Structural constraints and object similarity in analogical mapping and inference

Daniel C. Krawczyk
University of California, Berkeley, USA

Keith J. Holyoak and John E. Hummel
University of California, Los Angeles, USA

Theories of analogical reasoning have viewed relational structure as the dominant determinant of analogical mapping and inference, while assigning lesser importance to similarity between individual objects. An experiment is reported in which these two sources of constraints on analogy are placed in competition under conditions of high relational complexity. Results demonstrate equal importance for relational structure and object similarity, both in analogical mapping and in inference generation. The human data were successfully simulated using a computational analogy model (LISA) that treats both relational correspondences and object similarity as soft constraints that operate within a limited-capacity working memory; but not with a model (SME) that treats relational structure as pre-eminent.

Analogies provide a valuable source of new inferences and a means of expanding knowledge. Analogical reasoning is generally viewed as involving four major subprocesses: (1) retrieving an appropriate *source* analogue from long-term memory to compare with a novel *target* analogue, (2) mapping elements of the two analogues, (3) making inferences about the target as a function of its mapping to the source, and (4) using the source and target together to induce a more general schema or rule (e.g., Carbonell, 1983; Gentner, 1989; Holyoak, Novick & Melz, 1994). Finding a coherent mapping between the source and target is essential for generating useful

Correspondence should be sent to Daniel C. Krawczyk, Department of Psychology and the Helen Wills Neuroscience Institute, University of California, Berkeley, 3210 Tolman Hall, Berkeley, CA 94720, USA. Email: krawczyk@uclink.berkeley.edu

This project was supported by National Science Foundation grant SBR-9729023. We thank Wei-Chung Ooi, and Carla Webster for help running subjects and for the useful feedback. We thank Art Markman for running several SME simulations and for useful comments on an earlier draft, and Dedre Gentner for helpful discussions and comments on an earlier draft.

http://www.tandf.co.uk/journals/pp/13546783.html DOI: 10.1080/13546780342000043

inferences about the target, as inferences are necessarily based on the correspondences established during the mapping process.

Analogical mapping appears to be guided by three primary types of constraints (Holyoak & Thagard, 1989).

(1) The structural constraint of *isomorphism* implies (a) structural consistency, such that elements that correspond in one context should correspond in all contexts, and (b) one-to-one mapping, such that any element of one analogue should correspond to exactly one element in the other (Gentner, 1983).
(2) *Semantic similarity* implies that semantically similar elements (e.g., joint members of a taxonomic category) should tend to map to each other.
(3) *Pragmatic centrality* implies that mapping should give preference to elements that are especially important to goal attainment, and should try to maintain correspondences that can be presumed on the basis of prior knowledge.

Empirical evidence indicates that people's preferred analogical mappings tend to honour each of these constraints (see, e.g., Gentner & Toupin, 1986; Holyoak & Thagard, 1989; Keane, 1988; Markman & Gentner, 1993; Spellman & Holyoak, 1996), provided they can be jointly satisfied within the finite processing resources of working memory (Keane, Ledgeway & Duff, 1994; Kubose, Holyoak & Hummel, 2002).

However, the three types of constraint need not always favour the same set of mappings (e.g., Gentner & Toupin, 1986; Markman & Gentner, 1993; Ross, 1987, 1989; Spellman & Holyoak, 1996). For example, semantic constraints (more specifically, similarity of objects based on shared features, rather than corresponding roles) may favour one set of mappings, and structural or pragmatic constraints may favour another. In prior investigations of mapping conflict, story analogies have been used widely (Gentner & Toupin, 1986; Spellman & Holyoak, 1996), while other investigators have used picture analogies involving two analogous scenes with three to four objects included (Markman, 1996; Markman & Gentner, 1993). Most theories of analogical mapping tend to emphasise isomorphism as the most important constraint, especially as compared with object similarity. Indeed, object similarity is frequently viewed as a "foil", relative to which "true analogy" is measured. Most current analogy models are based on explicit graph matching of relational correspondences (e.g., the Structure Mapping Engine, SME, of Falkenhainer, Forbus & Gentner, 1989) or massively parallel constraint satisfaction (e.g., the ACME model of Holyoak & Thagard, 1989). In such models, structural similarity of relations will typically dominate local similarity of objects when these constraints conflict.

Indeed, the basic purpose of both graph matching and massively parallel constraint satisfaction is that such algorithms make it possible to find globally optimal structural mappings, even when those mappings are inconsistent with local constraints based on object similarity. In contrast, the LISA model (Learning and Inference with Schemas and Analogies; Hummel & Holyoak, 1997, 2003) operates within a limited-capacity working memory that will not always be sufficient to allow computation of the optimal structural mapping. LISA uses semantic similarity (of both predicates and objects) as an important initial guide to the identification of corresponding elements of the source and target. This aspect of the model will tend to allow for mappings based on attributes rather than relations. The working memory limitation follows from the model's synchronous activation of information, as there is a limit on the number of role-filler bindings that can be maintained out of synchrony with each other. It follows that for complex or large analogies the limitation on maintaining asynchrony of activation for possible role fillers will not allow an exhaustive consideration of all possible mappings and may leave certain aspects of the relational structure unsatisfied. Limited working memory for role bindings, coupled with the fact that object attributes are often an initial biasing factor in candidate mappings in LISA, will tend to lead the model to map based on attributes under some conditions, despite the fact that additional constraining relational information may have been available. Evidence from frontal-lobe damaged patients with working memory deficits supports this hypothesis, as these patients show a tendency to map based on attributes rather than relations compared to healthy control participants (Morrison et al., in press). In the present study the complexity of information within the problem is sufficient to lead to attribute mappings due to working memory limitations, despite the fact that the materials are available to the participants throughout the task.

The type of similarity shared by the objects is also relevant to the way in which mappings may be formed. Attribute similarity tends to be the focus of attention when similar objects, such as two cars or types of fruit, are compared, whereas more integrative processing may be emphasised when entities are not similar in attributes, but share some common thematic link (Wisniewski & Bassok, 1999). Empirical evidence suggests that comparison processes dominate when objects are similar, whereas thematic integration is emphasised when objects are dissimilar but are related through causal or relational roles. The tendency for both thematic and attribute comparisons to be made under different conditions (Bassok & Medin, 1997) may be related to the distinction between attribute and relational mappings. The present experiment includes similar objects (several human employees) that may encourage attribute comparisons, but these objects are interrelated by roles (cheating relations) that may encourage relational integration.

The present experiment sought to investigate whether there are circumstances under which local object similarity can override global relational similarity in determining reasoners' preferred mappings and their subsequent inferences. The experiment used a complex social analogy, in which relational similarity suggests that one set of elements will be involved in a subsequent analogical inference, whereas local similarity among the elements would potentially lead to the inclusion of a different set of elements in the inference. The materials allow a variety of inferences, some based on local object similarity and others on global relations, thereby making it possible to determine whether object similarity or relational similarity dominates as the preferred basis for inference when the constraints conflict. We predicted that due to their limited-capacity working memory, people will be at least equally likely to use object similarity as the basis for mapping in this particular analogy problem. Simulations of the problem using both the SME and LISA models are reported and compared to the results obtained with our human subjects.

METHOD

Participants

Participants were 48 students at the University of California, Los Angeles (one participant per condition, including counterbalancing). Participation in the experiment partially fulfilled a course requirement for an introductory psychology class. Subjects were run in groups ranging from one to five.

Materials, design and procedure

The materials were a pair of stories describing various people at a company, including a CEO and several employees, and situations in which one employee cheated another out of due credit for work. In order to ensure that people attended to and remembered the object attributes, we generated elements in which the traits of the characters were both rare and distinctive (e.g., being an ex-astronaut or ex-professional wrestler). The critical comparison was between the number of inferences based at least in part on object similarity and the number of inferences based purely on global relational structure.

The materials were presented in a 13-page booklet. The first page contained a brief overview of the contents, explaining that subjects would read two stories about corporations, followed by some tasks to be completed. Subjects were instructed that they should read the entire text and that they would be allowed to refer back to earlier pages in completing the task. The names of the characters and corporations, as well as the orderings of the texts, were counterbalanced to control for memorability and order effects (see Figure 1 for a schematic description of the analogy problem).

Story 1 appeared on the second page of the booklet (see Appendix A). The first paragraph described a mid-sized corporation (named Brightech or Offstar) that specialised either in the manufacture of light fixtures or office supplies, depending on which counterbalanced version the subject received. In either case the business was described as successful, with a good reputation, but suffering from frequent employee conflicts. The second paragraph described how the business was unveiling a new product that two employees had developed. Each employee was male with a common name and was described as being either experienced or inexperienced, having one positive trait, one negative trait, and an unusual former occupation. These descriptors were the relevant attributes that could be used for mapping

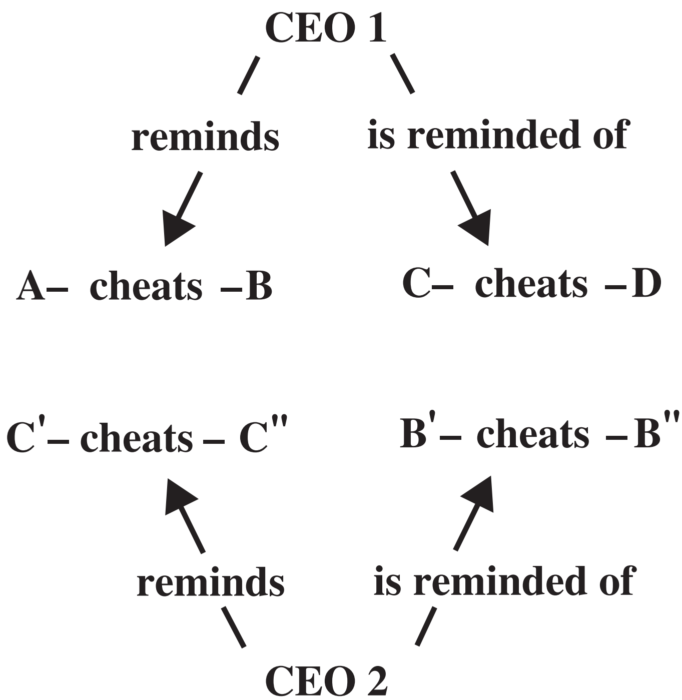

Figure 1. Structure of the stories. The letters stand for characters in the stories. All characters signified by the same letter shared identical object traits in the stories. The symbols (′) and (″) denote a new character with those same traits (e.g., C has the same traits as both C′ and C″). The “cheats” and “reminds” relations allow for structural mappings, while object similarities suggest competing, non-structural mappings.

based on semantic similarity. The story went on to describe how the new product was to be unveiled at an important meeting where one of the two employees responsible for the product cheats the other out of credit for it by claiming that it was all his idea. Although the company executives believe the claim of the cheating party, the CEO realises that the character claiming credit has cheated the other out of his fair share of acknowledgement.

The third paragraph of Story 1 describes how the CEO is reminded of a time he witnessed a similar incident, in which two other employees had done a good job on a company project, but one had cheated the other by taking full credit for the achievement.

In the final paragraph of Story 1, a company retreat is described, in which pairs of employees team up to work on new ideas for the company. The CEO makes sure that the individual who was cheated in the situation described in the first paragraph (the more recent situation in time) is teamed up with the cheater from the older situation (of which the CEO was reminded). The cheated individual in the new situation is described as being aware that his partner had taken credit for someone else's work in the past situation. The CEO then allows the cheated member of the pair to report their joint ideas at the retreat, with the intent that he will now be able to cheat the former cheater. The CEO intends to allow the cheated member of the pair to receive full credit for the ideas, as well as to punish the one who had taken more than his share of praise in the past. In this way the CEO is described as attempting to rectify inequities within the company ranks.

After reading Story 1, subjects completed a quiz on the contents of the story. This quiz consisted of seven multiple-choice questions, each of which contained four answer choices labelled *a* to *d*. Four of the questions listed a series of object traits and asked who had been associated with those traits. The answer choices for each of these trait questions included the correct answer as well as three incorrect choices of the same category as the correct answer (e.g., other employee names, but not the CEO's name). The remaining questions queried which character had cheated another named character from Story 1, and one question asked for the name of the CEO from the company. After participants completed the quiz, they were asked to correct their own quiz answers and could refer back to Story 1 in doing so.

Story 2 was designed to be highly similar to Story 1. However, its relational structure differed in several important ways (see Appendix A). In this story, the first paragraph described the company as having a good reputation, but with some disputes among the employees. Paragraph two described a new product that two employees had been responsible for developing. Critically, the surface semantic characteristics of these two employees were identical to those of the third character of Story 1 (i.e., the first employee described in the cheating situation of which the CEO was reminded). For example, if the third character of Story 1 had been an experienced manager who was friendly, careless, and formerly an astronaut, both of the Story 2 characters in this second paragraph shared these traits as well. Their names were the only descriptors that were not identical to that Story 1 character. As in Story 1, the third paragraph described how one of the characters cheats the other out of credit. Again, the board of directors believes him, but the CEO realises that he has cheated the other character.

The final paragraph of Story 2 stated that seeing the cheating situation over the new product reminded the CEO of another situation in which an employee had taken credit for work he had done with another employee (either cutting advertising costs, or landing a big contract). The two characters that had been involved in this prior cheating incident were identical (except for their names) to the second character from Story 1 (the second employee described in the cheating situation that appeared first in the text of Story 1) based on object similarity. The final part of Story 2 described the cheating that had occurred between these last two employees. Critically, no description was given of a retreat or any method used by the CEO to restore equality to the employees, as there had been in Story 1. This gap was left to be filled by inferences made by the participants.

The mapping task consisted of four pages (one per character to be mapped) that gave reasoners the opportunity to map each of the four Story 1 characters that had either cheated or been cheated (see Appendix B). Each page contained a heading that began with the word *Story 1* and the name of the character to be mapped. A reiteration of that character's traits was then provided (in order to boost the relevance of semantic information). The heading *Story 2* appeared beneath the Story 1 character description, and an inch to the right of it was the heading *Reason*. Beneath these headings were six lines in which participants were to write the names of the characters they wished to map to the Story 1 character listed at the top of the page, and give reasons to justify each mapping. The lines were spaced approximately one inch apart, and six were used in order to encourage the subject to make multiple mappings. Instructions included the statement that there might be no corresponding characters, or that there might be many, and that subjects should write down as many matches as they felt were plausible based on the contents of the two stories

Following the mapping task the subject was provided with a page for the inference task. In this task the page header read *Story 2 (the ending)*. Beneath the header, a paragraph stated that the company from Story 2 was now holding a retreat and teams of two employees would devise new strategies for the company. The CEO from Story 2 made sure that two particular employees are teamed up and that when it was time for ideas to be revealed, one character would have the opportunity to cheat the other. As in the ending to Story 1, the rationale for the CEO's behaviour was that he wished to balance the amount of credit that company employees were getting. Following the paragraph were six numbered inference spaces that read: "Who did (the CEO) pair together? ___ and ___ ." This format constrained the subject to provide only two names of characters in their inference, and was intended to prevent the subject from inferring irrelevant endings that did not use the analogy based on Story 1. Each of the six inference spaces also included a second statement reading: "Who will give the report and possibly cheat the other? ____". This question was critical in determining the degree to which the subject had used relational knowledge in making the inference, as a sufficiently detailed inference had to include information about who will cheat and who will be cheated. Subjects were instructed that they were to fill in as many answer spaces as they wished, but only those that they felt would provide plausible endings to Story 2. The final page of the booklet contained an instruction to rank each mapping in terms of how plausible it was thought to be in comparison to the others (whenever multiple inferences had been reported) by marking a number next to each inference, with 1 indicating the most highly plausible inference. If subjects had only provided one inference they were instructed to write a number 1 next to it.

RESULTS

Quizzes

Subjects obtained a mean proportion of .74 correct answers on the quiz initially; after they had corrected their quizzes a mean proportion of .99 answers were correct. These results indicate that subjects had either originally learned, or could review and identify, the correct answer to the majority of the questions.

Mapping task

The mapping task allowed subjects to report as many mappings as they wished, and each possible mapping fitted within one of several categories that corresponded to the inference task. All Story 1 characters allowed multiple plausible mappings due to the numerous surface traits that they shared with characters in Story 2, as well as the overlap in relational structure provided by the relations: *cheats, is cheated by*, and *reminds*. Mappings were divided into one of four categories, *object (O), relational (R), sensible other (SO),* and *unrelated (U)* mappings. O mappings were those that involved characters who would be relevant to the object inferences; e.g., those based primarily on object similarity relations among the objects. R mappings were those that included characters who were relevant to relational inferences; e.g., those based primarily on the relational role of the objects. SO mappings were based on both trait and relational information; e.g., "both characters were experienced managers and were cheated out of credit" and would fit Sensible Other inferences (i.e., those that were based on the story, but did not use relational or object similarity primarily as the basis for inferences, but rather a mix of both types of information; see below). Unrelated mappings did not make sense based on the stories. Examples of unrelated mappings include those in which a Story 1 character was mapped to another Story 1 character, and cases in which the reason given for the mapping was not factually correct in the stories.

Participants produced a mean of 2.36 mappings per character. For statistical analyses of types of mappings, means were tabulated based on proportions of each subject's total number of mappings.[1] A within-subjects repeated measures ANOVA indicated that there was a significant difference among mapping types, $F(3, 141) = 21.64$, $p < .01$. A Newmann-Keuls test (with alpha = .05) indicated that significantly more O mappings ($M = 0.33$), R mappings ($M = 0.30$), and SO mappings ($M = 0.32$) were made than U mappings ($M = 0.06$). No other differences were significant. It should be noted that there were 10 times as many possible SO mappings as possible O and R mappings (as each of the latter is unique); thus proportionally, O and R mappings were the leading categories.

Inference task

Inferences generated during this task were scored as belonging to different categories depending on the two names included, as well as on which of the two characters was said to give the report at the retreat and potentially cheat the other. Most combinations of names within an inference would uphold some aspects of the object similarities or relational structure; however, since few of these combinations are theoretically interesting, we focused the analysis on the two inferences that most captured the relational structure in one case and the surface similarity in the other. Table 1 describes these inference classifications.

The two most critical inferences, which we hypothesised would be reported with the highest frequency proportionally, are the pure object similarity inference O and the pure relational inference R. Inference O was

[1]The data were also analysed based on the raw numbers. This did not change the pattern of results in any way. With raw numbers the ANOVA for mappings was significant, $F(3, 141) = 34.13$, $p < .01$. Means were as follows: O ($M = 3.02$), R ($M = 2.79$), SO ($M = 2.92$), and U ($M = 0.46$). The ANOVA for inferences was significant as well, $F(3, 141) = 22.19$, $p < .01$. The means for raw numbers of inferences were: O ($M = 0.71$), R ($M = 0.65$), SO ($M = 11.60$), and U ($M = 0.08$). Proportions are presented in the Results section because they most clearly relate to the simulation results.

TABLE 1
Inference categories

O:	Employees C′ and B″ will be paired together and B″ will give the report and may cheat C′ out of credit for their work:
	(Preserves *object similarity* and *cheat* relation, partially violates *remind* relation)
R:	Employees C″ and B′ will be paired together and C″ will give the report and may cheat B′ out of credit for their work:
	(Preserves *cheat* relation and *remind* relation, violates *object similarity*)
SO:	Any of the other combinations of the Story 2 characters:
	(Each of the 10 possible inferences of this type preserves either some object similarity or relation)
U:	Any inference involving characters not mentioned in Story 2.

Note: The inference participant letters correspond to those used in Figure 1 to denote characters who will be invited and who will be giving the report.

that the first character mentioned in Story 2 (C′ from Figure 1) and the last character mentioned (character B″ from Figure 1) would be invited and character B″ would give the report. This inference preserves the object traits of the source story, while violating the relational correspondences between the two stories, as these character's roles were reversed from Story 1 with regard to the CEO reminding relation (i.e., in Story 1 the cheated character from the more recent situation will get a chance to cheat the cheater from the older situation described second). Inference R was that the second character mentioned in Story 2 (C″) would get paired with the first character from the second situation described in Story 2 (B′). This inference preserves the relational structure shown in Figure 1; however, it violates object similarity, as the characters B and C from Story 1 do not share similar traits with characters C″ and B′ from Story 2. A third relevant inference category, labelled SO (for "sensible other" inferences), included any other inferences that included characters from Story 2 in some arrangement that was sensible based on the stories. For example, the inference that characters C′ and C″ will be invited, with C′ cheating C′ in a sort of revenge scenario, would fit within the SO category. Any combination of characters from Story 2 that did not fit the S or O category were placed into the SO category. Finally, U (unrelated) inferences represented any inference that did not involve characters from Story 2.

For the purposes of analysis, we compared mean numbers of inferences of types O, R, SO, and U, after converting the numbers of inferences of each category into proportions of each subject's total number of reported inferences. A within-subjects repeated measures ANOVA indicated that there was a significant difference among inference types, $F(3, 141) = 18.78$, $p < .01$. A Newmann-Keuls test (alpha = .05) indicated that inferences O ($M = 0.28$) and R ($M = 0.23$) were reported more often than U inferences

(M = 0.02), and SO inferences (M = 0.50) were reported more often than O, R, and U inferences. It should be noted that there were 10 times as many inferences possible for the SO category than for either O or R; thus proportionally, O and R were the most favoured inferences overall.

Given that there was no difference between the O and R inferences, we performed an analysis on the rankings data to determine if one was favoured over the other. This analysis involved counting how often an inference was rated "1" (i.e., as most plausible of the inferences reported). As with the inference data, number 1 rankings for other inferences were calculated by summing all number 1 rankings not in the O or R category and converting to a proportion. Mean number 1 rankings were then compared using a within-subjects repeated measures ANOVA, which indicated a difference among the means, $F(2, 94) = 6.65$, $p < .01$. A Newmann-Keuls (alpha = .05) test indicated that inferences O (M = 0.33) and R (M = 0.23) received number 1 rankings on a significantly higher proportion of cases than did any of the other inferences (M = 0.04). Overall, the rankings results mirror the inference results, indicating that inferences O and R were essentially equivalent in being reported most often and receiving the most number 1 rankings.

Relationship between mappings and inferences

The reported mappings tended to correlate with the inferences that contained the mapped characters. Overall, a majority of the inferences of type O and type R were reported by subjects who had also fully mapped the characters that were involved in these inferences. For inference O, based on object similarity, 69% of the inferences were made by subjects who had reported mappings for all four characters involved in that inference. No inferences of this type were reported for which fewer than half of the characters had also been mapped previously. For inference R, based mostly on relational information, 56% of the inferences were made by subjects who had mapped all four characters involved in this inference; at least two characters had been mapped for 85% of this inference type. The difference between the total number of fully mapped inferences was not significantly different for inferences O and R.

Simulation results

A condensed representation of the analogy problem was run on both SME and LISA. These models were chosen as they are the most divergent with regard to the relative roles of relational and semantic constraints, with LISA being much more reliant on object similarity than is SME. Table 2 shows the numerical estimates of mapping goodness for each model compared to the

TABLE 2
Results of SME, LISA, and human reasoners

Type	*Proportion of total mappings by subjects*	*Proportion of total inferences by subjects*	*LISA mappings* (both based on a proportion of 50 runs)	*LISA inferences*	*SME (mapping score)*
Object (O)	.33	.28	.32	.26	1.01
Relational (R)	.30	.23	.32	.18	6.21
Sensible others (SO)	.32	.50	.40	.48	0.00
Unrelated (U)	.06	.02	.00	.00	0.00

The means for subjects are based on proportions of overall mappings and inferences for each subject.

data from the experiment. The models differ notably in the degree of flexibility they allow in the treatment of higher-order relations. SME rigidly enforces relational structure and especially higher-order relational structure (the "remind" relation in the current experiment). In contrast, LISA permits violations of relational structure due to its limited–capacity working memory, in turn based on its need to keep role-filler bindings out of synchrony. This aspect of the model places an inherent limit on the amount or complexity of the information that can be activated in the process of considering candidate mappings. A consequence of this limitation is a general tendency to map based on relations, but with exceptional circumstances (as created in the design of the present experiment) in which attribute matches may dominate for particularly large or complex analogies.

SME simulations. SME was run on this mapping task as follows (see Table 3). The source analogue stated that an employee, *A*, cheated another employee, *B*, that this caused *CEO-1* to be reminded that *C* cheated *D*, that *B* and *C* were invited to a retreat, and that *B* cheated *C* (at the retreat). The source also contained 18 additional propositions that listed the incidental properties of the principal characters (*A*–*D* and *CEO-1*). Specifically, *A* was characterised as a being a mean, intelligent, inexperienced ex-Navy SEAL, *B* as an experienced, mean, easy-going ex-magician, *C* as an experienced, easy-going, careless ex-astronaut, *D* as an inexperienced, intelligent, mean, ex-wrestler, and *CEO-1* as being a founder and good manager.

The target stated that because C' cheated C'', *CEO-2* was reminded that B' cheated B''. Eighteen additional propositions stated the attributes of the main characters. Specifically, both B' and B'' were described as experienced, mean, easy-going ex-magicians (to be similar to *B*), and both C' and C'' were described as being experienced, easy-going, careless

TABLE 3
Representations of source and target used in SME simulation

```
SOURCE
          (CAUSE
                    (CAUSE (CHEATS A B)
                              (REMINDED ceo-1
(CHEATS C D)))
                    (AND (INVITED B retreat-1)
                              (INVITED C retreat-1)))
(CHEATS b c))

(GOOD-MANAGER ceo-1)
(FOUNDER ceo-1)
(EX-WRESTLER D)
(MEAN D)
(INTELLIGENT D)
(INEXPERIENCED D)
(EXASTRONAUT C)
(CARELESS C)
(EASYGOING C)
(EXPERIENCED C)
(EX-MAGICIAN B)
(MEAN B)
(EXPERIENCED B)
(EASYGOING B)
(EX-NAVY-SEAL A)
(MEAN A)
(INTELLIGENT A)
(INEXPERIENCED A)

TARGET
(CAUSE (CHEATS C′ C″)
          (REMINDED ceo-2 (CHEATS B′ B″)))
(GOOD-MANAGER ceo-2)
(FOUNDER ceo-2)
(EX-MAGICIAN B″)
(MEAN B″)
(EASYGOING B″)
(EXPERIENCED B″)
(EX-MAGICIAN B′)
(MEAN B′)
(EASYGOING B′)
(EXPERIENCED B′)
(EXASTRONAUT C″)
(CARELESS C″)
(FRIENDLY C″)
(EXPERIENCED C″)
(EXASTRONAUT C′)
(CARELESS C′)
(FRIENDLY C′)
(EXPERIENCED C′)
```

ex-astronauts (to be similar to *C*), and *CEO-2* was described as being a founder and good manager (to be similar to *CEO-1*).

SME begins its operation by producing a subset of all possible mappings that occur within a problem (using the greedy-merge algorithm; Forbus & Oblinger, 1990). These candidate mappings are then given varying scores depending on how well they simultaneously satisfy all possible constraints. The net result of this process indicated that inference O (based most heavily on surface similarity, but also preserving the cheat relation) received a much lower structural evaluation (with an evaluation score of 1.01), than inference R (based most purely on relational structure) (with an evaluation score of 6.21) (see Table 2). As SME uses a deterministic algorithm, the model was only run once.

LISA simulations. We simulated this mapping task in LISA as follows (see Table 4). The source analogue stated that an employee, *A*, cheated another employee, *B* (proposition P1), that *C* cheated *D* (P2), and that P1 reminded the company's CEO of P2 (P3). The reminding (P3) caused the CEO (P5) to invite *B* (the cheated) and *C* (the cheater) to the company outing (P4). The source also contained four additional propositions (P6–P9) that listed the incidental properties of the four principal characters (*A*–*D*). (Specifically, *A* was characterised as a being an inexperienced, intelligent, mean ex-Navy SEAL, *B* as experienced, easy-going, mean ex-magician, *C* as an experienced, easy-going, careless ex-astronaut, and *D* as an inexperienced, intelligent, mean, ex-wrestler.) In the source, all propositions were given equal importance, and the causal statements supported one another with a weight of 5 (support and importance influence the likelihood with which propositions will be chosen to fire; see Hummel & Holyoak, 1997, 2003, for details).

The target stated that C' cheated C'' (P1), that B' cheated B'' (P2), and that P1 reminded the CEO of P2 (P3). Four additional propositions stated the attributes of the main characters. Specifically, both B' and B'' were described as experienced, easy-going, mean ex-magicians (to be similar to *B*), and both C' and C'' were described as being an experienced, easy-going, careless ex-astronauts (to be similar to *C*). As in the source, propositions in the target concerning the cheating and reminding relations supported one another with strength 5.

As LISA uses a stochastic algorithm, we ran the simulation 50 times and recorded the model's mappings and inferences on each run. Each run began with the target as the driver, firing six propositions chosen at random. Then the source was made the driver and the model again fired six propositions chosen at random. During the firing of these first 12 propositions, the model was not permitted to make analogical inferences. Analogical inference was then allowed, and the model fired an additional eight propositions in the

TABLE 4
Representations of source and target used in LISA simulation

SOURCE

Objects

A person employee A;
B person employee B;
C person employee C;
D person employee D;
CEO1 person boss clever CEO1;

Predicates

Cheats 2 action nasty evil cheats;
Reminded 3 state memory event reminder reminded
Invite 3 action offer deal invite;
Cause 2 action has effect cause;
TraitsA 1 inexperienced intelligent mean exnavyseal traitsa;
TraitsB 1 experienced easygoing mean exmagician traitsb;
TraitsC 1 experienced easygoing careless exastronaut traitsc;
TraitsD 1 inexperienced intelligent mean exwrestler traitsd;

Propositions

P1 (cheats A B);
P2 (cheats C D);
P3 (reminded CEO1 P1 P2); {
P4 (invite CEO1 B C);
P5 (cause P3 P4);
P6 (traitsa A);
P7 (traitsb B);
P8 (traitsc C);
P9 (traitsd D);

TARGET

Objects

C′ person employee C′;
C″ person employee C″;
B′ person employee B′;
B″ person employee B″;
CEO2 person boss clever CEO2;

Predicates

Cheats 2 action nasty evil cheats;
Reminded 3 state memory event reminder reminded;
Invite 3 action offer deal invite;
Cause 2 action has effect cause;
TraitsB 1 experienced easygoing mean exmagician traitsb;
TraitsC 1 experienced easygoing careless exastronaut traitsc;

Propositions

P1 (cheats C′ C″);
P2 (cheats B′ B″);
P3 (reminded CEO2 P1 P2);
P4 (traitsb B′);
P5 (traitsb B″);
P6 (traitsc C′);
P7 (traitsc C″);

source. Overall, LISA mapped the correct characters for the O inference on 24% of the runs and the R inference on 18% of the runs. On the remaining 48% of the runs, LISA generated inferences based on various mixtures of these character mappings, or on the S or O mappings but with only one of the two employees being invited to a company retreat. These remaining inferences roughly correspond to the SO category from the experiment (see Table 2).

DISCUSSION

The results indicate that object and relational overlap can have approximately equal impact on mapping and inference under conditions of high attribute similarity and high relational complexity, when these two constraints are placed in opposition. The mapping results indicate that subjects are likely to report multiple mappings under conditions in which different mappings can be derived from different types of information (cf. Markman, 1997; Spellman & Holyoak, 1992, 1996). The reasons subjects provided for their mappings demonstrated that prominent trait information is able to compete with shared relational roles as a strongly preferred basis for mapping. The inference results also support this conclusion, as subjects treated the two most preferred inferences, one based on object similarity and the other on relational structure, nearly equally. Not only did subjects generate object-based inferences as often as they generated relational inferences, but they also rated the object-based inferences as at least as plausible as the relational ones.

Our findings raise questions about the depth of understanding that underlies human analogical reasoning, even for our highly educated population of college students. Inferences based entirely on object similarity are typically suspect due to the local nature of the correspondences (Gentner & Toupin, 1986; Clement & Gentner, 1991). Inferences that honour object similarity at the cost of violating key structural aspects of the analogy run the risk of being coherent at only the local level, missing critical parts of the global structure that would lead to a more useful inference. Nonetheless, many subjects in our experiment preferred inferences that were based predominantly on object similarities, even when doing so violated a higher-order relation. The present findings suggest that human analogical reasoning is sensitive to correspondences based on object similarity—or at least exhibits a degree of flexibility in the relative contributions of object traits and relational constraints—which is not presently captured by the SME model. By contrast, the LISA model is able to simulate the human tendency to produce nearly equal numbers of object-based and relational mappings. LISA's tendency to generate object-based as well as relational inferences is due to its inherent working memory limits and the stochastic nature of its algorithm, which interact to produce different mappings on different runs. Our findings support analyses of the role of limited working memory in modulating analogical reasoning (cf. Hummel & Holyoak, 1997; 2003; Keane & Brayshaw, 1988; Kubose et al., 2002).

Manuscript received 29 January 2003
Revised manuscript received 14 August 2003

REFERENCES

Bassok, M., & Medin, D. L. (1997). Birds of a feather flock together: Similarity judgments with semantically rich stimuli. *Journal of Memory and Language*, *36*, 311–336.

Carbonell, J. G. (1983). Learning by analogy: Formulating and generalizing plans from past experience. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach* (pp. 137–161). Palo Alto, CA: Tioga.

Clement, C. A., & Gentner, D. (1991). Systematicity as a selection constraint in analogical mapping. *Cognitive Science*, *7*, 155–170.

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

Forbus, K. D., & Oblinger, D. (1990). Making SME greedy and pragmatic. In *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society* (pp. 61–68). Mahwah, NJ: Lawrence Erlbaum Associates Inc.

Gentner, D. (1983). Structure mapping: A theoretical framework for analogy. *Cognitive Science*, *7*, 155–170.

Gentner, D. (1989). The mechanisms of analogical learning. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 199–241). Cambridge: Cambridge University Press.

Gentner, D., & Toupin, C. (1986). Systematicity and surface similarity in the development of analogy. *Cognitive Science*, *10*, 277–300.

Holyoak, K. J., Novick, L. R., & Melz, E. R. (1994). Component processes in analogical transfer: Mapping, pattern completion, and adaptation. In K. J. Holyoak & J. A. Barnden (Eds.), *Advances in connectionist and neural computation theory. Vol. 2: Analogical connections* (pp. 113–180). Norwood, NJ: Ablex.

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

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

Hummel, J. E., & Holyoak, K. J. (2003). A symbolic-connectionist theory of relational inference and generalization. *Psychological Review*, *110*, 220–263.

Keane, M. T. (1988). *Analogical problem solving*. Chichester, UK: Ellis Horwood.

Keane, M. T., & Brayshaw, M. (1988). The Incremental Analogical Machine: A computational model of analogy. In D. Sleeman (Ed.), *European working session on learning* (pp. 53–62). London: Pitman.

Keane, M. T., Ledgeway, T., & Duff, S. (1994). Constraints on analogical mapping: A comparison of three models. *Cognitive Science*, *18*, 387–438.

Kubose, T. T., Holyoak, K. J., & Hummel, J. E. (2002). The role of textual coherence in incremental analogical mapping. *Journal of Memory and Language*, *47*, 407–435.

Markman, A. B. (1996). Structural alignment during similarity and difference judgments. *Psychonomic Bulletin and Review*, *3*, 227–230.

Markman, A. B. (1997). Constraints on analogical inference. *Cognitive Science*, *21*, 373–418.
Markman, A. B., & Gentner, D. (1993). Structural alignment during similarity comparisons. *Cognitive Psychology*, *23*, 431–467.
Morrison, R. G., Krawczyk, D. C., Holyoak, K. J., Hummel, J. E., Chow, T. W., Miller, B. L. et al. (in press). A neurocomputational model of analogical reasoning and its breakdown in Frontotemporal Lobar Degeneration. *Journal of Cognitive Neuroscience*.
Ross, B. H. (1987). This looks like that: The use of earlier problems and the separation of similarity effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *13*, 629–639.
Ross, B. H. (1989). Distinguishing types of superficial similarities: Different effects on the access and use of earlier problems. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*, 456–468.
Spellman, B. A., & Holyoak, K. J. (1992). If Saddam is Hitler then who is George Bush?: Analogical mapping between systems of social roles. *Journal of Personality and Social Psychology*, *62*, 913–933.
Spellman, B. A., & Holyoak, K. J. (1996). Pragmatics in analogical mapping. *Cognitive Psychology*, *31*, 307–346.
Wisniewski, E. J., & Bassok, M. (1999). What makes a man similar to a tie? Stimulus compatibility with comparison and integration. *Cognitive Psychology*, *39*, 208–238.

APPENDIX A

Example version of Story 1, the source analogue

Brightech Incorporated is a medium sized corporation specializing in the manufacture of light fixtures for large public buildings. They have been successful in their field for some time and have a favorable reputation among their major clients. Brightech is not without problems, though, particularly because of their frequent employee conflicts.

A recent project at Brightech was the development of the Mega-Watt 2000, an excellent new product that was expected to be the next revolution in industrial lighting. Harold and Robert were the primary employees involved in successfully completing the project. Harold had very little experience and had just joined Brightech within the last year as an assistant manager. Harold was quite intelligent; however, he was also mean-spirited. He had been a Navy SEAL prior to his career at Brightech. Robert had been with the company for over a decade and had reached the level of manager. Robert was easy-going, but also rather obnoxious. He had worked as a magician before his corporate life.

When time came for the Mega-Watt 2000 to be revealed at a shareholder's meeting, Harold managed to persuade everyone that he had done all the work on the brilliant new product and took all the credit for it. He managed to cheat Robert out of any benefits that might come his way for his part in developing Mega-Watt 2000. While the board of directors believed that it was indeed Harold's work, Steve, the CEO of the corporation, saw what had really happened and realized that Harold had cheated Robert.

Steve was one of the original founders of Brightech and had kept it competitive for many years. One reason for this was Steve's fantastic ability to keep the employees under control. Seeing Harold cheat Robert had reminded Steve of a similar situation that he had seen last month. In that case John and Tom were working together to cut advertising costs. John was an experienced manager who was known for being very friendly. He was also regarded as being a careless manager. John was an astronaut before coming to work at Brightech. Tom was new to Brightech and an assistant manager. While Tom was wise, he was also known to have a bad temper. Tom was once a pro wrestler before working as an executive. The pair had managed to cut the advertising costs in half. The board of directors was impressed with the accomplishment, but John had misled them to think that he had done it all himself. He had managed to cheat Tom out of the credit he deserved.

Later on a large company retreat was held where teams of two employees had to devise new strategies for the company. Steve made sure that Robert was teamed up with John. Robert was aware that John had unjustly taken credit for work in the past. When it came time to reveal the new ideas, Steve would have Robert report what they had come up with and this would give him the chance to get some credit himself and also the opportunity to deny John any credit for the work. If Robert decided to cheat John it would clear up some of the inequality in the employee ranks.

Example version of Story 2, the target analogue

The Offstar Corporation has been a producer of office supplies for quite a while. They have a good reputation and routinely have high sales. The only real problems at Offstar are that the employees do not always get along well.

A major project that Offstar had just completed was the Moto-chair, which was a deckchair that moved by motor power. Greg and Barry were the main people responsible for the completion of Moto-chair. Greg was an experienced manager who had been with the company for some time. Greg was friendly, but also quite careless. Greg had been an astronaut before becoming an executive. Barry was a long-time employee and was also a manager. Barry was friendly, but also rather careless. Barry had been an astronaut too before going to work in office supplies.

At the big office supply convention, the Moto-chair was a huge success and everyone wanted to start ordering them. Greg managed to take all of the credit for the Moto-chair and cheated Barry out of his share of the glory. The executive board praised Greg as the lone inventor of the Moto-chair, but Martin, the CEO of the Offstar Corporation, had seen that Greg had cheated Barry.

Martin was a well-respected CEO who had been in charge of the Offstar Corporation for years. He was especially skillful in employee management. Seeing Greg cheat Barry had reminded Martin of a recent incident between the employees Ken and Charles, who were working on getting a big contract for the company. Ken was a manager who had been with the company for a long time. He was easy-going, but had a tendency to be obnoxious. Ken was a former magician. Charles was also a manager who put in many years at the company. Charles was easy-going, but also known for obnoxious behavior. Charles had also been a magician before coming to the Offstar Corporation. Working together they had landed the big contract. The executive board was delighted, but Ken claimed to have done all the work himself and took all of their praise for himself. He had cheated Charles out of any credit.

APPENDIX B

Example question from Mapping questionnaire

STORY 1: Harold, who was inexperienced, intelligent, mean-spirited and an ex-Navy SEAL.

STORY 2	REASON
______________	____________________________
______________	____________________________
______________	____________________________
______________	____________________________
______________	____________________________
______________	____________________________