

4

DEPARTMENT OF COMPUTER SCIENCE
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

AD-A216 029



REPORT NO. UIUCDCS-R-89-1558

UILU-ENG-89-1777

SYSTEMATICITY AS A SELECTION CONSTRAINT
IN ANALOGICAL MAPPING

by

Catherine A. Clement
Dedre Gentner

DTIC

SELECTED
DEC 13 1989

S

CD

D

September 1989

89 12 15 0 5

This research is supported by the Office of Naval Research, Personnel
and Training Research Programs, Contract No. N00014-85-K-0559.

Reproduction in whole or part is permitted for any purpose of the
United States Government.

Approved for public release; distribution unlimited.

REPORT DOCUMENTATION PAGE				Form Approved OMB No 0704-0188	
1a REPORT SECURITY CLASSIFICATION Unclassified			1b RESTRICTIVE MARKINGS		
2a SECURITY CLASSIFICATION AUTHORITY			3 DISTRIBUTION / AVAILABILITY OF REPORT Approved for public release; Distribution unlimited		
2b DECLASSIFICATION / DOWNGRADING SCHEDULE					
4 PERFORMING ORGANIZATION REPORT NUMBER(S)			5 MONITORING ORGANIZATION REPORT NUMBER(S)		
6a NAME OF PERFORMING ORGANIZATION University of Illinois Department of Psychology		6b OFFICE SYMBOL (If applicable)	7a NAME OF MONITORING ORGANIZATION Cognitive Science (Code 1142CS) Office of Naval Research		
6c ADDRESS (City, State, and ZIP Code) 603 E. Daniel Champaign, IL 61820			7b ADDRESS (City, State, and ZIP Code) 800 N. Quincy Street Arlington, Va. 22217-5000		
8a NAME OF FUNDING / SPONSORING ORGANIZATION		8b OFFICE SYMBOL (If applicable)	9 PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER N00014-89-J-1272		
8c ADDRESS (City, State, and ZIP Code)			10 SOURCE OF FUNDING NUMBERS		
			PROGRAM ELEMENT NO 61153N	PROJECT NO	TASK NO
					WORK UNIT NR 442F-007
11 TITLE (Include Security Classification) Systematicity as a Selection Constraint in Analogical Mapping					
12 PERSONAL AUTHOR(S) Dr. Catherine A. Clement and Dr. Dedre Gentner					
13a TYPE OF REPORT Technical Report		13b TIME COVERED FROM 88-9-1 TO 92-10-31		14 DATE OF REPORT (Year, Month, Day) 89-9-15	
15 PAGE COUNT					
16 SUPPLEMENTARY NOTATION					
17 COSATI CODES			18 SUBJECT TERMS (Continue on reverse if necessary and identify by block number)		
FIELD	GROUP	SUB-GROUP			
05	10		Systematicity Analogy Structure-mapping		
			Constraints Mapping		
19 ABSTRACT (Continue on reverse if necessary and identify by block number)					
Analogy is often viewed as a partial similarity match between domains. But since between any two domains there are more partial similarities than good analogies, it follows that analogy is selective. Three experiments examined the selection constraint on which relations are mapped between a base and target in an analogy. In Experiment 1 subjects judged two matches to be included in an analogy: an isolated match, and a match embedded in a larger mapping system. Subjects preferred the embedded match. In Experiments 2 and 3 subjects made analogical predictions about a target domain. Subjects predicted information that followed from a causal system that matched the base domain, rather than the information that was equally plausible, but that created an isolated match with the base. Results support Gentner's (1983, 1989) structure-mapping theory that analogical mapping concerns systems and not individual predicates, and that attention to shared systematic structure constrains the selection of information to include in an analogy.					
20 DISTRIBUTION / AVAILABILITY OF ABSTRACT ☒ UNCLASSIFIED/UNLIMITED ☐ SAME AS RPT ☐ OTC USERS			21 ABSTRACT SECURITY CLASSIFICATION Unclassified		
22a NAME OF RESPONSIBLE INDIVIDUAL Dr. Susan Chipman			22b TELEPHONE (Include Area Code) 202-696-4318		22c OFFICE SYMBOL ONR 1142 PT

SYSTEMATICITY AS A SELECTION CONSTRAINT IN ANALOGICAL MAPPING

Catherine A. Clement
Eastern Kentucky University
and
Dedre Gentner
University of Illinois

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 Lise Dobrin, Robert Knippen, Alan Jeung, Maureen Clifford, Laura Kotovsky, Ronald Mawby, Robert Mitchell, Besty Perry, Mary Jo Rattermann, Brian Ross, and Robert Schumacher. Requests for reprints should be sent to Catherine A. Clement, Psychology Department, Eastern Kentucky University, Richmond, KY 40475-3108.

Send proofs to: Dr. Dedre Gentner
University of Illinois
Department of Psychology
603 East Daniel Street
Champaign, Illinois 61820

Short title: Systematicity as a selection constraint.

Approved for public release; Distribution is unlimited.

Approved for	
NTIS GRA&I	☒
DTIC TAB	☐
Unannounced	☐
Justification	
By _____	
Distribution	
Availability codes	
Dist	Availability for
A-1	Unlimited



Abstract

Analogy is often viewed as a partial similarity match between domains. But since between any two domains there are more partial similarities than good analogies, it follows that analogy is selective. Three experiments examined the selection constraints on which relations are mapped between a base and target in an analogy. In Experiment 1 subjects judged two matches to be included in an analogy: an isolated match, and a match embedded in a larger matching system. Subjects preferred the embedded match. In Experiments 2 and 3 subjects made analogical predictions about a target domain. Subjects predicted information that followed from a causal system that matched the base domain, rather than information that was equally plausible, but that created an isolated match with the base. Results support Gentner's (1983, 1989) structure-mapping theory that analogical mapping concerns systems and not individual predicates, and that attention to shared systematic structure constrains the selection of information to include in an analogy.

In an analogy a familiar domain is used to understand a novel domain: to highlight important similarities between the domains or to predict new features of the novel domain. For example, we use our knowledge about water flow to elucidate properties of electric circuitry. Such an analogy can lead to useful inferences, and reveal deep structural features about a domain. But how is an analogical mapping constructed?

Given that an analogy is a partial similarity, it is necessary to select which similarities count in an analogy. There are several levels of selection that need to be made. For example, most researchers agree that in an analogy structural relations are more important than object attributes (see Collins and Burstein, 1989, for a review). Thus, the analogy between a plumbing system and an electric circuit should not tell us that wires need to be hollow. But there is another kind of selection problem: How does a person decide which set of common relations should be preserved? Here, for example, we could preserve relations such as *...it is distributed to persons in a city; resources cost money...* or we could preserve *...degree of pressure determines flow rate*. The present paper concerns constraints on which relational mappings are included in an analogy.

Although researchers in AI have been concerned with how to reduce the space of possible mappings between a base and target domain (see Kedar-Cabelli, 1985a and Hall, 1988), research in psychology has not specifically addressed the question of selection constraints on mapping. Many psychological studies have either focused on simple analogies in which selection of information to map is a trivial problem, or have bypassed the selection problem in describing analogical processing. For example, much research has examined proportional analogies, or four-term analogy problems, e.g. $APPLE : EAT :: MILK : \underline{\hspace{1cm}}$ in which subjects must find a missing term(s) from a set of alternatives, e.g. (WHITE, DRINK, COW, SWEET). Sternberg (1977) has suggested a componential analysis identifying the sequence of steps used in solving such problems. However, he has not given an account of the crucial step of identifying a relation to map. This is chiefly because the selection issue does not arise in these kinds of analogies; they are carefully designed so that only one clear relation is meant to be mapped. Thus, accounts of these four-term analogy tasks have little to say about how

people construct complex, explanatory analogies that contain many mappings, and that may be a continuing source of hypotheses about the target domain.

Research on analogies used in problem solving and learning has also not addressed the selection problem. In most cases, the base domains given to subjects contain only the correct schemas to be mapped, so that again the selection problem does not exist (e.g. Gentner and Gentner, 1983; Gick and Holyoak, 1980, 1983). Typically subjects are able to map information successfully between analogs once a relevant analog is identified. But researchers have not been concerned with specifying mapping processes. Similarly, studies of analogy in learning programming languages (e.g. Rumelhart and Norman, 1981) have shown that subjects can modify known procedures to yield new procedures. But this research too has been criticized for failing to specify constraints on how familiar concepts are modified (Kedar-Cabelli, 1985a).

One of the few psychological models that is specific enough to address the problem of selection constraints is Ortony's (1979) salience imbalance theory of metaphor (analogy and metaphor are both non-literal similarity matches between domains). In Ortony's view, the interpretation of metaphor maps high-salience features of the base domain onto low-salience features of the target domain. Ortony, Vondruska, Foss and Jones (1985) have provided evidence that salience imbalance may be characteristic of metaphors. However, other research indicates that although subjects prefer metaphors that obey salience imbalance (possibly because of discourse conventions), salience imbalance does not specify a metaphor interpretation (Gentner and Clement, 1988). Interpretations of a metaphor cannot be predicted from the relative salience of features in subjects' prior representations of the constituent terms. Thus, salience imbalance may be part of the story, but it does not appear to provide an adequate selection rule for metaphor.

The present study, based on Gentner's (1980; 1983; 1989) structure-mapping account of analogy, investigates how the structure of analogous domains acts as a selection constraint on mapping. According to the structure-mapping view an analogy focuses on interrelated systems shared between two domains, and the relations included in an analogy are those that are embedded in a larger shared structure (e.g. a structure linked by causal relations). The goodness of a component match between two analogous

domains is not dependent just on properties of the match itself, but on the connection of the match to a shared system of matches. The structure-mapping account stems in part from the observation that useful analogies, such as those used in science or education, involve rich, inter-constraining systems of mappings between two domains rather than a set of independent correspondences (e.g. Black, 1962; Gentner, 1980, 1983; Holyoak, 1984; Kittay, 1987). Analogies that underlie our understanding of many everyday concepts also appear to have this coherent structure (Carbonell, 1983; Lakoff and Johnson, 1980). Gentner proposes that analogy is essentially a cognitive device for mapping systematic relational structures from a base to a target domain, and people's attention to relational structure guides the analogical mapping process.

Specifically, on the structure-mapping account, the construction and interpretation of analogy is guided by two kinds of implicit principles: structural principles (such as one-to-one object mapping and structural consistency) and selection principles. The first selection principle is that similarities in non-relational object attributes are irrelevant to the analogy. Analogy is a mapping of relations between objects rather than of object attributes which exist independently of other elements in a given domain. The second selection principle is the systematicity principle, which holds that the relations that are mapped are those that participate in a system of relations governed by higher-order relations that can also be mapped between the two domains. (Higher-order relations, such as causal relations, are relations among relations.) Thus, similarities in lower-order relations that are isolated, i.e. relations that are not linked to a larger matching system, are not part of the analogy. The present study is aimed at testing the use of the systematicity principle as a constraint on what information is included in an analogy.

Evidence for the Importance of Systematicity in Analogy

To our knowledge there is no direct evidence that systematicity can act as a selection constraint on mapping. However, some support for this hypothesis can be found in computational models of analogy, and in psychological research that provides evidence for the general role of systematic relational structure in analogical processing. Many computational models have utilized structural constraints on mapping (Burstein, 1983; Falkenhainer, Forbus and Gentner, in press; Thagard and Holyoak, 1988; Winston, 1980; see Hall, 1988

for others). The Structure-mapping Engine (SME) (Falkenhainer, Forbus and Gentner, 1988), which implements the specific principles of Gentner's structure mapping account described above, has produced interpretations of several analogies consistent with those of human reasoners (Skorstad, Falkenhainer, and Gentner, 1987).

Three lines of psychological research are also relevant. First, Gentner & Landers, 1985 and Rattermann & Gentner, 1987, have shown that judgments of the soundness of a match between two domains (that is, whether an analogy would yield justifiable inferences) are positively related to the systematicity of the match. Second, Gentner and Schumacher (1986) and Schumacher and Gentner (1988) found that systematicity can facilitate accurate analogical mapping. In their research, subjects were taught a device model and then asked to transfer their knowledge to an analogous device. Subjects were able to achieve accurate transfer in substantially fewer trials when their initial model possessed systematic structure than when it did not (even though the same procedures were used in both cases). Gentner and Toupin (1986) found similar results in a study of children's ability to map a story plot between two sets of characters. Finally, in a study of analogical problem solving, Holyoak and Koh (1987) varied the degree of structural correspondence between analogous problems. Consistent with the systematicity principle, they found that when the causal system describing the initial problem state differed between the base and target problems, subjects were less likely to spontaneously transfer the solution linked to the base causal system, even though the solution would have been adequate in the target problem.

In sum, previous research implicates systematicity in the evaluation of analogies and in the efficacy of transfer. However, this research has not provided a direct test of the claim that systematicity acts as a selection filter that guides the choice of information to map between domains. The issue of selection constraints in general has not been specifically addressed. That is, researchers have not manipulated alternative features to be mapped between complex domains, in order to isolate the basis for subjects' mapping choices.

Experiments to Address Systematicity as a Selection Constraint

We describe three experiments which examine systematicity as a selection constraint on analogical mapping. Our experiments looked

separately at two distinct processes that constitute mapping information from a base to a target domain: (a) *matching* existing information in the base and target and (b) *inferring* new information about the target that follows from the analogy with the base domain.

Our first experiment examined whether systematicity constrains the matching process. We constructed a set of novel analogies that were specifically designed to test whether mapping choices are guided by systematicity. In these analogies, subjects could choose which of two facts (lower-order relations) to map from a base to a target domain. In all cases, both facts were equally acceptable when considered as independent matches. But the two facts differed in whether they were part of a shared systematic relational structure (specifically, a shared causal structure). Under the systematicity hypothesis, subjects should prefer those matches that are embedded in a larger system of matches, and they should reject those matches that are relatively unconnected. Thus, mapping decisions should not be based exclusively on a matching fact itself. Decisions should be determined by a match's interdependence with other information shared by the two domains. With our materials, we could distinguish the effect of the lower-order facts themselves on mapping choices from the effect of the higher-order embedding of these facts.

Our second and third experiments examined whether systematicity constrains inferences carried over from the base domain to the target. Subjects were asked to make an analogical prediction about a target domain. Two candidate facts were present in the base domain that were equally plausible as inferences about the target. However, only one fact was linked to a causal system shared by the base and target. If a systematicity principle governs analogical inference, subjects should identify matching systems in the two domains, and map those relations that are present in the base system but missing in the target system. Subjects should not select just any base fact that could be plausibly inferred in the target, but should infer a fact that follows from a shared interdependent set of relations.

Before describing the structure of our materials in more detail, two further criteria in developing our tasks should be mentioned. First, to preserve a realistic degree of complexity, the situations described were fairly rich in information. Second, in the interests of generality, we did not wish to limit

ourselves to cases in which goal relevance was a possible selection constraint on mapping. Thus, subjects did not have to map information in order to solve problems or prove points in the target domain. Finally, it was important that subjects' answers be governed by use of the analogy rather than simply by prior beliefs about the target. Therefore we developed analogies between novel, fictional domains rather than using real world analogies.

Design of the materials. The analogies developed each consisted of a base and target passage describing novel objects or organisms on fictional planets. Each passage included two chief paragraphs. One paragraph described a causal structure that matched between the base and target, and the other described a causal structure that did not match. Subjects had to make mapping choices between key facts that were embedded either in the matching or in the non-matching causal structures. The key facts themselves always matched between the base and target.

To help understand the materials, we first describe in abbreviated prose the base passage for one analogy. This passage described creatures called *Tams* who live on a distant planet. (The actual passages were about one page long and were written in the style of an encyclopedia article.)

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 must relocate. At this time they stop using their underbellies.

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, a grown Tam's underbelly cannot function on a new rock.

Each paragraph describes a causal structure consisting of a *key fact* and a causal antecedent governing that key fact. The two key facts in this base domain are (1) *the Tams sometimes stop using their underbellies*, and (2) *the Tams' underbellies cannot function on new rocks*.

This base domain, and the analogous target domain called "The Robots", which describes robots who use probes to gather data from planets, are outlined in Table 1. The left column of the table shows the two causal structures of the base, with each key fact shown in italics, and the causal antecedents shown in bold face. The middle column of the table shows

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 •Exhaust minerals in one spot and must relocate on the rock •<i>So stops using underbelly</i>	•Gather data with probes •Exhaust data in one place and must relocate on the planet •<i>So stops using probes</i>	•Gather data with probes •Internal computers over-heat when gather a lot of data •<i>So stops using probes</i>
•Born with inefficient underbelly •Underbelly adapts and becomes specialized for one rock •<i>So underbelly can't function on new rock</i>	•Designed with delicate probes •Robots cannot pack probes to survive flight to a new planet •<i>So probes can't function on new planet</i>	•Designed with inefficient probes •Probes adapt and become specialized for one planet •<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.

Version 1 of The Robots which, like the base domain, contains two causal structures. Importantly, the key facts in each structure of the target match the key facts in the base: (1) *the robots sometimes stop using their probes* and (2) *the probes cannot function on new planets*. However, although both key facts in the target match the base domain, only the first key fact is linked to a causal system that also matches the base domain. We will call this the *shared-system* key fact. We predicted that subjects should prefer this shared-system fact in mapping. Although the other key fact matches the base domain, it is linked to a causal antecedent that does not match the base. (It should be noted that for ease of reporting, the matching causal structures in Table 1 are described in language that is similar at the surface level. In the actual passages the matching key paragraphs were written in more domain-specific language, and we tried to avoid extensive similarity in sentence structure within key paragraphs.)

To avoid confoundings with particular content there were two versions of a target domain, as shown in Table 1. In Version 1 of the target, the cause for the first key fact matches the base domain, but the cause for the second key fact does not. Version 2 of the target contains the same key facts but reverses which key fact is linked to a shared causal system -- in this version, the second key fact is the shared-system fact. This ensures that mapping preferences for shared-system 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 were not due to a preference for a particular fact in the context of a particular version of the target. In this task control subjects read only the target stories and judged which key fact was more important for each story. If these "target-only" subjects show no bias toward 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 shared-system key facts would be preferred over different-system key facts, even though both key facts in themselves match the base equally well. In Experiment 2 and 3, the materials were altered slightly and we examined subjects' inferences about the target domain, given its analogy with the base domain.

Experiment 1

Method

Overview

Four novel analogies were created each consisting of a base and target passage designed according to the structure described above. (The four analogies are presented in Appendix A.) Subjects in the Analogy group read both the base and target passage for each analogy, and subjects in a control group (Target-Only group) read only the target passage. For each analogy, after learning the passage(s) the two groups were given parallel tasks. These tasks required subjects to rate or choose among the two key facts in the target passage. As discussed above, each key fact matched the base passage, and each was embedded in a causal system. The manipulation was that this causal system matched between the base and target for the *shared-system fact* but did not match for the *different-system fact*. The Analogy subjects evaluated how well the two key facts contributed to the analogy. The Target-Only subjects evaluated the importance of the facts to the target passage. We predicted that, in their ratings and choices, Analogy subjects would prefer the shared-system fact, and Target-Only subjects would show no preference for either fact.

The experimental factors were Fact-Type (Shared-system v Different-system), a within-subjects variable, and Condition (Analogy v Target-Only), a between-groups variable. Additional factors were Passage (the four target passages), a within-subjects variable, and the counterbalancing factor Version-Set (two versions), a between-groups variable (see Table 1). "Version-Set refers to which version of the target passage was given. Within each Condition there were two subgroups of subjects, each received a different version of each of the four target passages.

Subjects

The subjects were 48 paid undergraduate students at the University of Illinois. Half were assigned to the Analogy Condition and half to the Target-Only Control Condition.

Procedure

Subjects were run in groups of three to six. The experimenter read aloud the instructions for each task. (The instructions were also given in writing.)

Subjects had as much time as they needed. Sessions for Target-Only subjects lasted approximately one hour. Sessions for Analogy subjects, who had more material to read, lasted approximately two hours.

Task for Analogy subjects. Our interest was in subjects' judgments of how well each key fact in the target contributed to the analogy with the base domain. We used a Rating Task and a Choice Task to assess these judgments. Since the materials are complex, we first gave subjects learning tasks to be sure that they understood the materials.

Learning tasks. The subjects' first task was to read the base passage carefully. Then, they were told that the next passage they were to read - the target passage - was analogous to the base passage. For the first analogy in the session, before reading the target, subjects were given an example of what we meant by "analogous". They were told "*We can say that plant stems and drinking straws are analogous. Even though they are two different things, they are both used as channels to bring liquid nutrients from below up to a living thing....*".

After reading the target, subjects were asked to match the objects that corresponded between the base and target domains. They were given a list of the approximately four or five central objects in the base passage. For example, in the analogy described in Table 1, the list included *Tams, rocks, underbellies, minerals, internal organs*. Subjects were to identify which objects in the target (if any) corresponded to each base object. Thus, the answers here were *Robots, planets, probes, data, and microcomputers*. The first correspondence, that between the subjects of the stories (e.g. the *Tams* and the *Robots*) was always filled in to illustrate the task. In order to avoid biasing subjects, the stories were designed so that there was no inconsistency in the object correspondences relevant to the shared and different-system facts.

Subjects then wrote out the ways in which the two domains were and were not analogous. The instructions were, for example: *Describe what is and is not analogous about the Tams and the Robots. When two things are analogous some things fit better than other things. Thus, some facts about the Robots contribute to making the analogy with the Tams a good analogy and some facts do not contribute to the analogy. Describe those facts that support the analogy and describe the facts that do not support the analogy.* Subjects

were encouraged to refer back to the passages if necessary. Subjects were given no feedback. The purpose of these preliminary learning tasks was simply to lead subjects to process the analogies thoroughly.

Rating Task. After the learning tasks, subjects were given two experimental tasks in which they evaluated how well the two key facts from the target contributed to the analogy. First, in a Rating Task, the two key facts were presented on a separate page and subjects rated the degree to which each fact supported the analogy on a scale of 1 to 7. For example, the italicized facts from the Target in Table 1 were presented, stated as in the passage:

(1) *Sometimes a robot must shut down its probes and*

(2) *Robots cannot go from one planet to another to gather data .*

The rating instructions, which were given in writing and read aloud were, for the Robots/Tams example: *When two things are analogous some things fit better than other things. Some facts support the analogy better than other facts. Assume you are in a debating match and you have to defend the claim that the Robots are analogous to the Tams. Below are two facts about the Robots. Please rate them on a scale of one to seven according to how well they support the claim.* Subjects were told to give a 7 to a statement that contributed well to making the analogy a good analogy and a 1 to a statement that did not contribute at all to the analogy.

Choice Task. Next subjects were given the two key facts again on a new page. Subjects were asked to choose between them and to explain their choice: *Still assuming that you are in a debating match, reread the two facts, they are printed again below. Which statement best supports the claim that the Robots and Tams are analogous?* Subjects were required to give a brief written explanation for their choice of one fact and their rejection of the other.

During both judgment tasks, subjects were encouraged to re-examine the passages as needed. Thus, there were no memory requirements.

Tasks for Target-Only Subjects. The Target-Only subjects read only the target passage. To ensure that these subjects thoroughly processed the passage, they were asked to summarize it. Subjects were told that their summaries should include discussion of the key objects from the target passage (these objects were listed for them). Following this learning task, Target-Only subjects were given the same two key facts given to Analogy subjects and asked to perform judgment tasks in the same order as the

Analogy subjects. First they *rated* on a scale of 1 to 7 how important each fact was to the target passage. Following this, they were presented with the facts again and were asked to *choose* which fact was most important to the target passage. Subjects were asked to explain their choices. (We were not interested in these justifications, but included this question in order to maintain similarity to the task given to Analogy subjects.)

Results

Rating Task

Table 2 shows the mean ratings for shared-system and different-system key facts for each target passage. As predicted, when asked to rate how well each fact contributed to the analogy, Analogy subjects rated shared-system key facts higher than different-system key facts (overall $M=6.28$ and $M=4.89$, respectively). Thus systematicity governed the Analogy subjects' preference. In contrast, the Target-Only subjects gave equivalent importance ratings to the two types of key facts, ($M=5.22$ and $M=5.31$), indicating that the Analogy group ratings are not due to differences in the *importance* of the key facts in the two versions of the target. To compare the ratings of each type of key fact in each condition a $2 \times 2 \times 4 \times 2$ repeated-measures ANOVA was conducted. The experimental factors were Fact-Type (Shared-system vs. Different-system) and Condition (Analogy group vs. Target-Only group). Additional factors were Passage (the four target passages) and the counterbalancing factor Version-Set (1 vs. 2). Both Subjects and Passage were treated as random variables in the analysis (see Clark, 1973).

There was a significant main effect for Passage, indicating that overall, ratings varied across passages, $F(3,132)=3.09$, $p < .05$. There were no other significant main effects. Significant two-way interactions were found between Condition and Passage, $F(3,132)=3.57$, $p < .05$, and Fact-Type and Passage, $F(3,132)=3.00$, $p < .05$. These interactions simply indicate that averaged over Fact Type, Condition had an effect that varied across Passage, and averaged over Condition, Fact-Type had an effect that varied across Passage.

The key prediction is an interaction between Condition and Fact-Type. This interaction was significant, $F(1,44)=13.99$, $p < .001$, indicating that the difference in ratings of shared-system and different-system key facts was different for the Analogy and Target-Only groups. When this interaction was analyzed, the simple effect of Fact-Type within the Analogy group confirmed

Table 2
Experiment 1. Results of Rating Task: Mean Ratings^a of Shared-system (correct)^b
and Different-system Key Facts for each Passage

Passage	Group							
	Analogy (n=24)				Target-Only (n=24)			
	Shared System		Different System		Shared System		Different System	
	M	SD	M	SD	M	SD	M	SD
1	6.12	.95	4.96	1.88	4.79	1.72	4.96	1.68
2	6.04	1.46	5.37	1.79	5.42	1.50	5.96	.99
3	6.42	1.10	4.33	1.86	5.62	1.28	5.42	1.06
4	6.54	.93	4.87	1.62	5.04	1.68	4.92	1.64
Overall	6.28	.96	4.88	1.32	5.22	.80	5.31	.57

^a Facts were rated on a 1-7 scale.

^b In this and the following tables we will use the term "correct" to remind the reader that it is the shared-system key fact that would be preferred by Analogy Subjects if they were reasoning according to the systematicity hypothesis.

that the mean rating for shared-system facts by Analogy subjects is significantly higher than the mean rating for different-system facts, $F(1,11)=10.44$, $p < .01$. This difference does not hold for the control group. No other two-way interactions were significant.


There is no evidence that the key interaction between Condition and Fact Type was specific to particular passages or to a particular Version-Set. That is, there was no triple interaction between Condition and Fact Type and either of these other factors. Finally, a significant triple interaction was found between Version-Set, Fact Type and Analogy, $F(3,132)=19.77$, $p < .001$. This simply indicates that, overall, a difference in ratings to the two fact types was found that depended on the version of particular analogies. No other effects were significant. In sum, the Rating Task responses support the prediction that systematicity would govern Analogy Subjects' preferences for particular matches between analogous domains.

Choice Task

The results for the Choice Task are consistent with those of the Rating Task. As predicted, Analogy subjects most often chose the shared-system fact as best contributing to the analogy. In contrast, Target-Only subjects showed no preference for this fact over the different-system fact in their importance judgments. To test the difference between groups, subjects were assigned a score for the number of choices of shared-system key facts across the four analogies (giving a possible score of 0-4). A 2×2 ANOVA (Condition $\times$ Version-Set) revealed a significant main effect of Condition, indicating that the mean for Analogy subjects (3.17) is significantly greater than the mean for Target-Only subjects (2.0), $F(1,44)=18.43$, $p < .001$. There was also a main effect of Version-Set, $F(1,44)=6.02$, $p < .05$, indicating that, overall, the number of choices of shared-system facts varied across Version-Set. As expected, the interaction between Condition and Version Set was not significant.

Figure 1 shows that the pattern held for each of the four passages. The Analogy subjects chose the shared-system fact 67% to 92% of the time, while Target-Only subjects chose it 42% to 54% of the time. Fisher exact tests reveal that the difference between groups is significant for three of the four passages: ($p < .05$, 1 tailed). In sum, these results show that the Analogy subjects viewed

Figure 1
Experiment 1 Choice Task: Percent of Subjects in each Condition Choosing the
Shared-system (correct) Fact for each Passage



the matching fact embedded in a shared causal system as better support for the analogy than an equally good match that was not so embedded.

The Rating and Choice Task results support the position that analogical processing is not a matter of concatenating independent matches, but of finding a connected system of matches. The goodness of a particular match is influenced by its neighboring matches which form a mutually constraining system. We next asked whether subjects were explicitly aware of the higher-order constraints that appear to govern their choices.

Justifications of Choices

In the Choice Task, Analogy subjects were asked to write a brief justification for their responses. They were to explain (a) why they thought their chosen fact supported the analogy best and (b) why the rejected fact did not support the analogy as well. The justifications for choices and rejections were separately scored. Our main interest was in whether subjects would show an explicit concern for the shared higher-order embedding of key facts. Thus, we scored justifications as to whether subjects referred to the causal systems supporting or not supporting a key match, or referred only to the lower-order match of the key fact itself. Responses of all subjects were coded, regardless of whether they chose the shared-system or the different-system key fact.

Two judges coded subjects' justifications for choosing a fact into four categories which can be ordered according to the level of focus:

1. Focus on similarity in causal structure (higher order similarity): Subjects stated that the cause of the chosen fact was similar in the base and target, or described the similar cause for the chosen fact. For example, if subjects chose the key fact *Robots cannot go from one planet to another to gather data*, they might give a justification such as, *Robots cannot go to another planet because they become adapted to one planet, and Tams cannot go to another rock because they are adapted to one rock*; or *The Robots cannot go to another planet for the same reasons that Tams cannot go to another rock*.

2. Focus on the similarity of chosen fact (lower order similarity): Subjects simply noted the direct similarity between the base and target key facts, e.g. *Robots cannot gather data on another planet and Tams cannot function on another rock*.

3. Focus on the chosen fact alone (importance of the target fact): Subjects asserted that the chosen fact is the most important fact.

4. Other: Gives any other response.

Subjects' justifications for rejecting a fact were classified into the same four categories, except that now justifications focused on: (1) *dissimilarities* in the cause for the rejected fact to the cause for the corresponding fact in the base, (2) the *dissimilarity* of the rejected fact to the corresponding fact in the base, or (3) the *lesser* importance of the rejected fact.

The important contrast is between the first and second categories. The first category reflects subjects' concern for the higher-order embedding of the matching fact, whereas the second category reflects only a concern for the lower-order match itself. The third and fourth types of responses could not be clearly categorized in this respect.

Table 3 shows the justification data. The left side of the table shows the distribution of responses when subjects made the predicted choice: that is when the shared-system fact was chosen (and thus the different-system fact was rejected). The right side shows responses when the different-system fact was chosen. The top half of the table shows justifications for choosing a fact, and the bottom half shows justifications for rejecting the other fact. Inter-rater agreement in classifying choice justifications was .86 before discussion, and agreement in classifying rejection justifications was .82 before discussion. (Data shown throughout the paper reflect final scoring decisions.)

As shown in Table 3, the results confirm that subjects who chose the shared-system fact did so because shared systematicity was important. Looking at the left side of the table, we see that the majority of these subjects' justifications both for their choice of shared-system fact, or their rejection of the different-system fact were concerned with shared causal information. Relatively few justifications were concerned only with the lower order match itself. Some examples of these justifications in which subjects explicitly mention the common causal constraints are shown in Table 4. (Note that these examples are all from subjects who gave the predicted response. The examples of choice and rejection justifications are not from the same subjects.)

Subjects did not always give higher-order justifications when choosing the shared-system fact. This suggests that people may sometimes operate on the basis of shared relational structure without necessarily being able to

Table 3

Experiment 1. Justifications by Analogy Subjects for Facts Selected in Choice Task:
Proportion and Frequency of Shared-system and Different-system Choices
(and Rejections)

<u>Justification Type:</u>	<u>Chosen Fact</u>			
	Shared-System Fact		Different-System Fact	
	(Correct)		(Incorrect)	
	<u>Prop.</u>	<u>Freq.</u>	<u>Prop.</u>	<u>Freq.</u>
Similar cause for chosen fact	.73	(56)	.35	(7)
Similarity of chosen fact	.16	(12)	.20	(4)
Importance of chosen fact	.04	(3)	.30	(6)
<u>Other</u>	<u>.07</u>	<u>(5)</u>	<u>.15</u>	<u>(3)</u>
Total	100	(76)	100	(20)

<u>Justification Type:</u>	<u>Rejected Fact</u>			
	Different-System Fact		Shared-System Fact	
	(Correct)		(Incorrect)	
	<u>Prop.</u>	<u>Freq.</u>	<u>Prop.</u>	<u>Freq.</u>
Dissimilar cause for rejected fact	.56	(43)	.05	(1)
Dissimilarity of rejected fact	.20	(15)	.30	(6)
Importance of rejected fact	.04	(3)	.20	(4)
<u>Other</u>	<u>.20</u>	<u>(15)</u>	<u>.45</u>	<u>(9)</u>
Total	100	(76)	100	(20)

Note. Out of 96 opportunities, subjects chose the shared-system fact (the predicted response) 76 times and the different-system fact 20 times. For each choice, justifications are shown for both the chosen fact and the rejected fact.

Table 4

Experiment 1. Examples of Justifications for Shared-system (correct) Choices^a

Justifications for choosing the shared system fact: *Robots cannot go from one planet to another to gather data.*

1. ...there is a direct comparison. Robots cannot go and be effective from one planet to another because they have adapted to one particular planet - Tams cannot move from one type of rock to another because they also have adapted.
2.the robots cannot move planets after it has developed probes and the Tam cannot move to a new rock once it has developed its underbelly.

Justifications for rejecting the different system fact: *Robots sometimes shut down their probes.*

1. ...the reason the robot shuts down is because it has received too much information whereas the reason the tam must stop sucking and move on is because it is no longer receiving enough minerals.
2. Although a robot must shut down its probes sometimes and a Tam's underbelly sometimes no longer functions, they are not as similar. The robot shuts down when it has too much of that which it is seeking (data) and the Tam shuts down when it does not have enough of what it is seeking (minerals).

^aResponses are for the Tams/Robot analogy, given target version 2.

articulate their reasoning. In this connection, it is interesting that people seemed to more often give higher-order justifications for acceptance of shared-system facts than for rejection of different-system facts, even though these represent the same set of ("correct") choice responses. This pattern is intelligible if we assume that there is some processing cost for negativity (e.g. Clark and Chase, 1972). Conceivably it is easier in the positive case for subjects to introspect and articulate the higher-order causal connections that govern their choice.

When subjects ("incorrectly") chose the different-system fact (right side of Table 3), their justifications were evenly distributed among the categories, focusing either on causal information, the similarity of the fact itself, or the importance of the chosen fact. Interestingly, when subjects mentioned causal information in justifying the different-system choice, they often invented or imputed a cause in the target that was similar to that in the base, even though this interpretation went beyond and/or was inconsistent with the information provided. For example, in Version 1 of the target domain *The Robots*, the robots are described as not specialized for a particular environment (see Table 1). However, some subjects read the passage as saying that the robots were specialized, thus creating a match in causal information with the base domain. In sum, subjects' Choice Task justifications indicate that links to a shared causal system guided which lower-order matches were judged to best belong to an analogy.

Discussion

As predicted, systematicity appears to constrain which similarities are selected for an analogy. In both their ratings and choices, given two equally good lower-order matches between the base and target, Analogy subjects preferred whichever matching fact was embedded in a matching causal system. Let us review this finding in more specific terms. Analogy subjects were given two lower-order relations, T1 and T2, which both matched key relations in the base domain. Subjects had to choose which match, B1 -> T1 or B2 -> T2, best belonged to the analogy. Note that each key fact was embedded in some causal system in both the base and target. Thus, the immediately dominating relation was cause in all cases. The manipulation was in whether the causal antecedents matched between domains: e.g.,

Base: CAUSE (X, B1) and CAUSE (Y, B2)

Target: CAUSE (X', T1) and CAUSE (Z, T2)

Analogy subjects selected the fact with like embedding - in this case, T1, whose causal antecedent, X' matches the antecedent (X) of B1 in the base.

Subjects' response justifications further confirmed their concern for the higher-order embedding of matching facts. Finally, the fact that control subjects, who saw only the target domain, showed no preference for the shared-system facts indicates that the Analogy responses are not due to the relative importance of the key facts in the target domain.

These results tell us that analogical matching is not merely a feature-by-feature decision. Analogical matching concerns systems of predicates, not individual predicates.

Experiment 2

Experiment 1 focused on matching known facts between two analogous domains. The results indicated that systematicity constrains which matching facts are selected as belonging to an analogy. Experiment 2 focuses on the use of analogies to predict new facts about a target domain. We asked whether systematicity would guide the inference process in which relations in the base domain are carried over to the target domain as candidate inferences about the target. In this experiment, instead of only asking subjects to choose among specified mapping possibilities, we included a task that allowed subjects themselves to find information to map from the base to the target. Specifically, after reading the base and target, subjects were asked to predict a fact about the target domain that followed from the analogy with the base. If systematicity guides analogical inference, then subjects' new predictions should center on facts that (if true in the target) would link into a causal structure shared by the two domains.

Method

Overview

Subjects read base and target passages that were similar to those of Experiment 1 and that again follow the design described in Table 1. However, for this experiment the two key facts were removed from a target passage; thus the italicized facts shown in the table were present only in the base passage. Of course, although these facts were removed from the target they were

plausible in that domain. The target domain still included the antecedent information that could potentially cause each key fact. For example, in the target shown in Table 1, in Version 2, the reader learns that the robots' computers overheat, and learns that the robots become adapted for one planet. As in Experiment 1, this antecedent information does not match the base domain in one case, but does match the base domain in the other case.¹

In this experiment subjects were asked to make a prediction (and later to judge given predictions) about the target domain. Our question was whether they would predict one of the two key facts omitted from the target, and more importantly, if so, which one. Thus, given the analogy in Table 1, we were interested in whether subjects would predict either that *the robots sometimes stop using their probes* or that *the robots fail to function on a new planet*. If subjects are simply trying to predict facts about the target that correspond to individual facts in the base, or that are plausible in the target, they should be equally likely to predict either of the omitted key facts. If, however, subjects are guided by systematicity in making predictions, they should predict the one key fact for which there was a matching antecedent in the base and target, i.e. the shared-system fact. Thus, given Version 2, subjects should predict that *robots' probes cannot function on a new planet*.

As before, in order to avoid confoundings with particular content, two versions of each target were used. Which key prediction would follow from a matching antecedent was counterbalanced across the two versions (see Table 1). (The two versions of the target passage for the Tams/Robots analogy are presented in Appendix B).

It was important that both the shared-system and different-system predictions be equally easy to construct in the target domain. We took several

¹ Aside from removing the key facts, some further modification of the materials used in Experiment 1 was sometimes necessary. The goal was to have the key facts in the target be plausible in the target and predictable given the analogy with the base domain. However we did not want these facts to be obvious or necessarily true given the target alone. Thus, to ensure that the key facts were not obvious consequences of the antecedent information in the target passages, we sometimes rewrote or removed some of the original antecedent information. However, enough antecedent remained to make the key facts plausible, and where desired, to allow a match with antecedent information in the base Passage. The specific content of some of the analogies was also modified for clarity.

steps to ensure this. First, given all the information in the base and target, subjects could easily identify the appropriate object correspondences and many relational correspondences between the two domains. Therefore, new analogous relations in the target corresponding to either key fact in the base could be easily created. (Inspection of the passages in Appendix B should make this clear). Second, as already mentioned, the target included antecedent information making both predictions plausible. Finally, to control for any differences in the ease or plausibility of constructing the two key facts in the target, the design was counterbalanced as discussed above, and, a Target-Only group created and judged predictions based only on reading the target passages.

In sum, in this experiment we asked Analogy subjects to make a prediction about the target domain based on the analogy with the base domain. Since much of the base and target passages already matched, subjects were constrained in their possible new predictions. However, they had two particularly plausible choices: there were always two facts presented in the base that were not present in the target for which parallel target facts could be straightforwardly constructed. Subjects could potentially carry over either of these key facts from the base domain. However, only one of them would follow from a causal system shared by the base and target.

Subjects

Subjects were 32 paid undergraduate students at the University of Illinois. Half were assigned to the Analogy Condition and half to the Target-Only Condition.

Procedure

Task for Analogy subjects. For each analogy, subjects first had learning tasks identical to those in Experiment 1. They then were given three experimental tasks requiring them first to make a prediction about the target domain and then to judge some possible predictions.

Prediction Task. After performing the learning tasks, 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 passage. The specific instructions (which were given in writing and read aloud) were, for example: *Because the Robots are analogous to the Tams, we might add some information to the story about the Robots. Aside from the*

information already stated about the Robots, the Tam story can suggest some predictions. Look again at the two stories; then in the space below predict one thing that might be true of the Robots that is suggested by analogy with the Tams. It was further emphasized that this should be a prediction about what might be true about the Robots and not something already written explicitly in the passage, and that this should be a prediction based on the analogy and not on the target Passage alone.

Prediction Rating Task. After making their own predictions, subjects rated possible predictions about the target according to how well they followed from the analogy. Subjects were told: Professor Zee answered the same question you just did. She said that the analogy with the Tams suggested two things about the Robots. She made the two predictions stated below. Neither of these were explicitly written about the Robots though both of these predictions are equally plausible. But, are these predictions suggested by the analogy with the Tams? Rate these predictions on a scale of 1 to 7 according to how well they follow from the analogy. (Note: one of her predictions may be the same as the one you just made).

The two key facts that appeared in the base but not in the target were then given to subjects as predictions about the target. For example, subjects were given:

(1) A Tam sometimes stops using its underbelly. Similarly, sometimes a robot shuts down its probes.

(2) Tams cannot go from one rock to another. Similarly, Robots cannot go from their current planet to another planet to gather data.

Prediction Choice Task. Finally, subjects were again presented with the two predictions on a new page and asked to Choose the prediction that is best suggested by the analogy. Subjects were also asked to write explanations for choosing one prediction and rejecting the other.

For all tasks subjects were told to refer back to the base and target passages while deciding on a response.

Task for Target-Only subjects. Subjects who only read the target domain performed Prediction, Rating and Choice tasks which paralleled the tasks given to Analogy subjects. Before completing these tasks subjects were asked to summarize the target by answering general questions that directed them to the essential parts of the Passage. e.g. *Explain what happens when the data*

gathered by the robots are no longer new. These questions were intended to ensure that subjects thoroughly attended to the parts of the passages describing the causal structures and key facts.

Prediction Task. Following the learning task subjects were asked to make a prediction about the target passage. Instructions were similar to those given to Analogy subjects, except that subjects were told to make a prediction about the target *given the information in the target Passage.*

Prediction-Rating Task. After making their own predictions, Target-Only subjects, like Analogy subjects, were told that a fictional person had made two predictions about the target Passage. Subjects were given the same two predictions given to Analogy subjects, except that there was no mention of the analogous facts from the base domain:

(1) *Sometimes a robot shuts down its probes.*

(2) *Robots cannot go from their current planet to another planet to gather data.*

They were asked to rate these according to *how well they followed from, or were predicted by*, information in the target passage.

Prediction-Choice Task. Finally, subjects were given the predictions again and asked to *choose which claim best follows from the information in the passage*, and to explain their choices. (As in Experiment 1 these explanations were not coded and will not be mentioned further.)

Results

Prediction Task

Two judges, who were blind to the condition of subjects being scored, grouped subjects' predictions into three categories²: (a) predictions of key facts that would follow from the shared causal system in the base and target (b) predictions of key facts that would follow from the different-system (c) all other predictions about the target. Inter-rater agreement before discussion was .88. Disagreements were readily resolved. A few subjects (.08 of the responses) predicted both the shared-system and different-system fact. These responses were discarded and omitted from further analysis.

As predicted, Analogy subjects most frequently predicted the shared-

²In some cases the content of subjects responses would indicate which condition the subject was in. However, one judge was blind to the experimental manipulation.

system rather than the different-system key fact. Table 5 shows that, across the four analogies, .53 of the Analogy subjects' responses were predictions of the shared-system fact, and only .12 were predictions of the different-system fact. As shown, the same pattern of results holds for each of the four passages used. Thus, although both predictions were clearly possible, Analogy subjects made the inference that was connected to a larger matching structure.

Two examples of shared-system predictions by Analogy subjects for the Tams/Robots analogy follow. These subjects were given Version 2 of the Robots passage, in which the shared-system prediction was that Robots cannot go from one planet to another to gather data.

1. *"I would predict that once the robots were specialized they would be unable to probe for data on other planets than what they were used to just as the Tams would not have the right textured underbelly for a new kind of rock."*
2. *"The robot may eventually be strictly unable to switch planets as the Tams cannot switch rock types."*

The frequency of shared-system predictions among Analogy subjects cannot be attributed to a bias in the materials. The responses of the Target-Only subjects indicate that the shared-system predictions were not highly salient or obvious predictions in the target domains. Not surprisingly, the most frequent response for the Target-Only subjects was to predict information other than the key facts. Some of these were rather creative. For example, two Target-Only subjects given Version 2 of the target predicted:

1. *"Since the robots are so sensitive to the different planets and they need to develop their own probes which related to that particular planet it may be predicted that the robots have trouble analyzing data and make incorrect assumptions and conclusions."*
2. *"Robots are able to control the spaceship which takes them from planet to planet."*

When Target-Only subjects did predict one of the key facts, their responses were evenly distributed between the two types of key facts.

To assess whether the preference for shared-system predictions among Analogy subjects was reliable, subjects were scored for the number of shared-system predictions minus the number of different-system predictions across the

Table 5
Experiment 2. Results of Prediction Task: Proportion of Subjects in each Group
Predicting Shared-system (Correct) and Different-system Key Facts for each Passage

Passage	Group					
	Analogy (n=16)			Target-Only (n=16)		
	Shared System	Different System	Other	Shared System	Different System	Other
1	.44	.06	.44	.06	.25	.69
2	.62	.06	.25	.06	0	.94
3	.44	.19	.25	.06	.12	.81
4	<u>.62</u>	<u>.19</u>	<u>.12</u>	<u>.31</u>	<u>.12</u>	<u>.56</u>
Overall	.53	.12	.27	.12	.12	.75

Note. A few Analogy subjects predicted both the shared-system and different-system fact (.08 of all responses). These responses are not shown here.

four analogies. (Note that shared system and different system predictions do not exhaust the possible responses.) A 2 x 2 ANOVA (the factors were Condition and the counterbalancing factor Version-Set) showed that the mean difference for Analogy subjects was significantly higher than the mean for Target-Only subjects ($M=1.63$ and $M=0$ respectively, $F(1,28)=20.05$, $p < .001$). The only other significant result was the main effect for Version-Set, $F(1, 28)=9.61$, $p < .01$, simply indicating that, overall, the number of shared-system predictions varied across Version-Set. No interactions were significant. Thus, a preference for shared-system predictions was found among Analogy subjects but not Target-Only subjects and this was true for each Version-Set.

Prediction-Rating Task

The mean ratings of each fact type for each target passage are shown in Table 6. As predicted, when asked to rate key facts according to how well they were predicted by the analogy, Analogy subjects gave higher ratings to shared-system than to different-system facts (overall, $M=6.03$ and $M=4.41$, respectively). The ratings of Target-Only subjects, who rated key facts according to how well they were predicted by the target Passage, showed no such difference (and if any showed the reverse preference) indicating that the materials were not biased in favor of the shared-system prediction (overall, $M=4.32$ and $M=5.22$ respectively for shared and different-system facts).

As in Experiment 1, a 2 x 2 x 4 x 2 repeated measures ANOVA was conducted to test the difference in ratings of shared-system and different-system facts in each condition. In addition to the experimental factors Condition and Fact Type, the analysis also included Passage and the counterbalancing factor Version-Set. Both subjects and Passage were treated as random variables in this analysis.

As expected, no main effects were significant. The predicted interaction between Condition and Fact-type was significant $F'(1,9)=10.99$, $p < .01$. The analysis of the simple effect of Fact-Type within the Analogy condition confirmed that Analogy subjects gave significantly higher ratings to shared-system facts than to different-system facts $F'(1,12)=7.86$, $p < .05$. This difference was not found for Target-Only subjects. No other effects tested by the overall ANOVA were significant.

Table 6
Experiment 2. Prediction-Rating Task: Mean Ratings^a of Shared-system (Correct)
and Different-system Predictions for each Passage

Passage	Group							
	Analogy (n=16)				Target-Only (n=16)			
	Shared System		Different System		Shared System		Different System	
	M	SD	M	SD	M	SD	M	SD
1	6.43	1.50	3.69	2.57	4.00	2.31	4.50	2.66
2	5.50	1.93	4.44	2.13	5.31	2.02	5.56	1.46
3	6.12	2.03	4.75	2.14	3.50	2.03	5.87	1.36
4	6.06	1.44	4.75	1.88	4.50	1.97	4.94	1.77
Overall	6.03	1.16	4.41	1.51	4.33	1.03	5.22	.76

^a Predictions were rated on a 1-7 scale.

Prediction-Choice Task

The results for subjects' choice of which key fact was the best prediction also support the systematicity hypothesis. Analogy subjects but not Target-Only subjects preferred the shared-system predictions. When subjects are scored for the number of choices of shared-system key facts across the four analogies (possible score is 0-4), a 2 x 2 (Condition x Version-Set) ANOVA revealed that the mean for Analogy subjects ($M=2.9$) is significantly greater than the mean for Target-Only subjects ($M=1.8$), $F(1,28)=16.43$, $p=.001$. As expected, no other effects were significant.

In contrast to the previous two tasks, here the results across the individual passages are somewhat varied. The pattern of results holds for three of the four passages, though the difference between the Analogy and Target-Only group reaches significance for only two passages. For these two passages, Analogy subjects chose the shared system fact 87% and 69% of the time, whereas Target-Only subjects chose this fact only 44% and 25% of the time, ($p < .05$, Fisher exact, one tailed tests.) For the remaining two passages, Analogy subjects chose the shared-system fact 62% and 75% of the time, and Target-Only subjects chose it 69% and 44% of the time; these differences between groups are not significant.³

Justifications of Choices

As in Experiment 1, Analogy subjects were asked to write a brief explanation for their Choice Task responses. They were to (a) explain why they thought their chosen prediction followed best from the analogy (b) explain why the rejected prediction did not follow from the analogy as well. The justifications for choices and rejections were separately scored. Our main

³ Note that for one passage the percentage of Target-Only subjects choosing the shared-system prediction is 69%. This gives a slight suggestion that this passage may be biased in favor of the systematicity predictions (though the percentage of Analogy subjects making the shared-system choice is smallest -69%- for this passage). The possibility that this passage was biased in favor of the systematicity predictions suggested that the analyses of the Prediction Task data should be reconsidered in the absence of this passage. The original findings still remain. That is, the ANOVA comparing the number of shared-system minus difference-system key facts predicted in each Condition still shows a significant effect for Condition even when predictions for this possibly biased passage are not included ($F(1,28) = 13.03$, $p < .001$).

interest was in whether subjects would refer to the shared or differing causal systems supporting a prediction in justifying their responses.

Two judges coded subjects' justifications for choosing a prediction into four categories similar to those used in Experiment 1. Responses in the first category show a concern for the higher-order embedding of the key prediction, and responses in the second simply show a concern for the similarity of the prediction to the base key fact. For responses in the final two categories it was either not clear whether subjects were responding on the basis of the analogy, or it was not clear how the analogy guided their responses. Thus, responses were categorized according to the level of focus:

1. Focus on the similar causal structure (higher order similarity). The subject described the cause for the chosen prediction that is similar between the base and target, or asserted *that* there is a similarity in cause.

2. Focus on the similarity of the prediction (lower order similarity). The subject simply noted the similarity or correspondence between the chosen prediction and the corresponding fact in the base.

3. Focus on the prediction alone -- on its plausibility (properties of the target). Subjects in this category did not refer to the base domain. They either stated that the prediction was likely because of the causal information in the target, or simply asserted that the fact was plausible in the target.

4. Other: Gives any other response.

Subjects' justifications for rejecting a fact were classified into the same four categories, except that now justifications focused on: *dissimilarities* in the cause of the rejected prediction to the cause for the corresponding fact in the base (some subjects who focused on the dissimilar cause also stated that the prediction was unlikely in the target.), the *dissimilarity* of the rejected prediction to the corresponding fact in the base, or the *implausibility* of the rejected prediction in the target.

Table 7 shows the justification data. As above, the left side of the table shows the distribution of responses when the shared-system prediction was chosen. The right side shows responses when the different-system prediction was chosen. The top half of the table shows justifications for choosing a prediction; the bottom half, for rejecting the other prediction. Inter-rater agreement for scoring of choice justifications according to the four categories

Table 7

Experiment 2. Justifications by Analogy Subjects for Predictions Selected in Choice Task: Proportion and Frequency of Shared-system and Different-system Choices (and Rejections)

<u>Justification Type:</u>	<u>Chosen Prediction</u>			
	<u>Shared-System</u>		<u>Different-System</u>	
	<u>(Correct)</u>		<u>(Incorrect)</u>	
	<u>Prop.</u>	<u>Freq.</u>	<u>Prop.</u>	<u>Freq.</u>
Similar cause for chosen fact	.77	(36)	.29	(5)
Similarity of chosen fact itself	.04	(2)	.18	(3)
Likelihood of chosen fact in target	.06	(3)	.29	(5)
<u>Other</u>	<u>.13</u>	<u>(6)</u>	<u>.23</u>	<u>(4)</u>
Total	100	(47)	100	(17)

<u>Justification Type:</u>	<u>Rejected Prediction</u>			
	<u>Different-System</u>		<u>Shared-System</u>	
	<u>(Correct)</u>		<u>(Incorrect)</u>	
	<u>Prop.</u>	<u>Freq.</u>	<u>Prop.</u>	<u>Freq.</u>
Dissimilar cause for rejected fact	.55	(26)	.29	(5)
Dissimilarity of rejected fact itself	.13	(6)	.23	(4)
Likelihood of rejected fact in target	.19	(9)	.29	(5)
<u>Other</u>	<u>.13</u>	<u>(6)</u>	<u>.18</u>	<u>(3)</u>
Total	100	(47)	100	(17)

Note. Out of 64 opportunities, subjects chose the shared-system prediction (the predicted response) 47 times and the different-system prediction 17 times. For each choice, justifications are shown for both the chosen prediction and the rejected prediction.

was .84 before discussion. Inter-rater agreement for scoring rejection justifications was .84 before discussion.

As shown in the left side of the table, when subjects chose the shared-system prediction, the majority of the justifications focused on the similar causal information in the base and target. Correspondingly, their justifications for rejecting the different-system prediction focused on the dissimilar causal information.

Examples of such causal justifications follow. Again examples are taken from subjects given Version 2 of the Robots passage in which the shared-system prediction is that robots cannot go from one planet to another to gather data, and the different-system prediction would be that the robots shut down their probes.

1. *"The second prediction [Robots cannot go from planet to planet] follows the analogy best. Robots will not be able to go from one planet to the next. We can assume this by looking at the story from the Tams. Since their underbellies (Tams) and filters and sensitivities (Robots) are specified they can't go from one place to the other. The first prediction [Robots shut down probes] doesn't follow the analogy because the Tams stop using its underbelly because there are no nutrients left. While the robot doesn't stop because there is no data, rather it will overheat if it doesn't."*
2. *"Two [Robots cannot go from planet to planet] is better because it relies on the analogy that both the Tams and the Robots specialize to the extent that are they are not transferable from rock to rock or planet to planet. One [Shut down probes] isn't good because when a Tam stops using its underbelly, it has exhausted the supply of minerals: whereas when a robot stops probing, it has been getting too much data and must shut down to avoid overheating."*

These comments reveal subjects' belief that predictions must be based on a shared higher-order structure. Few justifications were concerned only with the match between the target prediction itself and the corresponding base fact. Thus, the justifications for selecting predictions are consistent with the justifications for selecting matches found in Experiment 1. When subjects can make explicit their reasons for selecting information to map from a base to a

target domain, their inferences are guided by systematicity and shared structure.

Discussion

The results of Experiment 2 indicate that systematicity constrains analogical inference processes. That is, systematicity determines which predicates in the base domain will be imported as predictions into the target domain. In the Prediction Task, either the shared-system or different-system prediction was possible: each was a fact given in the base but not in the target, each could be easily constructed in the target, and each had antecedent conditions in the target. Yet subjects showed a strong preference for making the prediction that was supported by antecedent conditions that matched the base domain. The results of the Rating and Choice Tasks provide converging evidence that subjects prefer predictions sanctioned by systematicity. Overall, the predictions linked to a matching causal system were rated most highly and chosen most often as the predictions that follow well from the analogy. Furthermore, subjects explicitly focused on the matching causal structure in their choice justifications. Finally, the difference in performance between the Target-Only and Analogy subjects for each of the three tasks indicates that the results cannot be attributed to a bias in the materials in which the shared-system facts were inherently more salient or plausible in the target.

Convergent with the results of Experiment 1, the results of Experiment 2 indicate that analogical mapping concerns corresponding systems of predicates and not merely independent correspondences among individual predicates. Subjects ignored or rejected possible predictions that represented an isolated correspondence between the base and target, even though in themselves these predictions created a good match with the base. Rather, subjects generated analogical predictions that were supported by a larger systematic matching structure. It appears that in generating candidate inferences, just as in selecting predicates that belong to a match, people tacitly and sometimes explicitly seek connectivity and interdependency.

Experiment 3

In Experiment 1 and 2, subjects were always allowed to examine the base and target passages as they made mapping judgments and predictions. In the next experiment we asked whether the systematicity constraint would operate

when subjects had to rely on their memory representations of a base domain. In ordinary life, people frequently reason by analogy without the benefit of written material. Thus, our tasks so far could be allowing subjects to attain unrealistic levels of rigor in processing, or worse, could be somehow suggesting an unnatural strategy. In the more natural case, when people reason from a stored representation of a base domain, perhaps they are content with whatever correspondence they can most readily create with the target. Their selection of information to map may be less constrained by a concern for shared structure.⁴ We were somewhat reassured by evidence of various kinds that systematicity plays a role in natural analogizing: for example in scientific reasoning (Burstein, 1983; J. Clement, 1983; Gentner and Gentner, 1983; Gentner and Jeziorski, in press). However, we wanted to investigate the generality of the phenomenon. The present tasks check whether systematicity operates as a selection constraint even when subjects had to rely on a base domain represented in memory.

A second reason for conducting this experiment was to confirm that subjects' responses in the previous two experiments were guided by similarities in the underlying structures of the base and target passages, and not by uninteresting superficial features of the passages. In writing the analogies for all the experiments we attempted to avoid similarities in the surface form of sentences used to describe matching causal systems. However, replicating the previous experiments with memorized base domains would confirm that conceptual similarities in the absence of surface commonalities can support the systematicity constraint. (Note that we assume here that when material is committed to memory, the semantic content is better represented than the surface form of information.)

Thus, in Experiment 3, we again examined the effects of systematicity on the analogical predictions subjects make about a target domain. The basic method of this experiment was identical to that of Experiment 2, except that subjects first committed the base domains to memory.

⁴ Note that we are not talking here about how people access analogs in memory. This is generally agreed to be strongly influenced by surface similarities (e.g., Gentner and Landers, 1985; Holyoak and Koh, 1987; Ross, 1984; 1987). Rather we are concerned with how the analogical match is constructed given that a particular analog has been accessed.

Method

Subjects

Subjects were 48 paid undergraduate students at the University of Illinois. Half were assigned to the Analogy Condition and half to the Target-Only Condition.

Materials

Materials were designed identically to those used in Experiment 2. Three rather than four analogies were used because of the added length of the learning tasks. Two of the analogies were the same as those used in Experiment 2 (though modified slightly to facilitate comprehension) and one was a new analogy (derived from a previous analogy).

Procedure

Task for Analogy Subjects. Subjects first memorized the base domain. Then, given the target domain, subjects performed learning, prediction and judgment tasks identical to those in Experiment 2, except that subjects now could not refer back to the base passage.

Memorizing the base domain. Subjects studied a base domain for five minutes and then summarized it from memory. Next they reviewed the base passage and corrected or elaborated their summary by comparing it to the passage; no feedback was given by the experimenter. Finally, the base passage was removed and a multiple-choice test was given to assess subjects' understanding of the central events in the base passage. The test included questions about each key fact and causal system. Note that we tested knowledge of the causes for both key facts, and all subjects received the same test. Thus, the test did not bias subjects' attention toward a particular causal system. Subjects were given no feedback after their test.

To promote energetic performance on the memorization task, subjects were given a monetary motivation. That is, they were told that if they did well on their summaries and multiple-choice test, they would be paid an additional two dollars at the end of the session.

Learning tasks. After learning the base domain, subjects were given the target domain and, as in the previous experiments, worked out the object correspondences and described the ways in which the base and target were and were not analogous.

Prediction and Judgment Tasks. A Prediction Task, in which subjects made their own prediction about the target domain, as well as rating and choice tasks were given which were identical to those in Experiment 2. The instructions were altered slightly for clarification. Subjects were allowed to refer back to the target passage but not to the base passage.

Task for Target-Only Subjects. As in Experiment 2, in order to determine that the shared-system predictions were not simply the most plausible predictions in the target, a group of subjects made predictions and judgments based only on the information in the target passage.

Results

Prediction Task

Two judges who were blind to the condition of subjects being scored⁵ grouped predictions into the same categories as in the previous experiment: (a) predictions of key facts that would follow from the shared causal system in the base and target (b) predictions of key facts that would follow from the different-system, and (c) all other predictions about the target. (A few subjects, .07 of the responses, predicted both the shared-system and different-system fact. Again these responses were discarded and omitted from further analysis.) Inter-rater agreement before discussion was .96. Disagreements were readily resolved.

Table 8 shows the proportion of responses in each category. As predicted, Analogy subjects most frequently predicted the shared-system key fact. Furthermore, responses of the Target-Only subjects indicate that this prediction was not especially salient in the target passages. As before, to assess whether the preference for shared-system predictions among Analogy subjects was reliable, subjects were scored for the number of shared-system minus different-system facts predicted across the three passages. The mean for the Analogy group is .92 and the mean for the Target-Only group is .13. A 2 x 2 ANOVA (the factors were Condition and the counterbalancing factor Version-Set) showed that the predicted main effect for Condition was significant $F(1, 44)=5.15, p < .05$. Analogy but not Target-Only subjects

⁵Again, in some cases the content of subjects' responses would indicate which condition the subject was in.

Table 8
Experiment 3. Prediction Task: Proportion of Subjects in each Group Predicting
Shared-system and Different-system Key Facts for each Passage

Passage	Group					
	Analogy (n=24)			Target-Only (n=24)		
	Shared System	Different System	Other	Shared System	Different System	Other
1	.46	.21	.29	.04	.08	.88
2	.375	.375	.17	.04	0	.96
3	.75	.08	.08	.25	.125	.625
Overall	.53	.22	.18	.11	.07	.82

Note. A few Analogy subjects predicted both the shared-system and different-system fact (.07 of all responses). These responses are not shown here.

predicted more shared-system than different-system facts. The ANOVA also showed a significant main effect for Version-Set, simply indicating that overall the number of shared-system predictions varied across Version-Set, $F(1,44)=4.12, p < .05$. As expected, there was no interaction between Condition and Version-Set.

Table 3 shows that for one passage the data are not consistent with the systematicity predictions. For Passage 2, Analogy subjects predicted each fact type equally often. These inconsistent results appear simply to be due to some subjects' inaccurate representations of the causal information in the base passage for this analogy. As will be discussed below, several subjects incorrectly answered the multiple choice test questions about the base domain causal structure, and these errors were apparently related to responses on the prediction, as well as the rating and choice tasks.

In sum, consistent with the result of Experiment 2, Analogy subjects tended to make the prediction sanctioned by systematicity. Since in this experiment subjects had to memorize the base domains, the previous findings do not appear to be specific to the artificial situation of having a written base domain available for inspection. This experiment shows that analogical inferences are based on shared higher-order structure even when subjects must rely on their memory representations of the base domain.

Prediction-Rating Task

Table 9 shows the mean ratings for shared-system and different-system key predictions. Again as predicted, Analogy subjects rated shared-system facts more highly than different-system facts (overall $M=5.83$ and $M=4.89$ respectively) whereas Target-Only subjects showed the reverse pattern of ratings ($M=4.31$ and $M=5.39$ respectively).

A $2 \times 2 \times 3 \times 2$ (Condition, Fact-Type, Passage, and Version-Set) repeated measures ANOVA was conducted. (Passage was treated as a fixed-effects variable since the number of passages was small.) There was a significant main effect for Condition, $F(1,44)=4.37, p < .05$, simply indicating that overall, ratings varied across the two groups. No other main effects were significant.

As before, the the key prediction is an interaction between Condition and Fact-Type. This interaction was significant, $F(1,44)=10.82, p < .01$. Also, the simple effect of Fact-Type within the Analogy group showed that the difference

Table 9
Experiment 3. Prediction-Rating Task: Mean Ratings^a of Shared-system and
Different-system Predictions for each Passage

Passage	Group							
	Analogy (n=24)				Target-Only (n=24)			
	Shared System		Different System		Shared System		Different System	
	M	SD	M	SD	M	SD	M	SD
1	5.96	1.92	4.46	2.30	3.33	1.95	5.83	2.16
2	5.75	1.75	5.50	2.02	5.29	2.05	5.21	2.06
3	5.79	1.72	4.71	2.37	4.29	1.83	5.12	1.94
Overall	5.83	1.08	4.89	1.70	4.31	1.34	5.39	1.18

^a Predictions were rated on a 1-7 scale.

in ratings to shared-system and different-system facts by these subjects is significant, $F(1, 22)=5.25, p < .05$. Interestingly, the difference in ratings to the two fact types is in the opposite direction for Target-Only subjects (these subjects gave higher ratings to different-system facts) and this effect is also significant, $F(1,22)=5.58, p < .05$. This finding suggests that the materials were biased *against* the systematic choice.

Finally, the triple interaction between Condition, Fact-Type, and Passage was significant, $F(2,88)=5.41, p < .01$. As noted above, responses to Passage 2 do not follow the predicted pattern (again, apparently due to poor memory for the base passage), but the predicted interaction between Condition and Fact-Type holds for the other two passages. No other effects were significant.

Overall, the Rating Task results again confirm findings of the previous experiment. With the exception of one analogy, shared-system facts were rated higher than different-system facts. It is interesting that these results were found even when materials were apparently biased against the shared-system fact. The Analogy subjects valued a prediction which followed from the systematicity of the analogy even when an alternative prediction may have been more salient or plausible in the target domain.

Prediction-Choice Task

Results for the Choice Task also support the systematicity hypothesis. Analogy subjects but not Target-Only subjects preferred the shared-system prediction. Subjects were scored for the number of choices of shared-system predictions across the three passages (possible score is 0-3). A 2 x 2 ANOVA (Condition x Version-Set) showed that the mean number of shared-system predictions chosen in the Analogy group ($M=1.875$) was significantly greater than the number chosen in the Target-Only group ($M=1$), $F(1,44)=13.15, p = .001$. No other effects were significant. Considering the passages individually, again the predicted pattern does not hold for Passage 2 (the proportion of subjects choosing the shared-system fact is .54 and .58 in each group). However, the analogy and Target Only groups differ significantly for the other two passages. For each of these passages, .66 of the Analogy subjects chose the shared-system fact, in contrast to only .08 and .33 of the Target-Only subjects (Fisher exact, $p=.001$ and $p=.05$).

Choice Task Justifications

As in the previous two experiments we were interested in subjects' explicit reasons for their Choice Task responses. Justifications for choices and rejections of key predictions were coded in the same manner as in Experiment 2. Table 10 shows that the distribution of responses replicates the findings of the previous experiment: subjects who made the systematic selection focused on the similarity or dissimilarity of causal information for the key predictions. Few of these subjects were concerned only with whether the prediction itself corresponded to the base domain. Thus, subjects' explicit criteria for a good analogical prediction concern the connection of the prediction to a causal structure found in both domains.

Analogy Group Responses as Function of Comprehension of the Base Passage

After Analogy subjects read and memorized the base passage they answered multiple choice questions about the central events in the passage. Analysis of the relation between performance on this test and performance on the analogy tasks helps clarify the inconsistent findings for Passage 2. Performance on the multiple choice test for Base Passage 2 was poor relative to the other other passages. Specifically, on the test for Base Passage 2, 42% of the subjects made an error, but on the tests for Base Passages 1 and 3, only 8% and 29% of the subjects made an error. (Most subjects made only one error out of five to eight questions.) Furthermore, 12 of the 13 Passage 2 errors were to questions that specifically addressed the cause for one of the key facts. Consistent with this, Table 11 shows that the sub-group of subjects who were error-free on the multiple choice tests were more likely to respond to the analogy tasks according to the systematicity predictions than was the Analogy group as a whole. (Compare Tables 8 and 9.) This is especially true for Passage 2. These results must be interpreted with some caution, since with the removal of subjects Version-Set is no longer fully counterbalanced. However, the results indicate that the unexpected analogy task results discussed above for Passage 2 are due to some subjects' failure to either recall or understand the causal structure of the base domain.

Discussion

Overall, the Prediction, Rating, and Choice Task responses, and subjects' explicit justifications, replicate the findings of Experiment 2. Thus the previous

Table 10

Experiment 3. Justifications by Analogy Subjects for Predictions Selected in Choice Task: Proportion and Frequency of Shared-system and Different-system Choices (and Rejections)

<u>Justification:</u>	<u>Chosen Prediction</u>			
	Shared-System		Different-System	
	(Correct)		(Incorrect)	
	<u>Prop.</u>	<u>Freq.</u>	<u>Prop.</u>	<u>Freq.</u>
Similar cause for chosen fact	.69	(31)	.30	(8)
Similarity of chosen fact itself	.09	(4)	.30	(8)
Likelihood of chosen fact in target	.09	(4)	.18	(5)
<u>Other</u>	<u>.13</u>	<u>(6)</u>	<u>.22</u>	<u>(6)</u>
Total	100	(45)	100	(27)

<u>Justification:</u>	<u>Rejected Prediction</u>			
	Different-System		Shared-System	
	(Correct)		(Incorrect)	
	<u>Prop.</u>	<u>Freq.</u>	<u>Prop.</u>	<u>Freq.</u>
Dissimilar cause for rejected fact	.64	(29)	.11	(3)
Dissimilarity of rejected fact itself	.11	(5)	.15	(4)
Likelihood of rejected fact in target	.07	(3)	.30	(8)
<u>Other</u>	<u>.18</u>	<u>(8)</u>	<u>.44</u>	<u>(12)</u>
Total	100	(45)	100	(27)

Note. Out of 72 opportunities, subjects chose the shared-system prediction (the predicted response) 45 times and the different-system prediction 27 times. For each choice, justifications are shown for both the chosen prediction and the rejected prediction.

Table 11

Experiment 3. For each Analogy, Performance on the Prediction, Rating and Choice Tasks among Analogy Subjects who were Error-free on the Base Passage Test

Passage	Prediction Task: Proportion of Subjects Making Each Prediction		Rating Task: Mean Ratings to Each Prediction		Choice Task: Proportion of Subjects Choosing the Shared System Prediction
	Shared System	Different System	Shared System	Different System	
	prop.	prop.	M	M	prop.
1	.50	.18	6.22	4.27	.73
2	.50	.28	6.50	4.93	.71
3	.88	.06	5.59	4.35	.76

Note. The number of error-free subjects for each passage was 22 for Passage 1 (half in each Version Set); 14 for Passage 2 (8 in Version Set 1 and 6 in Version Set 2); and 17 for Passage 3 (7 in Version Set 1 and 10 in Version Set 2).

findings are apparently not the result of subjects' reliance on surface, textual features of the passages, rather than on the causal structure of the passages. Furthermore, subjects' use of systematicity as a selection constraint is apparently not restricted to the situation in which a written base domain is available for inspection. In the present experiment, as in typical cases of analogical reasoning, subjects made analogical inferences by drawing on their memory representations of the base domain. As before, the selection of analogical inferences appeared to be guided by a concern for the connection of these inferences to a shared causal structure. The effects are somewhat weaker than observed in the previous experiments, apparently because subjects did not always possess an accurate record of the base domain. However, in general subjects followed the same rules which constrain good analogical predictions to those that follow from a larger matching system.

General Discussion

Analogy is a selective form of comparison. Only certain kinds of similarities are important in mapping information between two analogous domains. The experiments presented showed that the choice of lower-order relations to include in an analogy is not determined just by the independent relations themselves, but also by whether they are embedded in a larger matching structure. This systematicity principle was found to constrain which *matches* between a base and target domain are selected for an analogy, and which *inferences* are drawn from an analogy. In what follows we discuss the implications of our findings, and how they link to several issues important to understanding analogy and similarity.

Systematicity Constraints on Selecting Matches.

Our first experiment required subjects to judge which of two possible matching facts contributed best to an analogy. Although both facts matched well between the base and target, subjects consistently preferred the fact that was linked to a causal system also found in both domains. Their mapping preferences and their explicit justifications indicate that subjects invoked a systematicity rule as a selection criterion. A match was given a high evaluation if it was linked by a higher-order relation to neighbors that also matched. Thus, even when subjects have no external purpose which might support particular

mapping preferences, they invoke a principle which concerns the place of a matching lower-order relation in the larger system of shared relations.

Subjects' mapping choices cannot be accounted for by other qualities of the matching facts, such as differences in how well the facts matched between domains, differences in the importance of the facts to goals internal to a domain, or differences in the validity of the facts. First, since two different versions of a target domain were used, the same fact was sometimes embedded in a matching causal structure, but other times embedded in a non-matching causal structure. Second, although a specific fact could have taken on more or less relevance or importance in its respective version of a target story, this did not appear to happen. Control subjects' responses indicated that shared-system and different-system facts were equally important in the target domains. Finally, the relative validity (truth value) of the facts in the target cannot have determined choices, since the analogies described fictional domains, and both facts were asserted as known facts about a domain.

One last alternative account of our findings can be considered. It might be argued that subjects preferred the shared-system match simply because it was preceded by other matching information, and the higher-order link between the key match and preceding matches was irrelevant. In contrast, on our account the higher order link is needed. That is, if we had preceded a matching key fact with totally unrelated information that also matched between domains, this would not have been sufficient to determine subjects' choices. For example, preceding the key matching fact, *Robots shut down their probes*, with *egg plants come in many varieties* would be unlikely to increase subjects preference for the key fact. Subjects' response justifications also show that the higher-order link between a matching fact and neighboring matches is important. That is, subjects often cited the similarity in causal information as the reason for their choice.

Although our results indicate that a systematic connection among matching relations is important in analogical mapping, an interesting question remains. That is, would a weaker kind of connection among matching information also constrain mapping? Perhaps shared structure would guide mapping choices, even when relations within a structure co-occur but are not interdependent (e.g. when they are regular conjunctions). Future research

should examine whether such weaker kinds of connections would support mapping choices.

Systematicity Constraints on Analogical Inference

When asked to draw an analogical inferences about a target domain, subjects appear to make those inferences that reflect the same relational interdependency in the base and target. In Experiments 2 and 3, subjects were given a base and target passage with many existing matches. However, two facts were always present in the base but absent in the target. Antecedents for both facts were present in the target, but the antecedent for only one fact matched the base domain. Subjects consistently predicted the fact whose antecedent matched that of the base system. Although the alternative prediction was equally possible, again the results indicate that it is the interdependence among predicates rather than isolated predicates themselves that govern mapping. Experiment 3 is important in showing that our findings are not unique to a relatively unnatural situation in which subjects have a written description of the base domain available for inspection. In Experiment 3, as in typical instances of reasoning by analogy, subjects had to rely on their memory representations of the base domain. Subjects' inferences continued to follow the systematicity principle.

These results show that subjects not only appreciated the aesthetics of analogical matching, but also appreciated the systematic use of analogy as a tool for learning about a target domain. Analogical mapping which follows the systematicity principle can be usefully predictive: if an implication between an antecedent and consequent is known to hold in a base domain, and the antecedent is also known in the target, then the base implication can be used to infer the consequent in the target. Consistent with this, in the present tasks, subjects did not simply select any base fact that could be plausibly transferred to the target. Rather, subjects formed a prediction that completed a matching causal system.

How Systematicity Operates as a Selection Constraint

Systematicity might function as a selection constraint in several different ways. For example, it might provide a criterion for evaluating a potential mapping already identified between domains, or, alternatively, it might enter into the search for potential mappings. In selecting matches between a base and target, either of these possibilities seems plausible. However, in the

inference process it is more plausible that systematicity constrains the search for predictions rather than operating after possible inferences have been identified (in this way systematicity would constrain the indefinite number of possible mappings). Thus, people would first find a matching structure and then derive a prediction from that structure.

The structure-mapping engine (SME) (Falkenhainer, Forbus, and Gentner, in press), the computer implementation of Gentner's structure-mapping theory, provides a possible model of the kind of processes people might follow. To arrive at a complete interpretation of an analogy, SME first begins locally, identifying individual relations (lower order relations) which match between a base and target domain. Local matches are then sorted into larger systematic sets according to the structural constraints of one-to-one mapping and structural consistency. Each set represents a different interpretation of the analogy. The matching relations within a set are given weights which reflect the match itself, and the extent to which it has higher-order links to other matching relations. From these weights, an evaluation of each systematic set is derived, and a final interpretation of the analogy is chosen. Importantly, SME also creates inferences about the target domain on the basis of matching structure. If there is a predicate connected to the base system, but not found in the target system, this predicate becomes a candidate inference in the target. Thus, in SME, systematicity enters into analogy construction in the evaluation of identified matches, and in the derivation of inferences from a matching structure. An attractive feature of this model is that the same processes that enter into forming matches also lead to analogical inferences. Candidate inferences can be drawn solely on the basis of shared structure. As mentioned earlier SME has been shown to be consistent with human reasoners in its interpretations of several analogies (Skorstad, Falkenhainer, and Gentner, 1987).

Structural Determinants of Similarity

The results of the present experiments show that whether a given or predicted match in two lower-order relations is considered a good analogical mapping, is not just determined by the match itself. Rather, a good mapping must also be connected to an interrelated system of matching relations. If we think of analogy as a kind of similarity, then these results conflict with many applications of Tversky's influential contrast model of similarity (Tversky, 1977),

that have assumed that the features entering in to similarity computations are independent of one another. In contrast, the present results indicate that it is precisely the non-independent features that matter most in an analogy. Additional evidence for the importance of interdependence among features is provided by recent work on perceptual similarity. Goldstone, Gentner and Medin (1989) demonstrated that judgments of similarity among patterns of geometric stimuli violated the independence assumption in that the importance of a relation depended on neighboring relations. Thus relational structure appears to be an important factor in judgments of literal similarity as well as in analogical mapping.

Ortony's (1979) salience imbalance model of metaphor, which extends Tversky's contrast model, takes the important step of noting that feature salience is context dependent. Also, Ortony notes that metaphorical domains may share a common schematic structure. However, Ortony's model, unlike structure-mapping, does not utilize interrelations among features in its computations. The present study demonstrated that such interrelations are critical to the computation of analogical similarity.

Relational correspondences based only on structure. The effects of systematicity may extend beyond those demonstrated in the present research. In Experiment 1 we found that embedding in a shared structure determined which of two *matching* relations contributed to an analogy. Future research should examine the extent to which shared structure may allow *non-matching* lower-order relations to be put into correspondence. That is, in the same way that dissimilar objects are mapped onto one another by virtue of their similar roles in a larger system, perhaps dissimilar lower-order relations may be mapped by virtue of their similar roles. For example, *differing* antecedent relations for the same consequent may be put into correspondence on the grounds that they lead to this same consequent (perhaps the relations are re-represented so that they have a similarity defined by their shared role). A meaningful analogy must include many matching relations. However, mismatching relations that are components of a much larger matching structure may be put into correspondence, despite their intrinsic dissimilarity, exclusively on structural grounds.

Interactions Between Structural Constraints and Goals

The present results show that people use interconnections among predicates to decide what matters to an analogy. But analogical reasoning often occurs in the context of a particular goal, and many interesting analogies involve goal structures. Several computational models of analogy have attempted to integrate structural and goal-relevance constraints on mapping (Burstein, 1983; Holyoak and Thagard, 1989; Kedar-Cabelli, 1985b; Thagard and Holyoak, 1988). What is the relation between goals and structural constraints in analogical processing?

It might be argued that in constructing an interpretation of an analogy, structural factors are insufficient to constrain mapping choices, and that mapping possibilities must be evaluated with respect to an extrinsic goal (for example, the reasoner's goal to solve a problem in the target domain, or to make a point about the target) (Burstein and Adelson, 1987; Burstein, 1983; Holyoak, 1985; Keane, 1988; Kedar-Cabelli, 1985b). Such a view is incompatible with the results of the present experiments which indicate that the relational structure intrinsic to the analogous domains can determine mapping even in the absence of an extrinsic goal. However, goal relevance and structural constraints may interact in several interesting ways when goals are present. One possibility is that goals directly influence the mapping process; for example, predicates may be weighted for goal relevance (Thagard and Holyoak, 1988). An alternative possibility is that extrinsic goals influence mapping indirectly. For example, Gentner (1989) proposes a before/after influence of goals on the mapping process: a person's goals influence the construal of the base and target domains that are input to the analogical mapping process, and after an interpretation is constructed, goals influence the evaluation of this interpretation, especially the inferences derived from it.

In many analogies, the systematicity and goal relevance of information is correlated. This is particularly true of problem-solving analogies and domains used in case based reasoning, (Gick and Holyoak, 1983; Holyoak, 1985; Kolodner, Simpson and Sycara, 1985; Kolodner, 1987; Schank, 1982). If we consider a goal schema as a particular case of a relational structure, then the same processes we have discussed here can operate to match the base and target in such analogies. This has the advantage of parsimony -- we can assume the same process across different contents and external contexts.

The question of how mapping processes will proceed with goal structures vs non-goal structures is an interesting one that future research should address. The present experiments provide the first clear demonstration of the more general point that structural matches determine mapping choices, and do so even in the absence of extrinsic goals. These findings allow us to understand how analogy can be a tool for discovery. Since an analogy can be constructed on the basis of the structures intrinsic to the analogous domains, analogical thinking is not restricted by a process that requires pre-specified goals to guide mapping. Thus, analogy can not only be a tool for instrumental reasoning, but also can be used to achieve the general human goals of explanation, and understanding, and the satisfaction of curiosity.

Conclusion

What factors constrain the choice of information to map in an analogy? Although one obvious factor is the individual feature matches between two domains, similarity alone is insufficient to determine whether a feature is mapped. Other factors are needed to determine which similarities are important. The present experiments have shown that when given a choice between component matches or predictions that are in themselves equally good or valid, subjects prefer those matches and make those predictions that maintain a highly systematic correspondence between the two analogous domains. There is prior evidence that systematicity affects the evaluation of analogies (Gentner & Landers, 1985; Rattermann & Gentner, 1987) and the efficacy of analogical mapping (Gentner and Toupin, 1986; Holyoak and Koh, 1987; Schumacher and Gentner, 1988). 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. Under the systematicity constraint matching low-order relations are not selected for an analogy unless they are connected to a larger matching system.

The present findings indicate that theories of analogy and similarity that focus on independent feature matches, and not on the higher-order embedding of the matches, are inadequate to account for analogical mapping. In deciding what to map, subjects invoked a principle that went beyond the intrinsic similarity of possible mapping choices, and that was independent of extrinsic goals. They selected information on the basis of its connection to a

larger matching structure. Thus, the systematicity of shared information appears to be a psychologically real constraint on analogical mapping.

References

- Black, M. (1962). Models and metaphors. Ithaca, N.Y.: Cornell University Press.
- Burstein, M. (1983). Concept formation by incremental analogical reasoning and debugging. In R. Michalski, J. Carbonell, & T. Mitchell (Eds.), Machine learning (Vol. 2). 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, 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.
- Carbonell, J. G. (1983). Metaphor and common-sense reasoning. Technical report #83-110. Carnegie-Mellon University.
- Clark, H. (1973). The Language-as-fixed-effect fallacy: A critique of language statistics in psychological research. Journal of Verbal Learning and Verbal Behavior, 12, 335-359.
- Clark, H. H. & Chase, W. G. (1972). On the process of comparing sentences against pictures. Cognitive Psychology, 3, 472-513.
- Clement, J. (1983). A conceptual model discussed by Galileo and used intuitively by physics students. In A. Stevens, and D. Gentner (Eds.) Mental models, Hillsdale, NJ: Erlbaum.
- Collins, A. & Burstein, M. (1989). A framework for a theory of mapping. In S. Vosniadou & A. Ortony (Eds.), Similarity, analogy, and thought. New York: Cambridge University Press.
- Falkenhainer, B., Forbus, K. D. & Gentner, D. (in press). The structure-mapping engine: an algorithm and examples. 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. (1989). The mechanisms of analogical learning. In S. Vosniadou & A. Ortony (Eds.), Similarity, analogy, and thought. New York: Cambridge University Press.
- Gentner, D. & Clement, C. A. (1988). Evidence for relational selectivity in the interpretation of analogy and metaphor. In G. H. Bower, (ed.). The Psychology of learning and motivation. Advances in research and theory. Vol.22, N.Y., New York: Academic Press.
- Gentner, D. & Gentner, D. (1983). Flowing waters or teeming crowds: mental models of electricity. In A. Stevens, and D. Gentner (Eds.) Mental models. Hillsdale, NJ: Erlbaum.
- Gentner, D. & Jeziorski, M. (in press). Historical shifts in the use of analogy in science. In B. Gholson, A. Houts, R. A. Neimeyer, & W.R. Shadish (Eds.) Psychology of Science.
- 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, Tucson, AZ.
- Gentner, D. & Schumacher, R. M. (1986). Use of structure-mapping theory for complex systems. Proceedings of the IEEE International Conference on Systems, Man and Cybernetics, Atlanta.
- Gentner, D. & Toupin, C. (1986) Systematicity and surface similarity in the development of analogy. Cognitive Science, 10, 277-300.
- Gick, M. & Holyoak, K. J. (1980). Analogical problem solving. Cognitive Psychology, 12, 306-355.
- Gick, M. & Holyoak, K. J. (1983). Schema induction and analogical transfer. Cognitive Psychology, 15, 1-38.
- Goldstone, R., Gentner, D. & Medin, D. (1989) Relations relating relations. To appear in Proceedings of the eleventh annual meeting of the Cognitive Science Society. Hillsdale, NJ: Erlbaum.
- Hall, R. P. (1988). Computational approaches to analogical reasoning: A comparative analysis. Artificial Intelligence, pp. 1-75.
- Holyoak, K. J. (1984). Analogical thinking and human intelligence. In R. J. Sternberg, (Ed.), Advances in the psychology of human intelligence, Vol. 2. Hillsdale, N. J.: Erlbaum.

- Holyoak, K. J. (1985). The pragmatics of analogical transfer. In G. H. Bower (Ed.), The Psychology of learning and motivation, (Vol. 1). New York: Academic Press.
- Holyoak, K. J. & Koh (1987). Surface and structural similarity in analogical transfer. Memory and Cognition, 15, 332-340.
- Holyoak, K. J. & Thagard, P. (1989). A computational model of analogical problem solving. In S. Vosniadou & A. Ortony (Eds.), Similarity, analogy, and thought. New York: Cambridge University Press.
- Keane, M. (1988). Analogical problem solving. Chichester, U. K.: Ellis Horwood.
- 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, Irvine California.
- Kittay, E. F. (1987). Metaphor: Its cognitive force and linguistic structure. Oxford: Carendon Press.
- Kolodner, J. , Simpson, R. L., & Sycara, K. (1985). A process model of case based reasoning and problem solving. Proceedings of the sixth annual International Joint Conference on Artificial Intelligence. Los Altos, CA: Morgan Kaufman.
- Kolodner, J. (1987). Capitalizing on failure through case-based inference. Proceedings of the ninth annual Conference of the Cognitive Science Society, Hilldale, NJ: Erlbaum.
- Lakoff, G. & Johnson, M. (1980). Metaphors we live by. Chicago: University of Chicago Press.
- Ortony, A. (1979). Beyond literal similarity. Psychological Review, 86, 161-180.
- Ortony, A. , Vondruska, R. J., Foss, M. A., & Jones, L. E. (1985). Salience, similies, and the asymmetry of similarity. Journal of Memory and Language, 24, 569-594.

- Rattermann, M. J. & Gentner, D. (1987). Analogy and similarity: Determinants of accessibility and inferential soundness. Proceedings of the ninth annual conference of the Cognitive Science Society, Hillsdale, NJ: Erlbaum.
- Rumelhart, D. E., & Norman, D. A. (1981). Analogical processes in learning. In J. R. Anderson (Ed.), Cognitive skills and their acquisition. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Schank, R. C. (1982). Dynamic memory. New York: Cambridge University Press.
- Schumacher, R. M., & Gentner, D. (1988). Transfer of training as analogical mapping. IEEE Transactions on Systems, Man, and Cybernetics. Vol 18 (4).
- Shannon, B. (1988). On the similarity of features. New Ideas in Psychology, 6, 306-321.
- Skorstad, J. , Falkenhainer, B., & Gentner, D. (1987). Analogical processing: A simulation and empirical corroboration. Proceedings of the meeting of the American Association of Artificial Intelligence, Seattle.
- Sternberg, R. J. (1977). Intelligence, information processing, and analogical reasoning: The componential analysis of human abilities. Hillsdale, NJ: Erlbaum.
- Thagard, P & Holyoak, K. (1988). Analogical problem solving: A constraint satisfaction approach. Proceedings of the Tenth Annual Conference of the Cognitive Science Society, Hillsdale, NJ: Erlbaum.
- Tversky, A. (1977). Features of similarity. Psychological Review, 84, 327-352.
- Winston, P. H. (1980). Learning and reasoning by analogy. Communications of the ACM, 23(12), pp. 689-703.

APPENDIX A
ANALOGIES USED IN EXPERIMENT 1

(Key facts in each target passage are shown in italics. Italics were not used in passages given to subjects.)

Analogy 1: Tams/Robots

Base Passage

The Tam

In a far away galaxy a life form was discovered called the "Tam". The Tam is a creature that lives off minerals in rock. It attaches itself to a rock and very slowly eats it. The Tam has a hard outer shell that looks like rock and an underbelly that has a surface like sandpaper. To eat rock the Tam wriggles its underbelly back and forth to loosen the top layer of the rock. Then the belly sucks in this loosened material. After the rock is sucked in, minerals in the rock are distributed to the Tam's internal organs and to the outer shell.

The Tam does not spend all its time sucking rock. When a Tam uses up all the valuable minerals in one spot on a rock it enters a "creeping" phase rather than a "sucking" phase. At this time the Tam's underbelly no longer functions and instead the Tam expends its energy creeping around the rock searching for a new spot. As the Tam creeps about the rock it will periodically stop moving and briefly puts its underbelly back into action to check the mineral supply in the new spot. If the supply is good the Tam stays at the new spot and its underbelly resumes constant activity until the mineral supply is exhausted. When the entire rock runs out of minerals the Tam dies.

Rocks and underbellies

There are several varieties of Tams. Each variety has a different underbelly. These differences result from the variety of rocks that Tams live on. The underbellies differ because the surface of different rocks is more or less difficult to loosen and suck in. For example, for hard rocks with large granules Tams must have an underbelly which is course-grained and highly permeable. In contrast, for softer rocks Tams must have an underbelly which is finer-grained and less permeable so that not too much rock is loosened and ingested at once.

Birth and adulthood

An amazing thing about Tams is that they are all born alike. They develop differences in their underbellies by adapting to a particular rock when they are young. When Tams are born they are tiny and winds can blow them from one kind of rock to another. However, within a few days of birth Tams get big enough to stay on one rock. They then adapt to that rock.

The underbelly of a newborn Tam functions in a clumsy way; it is poor at getting the required amount of minerals. But as the Tam adapts, its underbelly develops just the right texture and permeability for the particular rock the Tam lives on. Thus the Tam becomes specifically adapted to one kind of rock.

After a Tam has adapted to a particular rock it is very effective at bringing in minerals. However it is only effective on that one kind of rock. If a very fierce storm removed the older Tam to a new kind of rock it would not be able to adapt. Its underbelly could not develop the right texture and permeability to get the minerals needed. Fortunately such severe storms are rare.

Target Passage, Version 1

(Key facts here, but not in the passages given to subjects, are in italics)

The Robots

Highly sophisticated explorer robots are used to investigate planets throughout the galaxy. These robots gather and interpret information or data about the life forms on a planet. The robots have sensors or "probes" that act something like cameras. They take in all kinds of data about plants and animals. These data are eventually sent to one of several micro-computers inside the robot. These micro computers interpret the data.

Sometimes a robot must shut down its probes. When it no longer finds new information in a certain area on the planet it enters a period of travel and search. The robot turns off its probes and travels across the planet seeking a place where the data are different than data previously gathered. Every so often the Robot will turn on its probes to check the data in the current location. If the data are the same as old data the robot will turn off its probes and move on. If the data are new the robot will remain in that location and take up normal data gathering. When a robot finds no more new data it returns to the home planet.

Design

All of the explorer robots are built with "super-probes" that are extremely sophisticated. The robots' probes work efficiently under any lighting or atmospheric conditions. Different planets have different conditions that can affect the way in which information is detected by a probe. This had created problems for designers of probes. Since the designers could not know ahead of time what a robot's destination would be like they could not design probes that were specialized for the conditions on a particular planet. However, after years of research highly sophisticated super-probes were developed. These probes can function on any planet. Each explorer robot is equipped with these super-probes.

These super-probes are extremely sensitive and delicate. This means that the jostling of space travel can easily damage the probes during the flight to a planet. Consequently, elaborate packing is required to ship the robots safely to a planet. When a robot arrives at a planet the first thing it has to do is spend time carefully unpacking itself.

Because of this special packing a particular robot can only gather data on one planet before returning home. A robot is not able to repack its probes in the elaborate way needed to survive another space flight. When a robot returns to the home planet the probes have to be repaired before they are reused. *Thus, robots cannot go from one planet to another to gather data.*

Target Passage, Version 2

The Robots

Highly sophisticated explorer robots are used to investigate planets throughout the galaxy. These robots gather and interpret information or data about the life forms on a planet. Sensors or "probes" on the exterior of the Robots gather all kinds of data about plants and animals. These data are eventually sent to one of several micro-computers inside the robot. The micro-computers interpret the data.

Sometimes a robot must shut down its probes. This is because the microcomputers have a extremely sensitive electrical circuitry and can over-heat when they receive a great deal of complex data. When the probes are turned off the microcomputers can cool because they do not have to process

any new data. After the cooling the robot turns on its probes and again begins sending data to the microcomputers.

Design

The explorer robots differ from one another. Specifically, their probes differ depending on which planet they are sent to. The atmospheric and lighting conditions on a particular planet affect the ease and manner with which information can be detected by a probe. Thus probes must have different types of filters and degrees of sensitivity depending on the conditions existing on a planet.

Designers of robots cannot know ahead of time what a robot's planet will be like. Therefore the designers cannot know what the robots' probes should be like. A robot must develop the right probes after it arrives on the planet. A robot's probes are originally designed to work in a very general but inefficient way. However, once the robot gets started on a planet it develops the particular filters and sensitivity needed given the conditions on that planet. Thus each robot becomes specialized to operate very efficiently on its planet.

Once a robot's probes become specialized for one planet the robot cannot adjust its probes for the different conditions found on a different planet. If a robot was sent to a second planet its probes could not be changed to have the right filters and sensitivity needed to would work efficiently on the new planet. *Thus, robots cannot go from one planet to another to gather data.*

Analogy 2: Drilians/Veisal Plants

Base Passage

The Drilians

A large, four legged creature called a Drilian lives on the planet Munpar. Drilians are a highly intelligent and rather tense species.

Wumpums are relaxing

Drilians have a fascinating reproductive process. They rely on the presence of a lower animal for reproduction. The reproductive process works in the following way. Each fertile Drilian has a store of eggs (there are no "sexes" among Drilians). An egg begins to incubate when a Drilian is in the company of small fuzzy animals called Wumpums. Before incubation Drilian eggs are stored in a muscular sac inside the Drilian. When a Drilian is near a Wumpum the Drilian's body becomes extremely relaxed. When the egg sac

relaxes it releases an egg. The egg then moves to a special area inside the Drilian where it begins to incubate. The egg incubates in this area for exactly one year. Then it is laid and is hatched immediately. Thus, even though Wumpums do not pass any genes to a new Drilian, their relaxing effect on a Drilian is essential to the reproduction process.

Powerful minds

The planet Munpar where the Drilians live is a fertile environment hospitable to many infectious diseases. The Drilians have an unusual way of combating disease. They use mind control. Simply by thinking certain thoughts they can control their immune system. When a Drilian gets infected with a disease it can instruct its immune system to produce the correct antibodies so that the Drilian can resist the disease. Moreover, a Drilian can spread this resistance to other Drilians. When a Drilian is instructing its immune system other Drilians pick up on the brain waves emitted by the Drilian and resonate with them. These other Drilians then have the instructions needed for their immune systems to resist the disease. As a result the Drilians are seldom troubled by epidemics.

Target Passage, Version 1

The Veisel Plant

The Growth of the Veisel Plant

A plant called a Veisel plant grows in a rain forest on the planet Funl. Scientists have recently discovered how this plant generates new plants. The stalk of the veisel plant is covered with a mesh of fibrous material. Reproductive cells are held inside this fibrous mesh. Periodically, the rain that falls is a naturally occurring acidic rain. The acids in this rain cause the mesh of fibers to loosen. When this occurs the reproductive cells are carried off by wind and rain and they eventually fall to the ground. This ground provides the environment needed for the cells to germinate and finally develop into a new plant. Thus, Veisel plants are an interesting example of interaction among elements of nature. *Without the acidic rain new plants would not be generated.*

There are slight differences among the Veisel plants. For example some have large yellow flowers and others have smaller flowers that are bright blue in color. These differences apparently result from how much shade the plant

lives in. Plants in relatively open areas have the yellow flowers and plants that are well shaded by many surrounding trees have the blue flowers.

Dangerous Insects

Veisel plants have a useful way of resisting infestation by insects that can thrive on them. When reproductive cells are released from a parent plant the cells are typically carried by winds for a long time. Each cell ends up in some place distant from the parent and distant from other cells. Thus, the plants grow quite far apart from each other. A walker in the forest would be lucky to see more than one Veisel plant in a day. This growth pattern turns out to be crucial because it inhibits the spread of dangerous insects. Although an aphid or other destructive insect may invade an individual Veisel plant there will be no plants near by for the insects to spread to. *Because of this, large scale insect infestations across the general Veisel plant population are extremely rare.*

Target Passage, Version 2

The Veisel Plant

The Growth of the Veisel Plant

A plant called a "Veisel" plant grows in a rain forest that covers the planet Funl. Scientists have recently learned how new Veisel plants are generated. The scientists were amazed to discover that rain acts as one parent of a new plant. Periodically, the rain that falls is a naturally occurring acidic rain. The acids in this rain form a chemical bond with a Veisel plant's reproductive cells. The acids' special enzymes give a cell the additional substance needed to form a new plant. After the bond is fixed the cell begins to form into a new plant that is a product of the parent plant and the rain. When it reaches a certain size it is so heavy that it falls to the ground where it takes root. Thus, Veisel plants are an interesting example of the interaction among the elements of nature. *Without the acidic rain new plants would not be generated.*

The specific chemical content of some acidic rain is different from the content of other acidic rain. This causes slight differences among the Veisel plants. For example some rain contributes an enzyme that leads to large yellow flowers on the plant. Other rain creates a plant with smaller flowers that are bright blue in color.

Dangerous Insects

Veisel plants are equipped with an effective way of protecting themselves from infestation by insects that can thrive on them. When a plant is invaded by some destructive insect the plant's protection system develops a special powder that tastes terrible to the insect. Eventually the terrible taste will cause the insect to leave the plant. The wonderful thing is that the powder is blown by breezes from one plant to the other Veisel plants in the area. As a result these plants also become protected from this insect. Thus, one plant can save a whole population of Veisel plants from infestation. *Because of this, large scale insect infestations across the general Veisel plant population are extremely rare.*

Analogy 3: Space Ships/Zylots

Base Passage

The Space Ships

The planet Dolm builds space ships for inter-planetary mining. They recently built ships for flights to the planet Theal which is rich in valuable ores. To send ships to Theal the ship builders had to worry about two things: a dangerous gas and enemy ships.

Withstanding Gases

Interplanetary travel is very dangerous because of hostile gases that surround many planets. These gases contain "debinding" elements which sever molecular bonds of a space ship hull. Once debinding occurs the hull dissolves. Fortunately there is a way of building ships that can resist debinding. The ships can be made from a material that reorganizes itself in the presence of hostile gases before debinding can occur. This reorganization increases the strength of molecular bonds. The increased strength enables the material to withstand debinding actions.

The Trials of Two Spaceships. The planet Theal has many valuable ores. However, it turns out to be surrounded by a powerful gas called TGAS. TGAS is particularly dangerous because its debinding element acts extremely fast. This element is able to debind the molecular bonds of ship material before the material has time to reorganize and strengthen. Because of this, the first space vehicle sent to Theal was not able to withstand the TGAS. The material of the ship dissolved and the ship was destroyed.

Strangely enough, a second ship sent to Theal was able to withstand the TGAS. It turned out that this ship had stopped at another planet called Vando before going to Theal. Vando is surrounded by VNGAS. VNGAS is identical to TGAS with one exception: its debinding element does not act as quickly. Thus the ship material exposed to VNGAS was able to reorganize and increase its strength before the debinding element in VNGAS could work.

Since the ship material had reorganized itself in response to VNGAS it was able to withstand the similar TGAS. In other words, when the ship arrived at Theal the molecular bonds were already strong enough to resist the debinding element in TGAS. After seeing what happened to the two ships, the ship builders sent all ships to Vando on their way to Theal.

The Sronians

To get to Theal the space ship has to pass through territory owned by the planet Sron. Unfortunately Sronians are very aggressive and will try to seize travelers from Dolm. Thus designers of the space ships going to Theal had to take precautions. They devised a mechanism that would make the space ships extremely hard to detect by Sronians. When a space vehicle enters Sronian territory it can erect a giant screen which covers the ship. Special projectors behind the screen create an image on the front side that gives the appearance of space itself. The screen displays an image of blackness and stars. Also, the screen has a special texture so that it can absorb rather than reflect radar signals. Thus, the ship cannot be detected by sight or by radar. With this screen the space ships are able to travel safely through Sronian territory.

Target Passage, Version 1

The Zylots

Feeding Patterns

Small water creatures called "Zylots" are found in lakes and streams on the planet Crilo. Zylots feed on parasitic insects found in water. This is an unusual diet because these parasites are normally quite harmful to swallow. Their habit is to attach to larger creatures and suck their blood. Thus most creatures try hard to avoid these parasites. Only the Zylots can tolerate them. This is because they can typically digest the parasites before the parasites can infest and harm the Zylot.

Troublesome Food

There is one parasite that Zylots have trouble eating. This parasite is called a Dravit. Dravits live in a newly discovered lake on Crilo. The Dravit attaches itself to the stomach lining of other creatures and rapidly sucks large quantities of the creatures blood. When the host creature dies the Dravit moves on to a new host. Many Zylots die when they consume Dravits because they cannot digest the Dravit before it sucks a lot of blood. Therefore the population of Zylots in this lake is relatively small.

However, some Zylots can survive Dravits. It turns out that very large Zylots are able to digest Dravits. This is because these large Zylots have a large enough blood supply to be able to digest the Dravit before giving up a fatal amount of blood. Even these large Zylots have to rest after eating Dravits until their blood supply is restored. But, they at least can survive the Dravits.

The Giant Lorma

The Zylots on Crilo face a second problem. This is the Giant Lorma. The Giant Lorma are large shark-like predators that live in many lakes on Crilo. The Lorma have a voracious appetite for smaller creatures in the lake and Zylots are their favorite food. If a Giant Lorma sees a Zylot it will pursue it relentlessly and devour it. Fortunately nature has given Zylots a clever way of defending themselves from the Lorma. When a Zylot sees a Lorma a greenish substance exudes from its pores and covers its skin. Since the green color is the same as that of the plants in the lake the Zylot blends in with these plants. Further, to complete the disguise, the Zylot undulates in the water with movements similar to the movements of the water-plants. With this appearance a Zylot becomes very hard to distinguish from the surrounding vegetation. *Thus, Zylots can often avoid capture by the Giant Lorma.*

Target Passage, Version 2

The Zylots

Feeding Patterns

Small water creatures called "Zylots" are found in lakes and streams on the planet Crilo. Zylots feed on parasitic insects found in water. This is an unusual diet because these parasites are normally quite harmful to swallow. Their habit is to attach to larger creatures and suck their blood. Thus most creatures try hard to avoid these parasites. Only the Zylots can tolerate them.

This is because their digestive systems are extremely adaptable. They can often develop tolerance to whatever food is available in their environment. Thus the Zylot can adapt to a parasitic insect and digest it before the parasite can harm the Zylot.

Troublesome Food

There is one parasitic insect that Zylots have trouble eating. This parasite is called a Dravit. Dravits live in a newly discovered lake on Crilo. The Dravit swiftly attaches itself to the stomach lining of other creatures and rapidly sucks large quantities of the creatures blood. When the host creature dies the Dravit moves on to a new host. Many Zylots die when they consume Dravits because they cannot adapt their digestive system before the Dravit sucks a lot of blood. Therefore the population of Zylots in this lake is relatively small.

However, some Zylots can survive Dravits. These Zylots all come from another lake upstream where they eat another parasitic insect. This other parasite, called the Behgit, is structurally very similar to a Dravit but it does not suck large quantities of blood from its host. The Zylots can adapt their digestive systems the Behgits. Then, when these Zylots arrive at the lake containing Dravits they are already adapted to the structure of Dravits. They are able to digest Dravits without having much blood sucked.

The Giant Lorma

The Zylots on Crilo face a second problem. This is the Giant Lorma. The Giant Lorma are large shark-like predators that live in many lakes on Crilo. The Lorma have a voracious appetite for smaller creatures in the lake and Zylots are their favorite food. If a Giant Lorma sees a Zylot it will pursue it relentlessly and devour it. Fortunately / nature has given Zylots a valuable way of defending itself from the Lorma. A special adrenaline in Zylots gives them an amazing sprinting ability. For short periods of time Zylots can move at incredible speeds. This can enable them to get away to safety before capture by a Lorma. However, since the Zylots can only move this quickly for short periods they must see the Lorma coming in advance and must sprint to a safe place such as inside a cave or up on the water surface. If the Zylots are caught unaware in open water they may be doomed. Nevertheless, their sprinting ability is very valuable. *Zylots can often avoid capture by the Giant Lorma.*

Analogy 4: The Glora Pond/The Keelings

Base Passage

The Glora Pond

Inside a cave on a far-away planet was an interesting Glora pond. Glora is a liquid substance common on the planet that sometimes drips from the sides of cave walls. In one cave the Glora dripped to form a pond that had a perfectly stable flow of liquid into and out of it. Every time some Glora dripped into the pond an equal amount of Glora ran out a crack in the side of the pond.

The Unfortunate Sleps

Some newt-like creatures called Sleps relied on this unchanging level of Glora in the pond. The Sleps lived in a system of tunnels that were partially filled with Glora; the tunnels ran all along the sides of the pond just at the level of the Glora. Sleps are land creatures; they cannot breath in liquid. However, their skin requires repeated moistening with Glora. Thus these tunnels were a perfect home for the Sleps. However, it was essential that the Glora remain just at the right level to only partially fill the tunnels. Unfortunately, after a time the crack in the side of the pond became clogged with sediment and Glora flowed out of the pond at a much slower rate than it came in. As a result the Glora level increased and the tunnels were then completely filled with Glora. Sadly, all the Sleps living in these tunnels died.

No More Bubble and Sparkle

When the crack in the side of the pond became clogged other changes took place. In its normal state, when Glora could still flow out through the crack at its regular rate, Glora fizzled and bubbled because it was constantly in motion. However, when the crack became clogged by sediment Glora flowed much more slowly. Because of this the Glora thickened and became more and more sticky. The once sparkling Glora increasingly resembled thick syrup. To make matters worse the stickier the Glora got the harder it was for it to get through the crack so the slower it flowed out. This in turn made it more stickier and in turn slower and so on. It was a losing proposition. Finally the Glora became a still, stagnant mess.

Target Passage, Version 1**The Keelings**

A marsh on the planet Himory was populated by some flightless birds called Keelings. The Keelings were migratory and there was a constant movement of birds in and out of the marsh. Keelings liked to nest with several feet between nests. Therefore as new birds came into the marsh, other birds, whose nestlings were already grown, would travel out. The Keeling population was almost completely stable.

The Fate of the Juma Plant

A plant called the Juma grew in this marsh. Normally the Juma and Keelings lived side by side very comfortably. Though Keelings were plant eaters they normally lived off marsh grasses and did not eat Juma. Unfortunately, one day some Keelings that migrated into the marsh were carrying a disease. As the disease began to spread among the Keelings, the Keelings instinctively began to eat Juma. It turned out that the Juma had properties that cured the disease. Unfortunately for the Juma, the disease spread very rapidly and by the time all the Keelings were cured they had eaten all the Juma. *The population of Juma was wiped out.*

Keelings Change Their Behavior

Life in this marsh underwent a second major change. An earthquake damaged a mountain pass that migrating Keelings normally used when leaving the marsh. This had a devastating effect on the behavior of the Keelings.

Before the earthquake, when Keelings could migrate normally, the Keelings were very active. They worked constantly to reorganize themselves to be properly spaced. However, after the earthquake things changed. The earthquake made the pass out of the marsh very difficult to cross and Keelings were reluctant to try crossing it. As a result the rate of migration out of the marsh decreased dramatically. When this happened Keelings became extremely depressed and sluggish because they could not tolerate over-crowding. Moreover, in this sluggish state even fewer Keelings could make their way out of the marsh. This increased the over-crowding which increased the sluggishness which lead to further over-crowding. Eventually, the Keelings began to settle themselves into hovels in the ground and maintain

the minimum activity level to stay alive. *The birds became so sluggish they were almost completely dormant.*

Target Passage, Version 2

The Keelings

A marsh on the planet Himory was populated by some flightless birds called Keelings. The Keelings were migratory and there was a constant movement of birds in and out of the marsh. Keelings liked to nest with several feet between nests. Therefore as new birds came into the marsh, other birds, whose nestlings were already grown, would travel out. The Keeling population was almost completely stable.

The Fate of the Juma Plant

A plant growing in the marsh called the Juma Plant was highly dependent on this population stability of the Keelings. The Juma Plant requires an atmosphere which has a certain level of carbon dioxide in it. Since the Keelings breathe out carbon dioxide the number of Keelings in the marsh determined the amount of carbon dioxide in the atmosphere. With the stable Keeling population the carbon dioxide level was just right for the Juma. Unfortunately, one day the migration of Keelings out the marsh slowed down. An earthquake damaged a passageway out of the marsh so that now only a few Keelings could leave the marsh at one time. However, other Keelings continued to come into the marsh as usual. Consequently, the population of Keelings in the marsh increased significantly and the atmosphere became much richer in carbon dioxide. Unfortunately for the Marsh Juma, this atmosphere was intolerable. *The population of Juma was wiped out.*

Keelings Change Their Behavior

Life in this marsh underwent a second major change. A group of Keelings who migrated into the Marsh from a distant part of the planet had picked up a microbe in their travels. This microbe had a devastating effect on the Keelings' behavior. Normally the Keelings were very active creatures, always running too and fro as they built nests and organized their social groups. But as this microbe began to grow and spread among the Keeling population in the marsh the Keelings became extremely listless. They did not even bother to space their nests at appropriate distances from one another. Eventually, the Keelings began to settle themselves into hovels in the ground

and maintain the minimum activity level to stay alive. *The birds became so sluggish they were almost completely dormant.*

APPENDIX B
SAMPLE ANALOGY USED IN EXPERIMENT 2

Tams/Robots

Base Passage

The Tams

In a far away galaxy a life form was discovered called the "Tam". The Tam is a creature that lives off minerals in rock. It attaches itself to a rock and very slowly eats it. The Tam has a hard outer shell that looks like rock and an underbelly that has a surface like sandpaper. To eat rock the Tam wriggles its underbelly back and forth to loosen the top layer of the rock. Then the belly sucks in this loosened material. After the rock is sucked in, minerals in the rock are distributed to the Tam's internal organs and to the outer shell.

Although, this wriggling, sucking, underbelly is the most unusual trait of the Tam, the Tam sometimes stops using its underbelly. When the Tam uses up all the valuable minerals in one spot on a rock the Tam enters a "creeping" phase instead of a "sucking" phase. At this time the Tam must stop using its underbelly and instead must focus its attention on creeping around the rock. The Tam creeps about the rock in order to find a new spot. As the Tam creeps it will periodically stop moving and briefly puts its underbelly back into action to check the mineral supply in the new spot. If the supply is good the Tam stays at the new spot and its underbelly resumes activity.

Rocks and underbellies

There are several varieties of Tams. Each variety has a different underbelly. These differences result from the variety of rocks that Tams live on. The underbellies differ because the surface of different rocks is more or less difficult to loosen and suck in. For example, for hard rocks with large granules Tams must have an underbelly which is coarse-grained and highly permeable. In contrast, for softer rocks Tams must have an underbelly which is finer-grained and less permeable so that not too much rock is loosened and ingested at once.

Birth and adulthood

Tams are not born with these differences. They develop differences in their underbellies by becoming uniquely adapted to a particular rock when they are young. The underbelly of all newborn Tams functions in a clumsy way; it is poor at getting the required amount of minerals. But the young Tam

gradually develops an underbelly with just the right texture and permeability for the particular rock the Tam lives on. Thus the Tam becomes specifically fitted to one kind of rock.

After a Tam has become fitted to a particular rock it is very effective at bringing in minerals. However if a Tam moved to a new kind of rock it would not be able to get the minerals needed. This is because its underbelly had become suited to the kind of rock the Tam grew up on and would not have the right texture and permeability for a new rock. Thus, the Tam cannot go from one rock to another. Scientists have so far discovered over fifty different varieties of Tams living on fifty different kinds of rocks.

Target Passage, Version 1

The Robots

Highly sophisticated explorer robots are used to investigate planets throughout the galaxy. These robots gather and interpret information or data about the life forms on a planet. The robots have sensors or "probes" that take in all kinds of data about plants and animals. These data are eventually sent to one of several micro-computers inside the robot. These micro computers interpret the data.

It is useless for a robot to continue to gather data in a particular location on a planet when the data being gathered are no longer new. Therefore, when a robot no longer finds new information in a certain area it starts a journey period. The robot activates its mobility units and begins moving across the planet. Every so often the Robot will decide whether the data in the current location are the same as data previously gathered. If they are the same as old data the robot will move on. If the data are new the robot will remain in that location.

Design

All of the explorer robots are built with all-purpose "super-probes" that are extremely sophisticated. Any robot can automatically function efficiently on any planet it happens to be sent to by the home planet. The robots are equipped with super-probes that automatically can work efficiently under any lighting or atmospheric conditions. This means the robots never have to spend time adapting to a planet and they do not become specialized for one planet.

These super-probes are extremely sensitive and delicate. This means that the jostling of space travel can easily damage the probes during the flight to a planet. Consequently, elaborate packing is required to ship the robots safely to a planet. When a robot arrives at a planet it must carefully unpack itself. A robot is not able to repack its probes in the elaborate way needed to survive another space flight. When a robot returns to the home planet the probes are in a state of disrepair.

The use of explorer robots has dramatically increased the knowledge available about planets in the galaxy. The robots are able to go to places uninhabitable by humans.

Target Passage, Version 2

The Robots

Highly sophisticated explorer robots are used to investigate planets throughout the galaxy. These robots gather and interpret information or data about the life forms on a planet. Sensors or "probes" on the exterior of the Robots gather all kinds of data about plants and animals. After being gathered by the probes, these data are eventually sent to one of several micro-computers inside the robot. The micro-computers interpret the data.

The microcomputers inside the robot have extremely sensitive electrical circuitry . As a consequence, they can over-heat if they receive and process a great deal of complex data. This overheating can lead to serious and costly errors in the interpretation of the data; such errors could render all of the data gathered by the robot completely meaningless. Because of this possibility, procedures are instituted so that the microcomputers can stop receiving data for awhile and therefore temporarily stop interpreting the data. This period in which they don't have to process data allows them to cool and thus to maintain accurate functioning.

Design

The explorer robots differ from one another. Specifically, their probes differ depending on which planet they are sent to. The atmospheric and lighting conditions on a particular planet affect the ease and manner with which information can be detected by a probe. Thus probes must have different types of filters and degrees of sensitivity depending on the conditions existing on a planet.

A robot must generate the right filters and sensitivities of probes after it arrives on the planet. A robot's probes are originally designed to work in a very general way; as a consequence they are inefficient. However, the robots' probes eventually become highly efficient. Once the robot gets started on a planet the probes form the particular filters and sensitivity needed given the conditions on that planet. Thus each robot becomes specialized for its planet. The ability of these robots' probes to become specialized and efficient is a significant advance over earlier models of robots.

The use of the explorer robots has dramatically increased the knowledge available about planets in the galaxy. The robots are able to go to places uninhabitable by humans.

Distribution List [Illinois/Gentner] NR 887-551

Dr. Phillip L. Ackerman
University of Minnesota
Department of Psychology
75 East River Road
N218 Elliott Hall
Minneapolis, MN 55455

Dr. Beth Adelson
Department of Computer Science
Tufts University
Medford, MA 02155

AFOSR,
Life Sciences Directorate
Bolling Air Force Base
Washington, DC 20332

Dr. Robert Ahlers
Code N711
Human Factors Laboratory
Naval Training Systems Center
Orlando, FL 32813

Dr. John R. Anderson
Department of Psychology
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Dr. Stephen J. Andriole, Chairman
Department of Information Systems
and Systems Engineering
George Mason University
4400 University Drive
Fairfax, VA 22030

Technical Director, ARI
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Patricia Baggett
School of Education
610 E. University, Rm 1302D
University of Michigan
Ann Arbor, MI 48109-1250

Dr. Eva L. Baker
UCLA Center for the Study
of Evaluation
145 Moore Hall
University of California
Los Angeles, CA 90024

Dr. Meryl S. Baker
Navy Personnel R&D Center
San Diego, CA 92152-6800

prof. dott. Bruno G. Bara
Unità di ricerca di
intelligenza artificiale
Università di Milano
20122 Milano - via F. Sforza 23
ITALY

Dr. William M. Bart
University of Minnesota
Dept. of Educ. Psychology
330 Burton Hall
178 Pillsbury Dr., S.E.
Minneapolis, MN 55455

Leo Beltracchi
United States Nuclear
Regulatory Commission
Washington DC 20555

Dr. Gautam Biswas
Department of Computer Science
Box 1688, Station B
Vanderbilt University
Nashville, TN 37235

Dr. John Black
Teachers College, Box 8
Columbia University
525 West 120th Street
New York, NY 10027

Dr. Sue Bogner
Army Research Institute
ATTN: PERI-SF
5001 Eisenhower Avenue
Alexandria, VA 22333-5800

Dr. Jeff Bonar
Learning R&D Center
University of Pittsburgh
Pittsburgh, PA 15260

Dr. Gordon H. Bower
Department of Psychology
Stanford University
Stanford, CA 94306

Dr. Robert Breaux
Code 7B
Naval Training Systems Center
Orlando, FL 32813-7100

Dr. Ann Brown
Center for the Study of Reading
University of Illinois
51 Gerty Drive
Champaign, IL 61280

Dr. John S. Brown
XEROX Palo Alto Research
Center
3333 Coyote Road
Palo Alto, CA 94304

Dr. John T. Bruer
James S. McDonnell Foundation
Suite 1610
1034 South Brentwood Blvd.
St. Louis, MO 63117

Dr. Bruce Buchanan
Computer Science Department
Stanford University
Stanford, CA 94305

LT COL Hugh Burns
AFHRL/IDI
Brooks AFB, TX 78235

Dr. Joseph C. Campione
Center for the Study of Reading
University of Illinois
51 Gerty Drive
Champaign, IL 61820

Dr. Joanne Capper, Director
Center for Research into Practice
1718 Connecticut Ave., N.W.
Washington, DC 20009

Dr. Jaime G. Carbonell
Computer Science Department
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Dr. Susan Carey
Department of Cognitive
and Neural Science
MIT
Cambridge, MA 02139

Distribution List [Illinois/Gentner] NR 667-551

Dr. Pat Carpenter
Carnegie-Mellon University
Department of Psychology
Pittsburgh, PA 15213

CDR Robert Carter
Office of the Chief
of Naval Operations
OP-933D4
Washington, DC 20350-2000

Chair, Dept. of Psychology
College of Arts and Sciences
Catholic Univ. of America
Washington, DC 20064

Dr. Fred Chang
Pacific Bell
2600 Camino Ramon
Room 3S-450
San Ramon, CA 94583

Dr. Davida Charney
English Department
Penn State University
University Park, PA 16802

Dr. Micheline Chi
Learning R & D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15260

Professor Chu Tien-Chen
Mathematics Department
National Taiwan University
Taipei, TAIWAN

Dr. William Clancey
Institute for Research
on Learning
3333 Coyote Hill Road
Palo Alto, CA 94304

Dr. Charles Clifton
Tobia Hall
Department of Psychology
University of
Massachusetts
Amherst, MA 01003

Assistant Chief of Staff
for Research, Development,
Test, and Evaluation
Naval Education and
Training Command (N-5)
NAS Pensacola, FL 32508

Dr. Allan M. Collins
Bolt Beranek & Newman, Inc.
10 Moulton Street
Cambridge, MA 02238

Dr. Stanley Collyer
Office of Naval Technology
Code 222
800 N. Quincy Street
Arlington, VA 22217-5000

Brian Dallman
Training Technology Branch
3400 TCHTW/TTGXC
Lowry AFB, CO 80230-5000

Georg Delacote
Directeur de L'informatique
Scientifique et Technique
CNRS
15, Quai Anatole France
75706 Paris, FRANCE

Dr. Denise Dellarosa
Psychology Department
Box 11A, Yale Station
Yale University
New Haven, CT 06520-7447

Dr. Thomas E. DeZern
Project Engineer, AI
General Dynamics
PO Box 748/Mail Zone 2646
Fort Worth, TX 76101

Dr. Andrea di Sessa
University of California
School of Education
Tolman Hall
Berkeley, CA 94720

Dr. R. K. Dismukes
Associate Director for Life Sciences
AFOSR
Bolling AFB
Washington, DC 20332

Defense Technical
Information Center
Cameron Station, Bldg 5
Alexandria, VA 22314
Attn: TC
(12 Copies)

Dr. Thomas M. Duffy
Communications Design
Center, 160 BH
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Dr. Richard Duran
Graduate School of Education
University of California
Santa Barbara, CA 93106

Dr. John Ellis
Navy Personnel R&D Center
Code 51
San Diego, CA 92252

Dr. Susan Epstein
144 S. Mountain Avenue
Montclair, NJ 07042

ERIC Facility-Acquisitions
4350 East-West Hwy., Suite 1100
Bethesda, MD 20814-4475

Dr. Beatrice J. Farr
Army Research Institute
PERI-IC
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Marshall J. Farr, Consultant
Cognitive & Instructional Sciences
2520 North Vernon Street
Arlington, VA 22207

Dr. Paul Feltovich
Southern Illinois University
School of Medicine
Medical Education Department
P.O. Box 3926
Springfield, IL 62708

Distribution List [Illinois/Gentner] NR 887-551

Mr. Wallace Fearsig
Educational Technology
Bolt Beranek & Newman
10 Moulton St.
Cambridge, MA 02238

Chair, Department of
Psychology
Georgetown University
Washington, DC 20057

Dr. Dik Gregory
Admiralty Research
Establishment/AXB
Queens Road
Teddington
Middlesex, ENGLAND TW110LN

Dr. Gerhard Fischer
University of Colorado
Department of Computer Science
Boulder, CO 80309

Dr. Robert Glaser
Learning Research
& Development Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15260

Dr. Gerhard Grosseing
Atominstitut
Schuttelstrasse 115
Vienna
AUSTRIA A-1020

Dr. J. D. Fletcher
Institute for Defense Analyses
1801 N. Beauregard St.
Alexandria, VA 22311

Dr. Arthur M. Glenberg
University of Wisconsin
W. J. Brogden Psychology Bldg.
1202 W. Johnson Street
Madison, WI 53706

Prof. Edward Haertel
School of Education
Stanford University
Stanford, CA 94305

Dr. Linda Flower
Carnegie-Mellon University
Department of English
Pittsburgh, PA 15213

Dr. Sam Glucksberg
Department of Psychology
Princeton University
Princeton, NJ 08540

Dr. Henry M. Half
Half Resources, Inc.
4918 33rd Road, North
Arlington, VA 22207

Dr. Kenneth D. Forbus
University of Illinois
Department of Computer Science
1304 West Springfield Avenue
Urbana, IL 61801

Dr. Susan R. Goldman
Dept. of Education
University of California
Santa Barbara, CA 93106

Dr. Ronald K. Hambleton
University of Massachusetts
Laboratory of Psychometric
and Evaluative Research
Hills South, Room 152
Amherst, MA 01003

Dr. Barbara A. Fox
University of Colorado
Department of Linguistics
Boulder, CO 80309

Dr. Sherrie Gott
AFHRL/MOMJ
Brooks AFB, TX 78235-5601

Dr. Bruce W. Hamill
Research Center
The Johns Hopkins University
Applied Physics Laboratory
Johns Hopkins Road
Laurel, MD 20707

Dr. John R. Frederiksen
BBN Laboratories
10 Moulton Street
Cambridge, MA 02238

Dr. T. Govindaraj
Georgia Institute of
Technology
School of Industrial
and Systems Engineering
Atlanta, GA 30332-0205

Stevan Harnad
Editor, The Behavioral and
Brain Sciences
20 Nassau Street, Suite 240
Princeton, NJ 08542

Dr. Norman Frederiksen
Educational Testing Service
(05-R)
Princeton, NJ 08541

Dr. Wayne Gray
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Reid Hastie
Northwestern University
Department of Psychology
Evanston, IL 60208

Julie A. Gadsden
Information Technology
Applications Division
Admiralty Research Establishment
Portsmouth, Portsmouth PO6 4AA
UNITED KINGDOM

Dr. James G. Greeno
School of Education
Stanford University
Room 311
Stanford, CA 94305

Dr. John R. Hayes
Carnegie-Mellon University
Department of Psychology
Schenley Park
Pittsburgh, PA 15213

Dr. Dedre Gentner
University of Illinois
Department of Psychology
603 E. Daniel St.
Champaign, IL 61820

Distribution List [Illinois/Gentner] NR 067-551

Dr. Barbara Hayes-Roth
Knowledge Systems Laboratory
Stanford University
701 Welch Road
Palo Alto, CA 94304

Dr. Frederick Hayes-Roth
Teknowledge
P.O. Box 10119
1850 Embarcadero Rd.
Palo Alto, CA 94303

Dr. James D. Hollan
MCC
3500 W. Balcones Ctr. Dr.
Austin, TX 78759

Dr. Melissa Holland
Army Research Institute for the
Behavioral and Social Sciences
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Keith Holyoak
Department of Psychology
University of California
Los Angeles, CA 90024

Ms. Julia S. Hough
Lawrence Erlbaum Associates
110 W. Harvey Street
Philadelphia, PA 19144

Dr. Ed Hutchins
Intelligent Systems Group
Institute for
Cognitive Science (C-015)
UCSD
La Jolla, CA 92093

Dr. Barbara Hutson
Virginia Tech
Graduate Center
2990 Talenstar Ct.
Falls Church, VA 22042

Dr. Alice M. Isen
Department of Psychology
University of Maryland
Catonville, MD 21228

Dr. Janet Jackson
Rijksuniversiteit Groningen
Biologisch Centrum, Vleugel D
Kerklaan 30, 9751 NN Haren
The NETHERLANDS

Dr. Robert Jaanarone
Elec. and Computer Eng. Dept.
University of South Carolina
Columbia, SC 29208

Dr. Claude Janvier
Universite' du Quebec a Montreal
P.O. Box 8888, succ. A
Montreal, Quebec H3C 3P8
CANADA

Dr. Robin Jeffries
Hewlett-Packard Laboratories, 3L
P.O. Box 10490
Palo Alto, CA 94303-0971

Chair, Department of
Psychology
The Johns Hopkins University
Baltimore, MD 21218

Dr. Douglas H. Jones
Thatcher Jones Associates
P.O. Box 6646
16 Trafalgar Court
Lawrenceville, NJ 08648

Dr. Marcel Just
Carnegie-Mellon University
Department of Psychology
Schenley Park
Pittsburgh, PA 15213

Dr. Daniel Kahneman
Department of Psychology
University of California
Berkeley, CA 94720

Dr. Ruth Kanfer
University of Minnesota
Department of Psychology
Elliott Hall
75 E. River Road
Minneapolis, MN 55455

Dr. Milton S. Katz
European Science Coordination
Office
U.S. Army Research Institute
Box 65
FPO New York 09510-1500

Dr. Frank Keil
Department of Psychology
228 Uris Hall
Cornell University
Ithaca, NY 14850

Dr. Wendy Kellogg
IBM T. J. Watson Research Ctr.
P.O. Box 704
Yorktown Heights, NY 10598

Dr. Dennis Kibler
University of California
Department of Information
and Computer Science
Irvine, CA 92717

Dr. David Kieras
Technical Communication Program
TIDAL Bldg., 2360 Bonisteel Blvd.
University of Michigan
Ann Arbor, MI 48109-2108

Dr. J. Peter Kieras
Army Research Institute
Orlando Field Unit
c/o PM TRADE-E
Orlando, FL 32813

Dr. Walter Kintsch
Department of Psychology
University of Colorado
Boulder, CO 80309-0345

Dr. David Klahr
Carnegie-Mellon University
Department of Psychology
Schenley Park
Pittsburgh, PA 15213

Dr. Janet L. Kolodner
Georgia Institute of Technology
School of Information
& Computer Science
Atlanta, GA 30332

Distribution List [Illinois/Gentner] NR 887-551

Dr. Kenneth Kotovsky
Community College of
Allegheny County
808 Ridge Avenue
Pittsburgh, PA 15212

Dr. David H. Krantz
Department of Psychology
Columbia University
406 Schermerhorn Hall
New York, NY 10027

Dr. Benjamin Kuipers
University of Texas at Austin
Department of Computer Sciences
Taylor Hall 2.124
Austin, Texas 78712

Dr. David R. Lambert
Naval Ocean Systems Center
Code 772
271 Catalina Boulevard
San Diego, CA 92152-5000

Dr. Pat Langley
University of California
Department of Information
and Computer Science
Irvine, CA 92717

Dr. Marcy Lansman
University of North Carolina
The L. L. Thurstone Lab.
Davis Hall CB #3270
Chapel Hill, NC 27514

Dr. Jill Larkin
Carnegie-Mellon University
Department of Psychology
Pittsburgh, PA 15213

Dr. Jean Lave
Institute for Research
on Learning
3333 Coyote Hill Road
Palo Alto, CA 92304

Dr. Robert W. Lawler
Matthews 118
Purdue University
West Lafayette, IN 47907

Dr. Alan M. Lesgold
Learning R&D Center
University of Pittsburgh
Pittsburgh, PA 15260

Dr. Jim Levin
Department of
Educational Psychology
210 Education Building
1310 South Sixth Street
Champaign, IL 61820-6990

Dr. John Levine
Learning R&D Center
University of Pittsburgh
Pittsburgh, PA 15260

Dr. Michael Levine
Educational Psychology
210 Education Bldg.
University of Illinois
Champaign, IL 61801

Dr. Clayton Lewis
University of Colorado
Department of Computer Science
Campus Box 430
Boulder, CO 80309-0430
CE

Matt Lewis
Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213

Library
Naval Training Systems Center
Orlando, FL 32813

Library
Naval War College
Newport, RI 02940

Science and Technology Division
Library of Congress
Washington, DC 20540

Dr. Jane Malin
Mail Code EF5
NASA Johnson Space Center
Houston, TX 77058

Dr. William L. Maloy
Code 04
NETPMSA
Pensacola, FL 32509-5000

Dr. Elaine Marsh
Naval Center for Applied Research
in Artificial Intelligence
Naval Research Laboratory
Code 5510
Washington, DC 20376-5000

Dr. Sandra P. Marshall
Dept. of Psychology
San Diego State University
San Diego, CA 92182

Dr. Manton M. Matthews
Department of Computer Science
University of South Carolina
Columbia, SC 29208

Dr. Richard E. Mayer
Department of Psychology
University of California
Santa Barbara, CA 93106

Dr. Joseph C. McLachlan
Code 52
Navy Personnel R&D Center
San Diego, CA 92152-6800

Dr. James McMichael
Technical Director
Navy Personnel R&D Center
San Diego, CA 92152-6800

Dr. Barbara Means
SRI International
333 Ravenswood Avenue-
Menlo Park, CA 94025

Dr. Douglas L. Medin
Department of Psychology
University of Illinois
603 E. Daniel Street
Champaign, IL 61820

Dr. George A. Miller
Dept. of Psychology
Green Hall
Princeton University
Princeton, NJ 08540

Distribution List [Illinois/Gentner] NR 667-551

Dr. William Montague
NPRDC Code 13
San Diego, CA 92152-6800

Director, Manpower and Personnel
Laboratory,
NPRDC (Code 06)
San Diego, CA 92152-6800

Office of Naval Research,
Code 1142CS
800 N. Quincy Street
Arlington, VA 22217-5000
(6 Copies)

Dr. Randy Mumaw
Training Research Division
HumRRO
1100 S. Washington
Alexandria, VA 22314

Director, Human Factors
& Organizational Systems Lab,
NPRDC (Code 07)
San Diego, CA 92152-6800

Office of Naval Research,
Code 1142PS
800 N. Quincy Street
Arlington, VA 22217-5000

Dr. Allen Munro
Behavioral Technology
Laboratories - USC
1845 S. Elema Ave., 4th Floor
Redondo Beach, CA 90277

Library, NPRDC
Code P201L
San Diego, CA 92152-6800

Psychologist
Office of Naval Research
Branch Office, London
Box 39
FPO New York, NY 09510

Chair, Department of
Computer Science
U.S. Naval Academy
Annapolis, MD 21402

Technical Director
Navy Personnel R&D Center
San Diego, CA 92152-6800

Special Assistant for Marine
Corps Matters,
ONR Code 00MC
800 N. Quincy St.
Arlington, VA 22217-5000

Dr. Allen Newell
Department of Psychology
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Commanding Officer,
Naval Research Laboratory
Code 2627
Washington, DC 20390

Dr. Judith Orasanu
Basic Research Office
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Richard E. Nisbett
University of Michigan
Institute for Social Research
Room 5261
Ann Arbor, MI 48106

Dr. Harold F. O'Neil, Jr.
School of Education - WP21 801
Department of Educational
Psychology & Technology
University of Southern California
Los Angeles, CA 90089-0031

Dr. James Paulson
Department of Psychology
Portland State University
P.O. Box 751
Portland, OR 97207

Dr. A. F. Norcio
Code 5530
Naval Research Laboratory
Washington, DC 20375-5000

Dr. Stellan Ohlsson
Learning R & D Center
University of Pittsburgh
Pittsburgh, PA 15260

Military Assistant for Training and
Personnel Technology,
OUSD (R & E)
Room 3D129, The Pentagon
Washington, DC 20301-3080

Dr. Donald A. Norman
C-015
Institute for Cognitive Science
University of California
La Jolla, CA 92093

Office of Naval Research
Code 1133
800 North Quincy Street
Arlington, VA 22217-5000

Dr. David N. Perkins
Project Zero
Harvard Graduate School
of Education
7 Appian Way
Cambridge, MA 02138

Deputy Technical Director
NPRDC Code 01A
San Diego, CA 92152-6800

Office of Naval Research,
Code 114J
800 N. Quincy St.
Arlington, VA 22217-5000

Dr. Nancy N. Perry
Naval Education and Training
Program Support Activity
Code-047
Building 2435
Pensacola, FL 32509-5000

Director, Training Laboratory,
NPRDC (Code 05)
San Diego, CA 92152-6800

Office of Naval Research,
Code 1142BI
800 N. Quincy Street
Arlington, VA 22217-5000

Distribution List [Illinois/Gentner] NR 667-551

Department of Computer Science,
Naval Postgraduate School
Monterey, CA 93940

Dr. Steven Pinker
Department of Psychology
E10-018
MIT
Cambridge, MA 02139

Dr. Tjeerd Plomp
Twente University of Technology
Department of Education
P.O. Box 217
7500 AE ENSCHEDE
THE NETHERLANDS

Dr. Steven E. Poltrock
MCC
3500 West Balcones Center Dr.
Austin, TX 78759-6509

Dr. Harry E. Pople
University of Pittsburgh
Decision Systems Laboratory
1360 Scaife Hall
Pittsburgh, PA 15261

Dr. Mary C. Potter
Department of Brain and
Cognitive Sciences
MIT (E-10-039)
Cambridge, MA 02139

Dr. Joseph Paoika
ATTN: PERI-IC
Army Research Institute
5001 Eisenhower Ave.
Alexandria, VA 22333-5600

Dr. Lynne Rader
Department of Psychology
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Dr. Steve Rader
Northwest Regional
Educational Laboratory
400 Lindsay Bldg.
710 S.W. Second Ave.
Portland, OR 97204

Dr. James A. Reggia
University of Maryland
School of Medicine
Department of Neurology
22 South Greene Street
Baltimore, MD 21201

Dr. J. Wesley Regian
AFHRL/IDI
Brooks AFB, TX 78235

Dr. Fred Reif
Physics Department
University of California
Berkeley, CA 94720

Dr. Lauren Resnick
Learning R & D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15213

Dr. Gilbert Ricard
Mail Stop K02-14
Grumman Aircraft Systems
Bethpage, NY 11787

Dr. Linda G. Roberts
Science, Education, and
Transportation Program
Office of Technology Assessment
Congress of the United States
Washington, DC 20510

Dr. William B. Rouse
Search Technology, Inc.
4725 Peachtree Corners Circle
Suite 200
Norcross, GA 30092

Dr. Roger Schank
Yale University
Computer Science Department
P.O. Box 2158
New Haven, CT 06520

Dr. Alan H. Schoenfeld
University of California
Department of Education
Berkeley, CA 94720

Dr. Janet W. Schofield
816 LRDC Building
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15290

Dr. Judith W. Segal
OERI
555 New Jersey Ave., NW
Washington, DC 20208

Dr. Colleen M. Seifert
Institute for Cognitive Science
Mail Code C-015
University of California, San Diego
La Jolla, CA 92093

Dr. Ben Shneiderman
Dept. of Computer Science
University of Maryland
College Park, MD 20742

Dr. Lee S. Shulman
School of Education
507 Ceras
Stanford University
Stanford, CA 94305-3084

Dr. Robert S. Siegler
Carnegie-Mellon University
Department of Psychology
Schenley Park
Pittsburgh, PA 15213

Dr. Derek Sleeman
Computing Science Department
King's College
Old Aberdeen AB9 2UB
Scotland
UNITED KINGDOM

Dr. Richard E. Snow
School of Education
Stanford University
Stanford, CA 94305

Dr. Elliot Soloway
Yale University
Computer Science Department
P.O. Box 2158
New Haven, CT 06520

Distribution List [Illinois/Gentner] NR 667-551

Dr. Richard C. Sorensen
Navy Personnel R&D Center
San Diego, CA 92152-6800

Headquarters, U. S. Marine Corps
Code MPI-20
Washington, DC 20380

Dr. Wallace Wulfeck, III
Navy Personnel R&D Center
Code 51
San Diego, CA 92152-6800

Dr. Kathryn T. Spoehr
Brown University
Department of Psychology
Providence, RI 02912

Dr. Kurt Van Lehn
Department of Psychology
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Dr. Masoud Yazdani
Dept. of Computer Science
University of Exeter
Prince of Wales Road
Exeter EX44PT
ENGLAND

Dr. Robert J. Sternberg
Department of Psychology
Yale University
Box 11A, Yale Station
New Haven, CT 06520

Dr. Jerry Vogt
Navy Personnel R&D Center
Code 51
San Diego, CA 92152-6800

Mr. Carl York
System Development Foundation
1 Maritime Plaza, #1770
San Francisco, CA 94111

Dr. Thomas Sticht
Applied Behavioral and
Cognitive Sciences, Inc.
P. O. Box 6640
San Diego, CA 92106

Dr. Beth Warren
BBN Laboratories, Inc.
10 Moulton Street
Cambridge, MA 02238

Dr. Joseph L. Young
National Science Foundation
Room 320
1800 G Street, N.W.
Washington, DC 20550

Dr. John Tangney
AFOSR/NL, Bldg. 410
Bolling AFB, DC 20332-6448

Dr. Keith T. Wescoart
FMC Corporation
Central Engineering Labs
1205 Coleman Ave., Box 580
Santa Clara, CA 95052

Dr. Kikami Tatsuoka
CERL
252 Engineering Research
Laboratory
103 S. Mathews Avenue
Urbana, IL 61801

Dr. Douglas Wetsel
Code 51
Navy Personnel R&D Center
San Diego, CA 92152-6800

Dr. Perry W. Thorndyke
FMC Corporation
Central Engineering Labs
1205 Coleman Avenue, Box 580
Santa Clara, CA 95052

Dr. Barbara White
BBN Laboratories
10 Moulton Street
Cambridge, MA 02238

Dr. Martin A. Tolcott
3001 Veasey Terr., N.W.
Apt. 1817
Washington, DC 20008

Dr. Robert A. Wisher
U.S. Army Institute for the
Behavioral and Social Sciences
5001 Eisenhower Avenue
Alexandria, VA 22333-5600

Dr. Douglas Towne
Behavioral Technology Labs
University of Southern California
1845 S. Elena Ave.
Redondo Beach, CA 90277

Dr. Martin F. Wiskoff
Defense Manpower Data Center
550 Camino El Estero
Suite 200
Monterey, CA 93943-3231

Chair, Department of
Computer Science
Towson State University
Towson, MD 21204

Mr. John H. Wolfe
Navy Personnel R&D Center
San Diego, CA 92152-6800