

Systematicity as a Selection Constraint in Analogical Mapping
Catherine Clement and Dedre Gentner
Psychology Department, University of Illinois

In an analogy a person's knowledge about one domain is used to understand a second domain: to highlight important similarities between domains or to predict new features of the second domain. For example, we use our knowledge about water flow to elucidate properties of electric circuitry. A central fact about analogy is that it is selective. Not all commonalities are equally important in matching or prediction. It follows that a central problem in theories of analogical mapping is to identify the selection constraints that tell us which similarities and dissimilarities count in an analogy (Kedar-Cabelli, 1985a; Hall, 1988). Within the enormous space of possible similarities between domains, what constrains the choice of information to include in an analogy?

Two major classes of selection constraints have been utilized in computational models of analogy: goal-relevance constraints and structural constraints. The goal-relevance constraints direct analogical mapping toward information that is relevant to the problem at hand (Burstein, 1983; Holyoak, 1985; Kedar-Cabelli, 1985b). The structural constraints focus mapping on such properties of the matching information as structural consistency (Winston, 1980; Gentner, 1980, 1983; Falkenhainer, Forbus and Gentner, in press). For example, in her account of analogy Gentner proposes that people observe the structural constraints of 1 to 1 object correspondences and structural consistency. Further, she proposes that two implicit selection rules underlie the interpretation of analogy. First, people seek to map relations between objects rather than non-relational object attributes. Second, people use a systematicity principle: they seek to map relations that participate in a system of relations governed by higher-order relations - e.g. causal relations - that can also be mapped between the two domains.

In support of the systematicity constraint, Gentner points out that the most systematic interpretation of an analogy is generally the preferred interpretation. However, Holyoak (1985) notes that systematicity is typically correlated with goal relevance and therefore these observations do not allow us to conclude that systematicity by itself can act as a selection constraint. For example, if a person who is trying to solve a problem about heat flow invokes an analogy to water flow, the causal structure relating inequality in level to direction of flow is likely to be both the most relevant and the most systematic structure shared by the two analogs. In the present research we looked at analogical mapping outside of a problem-solving context in order to test whether systematicity can act as a selection constraint on mapping in the absence of goal relevance.

We describe two experiments that focus on different aspects of analogical mapping: (a) *matching* existing information in the base and target domains (seeing similarities or identities between aspects of the base and target domains); and (b) *carryover* of new information from the base to the target (making new inferences about the target domain). Since these are distinct processes in constructing an analogy it is important to determine whether both are governed by the same selection principles. Thus the two studies tested whether systematicity acts as a selection filter to constrain both which matches are made and which inferences are drawn.

For this research we created a set of novel analogies. The basic plan was to devise analogies in which subjects could choose which of two lower-order facts to map. Both facts matched the between the base and target equally well. However they differed in whether they were part of a shared systematic relational structure. If systematicity acts as a selection constraint then subjects should focus on the fact governed by the shared structure in their interpretations and predictions from the analogy. In developing these materials we were guided by two further criteria:

1. Goal-relevance should not be a possible selection constraint. Thus, subjects did not have to map information in order to solve problems or prove points in the target domain.
2. Subjects' prior factual knowledge about the domains or expectations about relevance

CLEMENT AND GENTNER

should not influence their mapping. Thus, we developed analogies between novel, fictional domains rather than using real world analogies. To preserve a realistic degree of complexity, the situations described were fairly rich in information.

Design of the materials

The analogies developed each consisted of a base and target passage describing novel objects or organisms on fictional planets. The passages were about one page long and were written in the style of an encyclopedia article. Each passage included two key paragraphs. Before describing the materials further, we give an abbreviated example of one base domain which described Tams, creatures who live on a distant planet. Key paragraph 1: The Tams live on rock and can grind and consume minerals from the rock through the constant action of their underbelly. However, periodically they run out of minerals in one spot on the rock, and therefore must stop using their underbellies while they relocate. Key paragraph 2: Although at birth the Tams have rather inefficient underbellies, eventually the underbellies adapt and develop a texture that is specially suited to the rock the Tam lives on. As a consequence, the grown Tam cannot function on a new rock.

Table 1 illustrates the relational structure of this base domain "The Tams" and an analogous target domain called "The Robots", which describes robots who use probes to gather data from planets. All subjects received the same base domain. However, as shown in Table 1, there are two versions of the target domain, each given to half the subjects for counterbalancing. The base passage contains two different causal structures (the two key paragraphs discussed above). Each structure consists of a *key fact* (shown in italics in the table) and a causal system governing that key fact. The two key facts in this base domain are (1) that the Tams sometimes stop using their underbellies, and (2) that their underbellies cannot function on new rocks. Each target domain also contains two causal structures, and in both versions of the target, the key facts in each structure match the key facts in the base: (1) the robots sometimes stop using their probes and (2) the probes cannot function on new planets.

Thus, at the level of the individual key facts, both versions of the target match the base domain perfectly. However, the causal system governing a key fact does not always match the base domain. That is, within each version of the target domain, both key facts match the base, but only one key fact is linked to a causal system that also matches the base domain. We will call this the *identically-governed* key fact. We predicted that subjects should focus on this fact in mapping.

Table 1. Relational Structure of the Base Domain "The Tams" and the Target Domain "The Robots"

Base - "The Tams"	Target - "The Robots"	
	Version 1	Version 2
Consume minerals with underbellies	Gather data with probes	Gather data with probes
<i>Exhaust minerals in one spot and must relocate on the rock</i>	Exhaust data in one place and must relocate on the planet	Internal computers over-heat when gather a lot of data
<i>So stops using underbelly</i>	<i>So stops using probes</i>	<i>So stops using probes</i>

Born with inefficient underbelly	Designed with delicate probes	Designed with inefficient probes
Underbelly adapts and becomes specialized for one rock	Robots cannot pack probes to survive flight to a new planet.	Probes adapt and become specialized for one planet
<i>So underbelly can't function on new rock</i>	<i>So probes can't function on new planet</i>	<i>So probes can't function on new planet</i>

Note. Key facts are shown in italics. Matching causal information is shown in bold face. In Experiment 2 italicized facts were removed from the target.

CLEMENT AND GENTNER

The materials were designed to avoid confoundings with particular content. It can be seen that in Target version 1, the cause for the first key fact matches the base domain, but the cause for the second key fact does not. Target version 2 simply reverses which key fact is identically-governed. This ensures that preferences for identically-governed key facts cannot be attributed to preferences for a particular key fact in itself. A further control task was included to ensure that subjects' responses cannot be attributed to preference for a particular fact in context. That is, a key fact might seem more or less salient depending on its context in a particular version of the target, independently of the analogy. To check for this possibility, control subjects, who read only the Target domains, judged which key fact was more important for each domain. If these control subjects show no bias towards one key fact or the other, then any preference among experimental subjects can be attributed to the specific effects of the match with the base system.

Experiment 1 examined the matches people include in an analogy. Subjects had to judge how well each key fact in the target contributed to the analogy with the base domain. We predicted that identically-governed key facts would be preferred over differently-governed key facts, even though in each target both key facts match the base equally well.

Experiment 1

Method

Subjects.

Subjects were 48 paid undergraduate students at the University of Illinois. Half were assigned to the experimental group and half to the control group. Subjects were run in groups of three to six people.

Task and Hypotheses.

Experimental subjects. The subjects' task was to study the base and target domains and then judge how well each key fact in the target contributed to the analogy. Since our materials are complex, subjects were given three *learning* tasks before judging the key facts to ensure that they processed the analogy thoroughly: (a) they read the base and target domains. (b) they identified which object in the target domain corresponded to each object from the base domain. (c) they described the ways in which the two domains were and were not analogous. Two judgment tasks were then given:

Rating Task. Subjects were given the two key facts from the target domain (e.g. the italicized facts from the Target in Table 1). Subjects rated on a scale of 1 to 7 how well each fact "supports the claim that the base and target were analogous; how well does each fact contribute to making the analogy a good analogy".

Choice Task. Subjects were again given the two key facts and chose which *best* supported the analogy.

During the learning and judgment tasks for each analogy subjects were encouraged to re-examine the descriptions of the two domains as needed. Thus, there were no memory requirements. We predicted that identically-governed key facts should receive higher ratings than differently-governed key facts, and should be more often chosen as the fact that best supports the analogy.

Control subjects. Control subjects read and summarized the target domain. Following this, subjects were given the two key facts and were asked to (a) rate on a scale of 1 to 7 how important each fact was to the target domain and (b) choose which of the key facts was most important. Assuming the materials are not biased, the control subjects should show no preference for either identically or differently-governed key facts.

Materials.

Subjects completed these tasks for four novel analogies each designed according to the structure described above. Analogies were given in four orders of presentation across subjects. For each analogy, all subjects received the same base version; there were two target versions which were each given to half the subjects in each group (see Table 1).

CLEMENT AND GENTNER

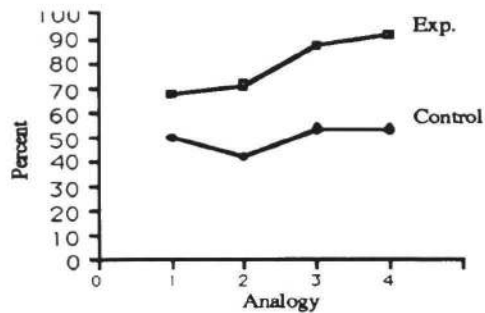


Figure 1. Exp.1 Choice Task: The Percentage of Subjects in Each Group Choosing the Identically Governed Key Fact for Each Analogy.

Results and Discussion

Rating Task

As predicted, experimental subjects rated identically-governed key facts significantly higher than differently-governed key facts. The mean ratings across the four analogies were 6.28 and 4.85 for each fact type respectively. The control subjects gave equivalent importance ratings to the two sets of key facts ($M = 5.2$ and 5.3) indicating that the experimental results are not due to differences in the importance of the key facts in the two versions of the target.¹

Choice Task

As shown in Figure 1 the results again support our prediction that the subjects should view the identically governed key fact as supporting the analogy. Across the four analogies, the experimental subjects chose the identically-governed fact 67% to 92% of the time, while control subjects chose it 42% to 54% of the time (i.e. at chance level).²

The results of Experiment 1 support the claim that systematicity acts as a selection constraint on which matches are selected as belonging to an analogy. Even though both key target facts matched the base domain equally well, the one that was linked to a larger causal structure that also matched the base domain was consistently judged to contribute best to the analogy. Further, the lack of preference in the control group indicates that the results are not due to the importance of a fact in the target domain.

Experiment 2

Experiment 1 focused on which of two matching facts would be selected as belonging to an analogy. The second experiment examined the carryover process in analogical mapping. We deleted certain key facts from the target passages in order to test whether systematicity would act as a selection constraint on subjects' predictions about a target domain.

Method

Subjects

Subjects were 32 (16 control and 16 experimental) paid undergraduate students at the University of Illinois.

Materials and Design

The base and target domains were identical to those of Experiment 1 except that the two key facts were removed from the target domains. Thus, in Table 1, the two italicized facts shown in each target were removed. This meant that the target domain included only the antecedent information that could potentially cause each key fact in that domain. As in Experiment 1, two versions of each target were used. In each version, one antecedent matched the base domain but the other antecedent did not.

CLEMENT AND GENTNER

Subjects were asked to make a prediction about the target domain. The question was whether they would predict one of the two key facts omitted from the target, and more importantly, which one they would predict. If subjects are simply trying to predict facts about the target that correspond to facts in the base, they should be equally likely to predict either of the omitted key facts. However, the systematicity prediction is that subjects should predict the one key fact for which there was a matching causal system, ie. identically-governed in the base and target.

Task

Experimental subjects. For each analogy subjects first had three learning tasks identical to those in Experiment 1. They were then given three experimental tasks:

Prediction Task. Subjects were asked to make one prediction about the Target domain. They were told to predict new information about the target "that was suggested by the analogy with the base domain".

Rating and Choice Tasks. After making their own predictions subjects were asked to judge possible predictions about the target. Subjects were given as predictions the two key facts that appeared in the base but not in the target. They rated each of these on a scale of 1 to 7 according to how well this prediction followed from the analogy. Finally, subjects were again given the two key facts and asked to choose which was best predicted by the analogy.

Control subjects. Control subjects made predictions after reading only the target domain. These subjects read and summarized the target and then performed Prediction, Rating and Choice tasks parallel to those given the experimental subjects. Control subjects predicted one new fact about the target "that might be true given the information about the target". Then, given the two key facts as predictions about the target, subjects rated each prediction "according to how well it followed from information in the target" and chose which prediction best followed from the information in the target.

Results and Discussion

Prediction Task

As predicted, experimental subjects tended to make the prediction sanctioned by systematicity. Figure 2 shows subjects' predictions about the target domain grouped into three categories: (a) predictions of key facts that would be identically-governed in the base and target (b) predictions of key facts that would be differently-governed (c) all other predictions about the target. As shown, experimental subjects most frequently predicted the identically-governed key fact. In contrast, the most frequent response for the control subjects was to predict other information, and their remaining responses were evenly distributed between the two key facts. The difference in performance between the control and experimental subjects indicates that the experimental results are not due to a bias in which one of the key facts was inherently more plausible in the target. More fundamentally, it indicates that the experimental subjects were indeed

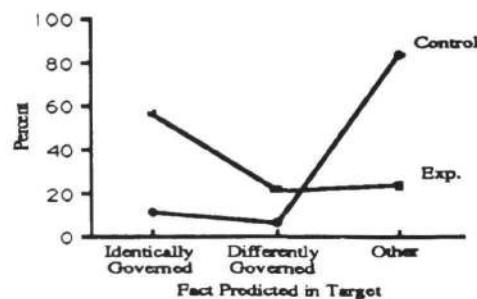


Figure 2. Exp. 2 Prediction Task: Percent of Predictions in Each Category in Each Group

guided by the analogy with the base.³

Rating and Choice Tasks

As hypothesized, experimental subjects rated identically-governed predictions significantly higher than differently-governed predictions ($M = 6.03$ and $M = 4.41$, respectively). Control subjects showed the reverse pattern, again indicating that materials were not biased in favor of the hypothesis ($M = 4.32$ and $M = 5.22$ respectively).⁴ The results for the choice task also support the systematicity hypotheses. The experimental subjects chose the identically governed fact 62% to 87% of the time, while control subjects chose it 25% to 66% of the time (they predicted at or below chance level -50%- for three of the four analogies).⁵ Thus, although for each analogy both potential predictions corresponded to facts given in the base domain, the one prediction linked to a matching causal system was judged by subjects to follow best from the analogy.

General Discussion

The results show that systematicity is a selection constraint on the choice of information to map in an analogy. In both experiments the subjects were provided with two possible matching facts or predictions, both of which matched well between the base and target, but only one of which participated in a governing higher-order system that also matched across domains. In Experiment 1 subjects chose which of these key matches contributed best to the analogy; they consistently preferred the fact linked to the larger matching system. In Experiment 2 subjects made predictions about the target domain. Their predictions, like their judgments of matches centered on the particular key facts that would follow from a matching causal antecedent.

Although the issue of selection constraints on analogy has attracted attention in AI (Winston, 1980; Carbonell, 1981), psychological research on this topic is just beginning. There is prior evidence that systematicity affects the subjective evaluation of the soundness of an analogy (Gentner & Landers, 1985; Rattermann & Gentner, 1987). However, to our knowledge the present research is the first test of the stronger claim that systematicity can act as a selection filter on mapping. This study provides evidence that systematicity can operate as a selection constraint in tasks in which no prior goal context exists, i.e. in which there is no problem to be solved so that relevance to existing goals cannot determine mapping choices.

It is important to note, however, that these results do not address whether goal-relevance can act as a filter instead of, or in addition to, structural constraints. Burstein and Adelson (1987) give protocol evidence that goal relevance constrained the selection of systematic information to map in an analogy used to solve computer programming problems. Thus goal relevance may be used to disambiguate analogous cases where the structural information is not sufficient, much as knowledge of the surrounding context is used to disambiguate sentence meanings in cases where syntactic information is insufficient, as in "flying planes can be dangerous".

Relevance may interact with structural constraints in several ways. For example, before analogical mapping begins task goals may govern the content of the domain representations that are inputs to the analogy process (Gentner, 1988). Or, as discussed above if multiple matching systematic structures are found, the choice among them may be based on relevance to current goals. These two constraints have been integrated to some extent in some computational models of analogy (Burstein, 1983; Kedar Cabelli, 1985). In any case our results reinforce the need to distinguish the structural coherence of information to be mapped from the relevance of this information to a current task goal. Although contextual relevance will certainly play a role in the interpretation of an analogy, the systematicity of shared information is itself a psychologically real constraint on mapping.

Footnotes:

1. A repeated measures ANOVA comparing the mean ratings of identically governed and differently governed key facts by each group revealed a significant interaction between group and fact -type $F'(1,34)=14.37, p < .001$. The analyses of the simple effects indicates that the difference between ratings of each fact type within the experimental group is significant, $F'(1,9)=10.86, p < .01$

CLEMENT AND GENTNER

2. The difference between groups is significant for three of the four analogies: ($p < .05$) Fisher Exact tests, 1 tailed. Another scoring method is to assign each subject a score for the number of choices of identically-governed key facts across the four analogies (possible score is 0-4), the mean for experimental subjects (3.17) is significantly greater than the mean for control subjects (2.0), $F(1,28) = 18.67, p < .0001$.

3. Agreement between these ratings and those of a second, blind judge, was .91. A repeated measures anova, comparing the mean number of identically governed and differently governed predictions by each group shows a significant interaction between group and fact type ($F(1,16) = 11.76, p < .01$). The analyses of the simple effects indicates that the difference between predictions of each fact type within the experimental group is significant $F(1,8) = 36, p < .001$. The difference within the control group is not significant.

4. The interaction between group and fact-type is significant $F(1,11) = 9.0, p < .025$. The analyses of the simple effects indicates that the difference between ratings of each fact type within the experimental group is significant, $F(1,11) = 6.05, p < .05$.

5. When subjects are scored for the number of choices of identically-governed key facts across the four analogies (possible score is 0-4), the mean for experimental subjects (2.9) is significantly different from the mean for control subjects (1.8), $F(1,16) = 16.2, p = .001$

References

- Burstein, M. (1983). Concept formation by incremental analogical reasoning and debugging. In R. Michalski, J. Carbonell, & T. Mitchell (Eds.), Machine Learning (Vol. 2, pp. 351-369). Los Altos, CA: Morgan Kaufman.
- Burstein, M., & Adelson, B. (1987). Analogical learning: Mapping and integrating partial mental models. In Program of the Ninth Annual Conference of the Cognitive Science Society (pp. 11-22). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Carbonell, J. G. (1981). Invariance hierarchies in metaphor interpretation, Proceedings of the Third Annual Meeting of the Cognitive Science Society, Berkeley, CA, pp. 292-295.
- Falkenhainer, B., Forbus, K. D. & Gentner, D. (in press). The structure-mapping engine Artificial Intelligence.
- Gentner, D. (1980). The structure of analogical models in science (Tech. Rep. No. 4451). Cambridge, MA: Bolt Beranek and Newman, Inc.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. Cognitive Science, 7 (2), pp. 155-170.
- Gentner, D. (1988). Mechanisms of analogical learning. To appear in S. Vosniadou and A. Ortony, (Eds.), Similarity and analogical reasoning.
- Gentner, D., & Landers, R. (1985). Analogical reminding: A good match is hard to find. In Proceedings of the International Conference on Systems, Man and Cybernetics, (pp.607-613). Tucson, AZ.
- Hall, R. P. (1988). Computational approaches to analogical reasoning: A comparative analysis. Artificial Intelligence, pp. 1-75.
- Holyoak, K. J. (1985). The pragmatics of analogical transfer. In G. H. Bower (Ed.), The Psychology of Learning and Motivation, (Vol. 1, pp. 59-87). New York: Academic Press.
- Kedar-Cabelli, S. (1985a). Analogy - from a unified perspective. (Tech. Rep. No. ML-TR-3). New Brunswick, New Jersey. Laboratory for Computer Science Research, Hill Center for the Mathematical Sciences, Rutgers University.
- Kedar-Cabelli, S. (1985b). Purpose-directed analogy. Proceedings of the Seventh Annual Conference of the Cognitive Science Society (pp. 150-159), Irvine California.
- Rattermann, M. J. & Gentner, D. (1987). Analogy and similarity: Determinants of accessibility and inferential soundness. In the Program of the Ninth Annual Conference of the Cognitive Science Society pp. 23-34.
- Winston, P. H. (1980). Learning and reasoning by analogy. Communications of the ACM, 23 (12), pp. 689-703.

Acknowledgements

This research was supported by a Child Health and Human Development Grant No. HD07205-08, and by the Office of Naval Research under Contract No. N00014-85-K-0559. We thank Laura Kotovsky, Ronald Mawby, Betsy Perry, Mary Jo Rattermann, Brian Ross, and Bob Schumacher.