

Adaptation as a Selection Constraint On Analogical Mapping

Mark T. Keane

Dept. of Computer Science,
University of Dublin, Trinity College,
Dublin 2, IRELAND
mark.keane@cs.tcd.ie

Abstract

In any given analogy, there are potentially a large number of possible mapping interpretations. One of the key issues in analogy research is how one of these mappings comes to be selected as optimal and used as the basis for the analogical comparison. It is well-established that structural factors, notably systematicity, can act as selection constraints on mapping. The present work tests to see if pragmatic and adaptation factors can also act as selection constraints on mapping. The selection of a mapping based on pragmatic factors proposes that people can exploit the higher-order, schematic structure of a domain to select one mapping over another. With respect to adaptation factors, the proposal is that a mapping will be selected if it is evaluated as being easily adapted relative to other competing mappings. Both of these predictions are tested in a novel, problem solving paradigm. The main finding is that adaptation factors do act as a selection constraint but that pragmatic factors do not. The implications of these results for computational models of analogy are discussed.

Introduction

Human thought excels in its capacity to deal with novelty and to overcome the unexpected and unfamiliar. In recent years, it has been demonstrated that the use of analogy is one strategy for dealing with novelty; novel problems can be solved by analogy to past experiences (see e.g., Gentner, 1983; Gick & Holyoak, 1980; Keane, 1988). Five main processes have been demonstrated to be important in analogical problem solving: representation, analogue retrieval, analogical mapping, adaptation and induction. Clearly, the way an analogue and a problem are mentally *represented* has an important impact on subsequent problem solving attempts by analogy (see e.g., Novick, 1988). The *retrieval* of an analogous case to solve a problem is harder if the domains are semantically-distant to one another (e.g., Gentner & Landers, 1985; Keane, 1987). *Analogical mapping* is the core process during which the analogy is drawn between the two domains. After an analogy is mapped it may have to be *adapted* in order to make the solution fit the target problem (see Novick & Holyoak, 1991; and the Artificial Intelligence literature on case-based reasoning, e.g., Kolodner, 1993; Smyth & Keane, 1993). With respect to *induction*, Gick & Holyoak (1983) have shown that subjects can induce a generalisation from the correspondences between two

analogous stories which facilitates subsequent problem solving.

In this study, we will be concerned with various factors involved in the mapping and adaptation stages of analogising. In *analogical mapping* people must map the conceptual structure of one domain (the *base* or *source domain*) to another domain (the *target domain*). During mapping, subjects *match* corresponding parts of the two domains. For example, in the water-flow/heat-flow analogy people might match the "the flow of water via a pipe between a beaker and a vial" with "the flow of heat via a bar between coffee and an ice-cube" (see Figure 1). Typically, the mapping will also involve the *transfer* or *carry-over* of knowledge from the base domain to the target; knowledge which constitutes a set of *candidate inferences* suggested by the analogy. In the present example, the causal relation between the pressure-difference and flow in the water-flow domain would be transferred as a candidate inference into the heat-flow domain to govern temperature-difference and heat-flow there.

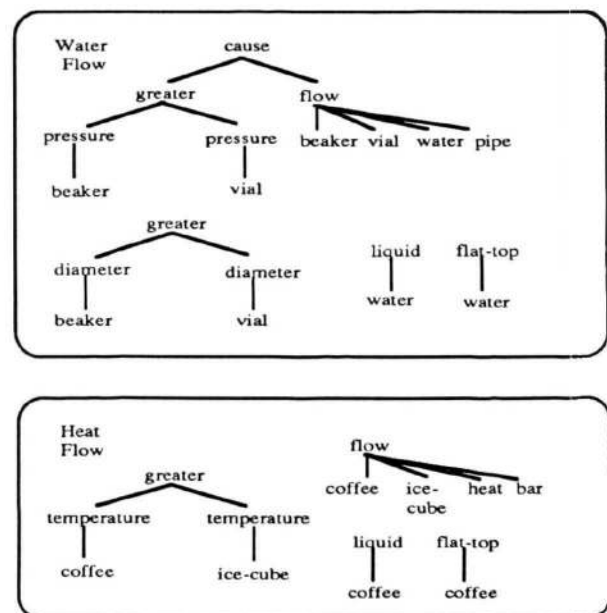


Figure 1 The Water-flow/Heat-Flow Analogy
(from Falkenhainer, et al., 1989)

No candidate inference by analogy is guaranteed to be true, it has to be verified and may have to be adapted to fit the target situation. Intuitively, the above causal, inference seems reasonable but experimentation could reveal that the dependency is only one of enablement, in which case the inference would have to be adapted (see Falkenhainer, 1987). Novick & Holyoak (1991) have demonstrated that adaptation occurs in solving mathematics problems by analogy to stories. But, very little is known about the nature of adaptation and how it influences the course of analogical problem solving.

The Theoretical Issue

One of the basic issues in analogical mapping is how one from a number of alternative mapping-interpretations is chosen for a given analogical comparison. For example, one mapping-interpretation for the above analogy involves the system of relations involving the causal dependency between pressure-difference/temperature-difference and flowing, another involves the mapping between the diameter-difference of the beaker and the vial and the temperature-difference between the coffee and ice-cube. People prefer and select the former mapping-interpretation rather than the latter. Why?

Gentner (1983, 1989) provides an account to explain the selection of this mapping in her structure-mapping theory and its associated model, the Structure-Mapping Engine (SME; Falkenhainer, Forbus & Gentner, 1986, 1989). Her account proposes that a structural factor -- systematicity -- acts as a selection constraint on alternative mappings. Stated simply, if one mapping involves a highly-interconnected set of matches (or is more systematic) it will be preferred over a mapping with less connectivity between its matches. This proposal has been supported repeatedly in the empirical literature (see e.g., Clement & Gentner, 1991; Gentner & Toupin, 1986; Keane, 1988). For example, Clement & Gentner (1991) have shown that when subjects were given a judgement task involving two systems of relations, they preferred a match embedded in a shared system of relations to one that was embedded in a different system of relations. Similar effects have been shown for the selection of candidate inferences. Systematicity is also the backbone of all computational models of analogy; including SME, the Analogical Constraint Mapping Engine (ACME; Holyoak & Thagard, 1989) and the Incremental Analogy Machine (IAM; Keane & Braysshaw, 1988; Keane et al., in press).

The present paper examines the previously-untested, role of two other factors on the selection of analogical mappings; pragmatic and adaptation factors. In the remainder of this paper, we first consider predictions based on these factors, then present an experiment testing these predictions, before considering the implications of the results found.

More Selection Constraints on Mapping

Consider another analogy, this time one involving Maier's (1931) "Two-String" problem and an analogous story about a fire (see Keane, 1988). The Two-String problem says:

Imagine you are in a room in which two strings are hanging from the ceiling. Your task is to tie the ends of the two strings together. However, when you take hold of one string and try to get the other string you find that it is too far away to touch (a diagram illustrating the problem was provided).

Traditionally, subjects propose a variety of solutions to this problem including a *stick* solution (use a stick to bring the string, that is furthest away, within reach) and a *swing* solution (swing one of the strings so that it comes within reach while holding the other string). Now, consider an analogous story made up of the setting information and plans 1 and 2 in Table 1 (where the text would indicate that one of the plans failed and the other succeeded). If we wish to solve the Two-String problem by analogy to this story, we are faced with two possible mappings either of which can be used to solve the problem. One mapping suggests the analogous solution of using a stick to reach the string (using the first plan) and the other mapping suggests the analogous solution of swinging one of the strings in order to bring it within reach (using the second plan).

Table 1 Summary of the Stories Used in the Experiment

Setting

People are stuck in a towering inferno and are trying to evacuate the building.

Plan 1:

A helicopter is positioned and a winch is thrown to the people, they grab hold of it and tie onto the winch, then swing from one building to the next.

Plan 2:

A helicopter is positioned and a winch is lowered to people, they reach for the winch, with a stick, tie on and are carried to the ground.

Plan 3:

Helicopter is positioned and a winch is lowered to people, they lasso the winch with a rope, tie on and are carried to the ground.

Which of these alternative mappings are people likely to select and on what is this selection likely to be based? Apart from the proven role of systematicity (which is held

constant for both plans by using equal numbers of actions), two other possibilities present themselves:

- pragmatic factors which assess which plan is likely to be the best, on the basis that "what was successful before is likely to be successful again", without further analysis of the plan's content
- adaptation factors which assess the specific content of the plan to see whether the solution it suggests is a good solution to the problem

A variety of different proposals have been made under the general heading of pragmatic factors (for some possibilities see Holyoak, 1985; Holyoak & Thagard, 1989; Keane, 1985, 1988; Keane et al, in press). The above proposal is just one specific pragmatic prediction made by Keane & Brayshaw (1988). This prediction hinges on the higher-order, schematic structure of the story; which can be cast as a "setting-->failed-plan-->successful plan" structure. Pragmatically, it might be a good idea to select the plan which succeeds on the assumption that, given an analogy between the story and problem, what succeeded before may succeed again. In an early version of IAM, Keane & Brayshaw proposed that the successful plan could be chosen in a top-down fashion using these schematic categories and then the contents of this plan would be mapped. ACME also has a mechanism for doing this sort of mapping; it can group parts of an analogue into "fields" (e.g., a SETTING field). The mapping is thus simplified because mappings are only established between information grouped in the same fields (see Holyoak & Thagard, 1989). In conclusion, a strong pragmatic prediction would be that the successful plan should always be selected to be mapped.

On the issue of adaptation, the key proposal made here is that adaptation enters into the selection of alternative mapping-interpretations. Specifically, if one analogous plan is evaluated as being more easily adapted to the problem than another, then it will be selected as the optimal mapping. In the Fire story, the stick plan is less adaptable than the swing plan. The stick plan proposes that a stick be used to reach one of the strings, a stick which is not present in the problem statement. Hence, the analogy suggests various actions of reaching and grabbing but also the creation of a new object in the Two-String problem. The swing plan simply suggests the new actions of swinging and grabbing, without the need to introduce a new object. As such, the analogous swing solution is easier to adapt. It may also be easier to adapt because, unlike the stick plan, it does not violate implicit constraints in the problem situation (i.e., that the problem is to be solved with the objects given).

If adaptation factors play a role in the selection of one mapping over another, then the more adaptable, analogous plan should tend to be selected when people are presented with alternatives.

Testing Pragmatic and Adaptation Factors

In this experiment, subjects were given a story to memorise and were then asked to use this story to solve a problem by analogy. We used the Two-String problem and four variants of the Fire story. The Fire story always had the same general structure involving setting information, an account of a failed rescue plan and one of successful rescue plan. In all versions of the story the setting information was the same, but the plans used were varied systematically. The setting information in the Fire story and the Two-String problem will be matched in forming the analogy; for instance, both domains involve propositions about getting from one location to another in order to achieve some end. The two plans constitute collections of candidate inferences (organized in a causal chain) which follow from matched propositions in the settings of both domains. So both plans are alternative mappings, one of which should be selected as the analogous solution to the problem.

Since each story has a failed and successful plan we can test the pragmatic predictions. If the pragmatic predictions are correct, overall, we should see more subjects using the successful plan rather than the failed plan. To control for primacy or recency effects between the story plans, half of the stories had a "setting --> successful plan --> failed plan" structure, while the other half had a "setting --> failed plan --> successful plan" structure.

Adaptability was manipulated by varying the relative-adaptability of one of the analogous-plans. The focus was on the stick plan which was always the successful plan in different versions of the story. In the low-adaptable versions of the story, the stick plan was paired with the swing plan to make it low-adaptable relative to the high-adaptable, swing plan. In the high-adaptable story-versions, the stick plan was deemed to be high-adaptable relative to a low-adaptable, lasso plan. The *lasso plan* is similar to the stick plan, with the exception that the rope is reached by lassoing it with another piece of rope. Intuitively, although this plan is very similar to the stick plan, it is a less adaptable solution to the problem because a rigid object (e.g., the stick) should be more effective than a non-rigid object (e.g., the lasso). If adaptation is a selection constraint then subjects should tend to prefer the stick plan when it is the high-adaptable plan relative to its paired-plan, than when it is low-adaptable relative to the paired-plan.

Finally, in order to substantiate our intuitions about the relative adaptability of plans, we asked subjects to rate how good the swing, stick and lasso solutions were as solutions to the Two-String problem. If our intuitions are correct, then the swing solution should be rated as being better as a solution than the stick solution and the stick solution should be rated above the lasso solution.

Method

Subjects. Sixty-three undergraduates in Computer Science at Trinity College Dublin took part in the experiment and were randomly assigned one of the story versions. Four of these subjects were excluded before data analysis because they had insufficient memory of the plans in the story.

Materials. The materials for the problem solving task consisted of four versions of the Fire story and a version of the Two-String problem.

The materials for the ratings task consisted of a booklet containing descriptions of the swing, stick and lasso solutions to the Two-String problem. Under each solution was a seven-point scale, ranging from 1 (very bad) to 7 (very good). Written instructions informed subjects that they were to rate how good they thought the solution was, as a solution to the problem. Each solution was printed on a separate sheet. The order of presentation of solutions in the rating task was randomised.

Procedure. The experiment had three distinct stages: the memory/recall task, the problem solving task and the ratings task. All the booklets handed out in each task were collected before those of the next stage were handed out. The memory/comprehension task involved subjects taking three minutes to read and memorise the Fire story. They were told that after the three minutes they would have to recall it in writing. Each subject in the experiment received one of the four versions of the Fire story. After subjects had recalled the story in writing and it was taken away, they were given the Two-String problem to read. They were asked "If you were to apply the story about the fire to this problem what answer would it suggest to you (irrespective of whether you think the answer is a good or a bad one)". We wanted a pure measure of the analogous plan used, so the caveat in parentheses was added to make sure that subjects did not use a solution by analogy and then reject it because they felt it was a poor solution. Finally, in the ratings task, subjects were asked to rate the goodness of a solution plan as a solution to the problem.

Subjects were tested in small groups of varying sizes and each experimental session took between 20 and 30 minutes.

Scoring. In analysing subjects' responses the interest was in the solution that was most immediately suggested to them by the story. Therefore, when subjects produced more than one solution, only the first solution produced was considered. In fact, 80% of subjects produced only one solution (20% produced two solutions). Subjects' solutions were categorised as *stick*, *swing*, *lasso solutions* or *other solutions* (which were based on none of the plans).

		Pragmatic		
		Successful	Failed	
Adaptability	Low	(Stick) 8	(Lasso) 3	11 (23%)
	High	(Stick) 20	(Swing) 17	37 (77%)
		28 (58%)	20 (42%)	

Figure 2 Frequency of Subjects' Solutions Classified According to the Variables Examined in the Experiment

Results

The results revealed the effects of adaptation factors but showed little evidence for pragmatic factors (see Figure 2). The pattern of goodness ratings confirmed our intuitive ranking of the adaptability of different plans.

Of the 59 subjects in the experiment, marginally more subjects used the successful plan (47%; 28 subjects) over the failed plan (34%, 20 subjects) while 19% (11) used neither plan [$\chi^2(2) = 7.34, p < .05$]. But, if we look at the 48 subjects who used one or other of the plans, we find that there is no reliable difference between subjects who selected the successful plan (58%, 28 out of 48), relative to those who selected the failed plan (42%, 20 out of 48; $\chi^2(1) = 1.2, p > .05$). Notable effects for adaptation factors were found. In low-adaptable versions of the story, subjects used the successful, stick plan less frequently (23%), than in the high-adaptable versions of the story (77%; $\chi^2(1) = 14.08; p < .001$). These results indicate that the adaptability of a plan has a strong effect on the selection of one mapping from amongst a number of alternatives.

Overall the goodness ratings reflected our intuitions about the relative adaptability of the solution plans. A one-way, repeated-measures analysis of variance of the ratings show a main effect of solution type $F(2, 116) = 50.07, p < .0001$. Planned comparisons, using Newman-Keuls tests, between solution types revealed that the swing solution ($m = 5.4$) is rated as being reliably different to the goodness of the stick solution ($m = 4.46, p < .01$) which in turn is rated more highly than the goodness of the lasso solution ($m = 2.71, p < .01$).

General Discussion

Two specific results have been uncovered by the present study which have important implications for computational models of analogy. First, it has been found that the

supposed pragmatic influence of exploiting higher-order categories in a domain (e.g., successful plan, failed plan) in order to select a mapping, has not been confirmed. Second, for the first time, we have shown that the ease of adaptation of a mapping can affect subjects' selection of a mapping in a crucial fashion.

With respect to pragmatic effects, ACME and IAM have the functionality to deal with higher-order structurings of a domain and to exploit this in drawing analogies. Holyoak & Thagard (1989) designed ACME with a facility for partitioning domains into different fields, and then just mapping the contents of like fields with like fields. In ACME the Fire story could be encoded as having the following fields: SETTING, SUCCESSFUL-PLAN, FAILED-PLAN. ACME has used such higher-order structurings in the mapping of other examples in the literature. However, there is no specific and unique empirical evidence to support the proposal that people exploit such higher-order categories in their analogising. Indeed, the present experiment provides negative evidence for the use of such categories. (It should also be noted in passing that ACME is really a matcher and cannot form the candidate inferences required by this task. It cannot form extended sets of candidate inferences without some additional mechanism or the use of implausible, dummy predicates for each inference in the target). Similar criticisms can be made of IAM. Early versions of this model had the functionality to deal with such higher-order categories like SUCCESSFUL-PLAN and FAILED-PLAN (see Keane & Brayshaw, 1988). Even if people form such categories, the evidence suggests that they do not exploit them when solving problems by analogy.

On adaptation, the present results suggest that any complete computational level account of analogical thinking should include *adaptation constraints* (see Keane et al., in press); constraints that evaluate the adaptability of alternative mappings. This evaluation could rely purely on local aspects of the candidate inferences being made: for instance, by noting whether the inferences suggest new objects which are absent from the original specification of the target-problem domain. However, such a local evaluation does not seem sufficient. It would result in evaluations which rated the swing solution as being better than the stick solution (because the latter suggests a new object) but would also evaluate the swing and lasso solutions as being equally adaptable (because both use string-like objects which are present in the problem). It seems more plausible that evaluation is performed by simulating the proposed solution in some way; the action of swinging a rope-like object is tested in the target problem or the action of reaching something with a stick is tested. Only a detailed simulation of this type will produce evaluations that correspond to the adaptability ordering supported by the present experiment (i.e., swing better than stick, and in turn stick better than lasso). Techniques for such evaluation have been widely explored in the area of analogy and case-based reasoning in

Artificial Intelligence (see e.g., Falkenhainer, 1987; Hammond, 1989; Kolodner, 1993; Smyth & Keane, 1993). Indeed, this study has interesting implications for case-based reasoning (see Keane, in press).

References

- Clement, C.A., & Gentner, D. (1991). Systematicity as a selection constraint in analogical mapping. *Cognitive Science*, 15, 89-132.
- Falkenhainer, B. (1987). An examination of the third stage in the analogy process: Verification-based analogical learning. In *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*. Los Altos: Morgan Kaufmann.
- Falkenhainer, B., Forbus, K.D., & Gentner, D. (1986). Structure-mapping engine. *Proceedings of the Annual Conference of the American Association for Artificial Intelligence*.
- Falkenhainer, B., Forbus, K.D., & Gentner, D. (1989). Structure-mapping engine. *Artificial Intelligence*, 41, 1-63.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7, 155-170.
- Gentner, D. (1989). Mechanisms of analogical learning. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning*. Cambridge: CUP.
- Gentner, D. & Landers, R. (1985). Analogical reminding: A good match is hard to find. *Proceedings of the International Conference on Systems, Man and Cybernetics*, Tucson, Arizona, November.
- Gentner, D. & Toupin, C. (1986). Systematicity and surface similarity in the development of analogy. *Cognitive Science*, 10, 227-300.
- Gick, M.L., & Holyoak, K.J. (1980). Analogical problem solving. *Cognitive Psychology*, 12, 306-355.
- Gick, M.L., & Holyoak, K.J. (1983). Schema induction in analogical transfer. *Cognitive Psychology*, 15, 1-38.
- Hammond, K. J. (1989). *Planning from memory*. New York: Academic Press.
- Holyoak, K.J. (1985). The pragmatics of analogical transfer. *The Psychology of Learning and Motivation*, 19, 59-87.
- Holyoak, K.J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, 13, 295-355.
- Keane, M. (1985). On drawing analogies when solving problems. *British Journal of Psychology*, 76, 449-458.
- Keane, M.T. (1987). On retrieving analogues when solving problems. *Quarterly Journal of Experimental Psychology*, 39A, 29-41.
- Keane, M.T. (1988). *Analogical problem solving*. Chichester: Ellis Horwood.
- Keane, M.T. (in press). Analogical asides on case-based reasoning. In K-D. Althoff, K. Richter, & S. Wess

- (Eds.), *First European Workshop on Case-Based Reasoning*. Amsterdam: Springer-Verlag.
- Keane, M.T., & Brayshaw, M. (1988). The Incremental Analogical Machine: A computational model of analogy. In D. Sleeman (Ed.), *European Working Session on Machine Learning*. London: Pitman.
- Keane, M.T., Ledgeway, T., & Duff, S. (in press). Constraints on analogical mapping: A comparison of three models. *Cognitive Science*, **18**.
- Kolodner, J. (1993). *Case-based reasoning*. Los Altos: Morgan Kaufmann.
- Maier, N.R.F. (1931). Reasoning in humans II: The solution of a problem and its appearance in consciousness. *Journal of Comparative Psychology*, *12*, 181-194.
- Novick, L.R. (1988). Analogical transfer, problem similarity, and expertise. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *14*, 510-520.
- Novick, L.R. & Holyoak, K.J. (1991). Mathematical problem solving by analogy. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *17*, 398-415.
- Smyth, B., & Keane, M.T. (in press). Retrieving adaptable cases. In K-D. Althoff, K. Richter, & S. Wess (Eds.), *First European Workshop on Case-Based Reasoning*. Amsterdam: Springer-Verlag.