Goal Similarity in Analogical Problem Solving

Bruce D. Burns

University of California, Los Angeles

Abstract

The role of goal similarity in analogical problem solving was investigated using highly simplified chess positions. Goal similarity was manipulated by instructing subjects to make an attacking or defensive move. Subjects received training positions, followed by a set of testings positions, each solvable by mapping to a training position. A normal chess position was also given. Testing positions maximally similar to training positions (including similarity of goals) were responded to most quickly, though this effect was not found for all positions. It was also found that when subjects had to avoid a fatal threat in a normal chess position, they were more likely to successfully defend against that move if they were told that they were losing than if told they were winning. The results indicate that goal similarity influences analogical problem solving.

Introduction

Seeing a new problem as analogous to an old problem can help develop a solution (Brown, Kane, & Echols, 1986; Carbonell, 1983; Gick & Holyoak, 1980, 1983). To do this requires retrieval of a source analogue and construction of a mapping between the source and the target analogue. How these analogical processes can be achieved remains uncertain. The goals of a cognitive system have often been proposed as important factors in analogy mapping, because the purpose of analogy mapping in problem solving is to accomplish the goals of the problem solving. This has led many artificial intelligence theorists to argue for the importance of goal accomplishment to retrieval of analogies (e.g., Carbonell, 1983, 1986; Hammond, 1989; Schank, 1982; Winston, 1980, 1982), as well as more psychologically orientated theorists (e.g., Holyoak and Thagard, 1989; Gentner, 1989) to argue that goal similarity has a role in analogical mapping.

This research was supported by Contract MDA 903-89-K-0179 from the Army Research Institute. I wish to thank Ramin Gabizadeh for data collection and Keith Holyoak for extensive comments. Correspondence should be sent to Bruce Burns, Department of Psychology, UCLA, Los Angeles, CA 90024, USA; or e-mail burns@cognet.ucla.edu

For the concept of goal similarity to be meaningful it needs to be distinguished from the more general influences of semantic similarity and structural consistency. The degree of semantic similarity relates to the similarity of the elements of a source and target analogy. Structural consistency refers to the degree to which the elements play similar roles in the analogues, which can be equated with a simple criterion for mapping: If two propositions are mapped, then their constituent predicates and arguments should also be mapped. There is a reasonable amount of psychological evidence that semantic similarity (e.g., Holyoak & Koh, 1987; Ross, 1987) and structural consistency (e.g., Gentner & Toupin, 1986; Holyoak & Koh, 1987; Ross, 1987, 1989) influence analogical mapping, although there is disagreement on exactly how these factors affect mapping. How, or whether, goal similarity is distinct from these factors is disputed. One reason for this dispute is that, despite the popularity of goal similarity in theories of analogical mapping and problem solving, there is very little psychological evidence that it plays a role. One suggestive study was that of Brown et al. (1986). They found that children who were directed to focus on the goal structure of a problem were better able to transfer solutions to analogous problems than were children not so directed. Despite this, both groups were able to recall the goals when later tested, suggesting that their poor performance was not due to a simple memory failure. However, it is possible that the directions acted as a hint, thus aiding retrieval.

The present study aimed to find evidence that goal similarity plays a role in analogical problem solving, and to develop a task suitable for the further study of the its effects. If goal similarity has been a slippery concept, this is partly due to the lack of a operational definition that might be provided by a task that can demonstrate its effects.

The Task

The problem solving task used in this study is novel and therefore it is necessary to explain its nature and rationale in some detail. The task involved highly simplified chess positions. Chess was chosen largely because it is possible to clearly define goals in terms of attacking and defensive moves. Different players given exactly the same position can have different goals (e.g., attack or defence) without requiring an overt change of the semantic or structural components of the problem.

The use of naive subjects dictated the use of a simple task. The positions consisted of two chess pieces on the board and one piece presented off the board, but waiting to be put on the board. For the attacking goal the task involved the subjects placing the piece that was off the board onto the board so as to guarantee that this piece would be able to capture an opposition piece on its next move. This next move occurs after one of those opposing pieces had had an opportunity to move. The task when the goal was defensive was the opposite of the attacking task: the subject must avoid the capture of a piece. The subject legally moved one of the white pieces already on the board in anticipation that the black piece off the board was about to be placed onto the board in an attempt to complete an attacking manoeuver similar to the one the subject used in the attacking task.

To help subjects achieve these goals they were trained on examples of two simple chess tactics known as *pins* and *forks*. These tactics are known by any experienced chess player, but they will not be familiar to novices. A fork manoeuver involves the attacking piece being placed so as to simultaneously attack two opposition pieces. Because the defending side's move consists of moving only one piece, only one piece will be able to escape the attack leaving the other to be captured. For a pin manoeuver the attacking piece is placed such that only one piece is directly attacked, but if this piece moves away then the one behind it can be captured. Hence a capture is still guaranteed. Examples of these training positions are presented in Figure 1 (an attacking bishop pin), Figure 2 (an example of a successful solution of the attacking bishop pin position), and Figure 3 (a successful solution to the defensive version of the bishop pin position).

After training on one of four attacking examples and on one of four defending examples, the subjects were given 40 testing trials consisting of similar simplified chess positions. Each testing position was a transformations of one of the eight training positions. Each transformation involved the same manoeuver (e.g., attacking bishop pin) but had the pieces changed (a change in semantic similarity), the relationship of the pieces to each other changed (structure inconsistency), or the goal changed (from attack to defence, or vice-versa). Because a subject received only two of the eight training positions (each involving a different manoeuver), transfer of a manoeuver was required when the subject was given one of the other training positions as a testing position. Therefore each testing position was classified by the type of transfer that would be required to get from the manoeuver used in one of the training position that the subject had been given, to the manoeuver in the training position of which the testing position was a transformation. The four transfer conditions were *similar*, *attack/defence*, *fork/pin*, and *bishop/rook*.

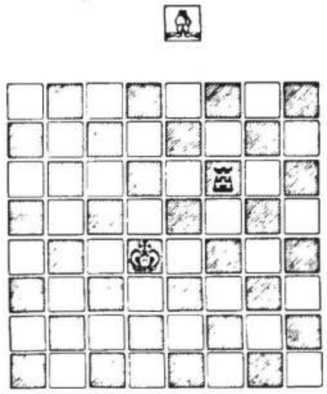

Figure 1. An attacking bishop pin position, as presented to a subject.

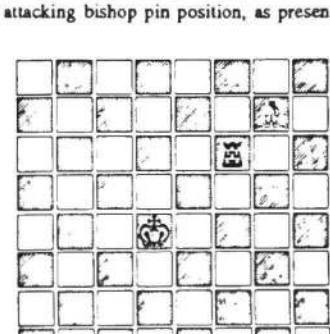

Figure 2. A solution to the attacking bishop pin position presented above.

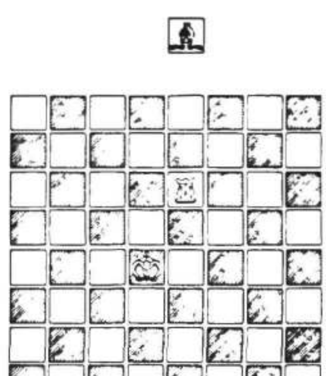

Figure 3. A solution to the defensive version of the bishop pin position presented above.

The similar condition involved the positions that used the same manoeuver (e.g., attacking rook fork) as that used in one of the subject's training positions and hence they differed the least from the training position. Of course, a subject's training position given as a testing trial would fall into this condition, though it could no longer be considered an analogue because it would actually be identical to the training position. The attack/defence condition involved a testing position that kept the rook/bishop and pin/fork characteristics of the manoeuver constant, but swapped attack for defence or vice-versa. This condition is considered to involve a change of goal. The fork/pin condition kept the attack/defence and bishop/rook characteristics of the required manoeuver the same, but swapped fork for pin, or vice-versa. This is considered to represent a structural change of the manoeuver. The bishop/rook condition involves positions that maintain the attack/defence and fork/pin characteristics but swap between rooks and bishops as the attacking piece. This is considered to primarily be a semantic change of the manoeuver, but because rooks and bishops move differently some structural change is unavoidable.

Due to the transformation of the training position, some form of transfer would be required to map even to a testing position that was a transformations of one the testing positions that the subject had been given. To control this it will be possible to directly compare positions that were constructed using the same type of transformation, but represent different transfer conditions.

Note that over subjects, each testing position appeared in each transfer condition, so that which condition a testing position was part of depended solely on what the subject's training positions were. Response time was used as the measure of performance because most of these positions are simple enough to be eventually solved by exhaustive search.

To investigate the effect of goal similarity, the comparison of the similar and the attack/defence conditions is most critical. If goal similarity is important for mapping or retrieval, then subjects should perform better on the positions in the similar condition than they do in the attack/defence condition. If subjects find the attack/defence positions harder then this will provide evidence that changing the goal makes problem solving harder and hence support the claim that goal similarity affects the analogical mapping and/or retrieval. It could be argued that differences in difficulty could be due simply to the attacking and defending tasks being different; there might be no transfer because training for one goal does not help with the other. However, there is in fact an intimate link between these two tasks. To solve a defending position requires knowledge of the attacking move. Subjects could not be sure that their solution to a defending position is correct unless they know what the threatened attack is, hence training on the attacking manoeuver should help the subject solve the defensive version of the manoeuver. Successful training on the defensive manoeuver involves becoming aware of the attacking form of that manoeuver.

Better performance should also be found in the similar condition as compared to the fork/pin and bishop/rook conditions, as previous work on problem solving tasks provide ample evidence that structure and semantics affect analogical mapping. The finding of such differences is important, however, because it would lend support to the assumption that solving the task involves mapping to the training positions.

A further attempt to investigate the effects of goals involved attempting to manipulate goals at the time of solving a problem. In a final position, subjects were given a full chess position and instructed to try to make the best move they could as though it were a normal chess position. This position involved a fatal move which was either available to the subject or that the subject had to identify so that they could defend against it. To manipulate goals, subjects were told either that they were winning or losing. This information is in fact irrelevant to selection of a correct move, but may have altered their immediate goal in the position. Subjects who were told they are winning would be more likely to make aggressive moves than those told they were losing. If goals are important then it would be predicted that telling a subject that they are winning should help them find the correct attacking move in the appropriate position, while telling subjects that they are losing should assist them in finding the correct defensive move when it is required.

Method

Subjects

A total of 81 subjects participated in the experiment. All subjects were from an introductory psychology subject pool at the University of California, Los Angeles. Seven subjects were eliminated because they failed to complete the task in the available time and two were eliminated because they never understood the requirements of the task. This left a sample of 72 subjects (53 male, 19 female).

Materials and Apparatus

All positions and instructions were presented by an Apollo series 4000 workstation with a 19 in. color monitor. Subjects made their responses using a mouse with three buttons. Response times were measured to an accuracy of one second.

Training positions

There were three factors each with two levels (fork or pin; bishop or rook; attack or defence) that were crossed, yielding eight training positions. However, only four different training conditions were used as each player was given two training positions that were the exact opposite of each other. For example, if one of the training positions was a bishop, pin, attack, then the other position would be a rook, fork, defence. In this way the subject was exposed to all possible elements. In particular, they became familiar with both the attacking and the defending task. It was especially important that the subject got a chance to clarify the nature of the defending task, since it is a more complex task than the attacking task.

Several constraints were applied to the construction of the eight training positions. The first constraint was that every attacking position must not only have a valid attacking solution, but must also have a defensive solution, because the same position was given as a defensive task. In effect, this means that there were only four different basic positions. The defensive versions of the training positions simply involved changing the goals and the colors of the pieces. Second, an attacking position should allow a pin solution, or a fork solution. Third, because of the restriction that an attacking piece cannot be placed onto a square where it could be immediately captured, it is not possible to create a fork position where one of the defending pieces is the same as the attacking piece. Hence one defending piece was always the opposite (in terms of rook and bishop) of the attacking piece and the other was a king. The use of a king made it easier to fulfill some of the other constraints. A fourth constraint was that it had to be taken into account that the training positions would be used as the basis for the construction of testing positions. The training positions therefore needed to be transformable in the appropriate ways.

Testing positions

Each testing position was based on one of the four basic training positions, but transformed in various ways. A total of 10 transformations were applied to each of the four basic training positions, yielding a total of 40 testing positions. Five of these transformations (*identical, rotation, change-piece, change-structure,* and *queen*) appeared for every training position. An identical transformation was the training position itself, which never appeared until at least the 22nd trial, by which time the subject had been exposed to all other possible manoeuvers. The rotation transformation was the training position rotated. The change-piece transformation was the same as the training position except that the defending piece that wasn't a king was replaced by a knight. The change-structure transformation replaced a defending piece and changed the relative position of the defenders (but an attacking maneuver of a similar type was still possible). The queen transformation used a queen as the attacking piece but allowed a similar attacking manoeuver to that used in the transformed training position. For fork training positions, five other types of more complicated transformations were used, but for pins (which are less flexible) these five were repeated so that each pin transformation was presented as both an attacking and a defending position.

As was the case for the training positions, there are many ways to construct the testing positions, but each must have both an attacking and a defending solution. The testing positions included 20 defending positions and 20 attacking. In order for all of the positions to be presented in both attacking and defending forms, two testing sets were formed. These two sets were identical, including the order of positions, except that every position that appeared as attacking in one set appeared as a defending position in the other set.

The order in which positions appeared was determined once. Order was determined by random selection, but simpler transformations had to appear before a more complex transformation of the same training position. Attacking and defending positions alternated, with odd numbered trials being attacking in one testing set and even numbered trials being attacking in the other. An equal number of pin and fork moves had to be attacking and defending.

Because the two training positions the subject received were maximally dissimilar, it was clear which transfer condition a testing position belongs to for each training condition. There was always one training position that was only one factor (attack/defence, fork/pin, or rook/bishop) different from the manoeuver used in the testing position, while the other was different by two factors (except for the similar condition where one training position has all three factors the same). Hence one training position was always clearly more similar to a given testing position than was the other training position.

It was this more similar training position that determined to which transfer condition a testing position belonged.

The final position was constructed so as to be a plausible chess position but with no worthwhile captures possible for either side. However, a fatal knight move was available to one side that would fork the other side's king and queen. Such a move is extremely dangerous because it guarantees the capture of the queen, which is the most powerful piece on the board. There was also a distractor move available, involving an attack on the other queen that is pointless at best, but is actually a losing move.

In the attacking version the subject had the opportunity to make the fatal move. In the defending version the identical position was used but with the colors reversed so that now white (the subject) must try to avoid this dangerous move, a task that could be easily accomplished if the subject realizes that it must be done. The attacking version of the position is presented in figure 4 with white about to move.

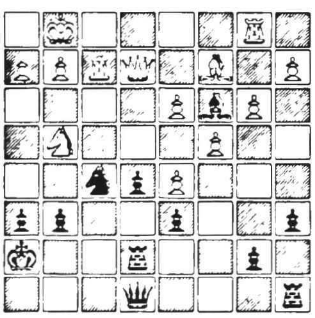

Figure 4. Final position (attacking version) with white about to move.

Procedure

Subjects were shown how the computer represented chess pieces and how to use the mouse to move pieces. They were told to follow the instructions that appeared on the screen, but that they could ask the experimenter questions if they were unclear with regard to any of the instructions.

The instructions first made clear that subjects would not see normal chess positions; in particular, kings were not special as they are in normal chess. Subjects were then given practice at detecting captures and were familiarized with using the mouse to move the pieces. Next, the subjects were presented with their attacking training position. They were to place the white piece on the board so as to guarantee the capture of a black piece no matter what move black subsequently made. It was emphasized that they could not place their piece on a square where it could immediately be captured. It was also made clear that it did not matter what happened to the white piece once it had captured a black piece, and therefore it was irrelevant if the piece they were to capture was protected by the other black piece.

The computer indicated if the subject's response was correct or incorrect. If an incorrect move was made, the computer displayed a correct response. This display lasted 10 s and showed the problem position with the white (attacking) piece placed on a correct square. To complete this stage subjects had to solve the position three times in a row. There were three different versions of the training position, each only differing from the others by a simple translation across the board.

The second set of practice trials presented subjects with their defending training position. Subjects were instructed that they were now to legally move one of the white pieces on the board so as to avoid the capture of a white piece. They were instructed that black was going to try to attack them in a similar way to that which the subject had used during the attacking training, but that they had a chance to pre-empt the placing of the attacking piece. If they made the right move at this stage then it would not matter where black placed their piece: black would not be able to force the capture of a white piece.

The link between attack and defence was made very explicit in the subjects instructions. They were advised to first think of where black was going to place their piece, and to then think how they could render that move harmless. It was also made clear that they had to be careful because even if they avoided one threat they might leave themselves open to another. Subjects were given similar feedback and the same completion criterion as they had for attacking training.

Subjects were then given the 40 testing trials from the randomly assigned testing set. No feedback was given at this stage, but subjects should have been able to recognize when they had made a correct move. After completing all testing trials subjects were told that they would be presented with a normal chess position and that they were to make the best move they could find. They were instructed that the relative value of pieces should be taken into account just as in normal chess (a factor that had been irrelevant up to this point). To guide them, a table giving the relative value of each piece was displayed. Whether the subject saw the attacking or the defending position was crossed with assignment to the winning or the losing condition. Subjects in the winning condition were instructed, "In your position, white is WINNING and is on the ATTACK". Subjects in the losing condition were instructed, "In your position, white is LOSING and is on the DEFENSIVE".

Results

The defending positions in which a queen was the attacking piece were inherently very difficult, as reflected in very low solution rates and very long response times. For this reason these moves were eliminated from the analysis of response time data, as were the attacking versions of these positions.

The other positions had a very high solution rate, particularity the attacking positions with 98% correct solutions. Defending positions were obviously harder but still had an 82% solution rate. As error rates were low, response times were analyzed. Analyses of response time data did not distinguish between correct and incorrect solutions. In an essentially self-terminating task such as this, response times will be greater for incorrect solutions than for correct ones. If response times for incorrect moves were eliminated, information about difficulty would be lost (as long as subjects were motivated).

Table 1
Means and standard deviations (in parenthesis) of response times for each transfer condition and for attacking and defending positions in seconds, collapsed across all transformation positions.

	similar	attack/ defence	fork/ pin	rook/ bishop	Total
Attacking	15.6 (10.5)	17.1 (10.0)	21.2 (23.5)	16.6 (13.9)	17.6
Defending	29.8 (30.5)	36.4 (36.8)	31.1 (29.3)	37.4 (39.0)	33.7
Overall	22.7 (23.9)	26.8 (28.6)	26.2 (27.0)	27.0 (31.0)	25.7

The mean response times for the attacking and defending positions in each of the four transfer conditions are presented in Table 1, collapsed over all transformations. A 2x4 between-subjects analysis of variance revealed a significant main effect of whether the position was attacking or defending, $F(1,2440) = 225$, $p<.01$, ($MS_e = 703$). This simply confirms the expectation that the defending positions are more difficult than the attacking positions. The predicted main effect of transfer condition was obtained, $F(3,2440) = 3.48$, $p<.05$, as was a significant interaction between the two variables, $F(3,2440) = 5.45$, $p<.01$.

The nature of the interaction was examined by looking at the attacking and defending positions separately (see Table 1 for means). For attacking positions there was a significant effect of transfer condition ($F[3,1220] = 7.88$, $p<.01$, $MS_e = 238$). However, planned comparisons revealed only that the fork/pin condition was significantly slower than the similar condition ($F[1,1220] = 20.34$, $p<.01$). There were no significant differences between the similar condition and the attack/defence condition ($F[1,1220] = 1.39$, $p<.25$) or the bishop/rook condition ($F[1,1220] = .67$, *n.s.*).

For defending positions there was also an effect of transfer condition ($F[3,1220] = 3.77$, $p<.025$, $MS_e = 1167$). The planned comparisons revealed that the similar condition was significantly faster than the bishop/rook condition ($F[1,1220] = 7.55$, $p<.01$), and the attack/defence ($F[1,1220] = 5.78$, $p<.01$), but not significantly faster than the pin/fork condition ($F[1,1220] = .21$, *n.s.*).

In order to investigate the differences for particular transformation positions, pin positions were examined. Because every subject received every pin position (which was not the case for forks) a more powerful repeated measures analysis could be used. The mean response times for each pin transformation position for each transfer condition are presented in Table 2.

Table 2
Mean response times (s) for pin positions for each transformation position in each transfer condition.

	Transformation			
	identical	rotation	change-piece	change-structure
similar	15.3	19.3	26.4	27.4
defence/attack	19.7	28.3	30.2	29.5
pin/fork	22.4	30.0	31.4	28.2
bishop/rook	21.2	29.2	28.1	23.3

For the identical transformation a multivariate analysis of variance revealed a significant effect of transfer, $F(3,69)=5.88$, $p<.01$. Planned univariate contrasts revealed significant differences between the similar condition and each of the other three transfer conditions: for similar with defence/attack, $F(1,71)=4.54$, $MS_e=160$, $p<.05$; with pin/fork, $F(1,71)=8.19$, $MS_e=226$, $p<.01$; and with bishop/rook, $F(1,71)=8.23$, $MS_e=155$, $p<.01$. Similar results were obtained for the rotation transformation in which a significant effect of transfer condition was obtained, $F(3,69)=5.49$, $p<.01$.

Planned univariate contrasts showed significant differences between the similar condition and each of the other three transfer conditions: for similar with defence/attack, $F(1,71)=4.65$, $MS_e=627$, $p<.05$; with pin/fork, $F(1,71)=6.51$, $MS_e=636$, $p<.05$; and with bishop/rook, $F(1,71)=7.84$, $MS_e=452$, $p<.01$. However there were no effects of transfer condition for the change-piece ($F[3,69]=.26$, *n.s.*) or the change-structure ($F[3,69]=.77$, *n.s.*) transformations.

Responses on the final position were analyzed by classifying all moves as correct or incorrect. A correct move in the attacking condition (when the subject had a chance to make the fatal move) was only recorded if they made the knight forking move. A correct move in the defending position (when the subject is threatened by the fatal move) is any move that renders the knight move harmless.

Table 3
Number of correct and incorrect solutions given to the final position in the attacking condition as a function of subject being given winning or losing manipulation.

	attacking position		defending position	
	winning	losing	winning	losing
Correct	5	9	0	4
Incorrect	6	8	16	9

The frequencies for correct and incorrect moves, for both attacking and defending positions, as a function of whether the subject was told that they were winning or losing, are presented in Table 3. There were no significant differences for attacking positions, $X^2(1) = .15$, *n.s.*, however the defending position yield a significant difference ($X^2[1] = 9.16$, $p<.005$) indicating that subjects were more likely to solve the problem, by recognizing the need for defence, when they were told that they were losing, rather than winning.

Discussion

These results provide support for the claim that goal similarity plays a role in analogical problem solving. In doing so they indicate that this task may be a useful one for further investigating the role of goal similarity in analogical processes. The overall findings, and those from the similar and rotation transformation positions, indicated that response times were faster if goals were kept consistent between the training and testing positions. This result is strengthened by the fact that the similar position appeared after many other attacking and defending positions had been seen and solved by the subjects. The fact that the conditions that we would expect to be hard to map, the fork/pin and bishop/rook conditions, were also generally slower then the similar condition supports the contention that mapping is being used to help solve positions.

The lack of a transfer position effect for the change-piece transformation positions is interesting. It appears to indicate that changing a piece is more than a simple semantic change. It may constitute a structural change, due to the fact that different pieces move differently. Hence changing a piece may change how the pieces relate to each other, even without changing their positions.

It is still not conclusively established, however that ease of analogical retrieval or mapping caused the faster responses in the similar conditions. Further research should try to establish directly how well subjects retrieve or map to these positions. Including a recognition task would be one way to do this. Another way to make the task more sensitive to mapping effects is to present ambiguous positions with multiple valid solutions and observe if subjects favor solutions that map most closely to their training positions

The finding that the goals at the time of test can influence solutions provides a powerful demonstration of the effects of influencing the subjects goals. This is particularly so since no other changes are made and the information that is provided is not directly relevant to determining the optimal move. However it is possible that this result could be explained in terms of demand characteristics. In future research this problem could be addressed by using less explicit instructions. It should also be noted that in this experiment the final position was not designed as an analogical task, as there is nothing provided as an explicit source analogue. Future studies could turn this into an analogical task by attempting to provide an explicit source. How well subjects could find such embedded analogues is unclear, but this ability could be examined by using a recognition task in which subjects would have to recognize an embedded analogy present in a full position (again goal similarity, structural, and semantic factors could be again manipulated).

This study did not examine skilled chess play; however the results have implications for chess skill, if they prove to be extendable to skilled chess players. If chess skill is based on analogical processes, then it would provide an explanation for the characteristic of chess masters that deGroot (1965) found that most distinguished them from less skilled players: that they simply selected better moves to examine in the first place. Because neither he nor the players were able to explain this, deGroot ascribed it to intuition, but its almost perceptual quality seems to fit quite well into an analogical framework.

In conclusion, this study provided support for the claim that goal similarity plays a role in problem solving, in that responses were made more quickly if consistent goals were given. This advantage may well be due to the improvement in the ease of analogical mapping that consistent goals provided, although further research is required to establish this with more confidence. The fact that this task is sensitive to such effects may mean it can be used to test different models of analogical mapping and problem solving. Such models often differ crucially in the ways they handle goals.

References

Brown, A. L., Kane, M. J., & Echols, C. H. (1986). Young children's mental models determine analogical transfer across problems with a common goal structure. *Cognitive Development, 1*, 103-121.

Carbonell, J. (1983) Learning by analogy: Formulating and generalizing plans from past experience. In R. Michalski, J. Carbonell, & T. Mitchell, (Eds.), *Machine learning: An artificial intelligence approach*. Los Altos: Tioga, 137-161.

Carbonell, J. (1986) Derivational analogy: A theory of reconstructive problem solving and expertise acquisition. In R. Michalski, J. Carbonell, & T. Mitchell (Eds.), *Machine learning: An artificial intelligence approach*, vol 2. Los Altos: Morgan Kaufman, 371-392.

deGroot, A. D. (1965). *Thought and choice in chess*. The Hague: Mouton.

Gentner, D. (1989). The mechanisms of analogical learning. In S. Vosniadou & A. Ortony (Eds.), *Analogy, similarity, and thought*. Cambridge, Mass.: Cambridge University Press.

Gentner, D., & Toupin, C. (1986). Systematicity and surface similarity in the development of analogy. *Cognitive Science, 10*, 277-300.

Gick, M. L., & Holyoak, K. J. (1980). Analogical problem solving. *Cognitive Psychology, 12*, 306-355.

Gick, M. L., & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology, 15*, 1-38.

Hammond, K. (1989). *Case-based planning*. Boston: Academic Press.

Holyoak, K. J., & Koh, K. (1987). Surface and structural similarity in analogical transfer. *Memory & Cognition, 15*, 332-340.

Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, 13, 295-355.

Schank, R. C. (1982). *Dynamic memory*. Cambridge: Cambridge University Press.

Ross, B. (1987). This is like that: The use of earlier problems and the separation of similarity effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13*, 629-639.

Ross, B. (1989). A further analysis of the access and use of earlier problems: Distinguishing different superficial similarity effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15*, 456-468.

Winston, P. H. (1980). Learning and reasoning by analogy. *Communications of the ACM, 23*, 689-703.

Winston, P. H. (1982). Learning new principles from precedents and exercises. *Artificial Intelligence, 19*, 321-350.