Transfer of Training as Analogical Mapping

ROBERT M. SCHUMACHER AND DEDRE GENTNER

Abstract—**Similarity is an important factor in transfer of training, but its precise role is not well understood. Is it overall fidelity or do some kinds of commonalities that matter more than others? It is proposed that the role of similarity in transfer can be clarified by a finer grained analysis. In this research, subjects learned a procedure for operating a simulated device and were asked to transfer the knowledge of the procedure to a new device. Two factors were varied: (1) the *systematicity* of the original device model (i.e. whether the subjects were given a coherent causal model or simply a set of operating procedures); and (2) the transparency of the mapping or degree of surface similarity between corresponding device components. The dependent measure was the number of trials to a criterion in the original and transfer devices. Results showed effects of both systematicity and transparency. Having a systematic mental model both facilitated learning of the initial device and promoted transfer to the target device. Transparency had strong effects on transfer: subjects learned the new device faster when corresponding pairs of components were similar than when noncorresponding pairs were similar. These results suggest that there are at least two separate factors to consider in transfer: the systematicity of the common domain model and the transparency of corresponding components.**

INTRODUCTION

In THIS RESEARCH we examine the determinants of transfer of training of a device model. It is widely accepted that similarity is a key determinant of transfer of training. But recent research in cognitive science indicates that the role of similarity in transfer may be complex. Novick [1] points out that different studies have led to different conclusions. Some studies indicate that similarity promotes positive transfer [2]; some that similarity often fails to create transfer (e.g., [3]–[5]); and some that similarity leads to negative transfer [2]. We believe these seemingly contradictory findings can be resolved by a more fine-grained account of similarity. In particular, we propose to analyze the transfer situation from the viewpoint of analogical mapping.

STRUCTURE-MAPPING THEORY

The structure-mapping theory of analogy [6]–[8] provides a framework for this analysis of transfer of training. In structure-mapping, analogy is defined as a mapping of knowledge from a familiar domain (the base) into another, usually less familiar, domain (the target). This mapping is constrained by a set of rules.[1] Objects in the base are placed in one-to-one correspondence with objects in the target:

$$M: b_i \rightarrow t_i.$$

Predicates are mapped from the base to the target according to the following mapping rules:

1) Attributes of objects are dropped: e.g., RED $(b_i) \nrightarrow$ [RED (t_i)].
2) Certain relations between objects in the base are mapped across: (e.g., COLLIDE $(b_i, b_j) \rightarrow$ COLLIDE (t_i, t_j)).
3) The particular relations mapped are determined by the *systematicity principle,* which states that a base predicate that belongs to a mappable system of mutually constraining interconnected relations is more likely to be imported to the target domain than is an isolated predicate. For example,

$$\text{CAUSE}\left[\text{PUSH}\left(b_i, b_j\right), \text{COLLIDE}\left(b_j, b_k\right)\right] \rightarrow \text{CAUSE}\left[\text{PUSH}\left(t_i, t_j\right), \text{COLLIDE}\left(t_j, t_k\right)\right].$$

Thus, in analogy the object correspondences are determined not by any intrinsic similarity in the objects themselves, but by their roles in the matching relational structures.[2]

The point of analogical mapping is to maximize overlap in relational structure: Fig. 1 shows an example analogy: the Rutherford analogy between the solar system and the hydrogen atom. To understand this analogy, a person must find the one-to-one correspondence between the objects of the solar system and the objects of the atom that gives a maximally systematic predicate match. Here, the most systematic set of matching predicates is CAUSE [MORE-MASSIVE-THAN (sun, planet), REVOLVE-AROUND (planet, sun)]. Thus the best object correspondences are sun $\rightarrow$ nucleus and planet $\rightarrow$ electron. Base relations belonging to the system, such as MORE-MASSIVE-THAN (sun, planet), are preserved; isolated relations, such as HOTTER-THAN (sun, planet) are discarded; and object attributes, like YELLOW (sun) and MASSIVE (sun) are dropped.

Manuscript received June 2, 1987; revised September 11, 1987. This work supported by the Office of Naval Research under contract no. N00014-85-K-0559, NR667-551. Portions of this work were presented at the Panel on Mental Models at the IEEE International Conference on Systems, Man, and Cybernetics, 1986, Atlanta, GA.

The authors are with the Department of Psychology, 603 E. Daniel St., University of Illinois, Champaign, IL 61820.

IEEE Log Number 8821880.

[1] We give a brief summary of structure-mapping theory here. A fuller description of the theory is given in [7]. For a computer-simulated process model, see [9].

[2] This distinguishes analogy from literal similarity. In literal similarity intrinsic object similarity is important as well as common relational structure.

0018-9472/88/0700-0592$01.00 ©1988 IEEE

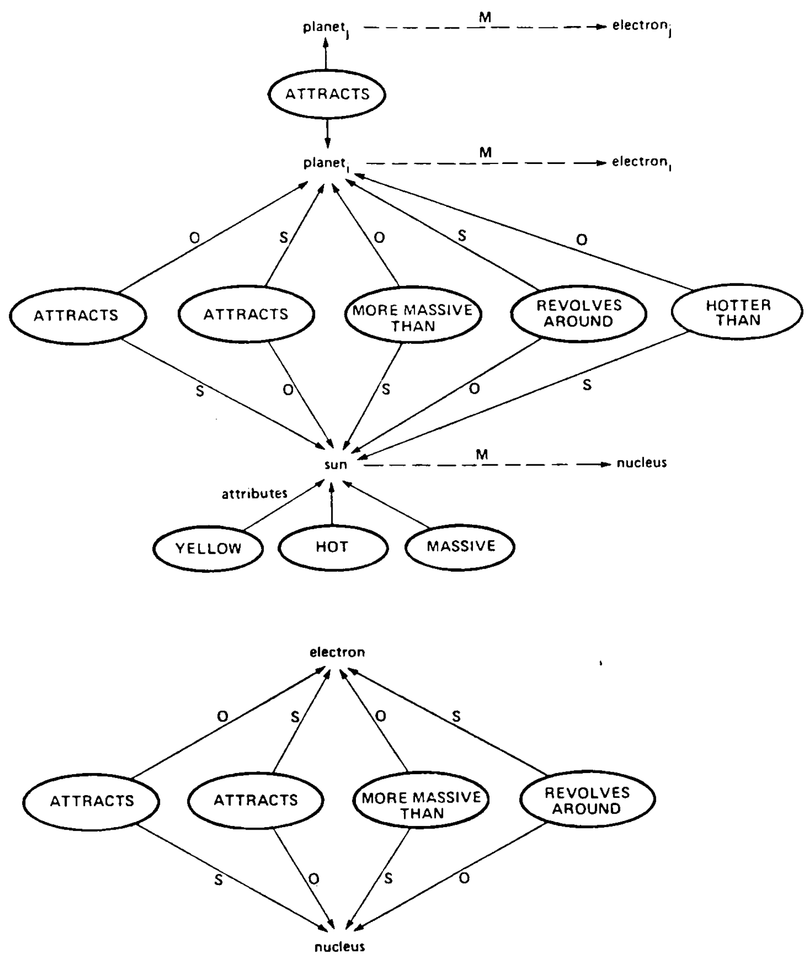

Fig. 1. Structure-mapping for Rutherford analogy: "The atom is like the solar system."

This kind of systematic match allows new predictions; predicates belonging to the base system but not present in the target can be mapped across as candidate inferences.

So far our focus has been on structure-mapping as a delineation of the essential distinctions that define an ideal analogical mapping. That is, we have been considering what Marr [10] called the computational level and what Palmer and Kimchi called the informational level of description [11], [12]. We turn now to considerations of the algorithmic or behavioral constraints. That is, we seek to go beyond the competence model to a performance model of how people actually carry out analogies under real-time constraints. Based on the discussion above, at least two separate factors should contribute to the ease of analogical transfer: the *systematicity* of the base model and the *transparency* of the correct object correspondences. We discuss each of these in turn.

Systematicity and Transfer

A good analogy conveys a coherent system of connected knowledge, not simply an assortment of independent facts. The systematicity principle reflects people's tacit preference for coherence and deductive power in analogy.[3] Because systematicity is a central notion here, we discuss it briefly before going on. A knowledge structure is systematic if it has higher-order relations that constrain its lower-order predicates. Thus when we speak of a *systematic domain model* we mean one in which there exist higher-order constraining relations (such as CAUSE or IMPLIES). When we speak of a preference for *systematicity* in analogy we mean that among the sets of common predicates that could serve as the interpretation of an analogy people prefer to choose sets with systematic structure.

Empirical studies have shown that adults focus on systematic relational structure in interpreting and evaluating analogy [8], [13], [4]. For example, given the analogy "A cloud is like a sponge" adults produce relational interpre-

[3]Note that the systematicity principle requires that the relational chain be *mappable* from the base to the target. Thus, a relational chain, such as a causal chain, in the base that matches a relational chain in the target constitutes support for including its subordinate members in the mapping.

tations like "Both can absorb water and hold it until some later time" over interpretations like "Both are round and fluffy." There is also evidence that systematicity may play an active role in guiding the mapping process [14]. Gentner and Toupin [14] had children and adults learn and then recreate story scenarios. They found that if the story possessed systematic causal or moral structure, subjects were more accurate in recreating the story using different characters. They suggested that the presence of higher-order relations may help constrain and guide the mapping of the lower-order plot events.

Related notions have been proposed in human factors and cognitive engineering. Norman [15], [16] argues that the learner's or user's mental model of a system should be coherent with respect to the system model to enable adequate performance. Norman's notion of coherence, as applied to an individual's mental model, has much in common with the notion of systematicity used here, in fact it seems to refer to a representation that is well-structured and internally consistent. A coherent mental model of a device should support prediction, explanation and diagnosis of system states over a less coherent model [17], [18]; by extension, it should also support transfer.

Based on this line of reasoning, we predict that if the user possesses a systematic model of the base domain this should increase the transfer accuracy between isomorphic devices.

Transparency and Transfer

Another factor that should be important in the mapping process is the *transparency* of the object correspondences. Transparency is defined as the degree of object-level similarity between corresponding component objects. For example, to the degree that base and target gauges that perform the same function look similar, transparency is high; to the degree that they look different (while maintaining the same function), transparency is low. As we have discussed, for an ideal analogizer, transparency should be irrelevant to determining the object correspondences. But in actual practice, surface similarities between corresponding objects may influence people's performance in achieving a correct analogical match. To take a very simple example, it may be easier to solve 2:4::20:_(40) than to solve 2:4::10:_(20). Even though both examples are equally valid analogies, it may be easier to match $4 \rightarrow 40$ then to match $4 \rightarrow 20$. Another reason the first analogy may be easier than the second is that the second invites a spurious match from $2 \rightarrow 20$—a *crossmapping* that may interfere with the correct mapping. That is, to the degree that one can easily determine how the objects in the base correspond with the objects in the target, the transfer of the predicate structure from base to target should be easier. Thus, we predict that transparency will have a strong effect on transfer accuracy.

The experiment reported here used a computer-simulated device. The device consisted of a set of interrelated gauges indicating system parameters such as the speed and engine temperature of a ship. Subjects learned how to operate the device and then transferred their operational procedures to a second device. In each case, we measured the number of trials subjects needed to reach a set criterion. In order to test whether having a systematic model of the original device would improve subjects' ability to transfer training to another device, we provided subjects with either a systematic or a nonsystematic model of how to operate the original device as discussed below. We predicted that the systematic model would enable subjects to better perform the transfer. A secondary prediction was that subjects with a systematic model would show faster learning of the original system [17].

The second factor varied was the transparency of the object correspondences between the base and target. We used two levels of transparency: 1) *high transparency:* the target components looked very similar to corresponding base components (e.g. the speed gauge in the target resembled the speed gauge in the base); and 2) *low transparency* (called the *crossmapped* condition): target components looked similar to noncorresponding base objects (e.g. the speed gauge in the target resembled the temperature gauge in the base). In the crossmapped condition a given target object looks like one of the base objects, but its role in the relational structure is different; thus the object similarities are in conflict with the functional similarities between the base and target domains. We predicted that crossmapping would greatly disrupt transfer since surface similarity is in conflict with structural similarity.

A third question was whether the effects of either systematicity and transparency in transfer would increase when the workload demands increased. To investigate this question, in a second series of trials, subjects were given an additional task (called the load task) to perform while operating the target device.

Method

Subjects

Subjects were 63 undergraduates at the University of Illinois with little or no background in the physical sciences, who were paid for their participation.

Materials and Procedure

Devices: Subjects learned to operate a computer simulated device panel like that shown in Fig. 2(b). For ease of presentation, we describe the systematic condition, pointing out differences with the nonsystematic condition when necessary.

Each panel consisted of six gauges. Two of these represented control switches: engine thrust and coolant valve. Four of the gauges were display gauges representing the following parameters: speed of the ship, the level of coolant in the engine, engine temperature, and distance travelled. These parameters were linked by a causal model shown in Fig. 3. In the systematic condition, subjects were given a scenario about a steamship that explained the operating

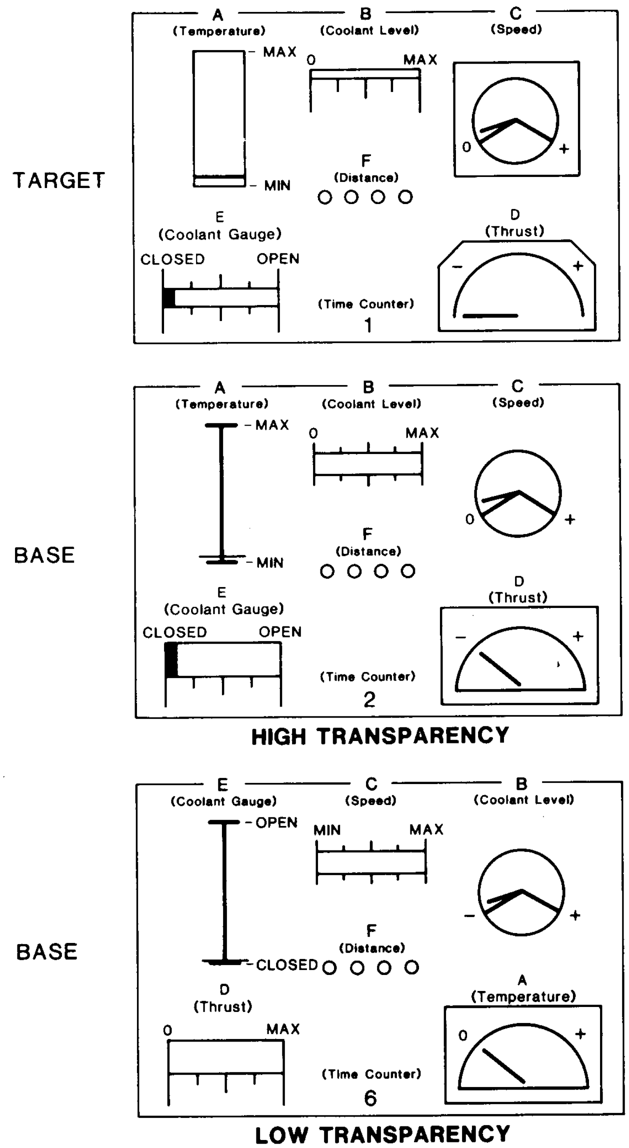

Fig. 2. Target and base device panels. Subjects first learned one of base devices and then transferred to target device.

procedures in terms of the causal model[4] (see also [17]). In the nonsystematic condition, subjects were given the same set of operating procedures, but no systematic causal model to organize their knowledge.

The subjects' task was to travel some distance in the given amount of time. This was operationalized as lighting all four lamps in the distance gauge before running out of time or entering a failure condition. Subjects had direct control of two of the five parameters displayed: 1) the engine thrust, which acted to increase the engine speed; and 2) the coolant valve opening, which acted to increase coolant level. Subjects controlled these parameters by pressing keys on a computer keyboard. Coolant level and engine speed worked opposite one another to influence the temperature: speed increased temperature, while coolant level held temperature down. Subjects had to achieve the optimal balance of speed and coolant level such that the ship would move fast enough to cover the required distance but would not overheat, and would maintain sufficient coolant level without coolant overflowing. Subjects in both systematic and nonsystematic conditions had the same task and the same operating procedures. They differed only in that the systematic group had a causal model that justified the procedures.

[4] No subjects were given the actual diagram shown in Fig. 3. Instead subjects in the systematic conditions read a description of the causal model.

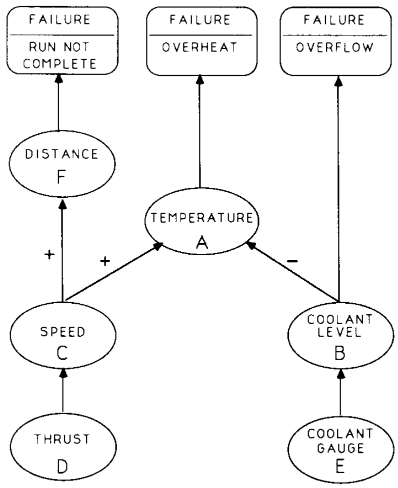

Fig. 3. Causal model of device operation—+'s and −'s indicate positive and negative relationships among the various gauges. Three failure conditions are shown at top.

Transparency conditions: There were two mapping conditions, which corresponded to high transparency (similar gauges—similar functions) and low transparency (cross-mapped condition): similar gauges—different functions). The variation in the mapping condition was achieved by varying the base device; the target device was identical for all subjects. (This was done to achieve maximum comparability across conditions in the transfer condition.) The two base devices are shown in Figs. 2(b) and 2(c), and the target device is shown in Fig. 2(a). To maximize the transparency effect, the appearance and location of the gauges were covaried, so that whenever two gauges looked similar they also occurred in the same relative location. Thus, in the high-transparency condition (Fig. 2(b)), each gauge in the target panel looked and functioned like one of the base gauges (and was placed in the same relative location as that gauge). In the low-transparency condition (Fig. 2(c)), cross-mapping occurred: each gauge in the target looked like one of the base gauges and was placed in the same relative location; but its function matched that of a different base gauge. Thus, any tendency to place the target gauge in correspondence with its lookalike counterpart in the base device would lead to error.

The design was a 2×2 factorial with two between-subjects factors: systematicity (systematic or nonsystematic device model) and mapping condition (high or low transparency), for a total of four treatment conditions.

Procedure: Subjects were randomly assigned to one of the four groups. They were given an introduction to the experiment and received a sheet of operating instructions on how to run the device.

Each gauge was described in terms of its operation and interactions with other gauges: what other gauges affect it and which gauges it affects. The operating instructions also set forth the possible run failures (described below). For both systematic and nonsystematic subjects a letter name was used to label each gauge, as shown in the devices in Fig. 2. Additionally, systematic subjects had semantic labels on their instruction sheets for each gauge (e.g., Speed (Gauge C)).[5] The operating instructions given to the nonsystematic subjects were identical to those given the systematic subjects except for the addition of semantic labels in the systematic condition. For example, the description of the temperature gauge for the systematic subjects was:

> Temperature (Gauge A) is controlled by Speed (Gauge C) and Coolant Level (Gauge B). Speed (Gauge C) acts to increase the level of Temperature (Gauge A). Coolant Level (Gauge B) acts to decrease or hold down the level of Temperature (Gauge A). In other words, Coolant Level (Gauge B) and Speed (Gauge C) work in opposite directions to control Temperature (Gauge A). Do not let Temperature go to its maximum.

For the nonsystematic subjects, the description of temperature was:

> Gauge A is controlled by Gauge C and Gauge B. Gauge C acts to increase the level of Gauge A. Gauge B acts to decrease or hold down the level of Gauge A. In other words, Gauge B and Gauge C work in opposite directions to control Gaüge A. Do not let Gauge A go to its maximum.

In addition to the operating instructions, subjects were also given a diagram of the base panel (Figs. 2(b) or 2(c), depending on the subject's mapping condition). They were allowed to study the operating instructions and the diagram while learning the device.

All subjects received the same procedural model of the device. The difference was that in the systematic conditions, subjects also received a causal model that explained the procedural steps, whereas in the nonsystematic conditions no additional causal justification was provided for the procedures. Because the systematic instructions provide an explicit set of higher-order causal constraints that motivate the operating procedures, subjects in the systematic condition should have a more robust procedural model than subjects receiving nonsystematic instructions.

Learning the base device: After subjects read the instructions, they learned to operate the base device. On each run subjects could either complete the run successfully or encounter one of three failure conditions (see Fig. 3): 1) *coolant overflow,* in which Gauge B (the coolant level) reached maximum; 2) *device overheating,* meaning that Gauge A (temperature gauge) reached maximum; or 3) *time out,* in which the time counter reached maximum before the required distance was covered. In a failure condition the subject saw a message indicating whether the failure was a Gauge A failure, a Gauge B failure, or a Time Out failure. As soon as one trial ended, the subject pressed

[5]The semantic labels did not appear on the diagram or on the screens. The device panels were identical for systematic and nonsystematic subjects and showed only the alphabetic labels.

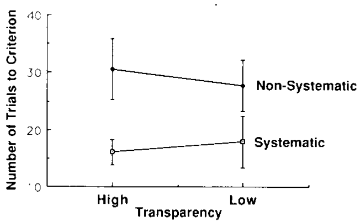

Fig. 4. Number of trials to criterion in base device. Standard error is indicated by bars.

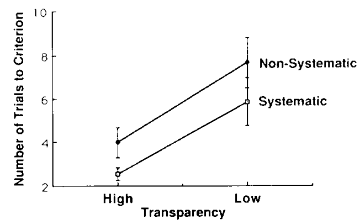

Fig. 5. Number of trials to criterion in target device. Standard error is indicated by bars.

a key to continue on to the next trial. Subjects were considered to have learned the base device when they reached a criterion of three correct trials out of five successive trials.

Transferring to the target: Upon reaching the learning criterion on the base device, subjects were immediately transferred to the target device. They were told that it was operated in the same way as the base device, but that they should study it carefully. The target gauges were always correctly labelled with the same letter labels as the base (e.g., C, D) according to functionality. Thus all subjects logically possessed enough information to allow perfect transfer of their knowledge of the base. The difference was that in the high-transparency condition this correct mapping was supported by the perceptual object matches, and in the low-transparency condition it was undermined by them. The criterion for learning the target was two correct out of four successive trials.

Target with load task: After reaching criterion on the target, subjects were given the load task. This was the same target task they had just learned, with an auxiliary load task requirement as follows. Approximately every five seconds a random number would appear at the same central location on the screen. The subject had to respond by pressing one of two keys within three time units (about 6 s). After two failures on the auxiliary task, the trial was aborted. Subjects then received a message indicating auxiliary task failure and a new trial began. Subjects were told that the load task was the secondary task and that controlling the device was the primary task. As in the target-only task device, the success criterion was two correct out of four successive trials.

Each subject was given an hour and fifteen minutes to learn the base device. Of the 63 subjects, 27 failed to reach criterion on the base device in the time allowed. They are not included in the analysis. Thus there were 36 subjects (nine in each cell).

Results

Learning the base: Systematicity had strong effects on the original earning of the base device. As shown in Fig. 4, subjects in the systematic condition required fewer trials to reach criterion than subjects in the nonsystematic condition. Using as the dependent measure the number of trials to criterion (3 correct runs of 5) a 2×2 (systematicity $\times$ mapping condition) analysis of variance (ANOVA) was performed.[6] The data were transformed using a logarithmic (base 10) function before performing the ANOVA. This analysis showed a main effect for Systematicity, $F(1,32)=8.398$, $p<0.01$. There was no effect of the mapping condition, $F(1,32)<1$, nor of the interaction of systematicity $\times$ mapping condition, $F(1,32)<1$. The significant effect of systematicity for the base device is evidence that giving a coherent causal model of a device helps initial learning of the device, as also found by Kieras and Bovair [17]. The lack of a main effect or interaction for mapping condition serves to confirm that there were no differences in difficulty of the base due to the transparency manipulation.

Transferring to the target: As predicted, both systematicity and mapping condition had strong effects on transfer difficulty (see Fig. 5). The more systematic the initial model and the more transparent the mapping of component correspondences, the more quickly subjects learned the target.

The dependent measure of difficulty of transfer was the number of trials to reach criterion (two correct out of four successive trials) on the target device. The data were transformed using a logarithmic (base 10) function before the ANOVA. A 2×2 ANOVA of systematicity $\times$ mapping condition showed a main effect of Systematicity, $F(1,32)=5.713$, $p<0.05$, and of mapping condition, $F(1,32)=21.616$, $p<0.01$. The main effect of Systematicity indicates that possessing a well-structured model for a device aids in transferring knowledge from one device to another. The main effect of mapping condition provides evidence for the importance of transparency: when object correspondences (and location cues) are crossmapped there is disruption of transfer even when the systems are absolutely identical in all other respects. There was no interaction between systematicity $\times$ mapping condition, $F(1,32)<1$.

Target with load: As in the previous task we predicted that it would be easier for subjects in the Systematic conditions and subjects in the high-transparency conditions to transfer to another device than for subjects in nonsystematic or low-transparency conditions. As Fig. 6 shows, the results lie in the predicted direction, but none

[6] Transparency should, of course, have no effect here. Mapping Condition was included as a check to be sure both base devices were equally difficult.

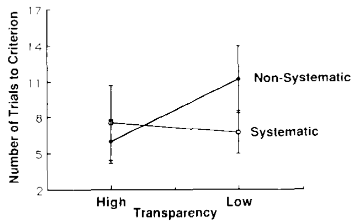

Fig. 6. Number of trials to criterion in load device. Standard error is indicated by bars.

of the effects is statistically significant. It seems plausible that the effects here were diminished by the target-with-load task always following the target-only task, so that substantial transfer had already occurred. As before, we used the number of trials to criterion (two correct out of four successive trials) as the dependent measure and performed an ANOVA on the log of the dependent variable. There were no significant main effects for Systematicity, $F(1,32) < 1$, or mapping condition, $F(1,32) = 2.402$, $p < 0.15$. The systematicity × mapping condition interaction was also non-significant, $F(1,32) = 1.803$, $p < 0.20$.

A second issue was whether the load task was more difficult than the target task alone. To look at the effect of the added workload, we compared performance between the target and target load tasks. For this analysis a systematicity × mapping condition × task ANOVA was performed, with task representing a within-subjects variable with two levels: target and target-load. Again, using the number of trials to criterion (two correct out of four successive trials) as the dependent measure, we carried out an ANOVA on the log of the dependent variable. There was a main effect of Task, $F(1,32) = 4.807$, $p < 0.05$, and of mapping condition, $F(1,32) = 11.026$, $p < 0.01$. No other effects were significant. The main effect for task indicates that, not surprisingly, the target load task was harder than the target task alone. The mapping condition main effect simply confirms the pattern discussed above: performance in high-transparency conditions is better than that in low-transparency conditions across the load and nonload target tasks.

Conclusion

In this research, we considered transfer of training from the viewpoint of analogical mapping. According to structure-mapping theory, the implicit goal of an analogical mapping process is to achieve maximal structural match between base and target. The mapping process involves setting up object correspondences and carrying across predicates. We predicted that two factors would influence the difficulty of the mapping process. The first factor is base systematicity—the presence of a coherent causal model of the base that can apply in the target [14]. The second factor is the transparency of the object correspondences: the greater the surface similarity between the corresponding objects in the base and target, the easier it should be to keep the mappings clear. The results of this study provide support for our two major predictions: both the systematicity of the base model and the transparency of the correspondences improve transfer of a device model.

Comparison with the Gentner and Toupin Study

On the whole these results parallel the findings of Gentner and Toupin [14] that systematicity and transparency affect transfer. However, there was one difference between the two studies. Gentner and Toupin found an interaction between systematicity and transparency: in the low-systematicity condition there was a large transparency effect, but in the high-systematicity condition there was no transparency effect: both transparency conditions did equally well. It was as though a systematic model insulated subjects from the misleading effects of surface mismatches. It seemed plausible that such an interaction might be seen in this task: i.e., that subjects in the systematic low-transparency condition might overcome the misleading object correspondences by relying on their understanding of the causal model. However, although the results of the load task (as shown in Fig. 6) suggest a trend towards this kind of interaction between systematicity and transparency, in neither the target task nor the target-load task was the interaction significant.

One possible reason that the systematicity × transparency interaction did not appear here is that there may have been a subject-selection effect. In the Gentner and Toupin task, all subjects completed the task. However, in the present study seventeen subjects in the nonsystematic conditions failed to complete the base task, as compared to only ten in the systematic conditions; therefore the nonsystematic conditions may have selected for better subjects. If such a selection effect obtained, any beneficial effects of systematicity on transfer performance could have been countermanded because the nonsystematic conditions would have had better subjects on the whole.

Levels of Analysis and Task Description

We turn now to the issue of levels of tasks. According to structure-mapping theory an ideal processor at the computational level would be able to ignore surface matches between objects and maximize the structural match. In pilot research, we find that subjects can achieve this kind of relational focus when they are explicitly told to carry out an analogy. The subjects are taught the systematic version of the base device, in the usual manner. They are then given the standard target diagram and told that the device is analogous to the original device. Our preliminary results suggest that, even in the crossmapped condition, subjects can explicate the analogy correctly: that is, they can focus on carrying over the systematic relational structure and ignore surface similarities.

In contrast, in the study reported here, performance was influenced by surface object similarities as well as systematic structure. This points up the importance of considering the algorithmic–behavioral level as well as the computational–informational level [10]–[12]. Different kinds of similarity may differentially influence performance depending on the task conditions—e.g., whether subjects are explicitly extracting commonalities or implicitly using them in transfer.

Structural and Surface Fidelity

Fidelity in the human factors literature behaves something like similarity in the general psychological literature. First there is widespread agreement that it is important but no consensus on exactly what it is [19]–[21]. Second, as with similarity, fidelity is used to refer to both structural and surface overlap. For example, a training simulator might be constructed (at considerable expense), so that it looks and behaves identical to a target device. This is an example of what we could call both surface and structural fidelity. In other cases fidelity is construed as "fidelity to the underlying principles of the system", which we might call structural fidelity. A major issue in training is deciding when and how to use these two types of fidelity. For instance, the STEAMER system is a successful Navy training device that emphasizes teaching principles rather than perfect surface fidelity [22].

Given that both high structural and surface fidelity are important, how should we incorporate these types of fidelity into training regimens? One way might be to train people on sets of a highly surface-similar devices first and then gradually decrease similarity so that they come to rely on the causal model and are not dependent upon the superficial features. This sequence is similar to a notion suggested by Ross [23] in learning problem schemas, and by Forbus and Gentner [24] in learning about physical domains. We believe that research separating surface and structural similarity will be important in achieving the optimal management of fidelity in training.

Acknowledgment

We thank Andrea Hoffman and Steve Clark for running subjects and Andy Lenzini for programming.

References

[1] L. R. Novick, "Transfer and expertise in problem solving: A conceptual analysis," Stanford University, unpublished manuscript, 1985.
[2] D. Gentner and D. R. Gentner, "Flowing waters or teeming crowds: Mental models of electricity," in *Mental Models*, D. Gentner and A. L. Stevens, Eds. Hillsdale, NJ: Erlbaum, 1983, pp. 99–129.
[3] M. L. Gick and K. J. Holyoak, "Schema induction and analogical transfer," *Cognitive Psych.*, vol. 15, no. 1, pp. 1–38, 1983.
[4] M. J. Rattermann and D. Gentner, "Analogy and similarity: Determinants of accessibility and inferential soundness," in *Proc. Ninth Annu. Meeting of the Cognitive Sci. Soc.*, Seattle, WA, July 1987.
[5] S. K. Reed, "A structure-mapping model for word problems," *J. Experimental Psych.: Learning, Memory and Cognition*, vol. 13, no. 1, pp. 124–139, 1987.
[6] D. Gentner, "The structure of analogical models in science," Bolt Beranek and Newman, Cambridge, MA, tech. rep. no. 4451, July 1980.
[7] ____, "Structure-mapping: A theoretical framework for analogy," *Cognitive Sci.*, vol. 7, no. 2, pp. 155–170, 1983.
[8] ____, "Mechanisms of analogical learning," to appear in *Similarity and Analogical Reasonging*, S. Vosniadou and A. Ortony, Eds. London: Cambridge Univ. Press, in press.
[9] B. Falkenhainer, K. Forbus, and D. Gentner, "The structure-mapping engine" in *Proc. Meeting of Int. Assoc. For Artificial Intell.*, Philadelphia, PA, Aug., 1986.
[10] D. Marr, *Vision*, San Francisco, CA: W. H. Freeman, 1982.
[11] S. E. Palmer, "Levels of description in information processing theories of analogy," *Similarity and Analogical Reasoning*, S. Vosniadou and A. Ortony, Eds. London: Cambridge Univ. Press, in press.
[12] S. E. Palmer and R. Kimchi, "The information processing approach to cognition," in *Approaches to Cognition: Contrasts and Controversies*, T. Knapp and L. C. Robertson, Eds., Hillsdale, NJ: Erlbaum, 1985.
[13] D. Gentner, "Evidence for a structure-mapping theory of analogy," University of Illinois, Dept. Comput. Sci., Urbana, IL, tech. rep. no. UIUCDCS-R-86-1316, Dec. 1986.
[14] D. Gentner and C. Toupin, "Systematicity and surface similarity in the development of analogy," *Cognitive Sci.*, vol. 10, no. 3, pp. 277–300, 1986.
[15] D. E. Norman, "Some observations on mental models," in *Mental Models*, D. Gentner and A. L. Stevens, Eds. Hillsdale, NJ: Erlbaum, 1983, pp. 7–14.
[16] ____, "Cognitive Engineering," in *User Centered System Design*, D. E. Norman and S. W. Draper, Eds. Hillsdale, NJ: Erlbaum, 1986, pp. 31–65.
[17] D. E. Kieras and S. Bovair, "The role of a mental model in learning to operate a device," *Cognitive Sci.*, vol. 8, pp. 255–273, 1984.
[18] W. B. Rouse and N. Morris, "On looking into the black box: Prospects and limits in the search for mental models," *Psych. Bull.*, vol. 100, no. 3, pp. 349–363, 1986.
[19] G. Lintern, "Context and dependency effects in skill acquisition," in *Trends in Ergonomics/Human Factors II*, R. E. Eberts and C. G. Eberts, Eds. New York, NY: Springer-Verlag, 1985, pp. 249–256.
[20] D. C. Wightman and G. Lintern, "Part-task training for tracking and manual control," *Human Factors*, vol. 27, no. 3, pp. 267–283, 1985.
[21] W. Schneider, "Training high-performance skills: Fallacies and guidelines," *Human Factors*, vol. 27, no. 3, pp. 285–300, 1985.
[22] A. Stevens and B. Roberts, "Quantitative and qualitative simulation in computer-based training," *J. Computer-Based Instruction*, vol. 10, pp. 16–19, 1983.
[23] B. H. Ross, "Remindings in learning and instruction," to appear in *Similarity and Analogical Learning*, S. Vosniadou and A. Ortony, Eds. London: Cambridge Univ. Press, in press.
[24] K. Forbus and D. Gentner, "Learning physical domains: Towards a theoretical framework," in *Machine Learning: An Artificial Intelligence Approach, Vol. II*, R. S. Michalski, J. Carbonell, and T. Mitchell, Eds. Los Altos, CA: Morgan Kaufmann, 1986, pp. 311–348.

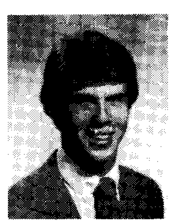

Robert M. Schumacher received his B.A. in 1981 from Eastern Illinois University, Charleston, IL, and his M.A. in 1987 from the University of Illinois, Urbana–Champaign. He received a Cognitive Science/Artificial Intelligence Fellowship from the University of Illinois, where he is a Ph.D. candidate in cognitive and experimental psychology.

His research interests include use of analogy and conceptual models in learning, human–computer interaction, and skill development.

Mr. Schumacher is a student member of the Human Factors Society and the Association for Computing Machinery.

Dedre Gentner received her B.A. in physics in 1967 from the University of California at San Diego and her Ph.D. in psychology in 1974 from the University of California at Berkeley.

She is an Associate Professor of Psychology at the University of Illinois and has been on the faculty of the University of Illinois since 1984. Prior to that she was a Senior Scientist at Bolt, Beranek and Newman, Inc. Currently she is on the Editorial Board for *Cognition* and for *Metaphor and Symbolic Activity*, and is a member of the governing board of the *Cognitive Science Society*. She edited the book *Mental Models*, along with Albert Stevens. Her research interests include analogy and similarity, representation of meaning and mental models of physical phenomena as well as their development of these.

Dr. Gentner is a member of the Psychonomic Society, the Cognitive Science Society, the Society for Philosophy and Psychology, and the Society for Research in Child Development, among others.