Retrieved May 3, 2004 from http://www.findarticles.com/cf_0/m2405/2_127/65913662/print.jhtml

Return to article page
This story was printed from LookSmart's FindArticles where you can search and read 3.5 million articles from over 700 publications.
http://www.findarticles.com

The Role of Context on Strategic Actions in Mastermind.

Journal of General Psychology, April, 2000, by John B. Best

ABSTRACT. When inexperienced solvers play consecutive Mastermind games, they begin to use a strategy in which specific hypothetical types are offered in response to specific feedback types, a strategy known as the modal hypothesis strategy. When participants who had learned that strategy were shifted onto Mastermind games whose code format was different from, and easier than, that they had experienced, their use of the strategy was degraded despite their familiarity and experience in dealing with the specific feedback types occurring in the new code. The results suggest that the use of a domain-specific strategy is at least somewhat influenced by the match between the context in which the strategy has been acquired and the context in which it is being put to use.

IN DISCUSSIONS of how problem-solving strategies come to be learned and maintained, a substantial body of research has focused on the relationship of domain knowledge to the production of specific strategic actions. According to this view, changes in domain knowledge may lead to the production of new and different strategic actions whose outcomes may, in turn, create situations that offer the solver a chance to discover more domain knowledge. For example, Simon and Reed (1976) demonstrated that a hint expressed in domain-specific terms produced a change in the moves of solvers on river-crossing problems such as missionaries--cannibals or Hobbits--Orcs. Such problems feature a number of people or creatures who are members of one of two categories, one of which is hostile to the other. The creatures must use a limited-capacity boat to transport their entire number across a river, under the constraint that if the hostiles ever outnumber the friendlies at any point in the crossing procedure, the hostiles will des troy the friendlies. Specifically, Simon and Reed advised solvers that, at some point in the crossing sequence, it was necessary to reach a state in which three of the five hostiles were alone on the far bank without the boat (in their problems, the boat's capacity was three people or creatures). Prior to being given the hint, the majority of participants produced a sequence of moves whose outcome was to keep the numbers of hostile entities approximately equal on both river banks (the balancing strategy). After receiving a hint, the participants produced a series of moves that created some temporary numerical imbalances. Simon and Reed described the hint as producing a shift in the participants' strategy, away from an intuitively appealing, but wrong, balancing strategy and toward a more correct means-end strategy.

Simon (1975) showed that the strings of moves chosen by the solvers depend in part on the structure of the task environment. This finding suggests that the sophistication of the solvers' strategies is influenced by the structure of the task environment. For example, the structure of the Tower of Hanoi problem enables solvers to create numerous subgoals, and this aspect of the problem's structure seems to enable solvers to plan long strings of moves. In tasks in which the structure does not permit the creation of subgoals, such as the missionaries-cannibals problem, solvers are compelled to rely on short sequences of moves whose functions are to balance the number of entities in the problem. The upshot of this view is that the effectiveness of the solver's

strategy in helping to discover aspects of the domain may be limited by the structure of the domain itself.

This view was amplified and extended in a study by Lewis and Anderson (1985). They showed that solvers in complex domains such as geometry learn to attend to particular features of a problem, namely, features that appear to be predictive of those problem-solving actions, or operators, that will lead eventually to a solution. In geometry problems, for example, Lewis and Anderson noted that the given facts that were provided in proof problems were frequently correlated with the preferred operator--in this case, the execution of an inference rule such as the side-angle-side postulate. From Lewis and Anderson's perspective, the solver, with increased experience in a domain, acquires schemata of problems, and these representations include recognition of those features that have been predictive of the success of certain operators. According to this view, the solver notes the features that are currently present in the problem, maps these features onto particular schemata and, on the basis of previous success, selec ts operators whose purpose is to move the solver closer to the goal. These basic findings have been supported in other research as well (Anzai & Simon, 1979; Karat, 1982; VanLehn, 1991).

The view that has emerged has concentrated on the role of the strategic action as a learned response to a specific situation that may occur in a given domain. As an increasing volume of domain-specific knowledge is acquired, the solver becomes correspondingly better able to recognize, categorize, and associate situations occurring in the domain with specific strategic actions. When this knowledge is modeled as a production system, the result is that the knowledge becomes compiled (Anderson, 1987; Anderson & Thompson, 1989) and proceduralized.

This position is based on the notion that the context for a strategic action is the entire domain-specific task environment in which the action has been learned--an assumption that seems warranted in the case of highly coherent domains such as geometry, or in the case of somewhat artificial problems such as river crossings. But, in other cases, the assumption that the context for a strategy is equivalent with the entire task domain may be less appropriate, perhaps because either other portions of the domain are organized differently or the structure of some portion of the domain is less homogenous or coherent than it is in other portions. There are some reasons to believe that the task domain of the game Mastermind possesses some of these latter-mentioned properties.

Strategic Actions in Mastermind

The commercial version of Mastermind is typically played by two people, one of whom acts as a "codemaker," creating a hidden left-to-right-ordered string of colored plastic buttons. The second player, who is the actual problem solver, functions as a "codebreaker," attempting to deduce the identity of each color and its place in the sequence. To accomplish this, the solver deploys a string of colored buttons as a hypothetical candidate for the hidden string, and then receives rather ambiguous feedback. Each unit of black feedback indicates that one of the colored buttons in the just-deployed hypothesis matches a code member in color and in location. Each unit of white feedback indicates that one color in the just-deployed hypothesis matches a color, but not a location in the code. An example of a code and hypothesis follows:

Previous work (Best, 1990) has shown that in games in which the participants are told that the code consists of four different colors drawn randomly from the pool of six colors, solvers typically use only the immediately preceding hypothesis and its feedback as a basis for their next hypothesis. That is, they tend to treat their just-deployed hypothesis as abase, which is then changed in predictable ways as a function of the feedback that hypothesis received. For example, in response to the foregoing hypothesis, the solver tends to make characteristic changes as follows:

Here, the solver has played two colors that did not appear in the base hypothesis (blue and green), and the white button has been moved to a different location in the string. This change is described (Best, 1990) as "two new colors and one location"; thus, Hypothesis 2 is an example of a "2C1L" type. It is important to note that each of the 13 types of feedback that typically occur in the four-different-colors code Mastermind game tends to be followed by a different type of hypothesis, using this characterization scheme. The most prevalently occurring hypothesis type for each feedback type is called its "modal" hypothesis.

When looked at in the context of play time of 1 hr by a novice player on such a four-different-colors code, modal hypotheses make up approximately 60% to 90% of all hypotheses advanced, depending in large part on which specific feedback type the solvers are responding to. I refer to the consistent deployment of modal hypotheses to each feedback type as it appears as a modal hypothesis strategy. It is relevant to note that the consistent deployment of modal hypotheses appears to be an effective way to search for the code in Mastermind. Other work (Best, 2000) has shown that a computer-generated Mastermind player who deployed modal hypotheses reduced the number of candidates for the code more quickly than did a plausible competitor strategy.

Given that modal hypotheses are driven by feedback, the opportunity to make particular modal hypotheses depends on the appearance of particular types of feedback. In turn, the distribution of feedback types varies somewhat as a function of the type of code that is being used in a particular Mastermind game. The 13 types of feedback do not necessarily occur with equal likelihood in Mastermind games using different types of codes. For example, in the type of game that has been described thus far, the four-different-colors code, or 1234 code, the feedback type lBk-lWh is awarded on approximately 11% of all hypotheses, but the feedback type "no feedback given" (i.e., the solver has played no correct colors at all) is very seldom awarded, usually less than 1% of the time. However, when the solver is playing a "two pairs" or 1122 code, lBk-lWh is awarded on approximately 25% of all hypotheses, and no feedback given is awarded 16% of the time.

At issue is the solver's response to these various feedback types as they occur in Mastermind games having different code formats (e.g., 1234 code vs. 1122 code). The claim is that playing Mastermind under different code formats is equivalent to a shift in the context in which the modal hypothesis strategy can be used, and the expectation is that such a shift in context should produce a decrement in the ability to deploy the modal hypothesis strategy. That is, even when problem solvers have acquired an amount of domain-specific knowledge sufficient to enable the production of particular strategic actions, that knowledge is limited by the context in which it was learned. In this case, the goal structure is an important element of the context. Thus, we would expect that a shift to a new, and previously unexperienced, goal structure (as when the code format in Mastermind is shifted) should produce a disruption in the use of the strategic action. Specifically, the contention is that in those situations in which the participants are quite familiar

with the feedback, when such participants are shifted to a new code format that is associated with different modal hypotheses, then the participants will deploy modal hypotheses with a lower frequency than participants who have not been shifted.

An example will help to clarify these expectations. Consider the feedback type 1Bk-1Wh: In games in which the code format is 1234, this feedback is given approximately 11% of the time, and its associated modal hypothesis, 2C1L, is played in response on 41% of those encounters with 1Bk-1Wh. In games where the code format is 1122, 1Bk-1Wh is awarded 25% of the time, and its associated modal hypothesis, 1C1L, is played in response on 50% of those encounters (note the change from the 1234 code format games). For novice solvers who have experience with only the 1234 code format, the expectation is that shifting to a 1122 code format will result in such solvers producing the new (to them) modal hypothesis significantly less than 50% of the time in response to 1Bk-1Wh.

It is expected that the effects of shifting context should be limited to just the case described previously. Thus, in those cases in which the shifted solvers are able to make the same modal response to feedback previously seen in the "old" code format (i.e., because the two code formats are followed by the same modal hypothesis type), then we should not expect a disruption in the deployment of modal hypotheses. Similarly, in cases in which a problem solver is shifted to a code format and encounters feedback types in the new code format that were previously unencountered in the old code format (and for which he or she thus has not acquired the knowledge to build the appropriate modal hypothesis), then we should not expect an impairment in the solver's ability to learn and deploy the appropriate modal hypotheses.

Method

Participants

The participants were 60 student volunteers in upper division psychology classes at Eastern Illinois University. They received a small (1.5% of their total grade) extra-credit bonus in their courses. The study's findings were explained to them at the end of the semester.

Design

The participants were randomly assigned to one of three experimental conditions. Twenty participants played four games with the code format 1122--that is, a two pairs code format (90 different possible codes); these participants are designated 1122 x 4. In addition, there were 20 participants in the "shift" condition: These participants played three games whose code format was 1234, followed by a single game whose code format was 1122 (1234 x 3, 1122 x 1). Finally, serving as a comparison group for the shift group, 20 participants played four Mastermind games with the code format 1234--that is, a four-different-colors-code format (360 possible codes), and these participants are designated 1234 x 4.

Procedure

The experimenter explained the rules of Mastermind and answered any questions the participants had. Then, each participant responded to a 5-point, labeled, Likert-type Mastermind familiarity

scale. Approximately 85% of the participants chose point 1 (I've never heard of this game before today). When the participant seemed comfortable with the task, the problem-solving session began. The task was administered using the standard, commercially available, plastic tokens. The experimenters were trained to avoid giving facial, verbal, or postural cues during their scoring of the participants' hypotheses. No participants indicated that an experimenter had provided any extra-task information. The participants were told that they would play four consecutive games of Mastermind or until a 1-hr time limit had elapsed. The experimenter announced the specific code format that would be used immediately prior to each game. Participants were permitted to produce up to 10 hypotheses per game; play was discontinued on any gam e that had not been solved at that point. Participants were not shown the code in games they did not solve.

Results

Hypothesis types for each feedback type were coded and tallied with the procedure that was described in the introduction. Table 1 contains the frequency of each hypothesis type in relation to preceding feedback, collapsed across all four games, for participants in the group that played four games of code format 1122. As has already been established for 1234 codes (Best, 1990), for each feedback type given in a 1122 game, a particular hypothesis was made with much greater frequency than any other individual hypothesis type, and that hypothesis is defined as the modal hypothesis for that feedback type.

The issue of strategic disruption was addressed by contrasting the proportion of modal hypotheses deployed by participants in the shift group at Game 4 (i.e., their first "easy" or 1122 game), with the proportion of modal hypotheses deployed at Game 4 by participants who had played three 1122 games prior to that point. For purposes of this contrast, several categories of feedback were defined. Some feedback types were "continuing" in the sense that they were awarded in both 1234 and 1122 games more than 5% of the time, as measured by the 1234 x 4 and 1122 x 4 groups. The 5% cut-off figure was arbitrary, but it was probably the lowest percentage that ensured that all participants had actually encountered that form of feedback at least one time in their first three games. There were four such feedback types: 2Bk, 2Wh, 1Bk-1Wh, and 2Bk-2Wh, collectively making up 30% of all feedback in the 1234 games and 69% of all feed-back in the 1122 games. Of these four types, three (2Bk, 2Wh, and lBk-1W) were associated wi th modal hypotheses in the 1122 games that differed from those in 1234 games, and so were categorized as "continuing-different modal." Of the four continuing feedback types, one (2Bk-2Wh) was associated with the same modal hypothesis in both 1122 and 1234 games, and so was categorized as "continuing-same modal." Finally, one feedback type ("no feedback given") was new in the sense that it occurred with a frequency greater than 5% in the 1122 games but less than that in the 1234 games, and thus was new to the shift-group participants, who were likely to encounter this feedback type for the first time when they shifted to the 1122 code in their final game. No feedback given was awarded less than 1% of the time in 1234 games, but it was awarded 16% of the time in 1122 games.

Table 2 contains the number of encounters with these feedback types that occurred for both the shift group and the 1122 x 4 group in their fourth games, the number of times that the modal was deployed in those encounters, and the resulting proportions of modal hypotheses. Also shown are the same figures for the 1122 x 4 group in their first three games, when these participants were presumably learning the modal hypotheses. The proportion of modal hypotheses made by shift

group in Game 4 in response to continuing-different modal feedback types (.39) was significantly lower than the proportion of modal hypotheses made by the 1122 x 4 group in their fourth game (.88, z = 4.8, p [less than] .001; Table 2), which is consistent with the expectation that the solvers' modal hypothesis strategy would be disrupted by a shift in context that required new hypotheses to be learned in response to feedback with which the shift participants were already familiar. It is worth noting that the percentage of continuing--differe nt modal hypotheses deployed by the shift group participants in Game 4 (i.e., .39) is not substantially different from the percentage of such hypotheses deployed by 1122 x 4 participants in their first three games (i.e., .48). This equivalence suggests that the experience gained by the shift-group participants, in playing their first three games using 1234 codes, did not confer any significant advantage on them with regard to finding the new correct modal hypotheses for continuing feedback in the 1122 code format.

Also, as expected, the proportion of modal hypotheses made by the shift group in Game 4 in response to the new feedback type (.84) was not significantly different from the proportion of modal hypotheses made by the 1122 x 4 group in their fourth game (.85). Finally, also as expected, the proportion of modal hypotheses made by the shift group in Game 4 in response to the continuing--same modal type (1.00) was not significantly different from the proportion of modal hypotheses made by the 1122 x 4 group in their fourth game (.92).

One possible objection to this interpretation is that the shift-group participants had simply tired of making the modal hypotheses after three games, and so their reduction in modal hypothesis production in Game 4 was the result of their interest in exploring other hypotheses. This interpretation can be countered by comparing the shift group's modal hypothesis production across Games 1 to 3, and Game 4 with that of the 1234 x 4 group. As Table 2 shows, modal hypothesis production for both groups was very similar in their first three games when both groups were dealing with 1234 codes. Moreover, there is no marked reduction in modal hypothesis production in the fourth game of the 1234 x 4 participants, compared with their own first three games, suggesting that they recognized the effectiveness of the modal hypothesis strategy.

Discussion

Consistent with my hypotheses, solvers' production of certain modal hypotheses was significantly impaired when they were shifted from a 1234 code to a 1122 code. Moreover, the disruption in modal hypothesis production was limited to just those cases in which the shift-group participants had to deal with producing a new modal hypothesis in response to a feedback type with which they were already familiar from their games in the 1234 code format (i.e., the continuing-different modal feedback types). The participants' modal hypothesis production for new feedback types was not disrupted, meaning, in this case, that the production of the appropriate modal was not significantly lower for the shift group in their final game than it was for the group that had already played three 1122 games. The absence of a disruptive effect here suggests that the disruption in strategic knowledge for shift-group participants was not produced by the appearance per se of a feedback type with which the shift group was unfamiliar. Thi s result, in turn, suggests that the shift to a different goal type (for the shift-group participants) did not impair the shift group's ability to learn new aspects of the modal hypothesis strategy. Similarly, participants in the shift group did not show a disruption in modal hypothesis production for the single feedback type, which occurred with high frequency in both 1234 and 1122 code formats and required the same

modal hypothesis in both formats (i.e., the continuing-same modal feedback type). The absence of a disruptive effect here suggests that the disruption in strategy is not necessarily produced by the shift to a new and previously unfamiliar goal type per se (i.e., 1122 vs. 1234), but rather by what the knowledge of this shift does to the solver's understanding of the relationship between specific feedback types and their specific associated modal hypotheses. This effect can be see in the following examples.

When solvers are told the format of the code in advance of their playing, as they were in this study, they are strongly inclined to advance hypotheses that could actually be the code. Thus, for example, the overwhelming majority of hypotheses advanced in 1234 games contain four different colors. It is important to note that modal hypotheses also conform to this principle, and each feedback type's associated modal hypothesis therefore results in a string that could actually be the code. For example, in a 1234 code format game, when solvers receive the feedback 2Bk-2Wh, they typically make a 2L hypothesis, and such a hypothesis will also have the same format that the code is known to have. This effect is seen in the following fictional but plausible couplet:

When participants are shifted to the 1122 code format, this particular connection is undisturbed in the sense that a 2L hypothesis in response to the feedback 2Bk-2Wh will result in a string that has the same format that the code is known to have in 1122 games as well as in 1234 games. This effect is seen in the following plausible couplet:

But for the three modal hypotheses that were disrupted by the shift to the 1122 goal context, this effect is not observed. That is, the modal response to the 1234 code types will not lead to the production of a hypothesis that has the structure that the code is known to have in a 1122 game. This difference is illustrated using the feedback type 2Bk. In the following plausible couplet the 2C hypothesis, which is the modal hypothesis in 1234 games, produces a string that does not have the structure that the code is known to have in the case of the 1122 games. And, presumably, this failure necessitates the disruptive search for a hypothesis that will have such a structure. As the following plausible couplet shows, a 1C hypothesis, which is the modal hypothesis for the feedback 2Bk in 1122 code format games, results in a match between the form of the hypothesis and the form that the code is known to have.

In some respects, these findings are similar to those of Luchins and Luchins (1991), who used a modified version of their water jars task (WJT; Luchins & Luchins, 1959). They found an illustration effect: Participants who saw example problems solved by pouring water from the central jar into the left-most or rightmost jar tended to follow the illustrated direction in solving problems on their own. Moreover, this effect was interleaved with an order of difficulty effect. Participants who started the solution series with the easiest problem were more influenced by the examples they had seen--that is, they were more likely to duplicate the directionality of the examples than were participants who were given a more complex initial problem to solve. These findings suggest that mental set might be understood as having its effects on learning: The ease or difficulty with which the solution algorithm is learned is influenced by the degree to which the initially presented cases facilitate a compilation of the product ion rules underlying successful solutions (Anderson & Lebiere, 1998).

Lovett and Anderson (1996) also observed Einstellung effects using a building sticks task (BST), which is isomorphic to the WJT. The goal in BST is to construct a stick of a given length, given

three stick lengths from which to build. Solvers faced a strategic choice: Using an overshoot strategy involved initially selecting a stick that was too long for the target and removing smaller stick lengths to arrive at the target length. Solvers who used the undershoot strategy built the desired stick by selecting an initial stick that was clearly too short and adding small stick lengths until the desired length was reached. Solvers were sensitive to both the history of the strategy's success in a global sense (some problems could not be solved with particular strategies) and to the current context of problem solving. That is, solvers adjusted their approach upon discerning that a particular problem may not be solvable with a particular strategy, even when that strategy had been successful overall. Lovett and Anders on concluded that the contributions of history and current context on operator choice were independent.

Despite some similarities to the Einstellung effects noted in Luchins and Luchins (1959) and Lovett and Anderson (1996), there are also some differences suggesting that a different interpretation of the current findings may be more productive. The WJT and BST have a common goal structure--the format of the problem's answer is always the same. In the present study, unlike the others described, what changed was not a specific goal that enabled the problem to be solved in fewer steps, but rather the goal structure itself changed for the participants in the shift group.

One implication of these findings is that knowledge of the code, or more generally, knowledge of the goal's structure, is part of the context in which strategic knowledge is assembled. Thus, these findings suggest a limit on the generality of a strategy. When solvers learn a strategy in a restricted portion of a problem's domain, this learning may not confer any strategic advantage on those solvers if and when they venture outside their previously restricted domain. With regard to Mastermind, this implication suggests that if solvers were not told about the format of the code prior to playing, they may not be decremented by later shifts to different goal structures, assuming that they are exposed to a variety of code formats during their training period. However, it is not clear that the modal hypothesis strategy would emerge at all under these circumstances, given that this strategy seems to be used in the service of moving the solver to a known goal.

This interpretation is consistent with that of Vollmeyer, Burns, and Holyoak (1996), who found that when people were provided with a specific goal at the outset of their learning experience in a rich, nonlinear domain, they tended to adopt a strongly goal-driven difference-reduction strategy. Although this strategy was shown to be helpful in moving the solver to the specific goal given at the outset of problem solving, such a strategy was also shown to be somewhat ineffective as a vehicle in learning about the totality of the problem space. Consequently, when such solvers were shifted to a different goal-state from the same problem space, they did not perform as well as solvers who had not been given a goal initially. Although it is difficult to think of playing Mastermind in any way except one that is goal-driven, it may well be the case that solvers who are left to discover the varieties of goal structures on their own (i.e., the varieties of code types) may have the opportunity to gather important informa tion about the effect of context on the implementation of a domain-specific strategy.

REFERENCES

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

Retrieved May 3, 2004 from http://www.findarticles.com/cf_0/m2405/2_127/65913662/print.jhtml

Anderson, J. R., & Lebiere, C. (1998). The atomic components of thought. Mahwah, NJ: Erlbaum.

Anderson, J. R., & Thompson, R. (1989). Use of analogy in a production system architecture. In S. Vosniadou & A. Ortony (Eds.), Similarity and analogical reasoning (pp. 267-297). Cambridge, UK: Cambridge University Press.

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

Best, J. B. (1990). Knowledge acquisition and strategic action in "Mastermind" problems. Memory & Cognition, 18, 54-64.

Best, J. B. (2000). Dissociations in problem solving and conditional reasoning in Mastermind. Manuscript submitted for publication.

Karat, J. (1982). A model of problem solving with incomplete constraint knowledge. Cognitive Psychology, 14, 538-559.

Lewis, M. W., & Anderson, J. R. (1985). Discrimination of operator schemata in problem solving: Learning from examples. Cognitive Psychology, 17, 26-65.

Lovett, M. C., & Anderson, J. R. (1996). History of success and current context in problem solving: Combined influences on operator selection. Cognitive Psychology, 31, 168-217.

Luchins, A. S., & Luchins, E. H. (1959). Rigidity of behavior: A variational approach to the effect of Einstellung. Eugene: University of Oregon Books.

Luchins, A. S., & Luchins, E. H. (1991). Task complexity and order effects in computer presentation of water jar problems. The Journal of General Psychology, 118, 45-72.

Simon, H. A. (1975). The functional equivalence of problem solving skills. Cognitive Psychology, 7, 268-288.

Simon, H. A., & Reed, S. K. (1976). Modeling strategy shifts in a problem-solving task. Cognitive Psychology, 8, 86-97.

VanLehn, K. (1991). Rule acquisition events in the discovery of problem-solving strategies. Cognitive Science, 15, 1-49.

Vollmeyer, R., Burns, B. D., & Holyoak, K. J. (1996). The impact of goal specificity on strategy use and the acquisition of problem structure. Cognitive Science, 20, 75-100.

COPYRIGHT 2000 Heldref Publications in association with The Gale Group and LookSmart.
COPYRIGHT 2001 Gale Group