

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Analogical Transfer of Non-Isomorphic Source Problems

Permalink

<https://escholarship.org/uc/item/9rz7v3nr>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 21(0)

Authors

Schmid, Ute

Wirth, Joachim

Polkehn, Knut

Publication Date

1999

Peer reviewed

Analogical Transfer of Non-Isomorphic Source Problems

Ute Schmid (schmid@cs.tu-berlin.de)

Department of Computer Science, Technische Universität Berlin, Franklinstr. 28, 10587 Berlin, Germany

Joachim Wirth (wirth@mpib-berlin.mpg.de)

Center for Educational Research, Max Planck Institute for Human Development, 14195 Berlin, Germany

Knut Polkehn (knut.polkehn@rz.hu-berlin.de)

Institute of Psychology, Humboldt Universität, 10099 Berlin, Germany

Abstract

In analogical problem solving, non-isomorphic source/target relations are typically only investigated in contrast to the ideal case of isomorphism. We propose to give a closer look to different types of non-isomorphic source/target relations and varying degrees of structural overlap. We introduce a measure of graph distance which captures the "size" of partial isomorphism between two structures and we present two experiments investigating the influence of different non-isomorphic relations on analogical transfer. In the first experiment we contrast transfer performance for isomorphic vs. source inclusive problems with high vs. low superficial similarity. In the second experiment we explore different types of partial isomorphisms: source inclusiveness, target exhaustiveness, and different degrees of source/target overlap. The results indicate that (1) transfer of isomorphs is not significantly influenced by superficial similarity but transfer of partial isomorphs is, and (2) partial isomorphs can be transferred successfully if the amount of structural overlap is at least as high as structural differences. The experiments were inspired by some open design questions for the analogy module of IPAL (a computational model integrating problem solving and learning).

Introduction

Analogical problem solving is commonly described by the component processes retrieval, mapping and transfer. The work presented in this paper focusses on *analogical transfer*. Transfer can be faulty or incomplete, even if retrieval and mapping were successful (Novick & Holyoak, 1991). We are especially interested in transfer of *non-isomorphic* sources – the standard case in everyday problem solving. Several studies (cf. Reed, Ackinclose, & Voss, 1990; Novick & Hmelo, 1994; Spellman & Holyoak, 1996; Gholson, Smither, Buhrman, Duncan, & Pierce, 1996) show that subjects *can* transfer non-isomorphic sources successfully – at least when retrieval and mapping information is given explicitly. In our experiments, we want to look closer at the influence of different types and degrees of structural source/target similarities on transfer success.

This question is interesting for several reasons: (1) In an educational context (cf. tutoring systems) the provided examples have to be carefully balanced to allow for generalization (learning). Presenting only isomorphs restricts learning to small problem classes, while too large a degree of structural dissimilarity can result in failure of transfer and thereby obstructs learning (Pirolli & Anderson, 1985). (2) A plausible cognitive model of analogical problem solving (cf. Falkenhainer, Forbus, & Gentner, 1989; Holyoak & Thagard, 1989; Hummel & Holyoak, 1997) should generate correct

transfer only for such source/target relations where human subjects perform successfully. (3) Computer systems which employ analogical or case-based reasoning techniques (Carbonell, 1986; Schmid & Wysotzki, 1998) should refrain from analogical transfer when there is a high probability of constructing faulty solutions. Thus, it can be avoided that system users have to check – and possibly debug – generated solutions. Information about conditions for successful transfer in human analogical problem solving can provide guidelines for implementing criteria for when analogical reasoning should be rejected in favor of other problem solving strategies.

Our experiments were mainly motivated by this last reason (Schmid, Mercy, & Wysotzki, 1998). We are well aware that analogical problem solving is strongly influenced by semantic and pragmatic aspects of the involved problems (Hummel & Holyoak, 1997). But we believe that there are still open questions with respect to the structural basis (Falkenhainer et al., 1989) of analogical transfer which are worthwhile to investigate (see also results on dominance of systematicity over pragmatic relevance in Markman & Sanchez, 1998).

In the following, we introduce our problem domain and describe how we constructed problems with different types and varying degrees of structural similarity. Afterwards, we first present an experiment contrasting the effect of superficial and structural similarity on transfer success; second we present an experiment contrasting target exhaustiveness, source inclusiveness, and different degrees of structural overlap between problems. Finally, we will describe how the experimental findings can be used to improve the performance of the analogy module of our problem solving and learning system IPAL.

Non-Isomorphic Variants in a Water Redistribution Domain

Water redistribution problems

Our material is based on a modification of the water jug domain (Luchins & Luchins, 1950). In contrast to the classical problems we investigate *redistribution* problems (Atwood & Polson, 1976), for example: Given three jugs with capacities $A = 36$, $B = 45$ and $C = 54$ liters and initial quantities $A = 16$, $B = 27$ and $C = 34$ liters, find a (minimal) sequence of operations *pour from jug x to y* so that the jugs contain $A = 25$, $B = 0$ and $C = 52$ liters. An example problem is given in figure 1.

The *pour*-operator is defined in the following way:

IF *not(empty(x))* and *not(filled(y))* THEN *pour(x,y)* resulting




Figure 1: A water redistribution problem

in:
 IF $current(x) \leq max(y) - current(y)$
 THEN $current(y) := current(y) + current(x)$, $current(x) := 0$
 ELSE $current(x) := current(x) - (max(y) - current(y))$, $current(y) := max(y)$.

Water can be poured only from a non-empty jug and only into a jug which is not completely filled. Pouring results in filling y up to its capacity (possibly leaving a rest of the water in x) or in emptying x (possibly leaving a free capacity in y).

Problem Analysis

We are especially interested in the influence of *structural* similarity between source and target on transfer success. For this reason, we use only problems from the same domain – water redistribution problems – and with identical goals – find a sequence of *pour* operations so that each jug contains the desired amount of water. That is, we keep semantic and pragmatic aspects (Holyoak & Thagard, 1989) constant.

To investigate analogical transfer, we want to make sure that subjects really refer to the source for solving the target problem. Therefore, the problems should be complex enough that the correct solution cannot be found by trial and error, and difficult enough that the abstract solution principle is not immediately inferable. We constructed redistribution problems for which exists only a single (for two problems two) shortest operation sequence – in problem spaces with over 1000 states and more than 50 cycle-free solution paths¹. The strategy for solving redistribution problems is to express goal quantities in terms of relations between initial and maximum quantities. For example, the goal quantity (25 liters) of the small jug (A) in figure 1 can be obtained in the medium jug by filling it up to its capacity (45) and pouring in the small jug ($25 = 45 - (36 - 16) = 45 - 36 + 16$). Abstracting from the given values, the relation between the goal quantity of jug A and given initial quantities (*start*) and capacities (*max*) is $goal(small-jug) = max(medium-jug) - max(small-jug) + start(small-jug)$ (see jug j_3 in fig. 2a). Even for three-jug problems, calculating the desired redistribution is quite complex².

Redistribution problems can be described by the following attributes, operations and relations:

¹The algorithm for generating the set of all solutions can be obtained from the authors.

²The rules for calculating redistribution sequences can be obtained from the authors.

- superficial features: names (A, B, C ...) and positions (left, right, middle ...) of jugs,
- relevant features: jug capacities ($max(j)$), initial quantities ($start(j)$), and goal quantities ($goal(j)$),
- relevant operations and relations: ordinal difference between jug capacities ($j_i < j_j$), relative differences between quantities (e.g. $goal(j_1) = (start(j_1) + start(j_2)) - (max(j_2) - max(j_3))$) for the largest jug, j_1 in fig. 2a).

The structure of the problem given in figure 1 is presented in figure 2a. Note, that we represent only the aspects of the problem structure which are *relevant* for calculating the solution, e.g. we do not represent the relation between $start(j_1)$ and $start(j_3)$ or $max(j_2)$ and $goal(j_2)$.

We do neither claim that human problem solvers without experience with this problem domain represent the relevant problem structure completely and correctly, nor do we make assumptions whether analogical problem solving is better modelled on a symbolic (Falkenhainer et al., 1989) or a sub-symbolic (Hummel & Holyoak, 1997) level. We constructed this “normatively complete” symbolic graph representation to explore the impact of different analytically given structural source/target relations on empirical observable transfer success.

Non-Isomorphic Source/Target Relations

In the following experiments, we are interested in a special kind of non-isomorphism – **partial isomorphism**. That is, we do not consider many-to-one (Spellman & Holyoak, 1996; i.e. epimorphisms, see Schmid et al., 1998) or one-to-many (Spellman & Holyoak, 1996; i.e. no morphism, see Schmid et al., 1998) mappings. Instead we investigate source/target relations which share a common substructure. There are different kinds of partial isomorphic source/target relations:

- **target exhaustiveness:** the target problem is completely contained in the source (Gentner, 1980),
- **source inclusiveness:** the source problem is completely contained in the target (cf. Reed et al., 1990),
- **source/target overlap:** source and target share a common part, but both problems have additional aspects (cf. Carbonell, 1986).




Figure 2: Structure of the source problem (a) and a partial isomorph (b)

For all three types of partial isomorphs, the *degree* of structural overlap can vary. For example, the common substructure can consist of only five nodes and their interrelations, while the target problem consists of twenty nodes vs. SIX nodes.

The degree of structural overlap is captured in measures of graph similarity, as for example:

$$d(G, H) = 1 - \frac{V_{GH} + N_{GH}}{\max(V_G, V_H) + \max(N_G, N_H)} \quad (1)$$

The distance between two graphs G and H is defined as the number of common arcs (V_{GH}) and nodes (N_{GH}) in relation to the number of nodes and arcs of the larger graph. This simple measure captures all information relevant for characterizing relations between redistribution problems. To capture mappings between different concepts (as *heat/water*; Falkenhainer et al., 1989) or relations (as $+/-$; Anderson & Thompson, 1989), the definition can be relaxed to bijective mappings between (similar) node and/or arc labels (Schädler & Wysotzki, 1998).

In the following experiments, we investigate a small subset of the possible variations of types and degrees of partial source/target isomorphisms.

Experiments

Experiment 1

In a first experiment, we explored the suitability of our domain for studying transfer of non-isomorphic sources. That is, we investigated (1) whether subjects can solve water redistribution problems by analogical transfer (problems are neither too difficult, resulting in a failure even to solve isomorphical problems, nor too easy resulting in ignoring the source problem for generating a correct solution), and (2) whether at least partial isomorphs with a moderate degree of structural dissimilarity can be transferred successfully, and (3) whether superficial similarity has an influence on transfer success.

The problem given in figure 1 was used as source problem. We investigated two structural variations: one target problem which is isomorphic to the source (deviating only in the absolute numbers for initial, maximum and goal quantities) and one which is a source inclusive partially isomorph with a high degree of structural overlap (one additional jug and one additional solution step, see fig. 2b and table 2). Additionally, we varied superficial similarity between source and target by renaming jugs and switching their positions. All variations (see table 1) were realized between subjects.

Subjects Subjects were 60 pupils of a Berlin gymnasium, aged between fourteen and nineteen (average 17.4), 31 male and 29 female.

Procedure The experiment was fully computer based. The overall time of an experimental session was about 45 minutes.

After general instruction, subjects learned how to use the program and were introduced to the problem domain by solving an initial problem with tutorial guidance. Jugs were represented graphically (see fig. 1); redistributions $pour(x, y)$ were performed by clicking first on x and then on y ; impossible moves (from empty or into full jugs) were rejected; all operations could be redone and the subjects could cancel their

current solution and start again. The tutor-module intervened if the solution path was longer than four, if the problem was two times restarted without solving it correctly, or after two minutes. The program only proceeded if the problem was correctly solved twice without tutorial help.

Next, subjects were informed about the general principle for finding a shortest solution path (i.e. by thinking about the goal quantities in terms of relations to initial and maximum quantities). Note, that the problems were too difficult to simply apply the general concept given in this instruction. Afterwards, subjects received the source problem – which was isomorphic to the initial problem – and were given the hint that the initial problem was similar to the current problem and that referring to its solution might help solving the new problem. Similarity was pointed out by explicitly presenting the mapping relations between the jugs of the problems (Novick & Holyoak, 1991; Novick & Hmelo, 1994). The initial solution could be retrieved by mouse click. To proceed with the current problem this window had to be closed. Tutorial support was identical to that for the initial problem. The program proceeded after the source problem was correctly solved.

We introduced an initial problem before the source problem for three reasons: Solving the source problem should not be disturbed by difficulties in interacting with the program; the handling of the recall of a prior solution could be introduced; and the subjects were “primed” to use analogical reasoning as solution strategy (in contrast to solving the problem by trial and error or guided search in the problem space).

After solving the source problem, each subject received one of a set of five versions of the target problem (see table 1). All problems had different absolute numbers for jug capacities, start and goal quantities than the source problem. Again, the similarity to the last problem (source) was pointed out by explicitly presenting the mapping relations. The target problem had to be solved without tutorial help, only by referring to the solution of the source (by the same procedure as given above). Time was restricted to ten minutes.

Finally, subjects were asked to give the mapping between the source and target jugs.

Results and Discussion To make sure that we investigate analogical transfer, we restricted the criterium “successfully solved” to subjects, who produced the correct shortest target solution in a single trial and gave the correct mapping between source and target when questioned afterwards. The main results are given in table 1.

Table 1: Target problems and results for experiment 1

	Target Problems		Results		χ^2
	structure	surface	success	no success	
1	isomorphism	same names	8	4	
2	isomorphism	one renaming (A ↔ B)	10	2	1 vs 2 0.67
3	isomorphism	two ren.s (A → B, B → C, C → A)	8	4	1+2 vs 3 0.22
4	partial isom. (additional jug)	no renaming	8	4	1+2+3 vs 4 0.11
5	partial isom.	one renaming (A ↔ B)	3	9	1+2+3+4 vs 5 8.09*

Kimball's $k \times 2$ test, $df = 1$, $\alpha > 0.4$ for constrasts 1 vs 2, 1+2 vs 3, 1+2+3 vs 4, $\alpha = 0.005$ for contrast 1+2+3+4 vs 5 (Bonferroni adjusted $\alpha = 0.025$)

There is a significant relation between source/target similarity and solution success (5×2 contingency table with $\chi^2 = 9.58$, $df = 4$, $\alpha < 0.05$). The crucial (only significant, see χ^2 in table 1) difference is between conditions 4 and 5. That is, partial isomorphism is a sufficient condition for successful analogical transfer *if* the surface similarity of problems is high. If partial isomorphs differ in surface features (naming and positioning of jugs), it seems too difficult to transfer the source solution, even if the mapping is given explicitly. The results are in correspondence with the findings of (Reed et al., 1990) in the domain of algebra word problems: In their second experiment they could show that transfer success was low for source inclusive problems sharing a common domain with a target (cf. travel rate) and high for isomorphic problems even if they differ in their domain (cf. travel rate vs. interest rate).

Experiment 1 suggest the following consequences for our further investigation of partial source/target isomorphs: (1) to investigate the influence of *structural* similarity, superficial similarity between source and target should be as high as possible, and (2) to demonstrate that there exists a degree of structural dissimilarity where a source is no longer relevant for generating a target solution we have to construct source/target relations with less structural overlap than problem 4.

Experiment 2

In the second experiment, we used again the problem given in figure 1 as the source, that is, a three jug problem which can be solved by a minimal sequence of four *pour* operations. We investigated the following source/target relations:

- Target exhaustiveness: a target problem in which the last operation of the source solution is not needed (i.e. a three jug, three operations problem; problem 1 in table 2),
- Source inclusiveness (2): a target problem, in which an additional operation is needed (i.e. a three jug problem solvable with five operations; problem 2 in table 2),
- Three degrees of source/target overlap: target problems consisting of four jugs and are solvable with five operations
 - the partial isomorphic target problem used in experiment 1 (called problem 4 in exp. 1, see fig. 2b; problem D1 in table 2),
 - target problems with progressively decreasing structural overlap to the source (problems D2 and D3 in table 2).

To control the degree of structural overlap, we represented all problem structures as graphs – using an extended version of the representation given in figure 2 (additionally, the solution sequences are explicitly coded³). We calculated the distances between source and target problems using formula (1). Problems 1 and 2 differ from the source only in the number of necessary operations. This high degree of structural overlap is reflected in the low source/target distances of 0.16 and 0.17 (see table 3).

³The complete representations for all problems can be obtained from the authors.

Table 2: Target problems for experiment 2

jug:	small	medium-small	medium-large	large
Problem 1 (target exhaustive)				
	–	j_3	j_2	j_1
max	–	48	60	72
start	–	21	36	45
goal	–	0	33	69
solution:		$j_1 \rightarrow j_2, j_2 \rightarrow j_3, j_3 \rightarrow j_1$		
Problem 2 (source inclusive)				
	–	j_3	j_2	j_1
max	–	48	60	72
start	–	21	36	45
goal	–	33	60	9
solution:		$j_1 \rightarrow j_2, j_2 \rightarrow j_3, j_3 \rightarrow j_1, j_2 \rightarrow j_3, j_1 \rightarrow j_2$		
Problem D1 (high overlap), see problem 4 in exp. 1				
	j_4	j_3	j_2	j_1
max	16	20	25	31
start	3	8	15	18
goal	0	13	3	28
solution:		$j_1 \rightarrow j_2, j_2 \rightarrow j_3, j_3 \rightarrow j_1, j_2 \rightarrow j_3, j_4 \rightarrow j_2$		
Problem D2 (medium overlap)				
	j_4	j_3	j_2	j_1
max	16	20	25	31
start	6	9	15	18
goal	14	14	0	20
solution:		$j_1 \rightarrow j_2, j_2 \rightarrow j_3, j_1 \rightarrow j_4, j_3 \rightarrow j_1, j_2 \rightarrow j_3$		
or		$j_1 \rightarrow j_2, j_1 \rightarrow j_4, j_2 \rightarrow j_3, j_3 \rightarrow j_1, j_2 \rightarrow j_3$		
Problem D3 (less overlap)				
	j_4	j_3	j_2	j_1
max	17	20	25	31
start	7	9	15	18
goal	16	5	0	28
solution:		$j_3 \rightarrow j_4, j_1 \rightarrow j_2, j_2 \rightarrow j_3, j_3 \rightarrow j_1, j_2 \rightarrow j_3$		
or		$j_1 \rightarrow j_2, j_3 \rightarrow j_4, j_2 \rightarrow j_3, j_3 \rightarrow j_1, j_2 \rightarrow j_3$		

Problems D1, D2 and D3 were constructed by changing and/or introducing relations between *goal*, *max* and *initial* quantity. On the operational level the modifications result in an additional operator after the solution sequence of the source problem (problem D1), in the middle (problem D2) and at the beginning (problem D3). The values for the source/target distances decrease from 0.37 over 0.55 to 0.59 (see table 3). Note, that the absolute values for graph distances are to some extent dependent on the way in which the graph representation is realized. If all problems are transformed in the same way into graphs, the *relative* source/target distances reflect the varying degrees of structural overlap in an uniform way.

The experimental comparison of the realized degrees of structural overlap was not guided by a specific hypothesis. Our general assumption was, that there exists a degree of structural overlap between source and target (smaller than for the source/D1 relation) which is insufficient for analogical transfer.

Subjects Subjects were 70 pupils of a Berlin gymnasium, aged between 16 and 17 (average 16.3), 18 male, 52 female.

Procedure The procedure was identical to experiment 1. Subjects were presented with one of the five target problems.

Results and Discussion To make sure that we really investigate the impact of source/target relations on analogical transfer, we excluded all subjects from the analysis who did not give the correct mapping of jugs after solving the target.

Table 3: Results for experiment 2

	1	2	D1	D2	D3
	target	source	decreasing overlap		
	exhaus.	incl.			
distance	0.16	0.17	0.37	0.55	0.59
success	6	7	10	5	1
no success	1	1	3	4	11

The main results are given in table 3.

Interestingly, there are no performance differences between source inclusive and target exhaustive source/target pairs, if problems differ only in the number of operations (problems 1 and 2, exact binomial test, $\alpha = 0.601$). Both problems could be successfully solved by most of the subjects.

The finding of experiment 1, that a source inclusive partial isomorph can be successfully solved, could be replicated (problem 4 in exp. 1 and problem D1 in exp. 2). There is a significant relation between structural source/target similarity and solution success (conditions D1, D2 and D3; exact 3×2 test, $\alpha = 0.002$). The crucial difference is between conditions D2 and D3 (exact binomial tests: D1 vs. D2 $\alpha = 0.1003$, D2 vs. D3 $\alpha = 0.0004$, D1 vs. D3 $\alpha = 0.0001$).

Partial isomorphism exists between a lot of problems, although most of them might not be in a source/target relation usually considered in analogical problem solving: Even if two problems share only a single node, formally there exists a partial isomorphism between them. Our results suggest that there is a degree of source/target dissimilarity when the source can be no longer considered as relevant for solving the target. Note, that we are not discussing retrieval of a source – which is much more restricted by semantic similarity (Reed et al., 1990) – but analogical transfer.

For the given problem domain and representation of problem graphs the results show, that source inclusive partial isomorphs can be good candidates for analogical problem solving as long as the structurally identical part of the problems (i.e. the common subgraph) is greater than the structural differences. This is reflected by a distance smaller than 0.5 calculated with the similarity metric given in formula (1). We hope to continue our experiments on conditions for source/target relations for transfer success, exploring different problem domains (cf. algebra word problems or geometry proofs) and further structural source/target relations (cf. one-to-many mappings).

Modeling Analogical Transfer in IPAL

IPAL is a prototype for integrating problem solving and learning based on the machine learning approach of inductive program synthesis (Schmid & Wysotzki, 1998). IPAL is primarily intended as an AI application and not as a cognitive model. It deals with programming problems (such as sorting of lists) and blocksworld problems or puzzles (such as Tower of Hanoi) in a uniform way. Currently, IPAL receives a problem description (initial states, goal, operators) as input, generates problem solutions by planning and generalizes these solutions to cyclic macro-operations (Shell & Carbonell, 1989). Cyclic macros represent solution strategies for

problem classes; for example, experience with sorting lists of three numbers can be generalized to a recursive program for sorting lists of arbitrary length, experience with solving a Tower of Hanoi problem with three discs can be generalized to a solution strategy for n disc problems.

We plan to integrate an analogy module in IPAL as a possible way to circumvent macro-generation from scratch. Currently our analogy module operates stand-alone and we are using it to explore conditions for successful analogical transfer based on structural information *alone*. Of course, for many real-world applications it is necessary to consider semantical aspects of problems. But we want to develop a (generic) adaptation algorithm which works context-free as far as possible. Thus, our strategy is, to try to extract as much information as possible from structural source/target relations (Schmid et al., 1998).

The analogy module works in the following way: When a new problem is solved, we check whether there already exists a cyclic macro under which the new solution can be subsumed. If a macro (source) can generate a solution sequence which is isomorphic to the current solution (target), the macro is transferred to the target domain and two new knowledge structures are committed to memory – the macro for the target domain and a macro generalizing over source and target. If there is no isomorphic source/target relation, IPAL has to decide whether (1) to generate a re-representation of the target which might result in an isomorphic source/target relation (currently done on the basis of rewrite-rules provided by the user), (2) to try adaptation nevertheless, or, (3) switch to macro-generation from scratch. If source and target are structurally too dissimilar, analogical transfer might require more effort than inductive inference and additionally has the danger of generating inadequate or erroneous solutions. This is the reason why we are investigating structural criteria for successful analogical transfer.

In our psychological experiments, we investigated the transfer of *problem solutions*; in IPAL we want to employ analogical transfer on the level of macros, i.e. *problem solving strategies*. But the decision *whether* a source-macro can be transferred to the new domain is determined on the basis of problem solutions (the new problem solving trace and a trace generated from the candidate source macro). Thus, information about conditions for successful transfer of problem solutions can give us valuable design hints for IPAL. Up to now, our analogy module eagerly adapts each source to the target but generates more than 50% erroneous solutions for non-isomorphic source/target relations (see Schmid et al., 1998 for the adaptation algorithm and test results for a variety of source/target pairs). On the basis of the experimental results, we plan to run new trials, comparing IPAL's performance when adapting sources with more vs. less than fifty percent overlap. Also based on the experimental results, we will not prefer target exhaustiveness (which involves deleting of information from the source) to source inclusiveness (which involves inserting additional information). In our current implementation deletion is preferred. To introduce new information we currently rely on structural constraints given by the partial source solution only. Another possibility might be to introduce a mechanism of internal analogy (Hickman & Lovett, 1991).

Conclusions

We reported two experiments investigating the influence of different types and degrees of non-isomorphic source/target relations on transfer success. For the water-jug redistribution domain we could show that partial isomorphic sources can be transferred successfully if source and target do not differ in surface features. Furthermore, we could demonstrate that there exists a degree of structural overlap between source and target where the source is no longer helpful for constructing a solution.

Cognitive models and other systems using analogical reasoning techniques usually make no restrictions with respect to the structural overlap between problems when retrieving a source. Retrieval (or at least its first stage) is usually guided by feature-based (i.e. superficial) similarity alone. Thus, a source which shares only a very small isomorphic substructure with the target will be treated in exactly the same way as a source with a high structural overlap. In this case, transfer might result in meaningless inferences or erroneous solutions. Of course, problems which share a greater amount of similar attributes might often also share a greater structural overlap, but this must not be true in general. For example, in the domains we are exploring with IPAL, there might be the following source candidates for solving the blockworld problem "build a tower of alphabetically ordered blocks": (1) another blockworld problem – "unstack a tower of alphabetically ordered blocks" – which shares a lot of attributes with the target, and (2) a list sorting problem which shares no attributes but can be solved with the same underlying strategy. We would prefer to retrieve the structurally rather than the superficially similar source problem. That is, we propose to consider not only attribute-based but also structural similarity between problems to minimize the risk of erroneous transfer.

Acknowledgements

This manuscript was written while the first author was on leave at the School of Computer Science, Carnegie Mellon University, supported by a DFG research scholarship. Thanks to Klaus Eyferth, Bruce Burns, Alex Petrov and Dario Salvucci for helpful discussions and again to Bruce Burns for comments on the manuscript.

References

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

Atwood, M. E., & Polson, P. G. (1976). A process model for water jug problems. *Cognitive Psychology*, 8, 191-216.

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, CA: Morgan Kaufmann Pub.

Falkenhainer, B., Forbus, K., & Gentner, D. (1989). The structure mapping engine: Algorithm and example. *Artificial Intelligence*, 41, 1-63.

Gentner, D. (1980). *The structure of analogical models in science*. Cambridge, MA: Bolt Beranek and Newman Inc.

Gholson, B., Smither, D., Buhrman, A., Duncan, M. K., & Pierce, K. (1996). The sources of children's reasoning errors during analogical problem solving. *Applied Cognitive Psychology, Special Issue: "Reasoning Processes"*, 10, 85-97.

Hickman, A. K., & Lovett, M. C. (1991). Partial match and search control via internal analogy. In *Proceedings of the 13th Annual Conference of the Cognitive Science Society* (p. 744-749). Hillsdale, NJ: Lawrence Erlbaum Ass.

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

Hummel, J., & Holyoak, K. (1997). Distributed representation of structure: A theory of analogical access and mapping. *Psychological Review*, 104(3), 427-466.

Luchins, A., & Luchins, E. (1950). New experimental attempts at preventing mechanization in problem solving. *Journal of General Psychology*, 42, 279-297.

Markman, A., & Sanchez, A. (1998). Structure and pragmatics in analogical inference. In *Advances in analogy research: Integration of theory and data from the cognitive, computational, and neural sciences* (p. 191-200). Sofia, Bulgaria.

Novick, L. R., & Hmelo, C. E. (1994). Transferring symbolic representations across nonisomorphic problems. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20(6), 1296-1321.

Novick, L. R., & Holyoak, K. J. (1991). Mathematical problem solving by analogy. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 17(3), 398-415.

Pirolli, P., & Anderson, J. (1985). The role of learning from examples in the acquisition of recursive programming skills. *Canadian Journal of Psychology*, 39, 240-272.

Reed, S. K., Ackinclose, C. C., & Voss, A. A. (1990). Selecting analogous problems: similarity versus inclusiveness. *Memory & Cognition*, 18(1), 83-98.

Schädler, K., & Wysotzki, F. (1998). Application of a neural net in classification and knowledge discovery. In *Neural Networks in Applications NN'98, Proceedings of the Third International Workshop* (p. 219-226). Magdeburg.

Schmid, U., Mercy, R., & Wysotzki, F. (1998). Programming by analogy: Retrieval, mapping, adaptation and generalization of recursive program schemes. In *Proc. of the Annual Meeting of the GI Machine Learning Group, FGML-98* (pp. 140-147). TU Berlin.

Schmid, U., & Wysotzki, F. (1998). Induction of recursive program schemes. In *Proceedings of the 10th European Conference on Machine Learning (ECML-98)* (p. 228-240). Springer.

Shell, P., & Carbonell, J. (1989). Towards a general framework for composing disjunctive and iterative macro-operators. In *Proceedings of the 11th IJCAI-89*. Detroit, MI.

Spellman, B. A., & Holyoak, K. (1996). Pragmatics in analogical mapping. *Cognitive Psychology*, 31, 307-346.