EUROPEAN JOURNAL OF COGNITIVE PSYCHOLOGY, 2003, *15* (3), 435–453

Is case-based reasoning a source of knowledge generalisation?

André Didierjean

Laboratoire de Psychologie Cognitive, Université de Provence–CNRS, Aix-en-Provence, France

This study examines the link between case-based reasoning and knowledge generalisation. Experiment 1 showed that case-based reasoning is not in itself a source of generalisation. In this experiment, subjects were able to adapt knowledge acquired during analysis of an example in order to successfully solve a very similar problem. But immediately afterwards, they failed on a problem symmetrical to the one they had just solved. Experiments 2A and 2B showed, in this same situation, that it is possible to lead subjects to generalise simply by stating that the practice problems were similar to the examples. These findings suggest that knowledge generalisation accompanying case-based reasoning is the result of a process that must be actively implemented by the subject, and not an automatic outcome of the source-to-target transfer.

A great deal of recent research has looked into the cognitive progress made during the analysis of examples. The theoretical aim of these studies is twofold: analyse the processes implemented in such situations, and examine the nature of the knowledge that is acquired. The experimental device is nearly always the same: In the first phase subjects analyse examples illustrating certain rules, and in the second phase they are asked to solve problems that require applying the rules present in the examples. This paradigm has been used in many areas, including reasoning by analogy (Reeves & Weisberg, 1994) and implicit learning (Redington & Chater, 1996).

Two types of processes appear to be at play in these situations, each associated with a different type of knowledge: (1) Subjects may store particular situations or exemplars in memory, and then later, when solving new problems, refer to these specific situations (Logan, 1988; Perruchet & Pacteau, 1990; Ross, 1984, 1987, 1989a, 1989b). (2) Subjects may extract the principles underlying

Requests for reprints should be addressed to A. Didierjean, LPC & CNRS, Université de Provence, 29 av. R. Schuman, 13621 Aix-en-Provence, France. Email: andre@up.univ-mrs.fr

This research was supported by a grant from the COGNITIQUE programme of the French Ministère de la Recherche.

http://www.tandf.co.uk/journals/pp/09541446.html DOI:10.1080/09541440244000247

the examples in order to apply them to future target situations (Catrambone & Holyoak, 1989; Cummins, 1992; Gick & Holyoak, 1983; Reber, 1989).

A number of studies have shown that novice subjects in a domain—who do not yet possess abstract rules—primarily use the first process (Reed, 1987, 1989; Reed & Bolstad, 1991). They store the specific features of the examples and attempt to adapt them to target problems. As subjects become more expert in a domain, the amount of domain-specific knowledge stored in memory increases, and part of their expertise lies in adapting that highly contextualised knowledge (e.g., Blessing & Ross, 1996; Chase & Simon, 1973; Gobet & Simon, 1996a, 1996b; Norman, Brooks, Coblentz, & Babcock, 1992; Vicente & Wang, 1998). At the same time, however, a number of studies suggest that another characteristic of experts is that they store more abstract knowledge than novices do (e.g. Chi, Feltovich, & Glaser, 1981; Chi, Glaser, & Rees, 1982; Cooke, Atlas, Lane, & Berger, 1993; Schank, 1982). A key question, then, concerns what processes subjects implement to generalise their knowledge based on the specific examples they encounter.

Some studies have shown that subjects build abstract knowledge during the example analysis phase, i.e., *prior to* the transfer. Accordingly, in the Gick and Holyoak (1983) study, subjects elaborated an abstract schema by comparing isomorphic source problems and detecting their similarities. This type of inductive process seems to be one means of generating abstract knowledge (see, for example, Cummins, 1992; Gick & Paterson, 1992). Certain findings suggest that subjects can even build a schema by analysing a single example (Ahn, Brewer, & Mooney, 1992). In this case, the subject's self-explanations seem to play a crucial role (Catrambone, 1994, 1995, 1996; Mooney, 1990).

The purpose of the present study was to look at another process, namely, knowledge generalisation *subsequent to* adaptation of a specific instance to a target problem. Ross and Kennedy (1990), who reused Ross's (1987, 1989a) paradigm, showed that subjects may build a schema after being reminded of the source. According to these authors, this generalisation mode is the one most naturally used by subjects (see also Novick & Holyoak, 1991). Since the Ross and Kennedy study, many authors have shown that generalisation is a ''natural'' consequence of the adaptation of specific situations (see, for example, Anderson, Fincham, & Douglass, 1997). However, few studies have taken an interest in the conditions under which this process occurs. In our minds, two hypotheses can be set forth:

(1) Unintentional generalisation. According to this hypothesis, generalisation is a natural consequence of the adaptation of a specific situation, and the subject does not have to deliberately engage in a target-to-source comparison process. Ross and Kennedy (1990, p. 52) defended this hypothesis, claiming that generalisation occurs concomitantly with analogical transfer: ''This learning results not from separate comparison process between solutions, but rather as by

a byproduct of the comparison that is needed to apply the earlier example and its solution to the current problem. Thus, the learning occurs from making the analogy, not from separate comparison process.'' As the specific situation is being adapted, the term-by-term comparison that takes place automatically leads to the creation of a schema. In this view, generalisation is solely dependent upon the source-target adaptation process. Reber (1989, 1997) used artificial grammar experiments to demonstrate the presence of such ''implicit'' generalisation processes. Despite the debates initiated by research on this issue, it is generally agreed that this type of process may exist (see Gomez & Schvaneveldt, 1994; Knowlton & Squire, 1996; but see Shanks & St John, 1994). In Reber's point of view, rule learning occurs in a passive manner (Reber, 1976), i.e., without the subject having to actively engage in a rule-extraction process. No statements will be made here about how conscious the subject might be during schema construction. In short, in the unintentional-generalisation hypothesis, subjects do not deliberately implement a generalisation process; it occurs in conjunction with the transfer.

(2) Deliberate generalisation. According to this hypothesis, knowledge generalisation is not an automatic outcome of the source-to-target transfer; it only takes place when the subject actually implements a generalisation process (such as detecting source and target similarities or making self-explanations of the problem). In the area of memory, it has long been known that memorisation is better and deeper when subjects deliberately try to memorise the information (e.g., Craik & Lockhart, 1972; Craik & Tulving, 1975; Goldman & Pellegrino, 1977). Similarly, generalisation concomitant to specific-knowledge adaptation may require the subject to engage in an active process. In line with this, certain authors found recently that the transfer observed in implicit learning tasks was linked to activities deliberately carried out by the subjects (Whittlesea & Wright, 1997; Wright & Whittlesea, 1998). Note that these authors also hypothesised that the transfer observed in this type of task, when it involved unknown letters (see, for example, Gomez, Gerken, & Schvaneveldt, 2000), was due to the adaptation of memorised exemplars actively compared to the test exemplars. Generalisation in this case is thought to be the outcome of this adaptation process. In short, in the deliberate-generalisation hypothesis, adaptation of a specific situation does not necessarily lead to knowledge generalisation. Any knowledge generalisation that follows an analogical transfer results from a process actively implemented by the subject.

The two experiments on chess playing presented here examine this generalisation process. They use the same experimental paradigm as in our previous study (Didierjean, Cauzinille-Marmèche, & Savina, 1999), where subjects were given chess problems and had to learn how to put the black king in checkmate in a few moves, no matter what the black player did. One of the advantages of this material is that it is relatively easy to see what moves have to be made: People who know the rules of chess but are novice players have no trouble understanding why such and such a sequence of moves always leads to checkmate. In this study, subjects first analysed examples of chess problems that illustrated different ways of achieving checkmate. Then they were given problems that were very much like the examples studied (same pieces to move and same corner of the chessboard), and problems that were unlike the examples. Some of the unlike problems were symmetrical to the like ones (for this material construction method, see also Gobet & Simon, 1996a). These unlike problems were identical to the like ones as to what pieces had to be moved, but differed in the orientation on the chessboard. The results showed that novices who studied sample chess problems built and applied relatively specific knowledge to solve the new problems: They managed to solve the example-like problems only when the king to put in checkmate was in the same corner of the chessboard. This result indicates that the subjects in this study used highly contextualised knowledge to reason by analogy. The two experiments presented below were based on the same experimental device, but this time our goal was to find out whether the mere process of adapting highly contextualised knowledge is a source of knowledge generalisation.

EXPERIMENT 1

Experiment 1 was aimed at determining whether subjects who transfer highly contextualised knowledge and successfully adapt it to a new problem, later generalise that knowledge. The ''unintentional-generalisation'' hypothesis leads to the prediction that this will happen systematically. The ''deliberate-generalisation'' hypothesis predicts that case-based reasoning can take place without subsequent problem schema construction.

The experimental setup was to have subjects study examples of chess problems and then test them on problems that were isomorphic to the examples but varied in how similar they looked to them. The order of the test problems was varied systematically: One group of subjects did very similar problems first (these problems are usually solved; see Didierjean et al., 1999) and then less-similar problems, including ones that were symmetrical to the similar problems (these problems are usually not solved; see Didierjean et al., 1999). Another group of subjects was tested on the same problems but in the opposite order (very similar problems in second place).

Based on the results obtained by Didierjean et al. (1999), the subjects in the second group can be expected to do well on the similar problems and poorly on the less-similar problems. For the first group, on the other hand, if adapting contextualised knowledge to a similar problem leads to generalisation, then we can expect subjects tested first on the similar problems to do better on the symmetrical problems presented after the similar ones. If analogical adaptation of examples is a genuine source of generalisation, then we can even expect the resulting generalised knowledge to enable subjects to correctly solve dissimilar problems. This experiment also included a control group. The control subjects analysed examples of problems that were not illustrative of the types of checkmate needed to win the game in the test phase.

Finally, in order to make sure the three groups were initially equivalent, subjects performed the memorisation task designed by Chase and Simon (1973). The many studies using this task have shown that the number of pieces recalled is a function of the subject's level of expertise. If the groups tested here were in fact matched on expertise, then the mean number of pieces recalled on each trial would be the same.

Method

Subjects

Ninety psychology students (mean age: 23 years, 7 months, *SD*: 7 months) participated in the experiment. All subjects judged themselves to be novices in chess (playing less than once a year) but familiar with the rules.

Materials

Pretest. The pretest materials consisted of a chessboard layout from a real game, after about ten moves.

Learning phase. The learning phase materials were two example problems presented on a chessboard, each one representing a different way of putting the opponent in checkmate.

1. *Checkmate-by-smothering*: The player sacrifices his/her queen by moving it next to the opponent's king. The situation is such that the opponent is forced to take the queen. The king is then ''smothered'', that is, he is completely surrounded by his own men. The opponent's knight is played and this puts the king in checkmate.
2. *Kiss-of-death-with-sacrifice*: The player sacrifices a piece by moving it next to the opponent's king. The situation is such that the king is forced to take the piece. The queen, protected by one of her own men, is moved next to the king, putting the king in checkmate.

Two example problems were added for the control group. They involved putting the black king in checkmate in some other way than checkmate-by-smothering or kiss-of-death.

Problem-solving phase. The material consisted of six chess problems. Three of the problems could be solved using the checkmate-by-smothering schema and the other three could be solved using the kiss-of-death-with-sacrifice schema. The three problems of each type varied as to how similar they were to the example (‘‘similar’’, ‘‘less-similar’’, and ‘‘dissimilar’’), but all required the same general schema.

The ‘‘similar’’ problem was very much like the example, in both the moves it required and its surface features: The pieces involved and their positions were nearly the same. The opponent’s king was located in the same corner of the chessboard, and the sequence of moves was very similar. The ‘‘less-similar’’ problem was symmetrical to the similar problem with respect to the axis that separated the chessboard into two equal halves for all pieces involved in the solution. This problem thus differed from the similar problem essentially by what corner of the chessboard the king was in. The ‘‘dissimilar’’ problem was based on the same solving principle, but the layout was considerably different from the example.

Figure 1 shows the examples and the similar, less-similar, and dissimilar problems for checkmate-by-smothering. The kiss-of-death problems are presented in the Appendix.

Procedure

The experiment was run in a single session lasting about 1 hour. Subjects were tested individually after being randomly assigned to one of three experimental groups: ‘‘control’’, ‘‘similar-first’’, and ‘‘similar-last’’.

Pretest. Subjects were seated facing a chessboard with the pieces placed next to it. Behind them was another chessboard showing the layout of a real game. On each trial, the subjects had to turn around and try to memorise the layout within 5 s. They were then given as much time as needed to reproduce the layout on the empty chessboard. There were five trials.

Learning phase. All subjects underwent a preliminary phase aimed at familiarising them with chess problems (explanations and an example). For the similar-first and similar-last groups, the subjects were shown two example chess problems (chessboard layouts), one after the other. The experimenter let the subjects try to solve the problems for three minutes, and then showed them the solution (if a subject found a correct solution to one of the examples, he/she was eliminated from the experiment). During the presentation of the solution, the experimenter did not offer a general explanation; he simply showed the sequence of moves and pointed out the immediate consequences of each move (e.g., ‘‘The whites move the rook to this square. This puts the king in check . . .’’). The same problems were then presented again so the subjects could reproduce the sequence of moves and explain them; this was repeated until the subject performed the task perfectly. The same procedure was used for the control group, but with the two problems requiring different types of checkmate.

Example problem

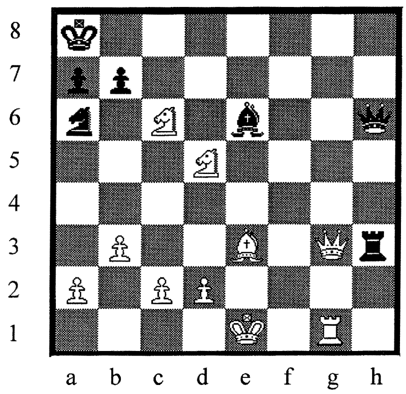

Solution:
1. Queen b8-Knight b8
2. Knight c7 checkmate

Similar problem

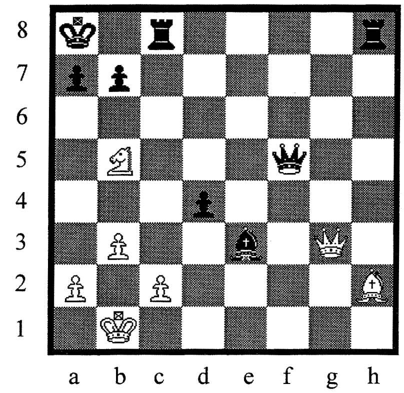

Solution:
1. Queen b8-Rook b8
2. Knight c7 checkmate

Less-similar problem

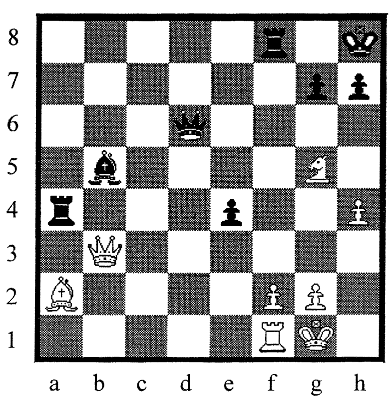

Solution:
1. Queen g8-Rook g8
2. Knight f7 checkmate

Dissimilar problem

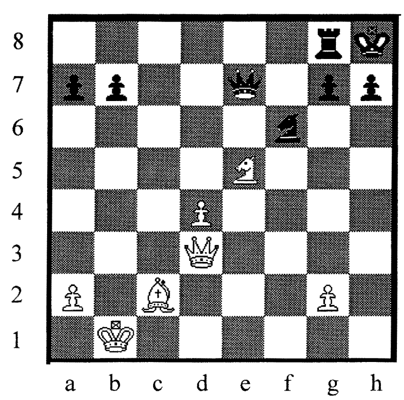

Solution:
1. Queen h7- Knight h7
2. Knight g6 checkmate

Figure 1. Checkmate by smothering: example, similar, less-similar, and dissimilar problems.

Post-test. The six problems to solve were given one at a time to the subjects. They had four minutes to find the solution. For the control group, a problem presentation order was defined and then counterbalanced (three subjects per group for each of the six orders). The two experimental groups differed as to the testing order: The similar-first group, as its name indicates, was given the two similar problems first and then the other four problems, and the similar-last group solved the less-similar problems and dissimilar problems first, followed by the similar problems.

Results

Pretest. No difference was found between the three groups as to the number of pieces recalled on the various trials, $F(12, 522) < 1$.

Post-test. The data were analysed using an ANOVA with three factors. The type of checkmate and the type of problem were within-subject factors, and the experimental group was a between-subject factor. Given that no difference was found between the two types of checkmate for the three groups, $F(1, 87) < 1$, the results presented below combine the two types of checkmate. Figure 2 shows the subjects' post-test performance on the three problem types.

There was an interaction between the subject group and the type of problem, $F(4, 174) = 6.0$, $MSE = 0.11$, $p < .001$. No significant difference was found between the similar-first and similar-last groups, $F(1, 87) = 1.61$, $MSE = 0.34$,

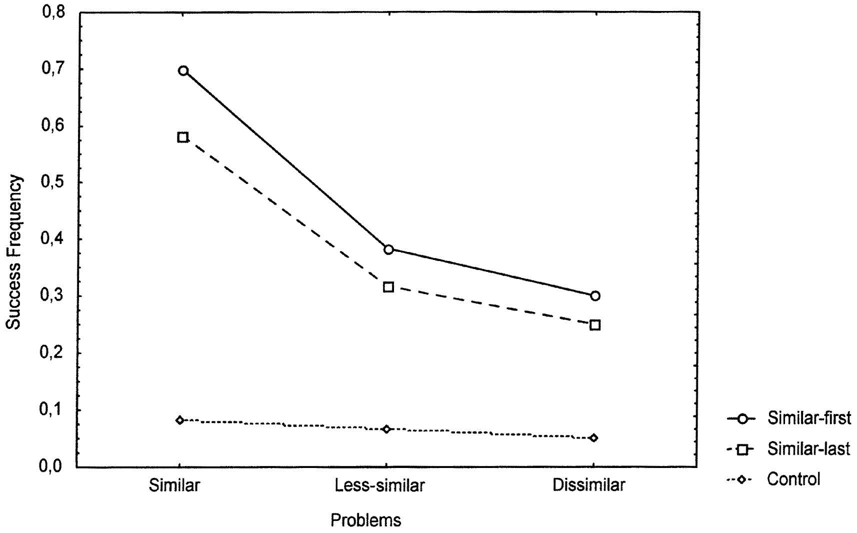

Figure 2. Results of Experiment 1. Success rate on similar, less-similar, and dissimilar problems for the subjects in the three groups.

$p = .21$. For both of these groups, similar problems were solved better than less-similar and dissimilar ones, $F(1, 87) = 86.29$, $MSE = 0.10$, $p < .0001$. No difference was observed between these two groups on the similar problems, $F(1, 87) = 2.12$, $MSE = 0.19$, $p = .15$, or on the less-similar problems, $F(1, 87) < 1$.

Discussion

The main purpose of Experiment 1 was to study the link between case-based reasoning and knowledge generalisation. The results showed first of all, that unlike the experimental group subjects, the controls, who were not given any appropriate training, failed massively on the post-test problems. This finding demonstrates that the experimental group subjects reasoned by analogy. Moreover, the difference in the success rates on similar and less-similar problems indicated that, on the average, the knowledge adapted by the subjects was contextualised. The findings show that retrieving a specific situation and adapting it to a similar problem does not necessarily lead to the construction of more abstract knowledge. In this experiment, the similar-first subjects solved the similar problems by analogy. Following the adaptation, this group as a whole apparently did not elaborate any reusable abstract knowledge. These subjects usually failed on problems that were symmetrical to the similar problems (ones that differed mainly by what corner of the chessboard the king was in). In particular, the lack of a difference between the profiles of these two experimental groups on less-similar problems raises the question of the link between analogical transfer and knowledge generalisation. This link was studied in greater depth in Experiment 2.

EXPERIMENT 2A

Two hypotheses were set forth here regarding the link between adaptation of a specific situation and the generalisation of knowledge.

Unintentional generalisation. The results of Experiment 1 showed that generalisation does not always follow analogical transfer. Why was there this discrepancy between our results and those obtained in many other studies? If the unintentional-generalisation hypothesis is valid nonetheless, then we can assume that the reason we obtained the results we did was that our subjects did little conceptualising during source problem storage, due to their lack of expertise in chess. They may have merely memorised the required sequence of moves without relying on an understanding of the concepts of the game. In the experiments by Ross (1987, 1989a; Ross & Kennedy, 1990), the subjects may have done more conceptualisation, although here, they did seem to know why the series of moves led to checkmate.

Deliberate generalisation. If the deliberate-generalisation hypothesis is correct, Experiment 1 subjects did not generalise their knowledge during the analogical transfer because the experimental situation did not lead them to implement such a process. Although they adapted an example to a similar problem, we can assume that they were not necessarily aware of the utility of extracting generalities from the situations presented.

In Experiment 2A, two groups were compared. One group was like the similar-first group in Experiment 1 and will hereafter be called the ''implicit-similarity'' group. The other group solved the problems in the same order as the first group. But for this group, the experimenter explicitly pointed out to the subjects that the similar problems were practice problems for learning the principles in the examples. This group will be called the ''explicit-similarity'' group.

If what the deliberate-generalisation hypothesis suggests is true, i.e., that Experiment 1 subjects did not generalise because they failed to implement the appropriate process, then we can expect an effect of the new instructions on the implementation of this process. If on the contrary, the unintentional-generalisation hypothesis is true, i.e., that what Experiment 1 subjects were lacking was sufficient conceptualisation of the domain, then we should not find any differences, no matter what instructions are given.

Method

Subjects

Sixty psychology students (mean age: 21 years, 8 months, *SD*: 19 months) participated in the experiment. All subjects said they were familiar with the rules, but judged themselves to be novices in chess (playing less than once a year).

Materials

The materials consisted of the experimental group problems used in Experiment 1: a board layout for the memory task, four checkmate-by-smothering problems, and four kiss-of-death-with-sacrifice problems. For both types of checkmate, the problems included an example, a similar problem, a less-similar problem (symmetrical to the similar problem), and a dissimilar problem.

Procedure

The experiment was run in a single session that lasted about 1 hour. It consisted of a pretest using the same memory task as in Experiment 1, followed by a learning phase and a post-test. The two experimental groups solved the problems in the same order as the similar-first group in Experiment 1. These two groups differed only by the instructions given for the various problems.

The implicit-similarity group heard the following instructions: ''I'm going to show you two examples of chess problems that illustrate two different ways of achieving checkmate. I'm going to let you try to solve the problems by yourself for a while, and then if you haven't found the answer, I'll show you. After the presentation of the examples, you will have to solve several problems, some of which are based on the same principles as the examples.'' After presenting the examples, the experimenter gave the subject the first problem and said, ''Now you're going to have to solve several problems, some of which are based on the same principles as the examples.''

For the explicit-similarity group, the instructions were as follows: ''First I'm going to show you two examples of chess problems that illustrate two different ways of achieving checkmate. I'm going to let you try to solve the problems by yourself for a while, and then if you haven't found the answer, I'll show you. After the presentation of the examples, you will have to solve two practice problems that are very similar to the examples you saw. This will help you learn how to achieve this type of checkmate. As a second step, you will have to solve several problems, some of which are based on the same principles as the examples.'' After presenting the examples, the experimenter gave the subject the first similar problem and said, ''This is the first of the two practice problems.'' When giving the second similar problem the experimenter said, ''Here is the second practice problem.'' After the two similar problems, the experimenter gave the next problem and said, ''Now you're going to have to solve several problems, some of which are based on the same principles as the examples.''

Results

Pretest. No difference was found between the two groups for the number of pieces recalled on the different trials, $F(6, 348) < 1$.

Post-test. The data were analysed using a three-factor ANOVA. The type of checkmate and the type of problem were within-subject factors, and the experimental group was a between-subject factor. No difference was found between the two types of checkmate for the two groups pooled, $F(1, 58) < 1$, so the data for the two types of checkmate was combined in the results presented below.

Figure 3 plots the subjects' success rate on the post-test for the three types of problems. We can see that the interaction between the subject group and problem similarity was not significant, $F(2, 116) = 2.13$, $MSE = 0.14$, $p = .12$. Planned comparisons indicated that the implicit-similarity group performed significantly better on similar problems than on less-similar ones, $F(1, 58) = 6.52$, $MSE = 0.13$, $p < .05$. For the explicit-similarity group, the success rate on these two problems was not significantly different, $F(1, 58) < 1$. No difference was observed between the two groups on the similar problems,

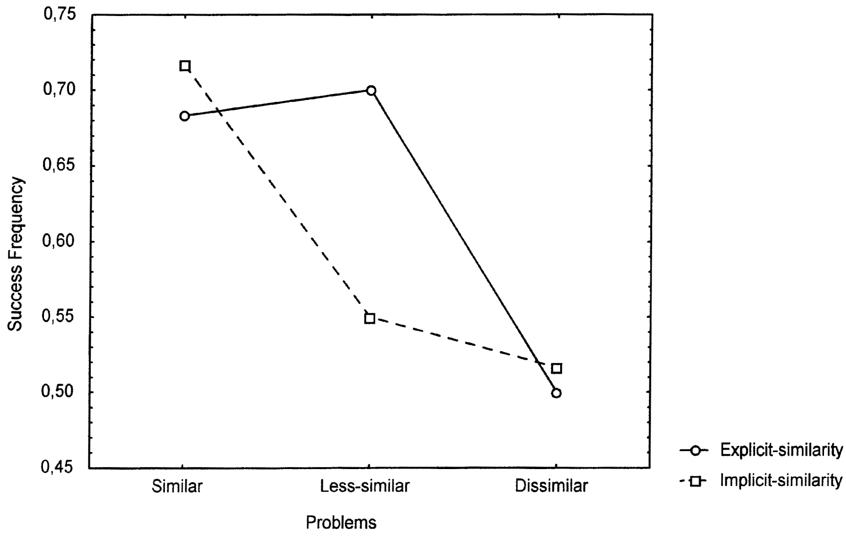

Figure 3. Results of Experiment 2A. Success rate on similar, less-similar, and dissimilar problems for the subjects in the two experimental groups.

$F(1, 58) < 1$, but a significant difference was found on the less-similar problems, $F(1, 58) = 4.46$, $MSE = 0.25$, $p < .05$. In addition, an analysis of the individual profiles showed that when the less-similar problems were correctly solved on either type of checkmate, 89% of the time it was after the similar problem had also been correctly solved.

Discussion

Note first of all that the implicit-similarity group had the same profile as the similar-first group. These subjects solved similar problems considerably better than less-similar and dissimilar ones.

The results were quite different for the explicit-similarity group. The performance of these subjects on less-similar problems was equivalent to their performance on similar problems. This finding means that simply changing the instructions led these subjects to generalise their knowledge more: They apparently constructed knowledge that did not depend on which corner of the chessboard was involved in the play. However, they did not generalise to an extent that allowed them to solve the dissimilar problems. This result suggests that what Experiment 1 subjects were lacking was not so much a conceptual understanding of the game as the use of a generalisation process. Note, however, that in Experiment 2A the two experimental groups also differed as to the prompting about using the examples. The instructions informed the explicit-similarity subjects that certain examples were very much like the problems they would solve *after* the similar problems. The implicit-similarity group was not given any further instructions between the similar problems and the other problems. Experiment 2B was aimed at making sure that the results obtained in Experiment 2A were not due to this difference in the instructions.

EXPERIMENT 2B

Experiment 2B was the same as Experiment 2A in every respect but one: The experimenter reminded the subjects in the implicit-similarity group that the problems to be solved were similar to the examples given before the less-similar and dissimilar problems.

Method

Subjects

Forty psychology students (mean age: 19 years, 5 months, *SD*: 17 months) participated in the experiment. All subjects said they were familiar with the rules, but judged themselves to be novices in chess (playing less than once a year).

Materials

The materials were the same as in Experiment 2A.

Procedure

The procedure was the same as in Experiment 2A with one exception: After the similar problems were solved, the experimenter said to the implicit-similarity group subjects, ‘‘Remember that among the problems you will have to solve, some are based on the same principles as the examples you just studied.’’

Results and discussion

Pretest. No difference was found between the two groups for the number of pieces recalled on the different trials, $F(6, 228) < 1$.

Post-test. The data were analysed using a three-factor ANOVA. The type of checkmate and the type of problem were within-subject factors, and the experimental group was a between-subject factor. No difference was found between the two types of checkmate for the two groups pooled, $F(1, 38) < 1$, so the data for the two types of checkmate were combined in the results presented below.

Figure 4 plots the subjects' success rate on the post-test for the three types of problems. We can see that the interaction between the subject group and the degree of problem similarity was significant, $F(2, 76) = 6.45$, $MSE = 0.15$, $p <$

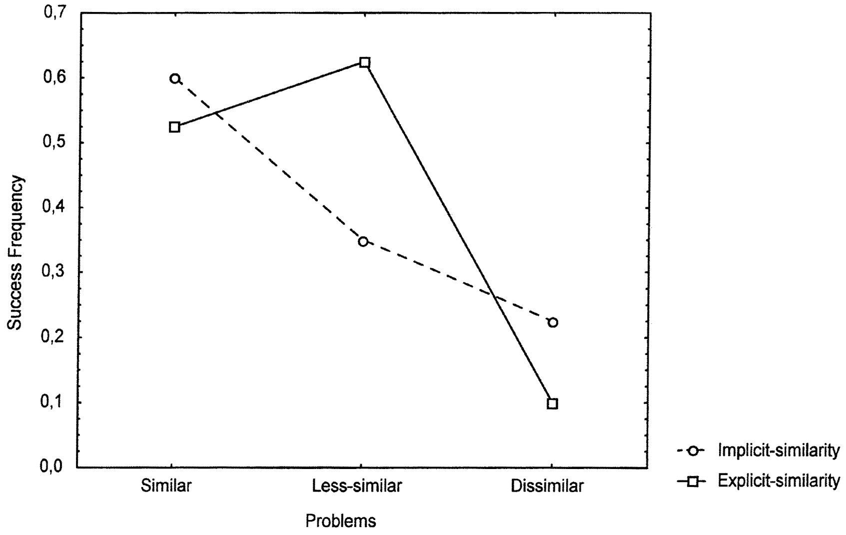

Figure 4. Results of Experiment 2B. Success rate on similar, less-similar, and dissimilar problems for the subjects in the two experimental groups.

.01. As in Experiment 2A, no difference was observed between the two groups on the similar problems, $F(1, 38) < 1$, but a significant difference was found on the less-similar problems, $F(1, 38) = 5.24$, $MSE = 0.29$, $p < .05$. In addition, an analysis of the individual profiles showed that when the less-similar problems were correctly solved on either type of checkmate, 81% of the time it was after the similar problem had also been correctly solved. These results confirm those obtained in Experiment 2A and show that the difference between the two groups was indeed due to whether or not they were prompted to generalise.

GENERAL DISCUSSION

The purpose of this study was to investigate the link between a particular form of reasoning by analogy, case-based reasoning, and the generalisation of knowledge. Following the pioneering work by Gick and Holyoak (1980, 1983), a great deal of research on reasoning by analogy has focused on the relationship between analogical transfer and knowledge generalisation, but most of the studies have looked at the conditions under which subjects generalise knowledge before making the analogical transfer (Catrambone & Holyoak, 1989; Cummins, 1992). Since these studies, knowledge generalisation has often been regarded as a necessary condition for analogical transfer. At the same time, other studies have shown that reasoning by analogy is influenced by the surface features shared by the source and target (Blessing & Ross, 1996; Heydenbluth & Hesse, 1996; Ross, 1987, 1989a). A widely accepted hypothesis is that reasoning by analogy can take on forms other than the adaptation of an abstract knowledge structure. For example, subjects may adapt a highly contextualised knowledge structure (Bernardo, 1994; Kolodner, 1993; Reed, 1989; Ross, 1987, 1989a; for a review, see Reeves & Weisberg, 1994), and in fact, this form of analogical reasoning is very common, not only in novices (Reed, 1987, 1989; Reed & Bolstad, 1991) but also in experts (Blessing & Ross, 1996; Norman et al., 1992). Our goal was to study the generalisation process that takes place after an analogical transfer.

Two hypotheses were set forth here. In the *unintentional-generalisation* hypothesis, knowledge generalisation is a ''natural'' consequence of reasoning by analogy; the analogical transfer itself is a generalisation process. In the *deliberate-generalisation* hypothesis, knowledge generalisation is not a natural consequence of reasoning by analogy; it requires the active implementation of a generalisation process by the subject.

Our results clearly support the latter hypothesis: The process of adapting a specific problem does not necessarily trigger knowledge generalisation. Experiment 1 showed that subjects can adapt highly specific, contextualised knowledge without engaging in a reasoning-by-analogy process that leads to the elaboration of less context-dependent knowledge. Experiments 2A and 2B showed that subjects' activity can be modulated solely by attracting their attention to the utility of generalising their knowledge during problem solving. Our subjects built knowledge that was sufficiently context-independent to allow them to succeed on the chess problems presented, regardless of where on the chessboard the opponent's king was located. In the study of the links between analogical transfer and knowledge generalisation, it would be worthwhile to conduct further research on the nature of the processes implemented by these subjects. Two interpretations can be proposed. The first is that giving a particular ''status'' to the target problems in Experiment 2 led the subjects to compare the source and target *after* the analogical transfer, in order to extract generalities from them. The second is that it was the nature of the knowledge transferred to the target that turned out to be different. The status given the targets in Experiment 2 may have triggered attempts to extract and adapt more abstract knowledge from the source.

Manuscript received July 2001
Revised manuscript received October 2002

REFERENCES

Ahn, W., Brewer, W. F., & Mooney, R. J. (1992). Schema acquisition from a single example. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*, 391–412.

Anderson, J. R., Fincham, J. M., & Douglass, S. (1997). The role of examples and rules in the acquisition of a cognitive skill. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *23*, 932–945.

Bernardo, A. B. I. (1994). Problem-specific information and the development of problem-type schemata. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 379–395.

Blessing, S. B., & Ross, B. H. (1996). Content effects in problem categorization and problem solving. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*, 792–810.

Catrambone, R. (1994). Improving examples to improve transfer to novel problems. *Memory and Cognition*, *22*, 606–615.

Catrambone, R. (1995). Aiding subgoal learning: Effects on transfer. *Journal of Educational Psychology*, *87*, 5–17.

Catrambone, R. (1996). Generalizing solution procedures learned from examples. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*, 1020–1031.

Catrambone, R., & Holyoak, K. J. (1989). Overcoming contextual limitations on problem-solving transfer. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*, 1147–1156.

Chase, W. G., & Simon, H. A. (1973). Perception in chess. *Cognitive Psychology*, *4*, 55–81.

Chi, M. T. H., Feltovich, P. J., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science*, *5*, 121–152.

Chi, M. T. H., Glaser, R. G., & Rees, E. T. (1982). Expertise in problem solving. In R. Sternberg (Eds.), *Advances in the psychology of human intelligence*. Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Cooke, N. J., Atlas, R. S., Lane, D. M., & Berger, R. C. (1993). Role of high-level knowledge in memory for chess positions. *American Journal of Psychology*, *106*, 321–351.

Craik, F. I. M., & Lockhart, R. S. (1972). Level of processing: A framework for memory research. *Journal of Verbal Learning and Verbal Behavior*, *11*, 671–684.

Craik, F. I. M., & Tulving, E. (1975). Depth of processing and the retention of words in episodic memory. *Journal of Experimental Psychology: General*, *104*, 268–294.

Cummins, D. (1992). Role of analogical reasoning in the induction of problem categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *5*, 1103–1124.

Didierjean, A., Cauzinille-Marmèche, E., & Savina, Y. (1999). Learning from examples: Case-based reasoning in chess for novices. *Current Psychology of Cognition*, *18*, 337–361.

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

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

Gick, M. L., & Paterson, K. (1992). Do contrasting examples facilitate schema acquisition and analogical transfer? *Canadian Journal of Psychology*, *46*, 539–550.

Gobet, F., & Simon, H. A. (1996a). Recall of random and distorted chess positions: Implications for the theory of expertise. *Memory and Cognition*, *24*, 493–503.

Gobet, F., & Simon, H. A. (1996b). Templates in chess memory: A mechanism for recalling several boards. *Cognitive Psychology*, *31*, 1–40.

Goldman, S. R., & Pellegrino, J. W. (1977). Processing domain, encoding elaboration and memory trace strength. *Journal of Verbal Learning and Verbal Behavior*, *16*, 29–43.

Gomez, R. L., Gerken, L., & Schvaneveldt, R. W. (2000). The basis of transfer in artificial grammar learning. *Memory and Cognition*, *28*, 253–263.

Gomez, R. L., & Schvaneveldt, R. W. (1994). What is learned from artificial grammars? Transfer tests of simple association. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 396–410.

Heydenbluth, C., & Hesse, F. W. (1996). Impact of superficial similarity in the application phase of analogical problem solving. *American Journal of Psychology*, *109*, 37–57.

Kolodner, J. (1993). *Case-based reasoning*. San Mateo, CA: Morgan Kaufman Publishers.

Knowlton, B. J., & Squire, L. R. (1996). Artificial grammar learning depends on implicit acquisition of both abstract and exemplar-specific information. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*, 169–181.

Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, *9*, 492–527.

Mooney, R. J. (1990). Learning plan schemata from observation: Explanation-based learning from plan recognition. *Cognitive Science*, *14*, 483–509.

Norman, G. R, Brooks, L. R, Coblentz, C. L., & Babcook, C.J. (1992). The correlation of feature identification and category judgments in diagnostic radiology. *Memory and Cognition*, *20*, 344–355.

Novick, L. R., & Holyoak, K. J. (1991). Mathematical problem solving by analogy. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*, 398–415.

Perruchet, P., & Pacteau, C. (1990). Synthetic grammar learning: Implicit rule abstraction or explicit fragmentary knowledge? *Journal of Experimental Psychology: General*, *119*, 264–275.

Reber, A. S. (1976). Implicit learning of synthetic languages: The role of instructional set. *Journal of Experimental Psychology: Human Learning and Memory*, *2*, 88–94.

Reber, A. S. (1989). Implicit learning and tacit knowledge. *Journal of Experimental Psychology: General*, *118*, 219–235.

Reber, A. S. (1997). How to differentiate implicit and explicit modes of acquisition. In J. D. Cohen & J. W. Schooler (Eds.), *Scientific approaches to consciousness*. Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Redington, M., & Chater, N. (1996). Transfer in artificial grammar learning: A reevaluation. *Journal of Experimental Psychology: General*, *125*, 123–138.

Reed, S. K. (1987). A structure-mapping model for word problems. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *13*, 124–139.

Reed, S. K. (1989). Constraints on the abstraction of solutions. *Journal of Educational Psychology*, *81*, 532–540.

Reed, S. K., & Bolstad, C. A. (1991). Use of examples and procedures in problem solving. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*, 753–766.

Reeves, L. M., & Weisberg, R. W. (1994). The role of content and abstract information in analogical transfer. *Psychological Bulletin*, *115*, 381–400.

Ross, B. H. (1984). Remindings and their effects in learning a cognitive skill. *Cognitive Psychology*, *16*, 371–416.

Ross, B. H. (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. H. (1989a). Distinguishing types of superficial similarities: Different effects on the access and use of earlier problems. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*, 456–468.

Ross, B. H. (1989b). Reminding in learning and instruction. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 438–469). Cambridge, UK: Cambridge University Press.

Ross, B. H., & Kennedy, P. T. (1990). Generalizing from the use of earlier examples in problem solving. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *16*, 42–55.

Schank, R. C. (1982). *Dynamic memory: A theory of reminding and learning in computers and people*. Cambridge, UK: Cambridge University Press.

Shanks, D. R., & St John, M. F. (1994). Characteristics of dissociable human learning systems. *Behavioral and Brain Sciences*, *17*, 367–447.

Vicente, K. J., & Wang, J. H. (1998). An ecological theory of expertise effects in memory recall. *Psychological Review*, *105*, 33–57.

Whittlesea, B. W. A., & Wright, R. L. (1997). Implicit (and explicit) learning: Acting adaptively without knowing the consequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *23*, 181–200.

Wright, R. L., & Whittlesea, B. W. A. (1998). Implicit learning of complex structures: Active adaptation and selective processing in acquisition and application. *Memory and Cognition*, *26*, 402–420.

APPENDIX

"Kiss of death" example

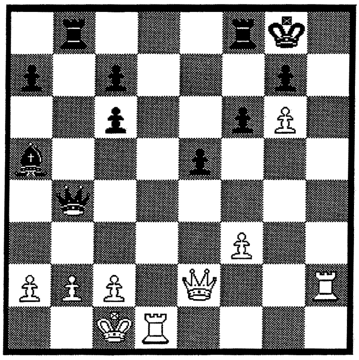

Solution: 1. Rook h8-King h8 2. Queen h2-King g8 3. Queen h7 checkmate

"Kiss of death" similar problem

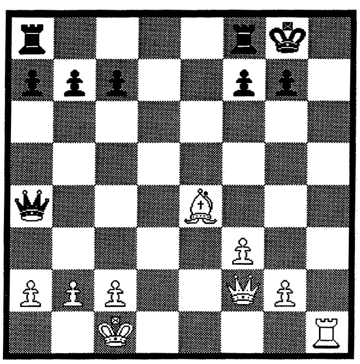

Solution: 1. Rook h8-King h8 2.Queen h4-King g8 3. Queen h7 checkmate

"Kiss of death" less-similar problem

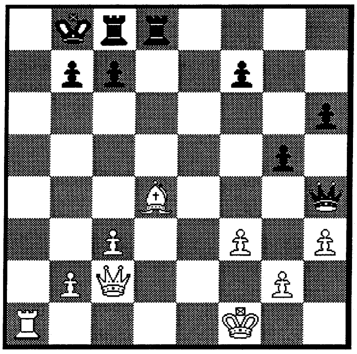

Solution: 1. Rook a8-King a8 2. Queen a4-King b8 3. Queen a7 checkmate

"Kiss of death" dissimilar problem

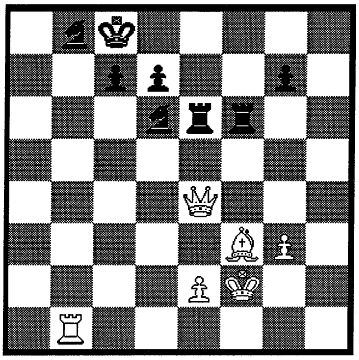

Solution: 1. Rook b8-King b8 2. Queen a8 checkmate