistor be *higher*, *lower*, or *equal to* the measurement after the resistor? The questions with circuits were similar, but referred to actual points on the circuits.

Posttests. I developed a four-part criterion test battery to measure the breadth and depth of knowledge and skills acquired from the tutor. The four-part battery was administered on-line at the end of the tutor. The first two tests in this learning outcome battery were identical to the two pretests discussed earlier (i.e., declarative knowledge and conceptual understanding).

The third posttest assessed the degree to which procedural skills were acquired. This test involved the **application** of Ohm's and Kirchhoff's laws in the solution of different problems. These questions did require computations in order to solve them. An on-screen calculator was provided to help solve these items. There were two types of questions, half with accompanying pictures of circuits and the others without pictures. Each question corresponded to a principle of Ohm's or Kirchhoff's laws. Problems with pictures displayed a circuit and the subject was required to compute what the reading was at some point for some component. The subject was required to apply the correct formula (e.g., V = I*R). Two of the three values were given and the solution required computing the unknown value. An example test item was: If the resistance in a circuit is 16 ohms and the current is 30 amps, then what is the voltage?

The fourth posttest in the criterion battery measured a subject's ability to **generalize knowledge and skills** beyond what was explicitly instructed by the tutor. The subject was required to generate or design circuits to do specific things. Thus, the test required not only a functional understanding of the laws and principles, but also the ability to compute solutions to novel problems. An example item from this test is included in Fig. 2.6.

In summary, the four tests were designed to measure different aspects of electronics knowledge and skill acquisition, from declarative knowledge understanding to quantitative understanding and ability to apply and transfer Ohm's and Kirchhoff's laws.

Learning Efficiency. I defined two learning efficiency measures. Because instruction in this tutor was self-paced, subjects could take as long as they needed to complete the curriculum. Some subjects were faster acquiring the new material, and others were slower. So the first index was defined as total *time on tutor*. The tutor was also open-ended as far as the number of problems generated per principle. That is, the number of problems a person received was a function of how many problems were needed to reach the mastery criterion (i.e., correctly solving three consecutive problems) per principle. Thus the minimum number of problems that would be created for a given principle was three. So the second learning efficiency index was defined as the total *number of problems* received. Although these two efficiency measures are somewhat related (i.e., it generally takes longer to complete the tutor if there are more problems to solve) they measure slightly different aspects of learning efficiency: speed and accuracy.

Hypotheses. In an earlier study, Robert Glaser and I (Shute & Glaser, 1990) found that individuals demonstrating systematic, exploratory behaviors (e.g.,

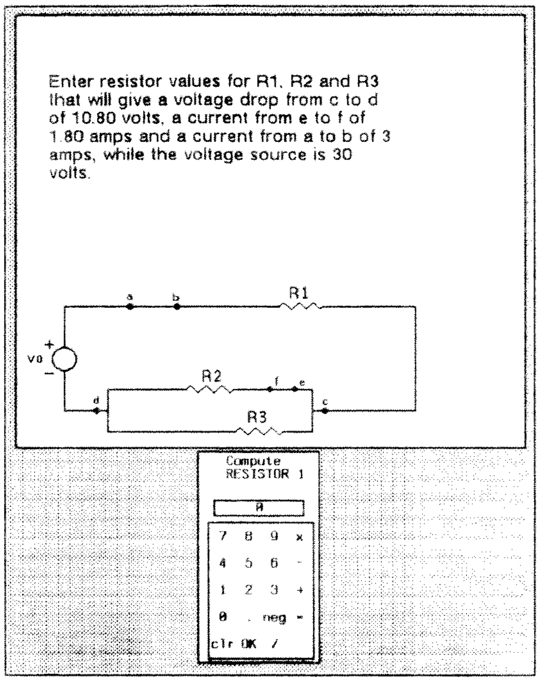

FIG. 2.6. Example test item from posttest 4: Generalization of knowledge and skills.

recording baseline data before any changes were made, limiting the number of variables changed in an experiment) were significantly more successful in a discovery microworld than those individuals evidencing less systematic behaviors. On the basis of this finding, I believed there would be a main effect of exploratory behavior on outcome where "more exploratory" would be associated with "better outcome" across criterion measures. Furthermore, I hypothesized this main effect of exploratory behavior to be even more pronounced in the rule-induction environment which supported inductive activities. In other words, I predicted that learners evidencing a lot of exploratory behaviors (procedural and declarative) should perform better on the outcome measures if they learned from the inductive environment than if they learned from the application environment. Conversely, less exploratory learners would benefit from the structured, application environment rather than the inductive environment.

RESULTS

Learning Outcome. The first criterion I investigated was learning outcome, defined as the percent correct scores on the four posttests. Although I originally created these tests to measure different facets of knowledge and skill acquisition, they turned out to be significantly correlated with one another: Posttests 1 and 2 ($r = .33$), 1 and 3 ($r = .76$), 1 and 4 ($r = .58$), 2 and 3 ($r = .41$), 2 and 4 ($r = .44$), and 3 and 4 ($r = .66$). Because of this interdependence among the test data, as well as a desire to keep analyses fairly simple, I computed a factor analysis (principal components) on the four posttest scores and a single factor was extracted, accounting for 65.1% of the posttest variance. The factor scores were saved for each individual (postfac) with loadings per test as follows: Posttest 1 (.85), Posttest 2 (.62), Posttest 3 (.90), and Posttest 4 (.84). Similarly, I computed a factor analysis (principle components) on the pretest data and one factor was extracted (prefac) accounting for 60.2% of the pretest data. Factor loadings for the pretests on this factor were both .78.

The composite learning outcome measure, postfac, was then examined as a function of learning environment and exploratory behavior—declarative and procedural. In addition, I wanted to look at the results of *just* the exploratory behaviors and environment on learning outcome without confounding the results with incoming knowledge (prefac correlates highly, .61, with postfac). By holding incoming knowledge constant (i.e., included as an independent variable in the regression equation), I can isolate the influence of specific behaviors on outcome.

I computed a multiple regression analysis using postfac as the dependent variable. The independent variables included: prefac, learning environment, procedural exploratory behavior (i.e., the proportion of time spent using the meter and changing components in relation to the total time on the tutor) and declarative exploratory behavior (i.e., proportion of time spent viewing definitions in relation to the total time on the tutor). Also the two interactions between the exploratory behaviors and environment were tested.

Results from this analysis showed that 42% of the variance of the outcome factor could be accounted for by these few variables (multiple R = .65). Not surprising, there was a main effect of prefac whereby individuals with more incoming domain-specific knowledge performed better on the outcome measures than those with less incoming knowledge: $t_{(1,299)} = 13.03$, $p < .001$. But there was no significant main effect of learning environment on learning outcome ($t_{(1,299)} = -1.83$, $p = .07$). As seen in Table 2.1, the pretest and posttest factor scores were similar (i.e., close to the mean of 0) in the two learning environments so neither environment showed a distinct learning advantage. But there was a slight advantage of the rule-application environment over the induction environment. There was a significant main effect of procedural exploratory behavior predicting the outcome factor: $t_{(1,299)} = -2.16$, $p < .05$. In this case, high procedural exploratory behaviors were associated with poor outcomes. There also was a significant main effect of declarative exploratory behavior on outcome: $t_{(1,299)} = 3.57$, $p < .001$. But here, the proportion of time allocated to reading definitions was a positive predictor of learning outcome. Finally, and of greatest interest, there was a significant interaction involving procedural (but not

TABLE 2.1
Summary Statistics of Learning and Behavioral Indicators by Environment

Variable	*Mean*	*SD*	*Min*	*Max*
Rule-Application (N = 152)				
Behaviors				
METER (minutes)	74.14	46.54	4.23	208.13
CHANGE (minutes)	9.45	9.39	0.00	40.00
DEFINITIONS (minutes)	13.21	9.43	0.68	46.62
TIME (MINUTES)	656.21	222.74	311.38	1230.58
METER + CHANGE (minutes)	83.59	51.03	4.23	220.63
PREFAC (factor score)	0.08	1.06	-1.82	3.41
Proportions				
METER = CHANGE / TIME	0.12	0.05	0.01	0.29
DEFINITIONS /TIME	0.02	0.01	0.0	0.08
Criteria				
POSTFAC (factor score)	0.08	1.06	-1.81	2.75
TOTAL PROBLEMS	140.16	45.95	79.00	291.00
TIME (hours)	10.93	3.71	5.19	20.51
Rule-Induction (N = 154)				
Behaviors				
METER (minutes)	82.14	52.81	4.30	243.48
CHANGE (minutes)	9.41	9.07	0.00	51.43
DEFINITIONS (minutes)	12.96	8.53	0.00	43.60
TIME (minutes)	687.60	201.74	382.52	1219.33
METER + CHANGE (minutes)	91.55	55.94	10.47	243.48
PREFAC (factor score)	-0.07	0.93	-1.57	2.33
Proportions				
METER + CHANGE / TIME	0.13	0.07	0.01	0.31
DEFINITIONS /TIME	0.02	0.01	0.00	0.06
Criteria				
POSTFAC (factor score)	-0.08	0.93	-1.81	2.58
TOTAL PROBLEMS	151.73	51.18	89.00	337.00
TIME (hours)	11.47	3.37	6.38	20.32

declarative) exploratory behavior and learning environment predicting learning outcome: $t_{(1,299)} = 2.44$, $p < .02$.

To illustrate this interaction, expected values were computed from the regression equation for four groups of subjects: Individuals one standard deviation above and below the average "procedural exploration" score in each of the two learning environments. These results can be seen in Fig. 2.7. Error bars are included in the plots of these expected values—approximate standard error measures for each group (i.e., square root of mean-square error divided by N). As can be seen in the figure, subjects who spent a large proportion of time engaged in procedural exploratory behaviors performed much better on the posttests (postfac) if they had been assigned to the rule-induction environment than the rule-application environment. But subjects showing fewer exploratory behaviors

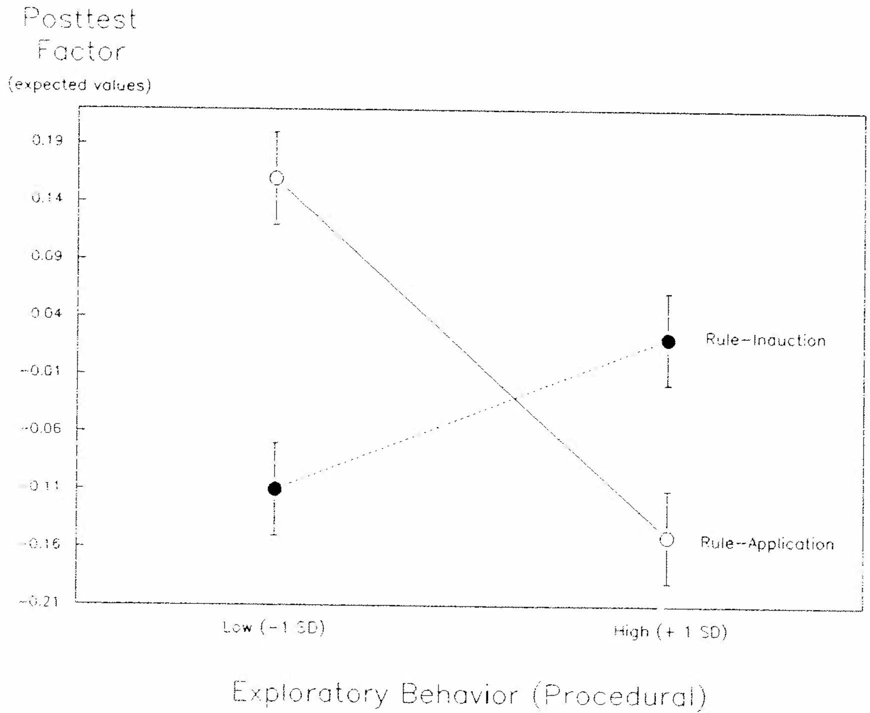

FIG. 2.7. Interaction of procedural exploratory behavior and learning environment on posttest factor score (expected values).

learned much more if they had been in the rule-application environment rather than the rule-induction environment.

Learning Efficiency: Time on Tutor. Similar to the preceding analyses with the outcome data, I computed a multiple regression analysis using time on tutor as the dependent variable and the same predictor variables as above (viz., prefac, learning environment, procedural and declarative exploratory behaviors, and the two interactions between behaviors and environment), and accounted for 26% of the efficiency variance (multiple R = .51). Again, prefac was included in the equation to control for differences in incoming knowledge that might impact learning rate. There was a significant main effect due to prefac: $t_{(1,299)} = -8.38$, $p < .001$ (i.e., more incoming knowledge associated with less time on tutor). There also was a significant main effect of environment on efficiency $t_{(1,299)} = 2.93$, $p < .005$. Individuals in the rule-application environment completed the tutor in less time than did those in the rule-induction environment. There also was a significant main effect of procedural (but not declarative) exploratory behavior on learning efficiency: $t_{(1,299)} = 4.34$, $p < .001$. In this case, using the

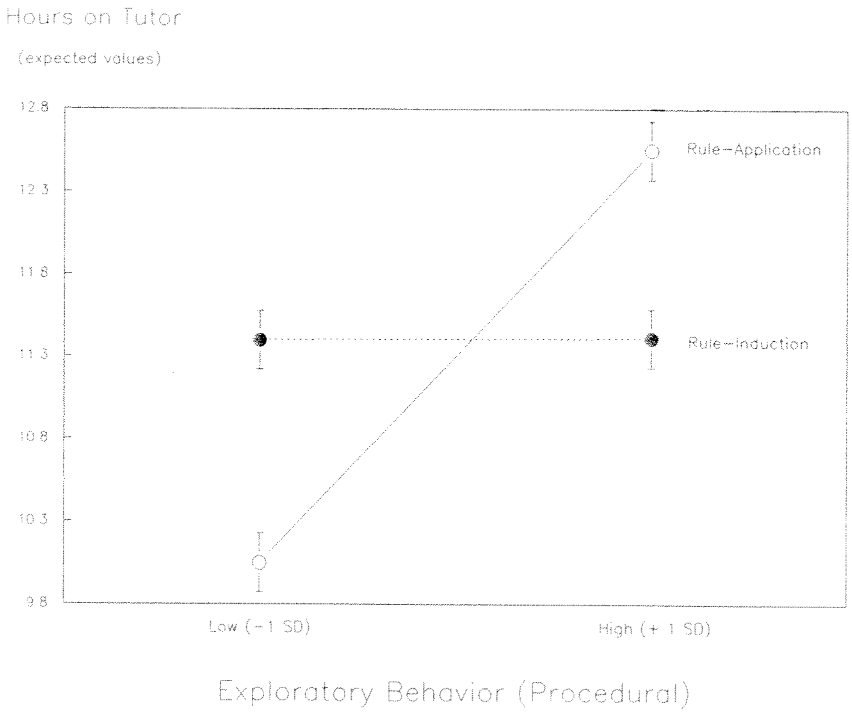

FIG. 2.8. Interaction of procedural exploratory behavior and learning environment on total time on tutor (expected values).

on-line exploratory tools was costly in terms of tutor completion time, and despite this increase in time, there was no payoff in increased outcome. On the contrary, from the results with postfac, above, we see that procedural tool usage was associated inversely with learning outcome factor. Finally, there was a significant interaction involving procedural behaviors and learning environment on learning efficiency $t_{(1,299)} = -3.43$, $p < .001$.

To illustrate this interaction, I computed expected values from the regression equation for four groups of subjects: Individuals one standard deviation above and below the mean procedural exploration score in each of the two learning environments. The results seen in Fig. 2.8 were as follows: Procedural exploratory behaviors were unrelated to hours on the tutor for individuals in the rule-induction environment, but positively related for the rule-application environment (where more behaviors = more time on tutor).

Learning Efficiency: Total Number of Problems Required. A final regression analysis was computed using total number of problems as the dependent variable and the same set of predictor variables as used above. About one third

(33%) of the variance was accounted for by the set of independent variables (Multiple R = .58). There was a significant main effect due to prefac: $t_{(1,299)} = -8.61$, $p < .001$. More incoming knowledge, again, was associated with fewer problems required to reach mastery criterion. There was a significant main effect of environment on efficiency $t_{(1,299)} = 3.63$, $p < .001$. Similar to the findings using time on tutor as the dependent variable, individuals in the rule-application environment required fewer problems, overall, compared to individuals in the rule-induction environment. And there was a significant main effect of declarative (but not procedural) exploratory behavior on number of problems: $t_{(1,299)} = -2.89$, $p < .005$. People who looked at many definitions required fewer problems to reach criterion, so it was a facilitative activity. There was also a significant interaction involving procedural behaviors and learning environment on learning efficiency $t_{(1,299)} = -3.15$, $p < .005$.

Expected values from the regression equation for four groups of subjects were computed: Individuals one standard deviation above and below the mean procedural explore score in each of the two learning environments. These results can be seen in Fig. 2.9. The depicted interaction shows that procedural exploratory

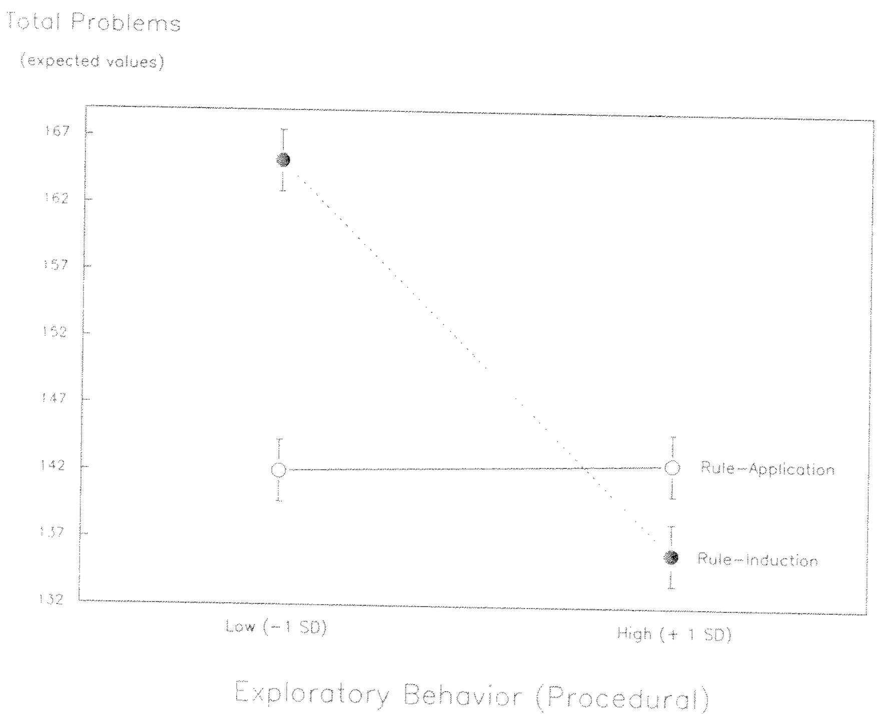

FIG. 2.9. Interaction of procedural exploratory behavior and learning environment on total number of problems needed (expected values).

behaviors were unrelated to number of problems required for individuals in the rule-application environment, but was significantly related for the rule-induction environment (i.e., more behaviors associated with a reduction in number of problems needed to reach mastery).

Within-tutor Analyses

The procedural exploratory index used in the foregoing analyses was computed as total time spent metering and changing component values divided by total time on the tutor. But sometimes total measures can be misleading (e.g., see Shute, 1989). A more refined way of looking at these data is to examine them across similar problem types or time (see Shute & Kyllonen, 1990). To accomplish this goal, I defined 26 new proportions corresponding to each of the 26 principles in the curriculum (rather than just the single proportion). These were computed as the amount of time spent metering and changing values divided by the amount of time spent on each principle.

I computed a factor analysis with varimax rotation on the 26 by 26 covariance matrix of proportions (i.e., time metering plus time changing values divided by total time for each principle). The varimax rotation converged in four iterations yielding a two factor solution. The two extracted factors accounted for 94.2% of the variance of these proportions. Table 2.2 shows the descriptions of the factors, along with associated principles and respective factor loadings.

Factor scores for the two factors were saved for each individual and then used in subsequent analyses. Relationships between factor scores and the criterion measures can be seen in Table 2.3, separated by environment.

In the rule-application environment, the data suggested that early on during the course of learning (factor 1 data, principles 1–9), the proportion of time spent engaging in procedural explorations was not significantly correlated with either posttest factor score or time on tutor.[1] During later learning of the more difficult concepts (factor 2 data, principles 10–26), higher proportions of procedural explorations were negatively correlated with the outcome and efficiency measures (i.e., lower posttest scores and longer time on tutor).

On the other hand, in the rule-induction environment, we see a different pattern of correlations. Early learning (factor 1 data) showed that higher proportions of procedural behaviors were *positively* correlated with outcome and efficiency measures (i.e., higher posttest scores, less time on tutor, and fewer problems to reach mastery). But later on, there was no correlation among the proportions and the learning measures.

The last analysis examines whether an individual's *initial* exploratory data can ultimately be used to predict learning outcome and efficiency measures differen-

---

[1]There was, however, a significant correlation between this proportion and number of problems required where higher proportions were associated with fewer problems.

TABLE 2.2

Factor Analysis Solution With Descriptions and Loadings for Each of the Two Factors Underlying Procedural Exploratory Behaviors

Factor 1: These are the first nine principles in the curriculum—simple Kirchhoff's problems involving current flow and voltage drop in series and parallel circuits.

Principle 3 (loading = .927):	The current is the same before and after a resistor.
Principle 4 (loading = .926):	The current before a resistor is equal to the current after a resistor in a parallel net.
Principle 5 (loading = .905):	The current in the branches of the parallel net sums to the current in the entire net.
Principle 6 (loading = .886):	The current in a component is lower than the current for the entire net.
Principle 2 (loading = .881):	The current is the same before and after a voltage source.
Principle 7 (loading = .852):	Voltage drop is lower across any single component of a series net than across the whole net.
Principle 8 (loading = .790):	Voltage drops across components of a series net sum up to the voltage drop across a whole net.
Principle 9 (loading = .730):	Voltage drop is the same across parallel components.
Principle 1 (loading = .683):	The current at one point in an uninterrupted wire is equal to the current at another point in an uninterrupted piece of wire.

Factor 2: This factor is characterized by principles representing later, more difficult problems: Ohm's law (i.e., the interrelationship among voltage, current, and resistance) and the integration of Kirchhoff's and Ohm's laws.

Principle 26 (loading = .958):	Voltage drop is the same across any component as it is across the whole parallel net.
Principle 25 (loading = .958):	Voltage drop is the same across parallel components.
Principle 24 (loading = .957):	The current in a component is lower than the current for the entire net.
Principle 23 (loading = .957):	The current in the branches of a parallel net sums to the current in the entire net.
Principle 22 (loading = .952):	The current before a resistor is equal to the current after a resistor in a parallel net.
Principle 21 (loading = .947):	Voltage drop is the same across any component as it is across the whole parallel net.
Principle 20 (loading = .938):	Voltage drop is the same across parallel components.
Principle 19 (loading = .926):	The current in a component is lower than the current for the entire net.
Principle 18 (loading = .918):	Current is the same across a resistor.
Principle 17 (loading = .905):	Current in the branches of a parallel net sums to the current in the entire net.
Principle 16 (loading = .886):	The current before a resistor is equal to the current after a resistor in a parallel net.
Principle 15 (loading = .867):	Current is the same before and after a resistor.
Principle 14 (loading = .850):	If the voltage goies up or down and the resistance stays the same, this implies that the current will go up or down with the voltage.
Principle 13 (loading = .833):	Current is equal to voltage divided by resistance ($I = V/R$).
Principle 12 (loading = .805):	When the current goes up or down and resistance stays the same, this implies that the voltage should also go up or down.
Principle 11 (loading = .761):	Voltage is equal to current multiplied by resistance ($V = I^*R$).
Principle 10 (loading = .705):	Voltage drop is the same across any component as across the whole parallel net.

TABLE 2.3
Correlations Among Procedural Factor Scores and Criterion Measures, Separated by Learning Environment

	Postfac	*Time*	*Problems*
Rule-Application Environment (N = 152)			
Factor 1	.17	-.13	-.28**
Factor 2	.23*	.35**	.14
Rule-Induction Environment (N = 154)			
Factor 1	.37**	.21*	-.38**
Factor 2	.07	.12	-.18

Notes. * $p < .01$; ** $p < .001$. Factor 1 = Early problems (principles 1-9) in the curriculum dealing with Kirchhoff's law, and Factor 2 = More difficult problems (principles 10-26) involving Ohm's and Kirchhoff's laws.

tially by environment. This has implications for generating decision rules for matching learners to environments. Rather than using factor 1 data (which consisted of the first nine principles in the curriculum), I was interested in testing whether exploratory behaviors, evidenced during learning the first principle, by itself, could predict any outcome or efficiency measures. The data used in this analysis included the amount of time a person spent in procedural explorations while learning principle 1 divided by the total time spent learning principle 1 (PEB1). The other independent variables included in the regression equation were: prefac, learning environment and PEB1 by environment interaction.

Results showed that these independent variables significantly predicted postfac (Multiple R = .62). There was a significant main effect due to prefac ($t = 13.5$; $p < .001$), where more incoming knowledge was a positive predictor of posttest performance. There was also a significant main effect due to learning environment ($t = -2.3$; $p < .05$) where the rule-application environment was associated with higher outcome performance. Finally, and of most interest, the interaction between exploratory behavior and environment was significant ($t = 2.2$, $p < .05$). Higher procedural proportions were associated with greater outcomes in the rule-induction environment, but not the rule-application environment. A graph of this interaction may be seen in Fig. 2.10.

The interaction term did not significantly predict time on tutor, but did predict total number of problems ($t = -2.3$, $p < .05$). A graph of this interaction may be seen in Fig. 2.11. Thus, the interaction between very early exploratory behaviors and learning environment may be used as a valid predictor of learning outcome and efficiency.

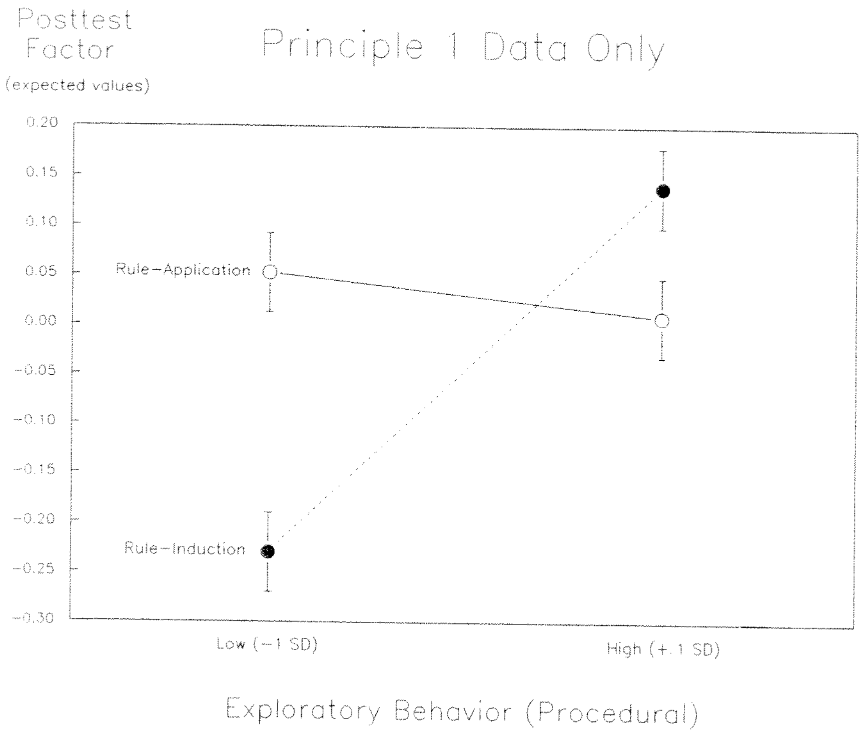

FIG. 2.10. Principle 1 data only—Interaction of procedural exploratory behavior and learning environment on posttest factor score (expected values).

DISCUSSION

In summary, I used an intelligent tutoring system with two different learning environments as a complex but controlled learning task to investigate possible learner style by treatment interactions. This represents a new generation of ATI research, more rigorously controlled than ATI research conducted during the 1960s and 70s. The learning environments (or treatments) in this study were identical, differing only in the feedback provided to the learner. After problem solution, whether correctly or incorrectly answered, one environment directly stated the relevant principle and the learner applied it in the solution of related problems; the other environment required the learner to induce the relevant principle, providing only the variables involved in the rule, but not their relationship(s). Findings showed that when learner styles (exploratory behaviors) were matched to environment, learning was superior compared with mismatched

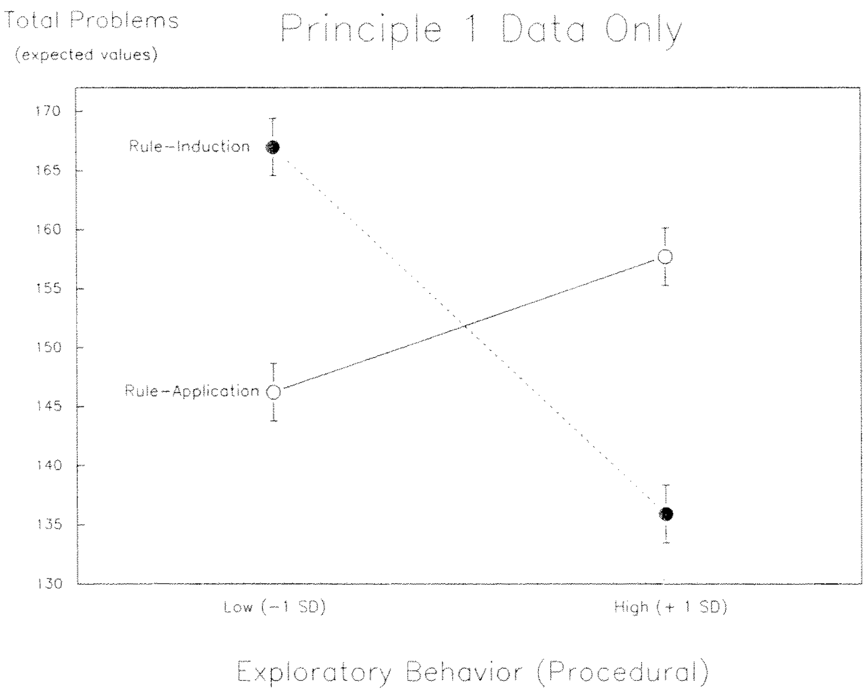

FIG. 2.11. Principle 1 data only—Interaction of procedural exploratory behavior and learning environment on total number of problems needed (expected values).

conditions. This suggests a new approach to student modeling using "new" ATI methodology (computer-administered learning tasks) and focusing on cognitive tool use as the behavior to model. This contrasts with, for example, model-tracing which records and diagnoses low-level productions underlying the learning process.

The first environment, rule-application, was straightforward and clear—all information necessary to solve a problem was presented to the learner. Subjects in this environment had no tenable need to engage in exploratory, extracurricular behaviors. On the other hand, the rule-induction environment required active participation in the learning process because the tutor only gave learners parts of a principle. Subjects had to come up with the conceptual glue (functional relationships) themselves, by any means they could. Thus, it was believed that the inductive environment would support (if not actively promote) the use of exploratory behaviors so that learners could obtain information needed to solve the problems. But results showed no significant differences between environments for either procedural or declarative exploratory learning behaviors (see Table 2.1). The mean procedural proportions were .12 and .13 in the rule-application and rule-induction environments, respectively. And the mean declarative proportions were .02 and .02 in the rule-application and rule-induction environments, respectively. Thus the two learning environments did not produce different profiles of exploratory behaviors.

In either environment, several different reasons can possibly explain an individual's exploratory behaviors. First, a learner unable to solve the problem being worked on may grope for something that he or she can do instead ("floundering" basis for the behavior). Another person may employ the tools after carefully designing an experiment involving the systematic manipulation of a circuit and taking controlled meter readings. This use of tools may ultimately supplement current understanding and yield other valuable insights ("methodic search for knowledge" reason for behavior). And finally, another person may simply use the on-line tools for fun and diversion ("playful curiosity" basis for behavior). The floundering basis may be associated with cognitive deficits while the methodic search basis may be associated with cognitive surfeits. Playful curiosity could be associated with either/neither.

If exploratory behavior simply reflects cognitive ability, in the rule-application environment, where there was no actual need to explore (the system presented the rule to learners), we would expect to see negative correlations between tool usage and cognitive ability. But in the rule-induction environment, applying exploratory behaviors may denote methodic (and perhaps necessary) knowledge searches. If that were the case, we would expect positive correlations between explorations and cognitive ability. Although not reported, I did compute these correlations and found no significant correlations between procedural exploratory behaviors and cognitive ability—overall and when the data were separated by environment.[2] Moreover, there was no significant correlation between tool usage and prefac (domain-specific incoming knowledge), overall and separated by environment. The degree to which an individual engages in exploratory behavior seems to be unrelated to cognitive ability and unrelated to incoming knowledge.

Although this chapter focuses on investigating possible interactions between exploratory behaviors and environment on outcome measures, some of the main effects turn out to be significant and illuminating; they will be discussed first.

Learning Environments. Was one environment more successful than the other in promoting knowledge and skill acquisition for the subject matter of basic electricity? Individuals in the rule-application environment, overall, took significantly less time to complete the curriculum and required significantly fewer problems to reach mastery compared to subjects in the rule-induction environment. But sometimes a large investment of time may actually result in greater

[2]The cognitive knowledge and skill measures that I examined in relation to exploratory behavior (procedural) included: working memory capacity, information processing speed, associative learning skill, inductive reasoning skills, procedural learning skills, and general knowledge.

gains or outcomes (e.g., the race between the tortoise and the hare). This was not the case in the present study. As a matter of fact, just the opposite was found. Not only did subjects in the rule-application environment learn the material more efficiently (i.e., take less time and require fewer problems to complete the tutor), they also had slightly higher (albeit nonsignificant) posttest factor scores than subjects learning in the rule-induction environment.

Tool Usage (exploratory behaviors). The next main effect involves exploratory behaviors and their impact on learning. A considerable amount of effort is expended by ITS designers and programmers creating multifarious "bells and whistles" in their systems. The point, of course, is to entice learners (as well as teachers, fellow researchers, and so on) with alterative and entertaining ways to learn. This tutor was no exception. Some very impressive features and capabilities were built into the tutor. For instance, an individual could use an on-line meter (voltmeter or ammeter) to obtain readings from different parts of a circuit. One could change a component's value (e.g., voltage source) and see the ramifications on the circuit. Finally, a person was free to peruse the on-line hypertext dictionary of terms. Which, if any, of these "bells and whistles" were important to learning? Did using these tools (and consequently engaging in exploratory behaviors) actually help or hinder learning? A significant negative main effect was found for procedural exploratory behavior in relation to posttest factor scores and time on tutor (where more procedural tool usage was associated with lower posttest factor scores and more time to complete the tutor). However a positive (facilitative) main effect was found between declarative exploratory behavior and both posttest factor scores and total number of problems needed to reach mastery. These data imply that using the on-line dictionary was a positive behavior but using the fancy meters and changing circuit values were, in general, negative behaviors.

There are several possible explanations to account for these findings: (1) Disruption of procedural skill acquisition; (2) Problems associated with using gross indicators in data analyses (e.g., overall procedural exploratory behavior proportion); and (3) Need to additionally consider other variables in the equation (e.g., learning environment, degree of tool usage). Each of these are discussed in turn.

Disruption of Proceduralization. Many cognitive psychologists have shown that successful skill acquisition depends on sustained and consistent practice opportunities (e.g., Ackerman, 1988; Anderson, 1987; Schneider & Shiffrin, 1977). If a person focuses on problem solution, then proceduralization is facilitated. But when that person departs from problem-solving activities and goes off to, for example, engage the on-line tools (for whatever reasons), that detracts from, and thus disrupts the compilation process. Referencing the dictionary may be an exception to this disruption because information found in the dictionary directly relates to relevant variables and their relations. Furthermore, there are limited garden paths available to traverse with the dictionary (18 terms defined in all). On the other hand, there are unlimited ways of manipulating circuits (e.g., successively increase a resistor value by one ohm).

Problems Using Total Counts, or "Gross Indicators" in Analyses. In dynamic learning situations spanning a duration of time, examining one variable defined as the sum of actions can be deceptive, especially when viewed in relation to other variables. For instance, Shute (1991) reported findings from a study employing a Pascal ITS as the learning task. One variable defined the total number of hints a person requested from the tutor. This gross indicator correlated with learning outcome ($r = -.64$), implying that hint-asking was, overall, a very unsuccessful behavior. But this was disturbing because one main feature of ITS's is their ability to provide individualized help when needed. When these data were analyzed across time (rather than using the gross count), asking for hints had much higher negative correlations with outcome during latter stages than the earlier stages of learning.[3]

A factor analysis computed on the data from the current study showed a clean two factor solution (i.e., Factor 1 = first nine principles, and Factor 2 = remaining principles). This breakdown allowed the data to be globally examined across time (i.e., early vs. later learning). Like the data from the Pascal study discussed earlier, findings with these separated data showed that, in fact, procedural exploratory behaviors were positively correlated with outcome measures early on (significantly so within the rule-induction environment, and the trend present in the rule-application environment). But later usage of these same tools was negatively correlated with outcome and efficiency measures, only within the rule-application environment, however. So, the simple main effects involving exploratory behaviors should be qualified (e.g., by time data).

Need to Additionally Consider Other Variables in the Equation (e.g., learning environment, degree of tool usage). The interaction hypothesis tested was whether individuals with above average exploratory behaviors would perform better in the rule-induction environment than the rule-application environment. Conversely, less exploratory individuals were believed to do better in the more didactic, rule-application environment than the more taxing rule-induction environment. The basis for this belief is that when learning environment is matched to certain characteristics of the learner, then performance is optimized (e.g., Pask & Scott, 1972). In fact, all three dependent measures (i.e., posttest factor score, time on tutor, and total number of problems required), showed significant learner

[3]This pattern of correlations between hints and outcome over time was seen even after cognitive process measures were partialled out of the hint-asking variable (e.g., working memory capacity, processing speed, general knowledge).

style by environment interactions. Each told the same basic story, but there were some subtle differences.

Posttest Factor Scores. This disordinal interaction was straightforward: Two opposite trends defined the correlations between exploratory behavior and posttest score. A positive linear trend expressed the relationship between exploratory behavior and outcome in the rule-induction environment (more is better), while a strong negative trend defined the relationship between exploratory behavior and outcome in the rule-application environment (more is worse). On the basis of these results, active explorers would do better on the outcome tests if learning from the inductive environment. But less exploratory folks should, unequivocally, be assigned to the straightforward application environment (see Fig. 2.7) to achieve their best posttest scores.

Time on Tutor. The significant interaction depicted in Fig. 2.8 showed that high explorers progressed through the curriculum in significantly less time if assigned to the rule-induction environment (again signifying a match between learner and environment). And low explorers completed the tutor much faster if assigned to the rule-application environment (another match). Now consider the slopes of the regression lines. High and low explorers in the rule-induction environment spent approximately the same amount of time on the tutor (11.4 hours, flat slope). Within the rule-application environment, though, a person's exploratory level really influenced learning efficiency (steep slope). A low explorer appropriately placed in the rule-application environment completed the tutor, on average, about 2.5 hours faster than a high explorer inappropriately assigned to the rule-application environment.

Total Number of Problems. The significant interaction shown in Fig. 2.9 supports the previous findings that low exploratory subjects assigned to the rule-application environment perform better on the tutor than low exploratory subjects assigned to the induction environment (i.e., require fewer problems to reach mastery). And high explorers in the inductive environment require fewer problems to complete the tutor compared to high explorers in the application environment. For this criterion measure, the rule-application environment showed no difference between high vs. low explorers in terms of the number of problems required (about 142 per group, flat slope). But the influence of exploratory behavior on number of problems was particularly striking within the rule-induction group. A difference of 30 problems separated subjects due to matched or mismatched condition. That is, low explorers who found themselves stuck in the rule-induction environment required 166 problems, on average, to complete the curriculum while high explorers, appropriately assigned to the inductive environment, required only 136 problems to reach mastery.

The main conclusion from these findings is that learning outcome and efficiency may be optimized by considering an individual's learning style in the assignment of person to learning environment. But here is the catch: We would like to develop some decision rule(s) for optimal placement of individual to environment. We have seen that a person's exploratory level impacts outcome performance differentially by learning environment. And we can obtain data about a person's exploratory level during tutor interactions. Then how can we make a priori decisions regarding placement? One solution is to *not* make *a priori* decisions. Instead, we could use early tutor data in the decision rule, providing these data showed some predictive validity. In fact, exploratory behavior data, tallied during the initial learning phase (principle 1 data only), *were* shown to be significant predictors of learning data in this study. The early behavior by environment interactions were shown to be significant predictors of learning outcome and efficiency (postfac and total number of problems).

In practice, the learner-to-environment assignment would work as follows: All individuals would initially be assigned to a default learning environment. Results from the study reported in this chapter suggest that, for instructing basic principles of electricity, the default environment should be rule-application because it displayed a distinct advantage over the inductive environment in terms of learning time and number of problems needed for attaining mastery (as well as a marginal advantage of posttest factor scores). Persons would then proceed through the tutor, and information on their explorations would be tallied in real-time. After the first principle was mastered, they would either be switched to the rule-induction environment if exploratory behavioral level was greater than average, else they would remain in the rule-application environment. Decision rules can, of course, be made even more comprehensive with the inclusion of additional conditionals. For example, some other ATI results reported by Shute (1992) suggest that considering an individual's associative learning abilities can inform decisions about which learning environment is the more suitable.

These findings have a direct implication for instruction (e.g., ITS design issues). As psychologists and educators concerned with instruction, our goal should be to maximize learning for as many individuals as possible. Results from this research provide information about which learning environments are more suitable for which learners, and why. In this study, we saw that low exploratory individuals learned efficiently from structured learning environments (rule-application) while high exploratory individuals learned best from freer learning environments (rule-induction). The reason "why" is due to the match between learner and environment characteristics. Furthermore, exploratory behavior does not appear to be simply an artifact of aptitude: the correlations between this learning style measure and various cognitive process measures were zero.

This study also addressed the issue of the utility of various "bells and whistles." Preliminary evidence suggested that, for many learners, all that glitters is not gold. In other words, simply having many and dazzling on-line tools in the environment without requirements for their use may be a wasted effort. Directed tool use may actually have positive effects on learning outcome and efficiency, but that was not tested in the current study. In conclusion, an ITS can potentially increase its effectiveness and progress toward the goal of optimizing learning by adapting to an individual's particular learning style. Learning environments are easily modified while learner attributes (e.g., styles, aptitudes) are less easily altered.[4] However, comparing the relative flexibility of styles to aptitudes, Baron (1985) argues that learning styles are considerably more modifiable than aptitudes (processing components). So, these data can provide a point of departure for building more adaptive learning environments.

ACKNOWLEDGMENTS

I sincerely thank the many people who worked with me on this project. Pat Kyllonen, Ray Christal, Bill Tirre, Scott Chaiken, and Lisa Gawlick-Grendell provided me with many valuable suggestions, from conceptual to analytical, during the conduct of this research. Roy Chollman, Rich Walker, Janice Hereford, Wayne Crone, and Linda Robertson-Shule provided excellent support and assistance in different aspects of the massive data collection job. Finally, I would like to thank the creative and industrious programmers on this project: Tony Beauregard and Kym Costa, as well as the original programmers at LRDC.

The research reported in this paper was conducted as part of the Armstrong Laboratory, Human Resources Directorate—Learning Abilities Measurement Program (LAMP). This study represents basic research funded by the Air Force Office of Scientific Research. The opinions expressed in this article are those of the author and do not necessarily reflect those of the Air Force.

Correspondence concerning this paper should be addressed to Valerie J. Shute, Armstrong Laboratory, Human Resources Directorate, Cognitive Skills Assessment Branch, Brooks Air Force Base, Texas, 78235.

REFERENCES

Ackerman, P. L. (1988). Determinants of individual differences during skill acquisition: Cognitive abilities and information processing. *Journal of Experimental Psychology: General, 117*, 288–318.

Anderson, J. R. (1987). Skill acquisition: Compilation of weak-method problem solutions. *Psychological Review, 94*, 192–210.

Anderson, J. R., Boyle, C. F., & Reiser, B. J. (1985). Intelligent tutoring systems. *Science, 228*, 456–462.

---

[4]However, I am not at this time claiming that exploratory behavior is a trait (i.e., relatively stable individual characteristic). It is probably more a "propensity" and potentially modifiable. More research is needed in this area.

Baron, J. (1985). What kinds of intelligence components are fundamental? In S. F. Chipman, J. W. Segal, & R. Glaser (Eds.), *Thinking and learning skills* (Vol. 2, pp. 365–390). Hillsdale, NJ: Lawrence Erlbaum Associates.

Collins, A., & Brown, J. S. (1988). The computer as a tool for learning through reflection. In H. Mandl & A. Lesgold (Eds.), *Learning issues for intelligent tutoring systems* (pp. 1–18). New York: Springer-Verlag.

Corbett, A. T., & Anderson, J. R. (1989). Feedback timing and student control in the Lisp intelligent tutoring system. In D. Bierman, J. Brueker, & J. Sandberg (Eds.), *Proceedings of the Fourth International Conference on Artificial Intelligence and Education* (pp. 64–72). Springfield, VA: IOS.

Cronbach, L. J., & Snow, R. E. (1977). *Aptitudes and instructional methods: A handbook for research on interactions.* New York: Irvington.

Glushko, R. J., & Cooper, L. A. (1978). Spatial comprehension and comparison processes in verification tasks. *Cognitive Psychology, 10,* 391–421.

Hunt, E., & MacLeod, C. M. (1979). The sentence-verification paradigm: A case study for individual differences. In R. J. Sternberg & D. K. Detterman (Eds.), *Human intelligence: Perspectives on its theory and measurement* (pp. 89–104). Norwood, NJ: Ablex.

Kyllonen, P. C., & Shute, V. J. (1989). A taxonomy of learning skills. In P. L. Ackerman, R. J. Sternberg, & R. Glaser (Eds.), *Learning and Individual Differences* (pp. 117–163). New York: W. H. Freeman.

Lesgold, A. M., Bonar, J., Ivill, J., & Bowen, A. (1989). An intelligent tutoring system for electronics troubleshooting: DC-circuit understanding. In L. Resnick (Ed.), *Knowing and learning: Issues for the cognitive psychology of instruction.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Pask, G., & Scott, B. C. E. (1972). Learning strategies and individual competence. *International Journal of Man-Machine Studies, 4,* 217–253.

Pellegrino, J. W., Mumaw, R. J., & Shute, V. J. (1985). Analyses of spatial aptitude and expertise. In S. Embretson (Ed.), *Test Design: Contributions from psychology, education, and psychometrics* (pp. 45–76). New York: Academic Press.

Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review, 84,* 1–66.

Shute, V. J. (1989). Individual differences in learning from an intelligent tutoring system. *Technology and Learning, 3,* 7–11.

Shute, V. J. (1991). Who is likely to acquire programming skills? *Journal of Educational Computing Research, 7,* 1–24.

Shute, V. J. (1992). Aptitude-treatment interactions and cognitive skill diagnosis. In J. W. Regian & V. J. Shute (Eds.), *Cognitive diagnosis and automated instruction.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Shute, V. J., & Glaser, R. (1990). A large-scale evaluation of an intelligent discovery world: Smithtown. *Interactive Learning Environments, 1,* 51–76.

Shute, V. J., Glaser, R., & Raghavan, K. (1989). Inference and discovery in an exploratory laboratory. In P. L. Ackerman, R. J. Sternberg, & R. Glaser (Eds.), *Learning and individual differences* (pp. 279–326). New York: W. H. Freeman.

Shute, V. J., & Kyllonen, P. C. (1990). *Modeling individual differences in programming skill acquisition* (Tech. report no. AFHRL-TP-90-76). Air Force Systems Command, Brooks Air Force Base, TX 78235-5601.

Sleeman, D., Kelly, A. E., Martinak, R., Ward, R. D., & Moore, J. L. (1989). Studies of diagnosis and remediation with high school algebra students. *Cognitive Science, 13*(4), 551–568.

White, B. Y., & Horowitz, P. (1987). *Thinker tools: Enabling children to understand physical laws* (Report No. 6470). Cambridge, MA: Bolt, Beranek, and Newman.