

The One-to-One Constraint in Analogical Mapping and Inference

Daniel C. Krawczyk^a, Keith J. Holyoak^b, John E. Hummel^b

^a*University of California, Berkeley*

^b*University of California, Los Angeles*

Received 20 January 2004; received in revised form 29 November 2004; accepted 17 December 2004

Abstract

Theories of analogical reasoning have assumed that a 1-to-1 constraint discourages reasoners from mapping a single element in 1 analog to multiple elements in another. Empirical evidence suggests that reasoners sometimes appear to violate the 1-to-1 constraint when asked to generate mappings, yet virtually never violate it when asked to generate analogical inferences. However, few studies have examined analogical inferences based on nonisomorphic analogs, and their conclusions are suspect due to methodological problems. We sought to elicit mixed inferences that could result from combining information from 2 possible mappings. Participants generated 2-to-1 correspondences when asked for explicit mappings, but did not produce mixed inferences. Multiple correspondences appear to arise from multiple isomorphic mappings, rather than from a single homomorphic mapping.

Keywords: Analogy; Reasoning; Problem solving; Human experimentation; Knowledge representation

Reasoning by analogy involves a series of steps that generate novel inferences about an unfamiliar target domain based on prior knowledge of a familiar source domain. The source and target analogs are analyzed for similar attributes, relations, and systems of connected relations, and this information is used to map the two analogs and identify corresponding elements. These correspondences can then be used to make inferences about the target by “copying with substitution” from the source to the target; that is, creating new target propositions by substituting corresponding target elements for source elements in initially unmapped source propositions. For example, when an analogy is formed between the solar system and an atom (Gentner, 1983), a reasoner might first notice that the nucleus corresponds to the sun, that the planets correspond to electrons, and that planets revolve around the sun. Using these mappings, the reasoner may then copy the shared relation (“revolve around”) and substitute the corresponding elements to infer that the electrons revolve around the nucleus. (For examples of

copy-with-substitution algorithms embodied in computational models, see Falkenhainer, Forbus, & Gentner, 1989; Holyoak, Novick, & Melz, 1994; Hummel & Holyoak, 2003.)

The fundamental purpose of analogy is to generate plausible and useful inferences. It follows that analogical mapping and inference, although inevitably fallible, must be governed by constraints that favor inferences likely to prove true and goal relevant. A number of interrelated constraints have been proposed. Gentner (1983) emphasized that analogy is guided by structural constraints based on predicate–argument structure, and that analogical inferences will preferentially follow from mappings based on higher order relations (the *systematicity* principle). Gentner also proposed the *one-to-one constraint*, which implies that each correspondence should be unique (Falkenhainer et al., 1989; Gentner, 1982, 1983). The one-to-one constraint was also included in Holyoak and Thagard's (1989) formulation of analogical constraints. In their multiconstraint theory, the basic structural constraint is *isomorphism*, which is a pressure to find mappings that are structurally consistent (i.e., if a source and target element correspond in the context of one relational role, these same elements should correspond in all relational roles) and map elements as one-to-one correspondences.

This article focuses on further clarifying the nature of the one-to-one constraint. In messy, realistic analogies, the structure of the source and target will typically fall short of a pristine isomorphism. When multiple mappings are plausible for a single element, the one-to-one constraint encourages selection of the best fitting mapping (as determined by other constraints, such as semantic similarity and maximal match within the overall relational structure), rather than maintaining multiple mappings for a single element. The one-to-one constraint is especially important in analogical inference because inferences derived using a mix of incompatible element mappings are likely to be incoherent (see Falkenhainer et al., 1989; Gentner, 1982; Markman, 1997). For example, suppose the source analog “Iraq invaded Kuwait but was driven out” were mapped onto the target “Russia invaded Afghanistan; Argentina invaded the Falkland Islands.” If the one-to-one constraint on mapping does not prevail over the structural evidence that Iraq can map to both Russia and Argentina, and Kuwait to both Afghanistan and the Falklands, then a copy-with-substitution algorithm might form the target inference, “Russia was driven out of the Falkland Islands.” For this example prior knowledge shows that this “mixed” inference is false. The more basic problem, however, is that the inference is incoherent—the target analog gives no basis for inferring that Russia had anything to do with the Falklands. This lapse in inferential coherence arises from the use of multiple mappings for two source elements (here, Iraq and Kuwait), coupled with a failure to maintain structural consistency (i.e., *if* Iraq maps to Russia, then Kuwait should map to Afghanistan, not the Falklands).

A controversy has existed in the literature as to how the one-to-one constraint is implemented. Computational models of analogical mapping have universally included some version of a one-to-one constraint, but empirical evidence concerning whether and when people might violate the one-to-one constraint when they form mappings or make analogical inferences is quite limited. Current models of analogy vary in the way this constraint is enforced, with some models strictly enforcing one-to-one mappings, whereas others treat it as a looser constraint. Representative models for these two positions are the Structure Mapping Engine (SME; Falkenhainer et al., 1989) and the Analogical Constraint Mapping Engine (ACME; Holyoak & Thagard, 1989). SME has a strong implementation of this constraint and only allows one-to-one mappings. The model nonetheless allows for multiple one-to-one mappings to oc-

cur in the context of multiple isomorphic mappings, thereby accounting for the fact that there may be more than one possible mapping for a given element. By contrast, ACME used a looser one-to-one constraint that could be overridden under certain mapping situations by the need to maximally satisfy pragmatic and semantic constraints in addition to structural constraints, thus allowing the model to form true many-to-one mappings in such cases. The critical difference is that the strong constraint always keeps elements in separate one-to-one mappings, whereas the looser constraint allows the possibility of multiple elements mapping to a single element within the same representation.

These differences in computational implementation prompted a series of prior studies that empirically addressed the one-to-one constraint. Spellman and Holyoak (1992, 1996) found that participants tended to generate occasional two-to-one mappings for ambiguous elements; nonetheless, they did not generate mixed inferences based on inconsistent mappings, as might occur if the constraint were loosely enforced, allowing a many-to-one representation. They suggested two possible explanations for their finding of occasional two-to-one mappings. One possibility is that people sometimes form a single set of homomorphic correspondences in which one source element maps to two target analogs (for a discussion of homomorphisms, see Krantz, Luce, Suppes, & Tversky, 1971). Alternatively, people may construct two separate isomorphic mappings from source to target and then report the results of both on a mapping test (consistent with a strong one-to-one implementation). The latter explanation more readily accounts for the apparent absence of mixed inferences. If inferences were derived from a single homomorphic mapping, then mixed inferences would be expected to sometimes occur; but if inferences are derived from one isomorphic mapping (perhaps the more dominant of two), then mixed inferences would be prevented. These results provided some support for a strict one-to-one implementation, but did not end the controversy, as the many-to-one representation cannot be completely ruled out based on the findings of Spellman and Holyoak (1992, 1996).

In an effort to disambiguate the source of one-to-two mappings, Markman (1997) ran a series of experiments that allowed for the construction of possible homomorphic mappings. The experiments included explicit mapping and inference tasks. To determine whether participants were forming unified homomorphic representations, or multiple isomorphic representations, Markman looked for inconsistent object substitutions in analogical inferences. For example, in Markman's Experiment 1, an inconsistent condition was included in which an English department and a biology department in a source University each corresponded to a computer science department and a music department in a target university. The analogs allowed both one-to-many (biology to music and computer science) and many-to-one (English and biology to music) mappings. After a correspondence task, in which participants mapped the departments, they were asked to infer new facts about the target department. Across a series of four similar experiments, Markman found that people would sometimes report many-to-one mappings, but virtually never produced mixed inferences. The results were interpreted as evidence that people may compute multiple isomorphic mappings, but do not compute a single homomorphic mapping.

However, there are reasons for caution in accepting this conclusion based on Markman's (1997) findings, as the design of his experiments may have discouraged the reporting of inferences that could have demonstrated violations of the one-to-one constraint. Although the universities in the materials were fictional, the departments and their characteristics were realistic

examples of occurrences at universities in general. Prior knowledge may have led participants to generate inferences based on their personal knowledge of what might occur at universities under the specified conditions, rather than actually reasoning by analogy. Markman classified inferences of this type as “other,” which occurred with high frequencies in all experiments.

Prior knowledge may also have created a confounding of analogical consistency with a priori plausibility of inferences. In Markman’s (1997) Experiment 1, inconsistent inferences necessarily included different departments in “cause” and “effect” roles (e.g., “Because the Computer Science faculty obtained grants, the Music department hired research assistants”), whereas consistent inferences included a single department in both roles (e.g., “Because the Computer Science faculty obtained grants, the Computer Science department hired research assistants”). It seems very likely that college students would find statements in which different departments played the cause and effect roles implausible, even in the complete absence of a source analog.

These concerns motivated a pilot study to determine whether a confounding with plausibility may have affected the likelihood of inferences based on homomorphic mappings in the Markman (1997) study. Thirty students at the University of California, Los Angeles rated the plausibility of the inferences in the university materials developed by Markman (1997). The survey included consistent, inconsistent, and “other” inferences, the three types that were relevant to one-to-many mappings in Markman’s study. In the absence of any analogical context, participants were asked to report how likely each inference type appeared based on their intuitive feelings, using a rating scale from 1 (*very unlikely*) to 10 (*very likely*). Participants rated consistent inferences ($M = 6.72$) to be more plausible than inconsistent inferences ($M = 3.22$). A dependent measures t test revealed that these means were significantly different, $t(59) = 5.40$, $p < .01$. Participants also rated consistent inferences more plausible than irrelevant inferences ($M = 3.52$), $t(59) = 5.70$, $p < .01$. Such plausibility differences may have biased participants against producing inconsistent inferences, making it unlikely that evidence would be found for homomorphic mappings. Thus the controversy was not fully settled based on Markman’s (1997) findings.

In this experiment, the one-to-one constraint was investigated under conditions in which inconsistent and consistent inferences were equated in a priori plausibility. We counterbalanced inference types to equalize plausibility so that a statement that represented an inconsistent inference in one version represented a consistent inference in another due to rotation of the various relations described in the materials. This experiment examined whether the one-to-one mapping constraint is subject to violation for analogies based on novel materials that suggest potential two-to-one mappings based on the physical traits of fictional organisms.

1. Method

1.1. Participants

Participants were 32 University of California, Los Angeles students fulfilling a course requirement for an introductory psychology class. Participants were run in groups ranging from 1 to 5.

1.2. Materials, design, and procedure

The materials were presented in a 21-page booklet and were based on a cover story indicating the participant was to play the role of a futuristic biologist on an outer-space mission. The story included background information about the purpose of the mission (to explore life on a distant planet and take note of species traits and interactions). A description of four primary organisms from a source planet followed, with each organism described in one paragraph. These descriptions included a nonword species name for each primary organism, along with two of its unique physical traits. For example, one species had long flexible antennae and tails that were curled and translucent. Following the trait description, two other organisms were introduced. One was a neighbor (described as a species that lived near the primary organism), and the other was a competitor of the primary organism.

The third page included text that described the interactions of the organisms on the source planet. Separate paragraphs were used to describe each trio of organisms (the primary organism along with its neighbor and competitor). In each paragraph, the primary organism was described as obtaining food in a specific way (e.g., hunting for it). The rest of the paragraph explained that the other two organisms in the trio reacted to the primary organism's feeding behavior in a specific way. Each neighbor was described as performing an innocuous behavior (e.g., napping) to remain out of the way of the primary organism. Each competitor was described as trying to intimidate the primary organism (e.g., by demonstrating speed). Lastly, the neighbor and competitor were described as staying out of each other's way, so they only perform their particular action when the other is performing its action as well. This relation ensured that each organism in the trio performs its particular action *only* in combination with the actions of its fellow cohorts, thus tying together each group of actions as a distinct set linked by co-occurrence in time and place. Each paragraph differed slightly with regard to exact phrasing and particular filler information to ensure that the descriptions were distinct from one another.

Three quizzes were formulated to ensure that participants attended to the relevant traits and actions that they would later need to map and make inferences about the organisms. The content of each quiz question described a fact (a trait or interaction of the organisms) that would be needed for the later mapping and inference tasks. On the page following each quiz, participants were instructed to correct the previous quiz by reading back through the materials and writing down the correct letter choice to the left of the question number (on the original quiz page) for any missed questions. This procedure ensured that participants had attended to and written the correct answer to each of the questions by the end of the quiz portion.

A mapping task appeared after the third quiz. It included instructions set within the context of the cover story indicating that the participant had met Dr. Krick, a space biologist who had some incomplete field notes about another remote planet (the target planet) that he had previously explored. The instructions stated Krick believed that the organisms on the target planet were similar to those on the source planet and that the participant should write down the names of any organisms from the source planet that corresponded to the one on the target planet that was named in the materials. The field notes about the target planet began with the name that Dr. Krick had assigned to the primary organism, followed by a description of the unique pair of the primary organism's traits. There were descriptions of four primary organisms and their traits in total, presented on four separate pages. The descriptions of two out of the four primary organ-

isms allowed for a two-to-one mapping to the primary organisms on the source planet. Multiple potential mappings were achieved by having each of the primary organism's traits be identical to one of the traits from two different primary organisms from the source planet; thus two target organisms matched one source organism, each on the basis of sharing one trait. These possible two-to-one correspondences could be represented within either an integrated homomorphic mapping (Fig. 1A) or within two separate isomorphic mappings (Fig. 1B).

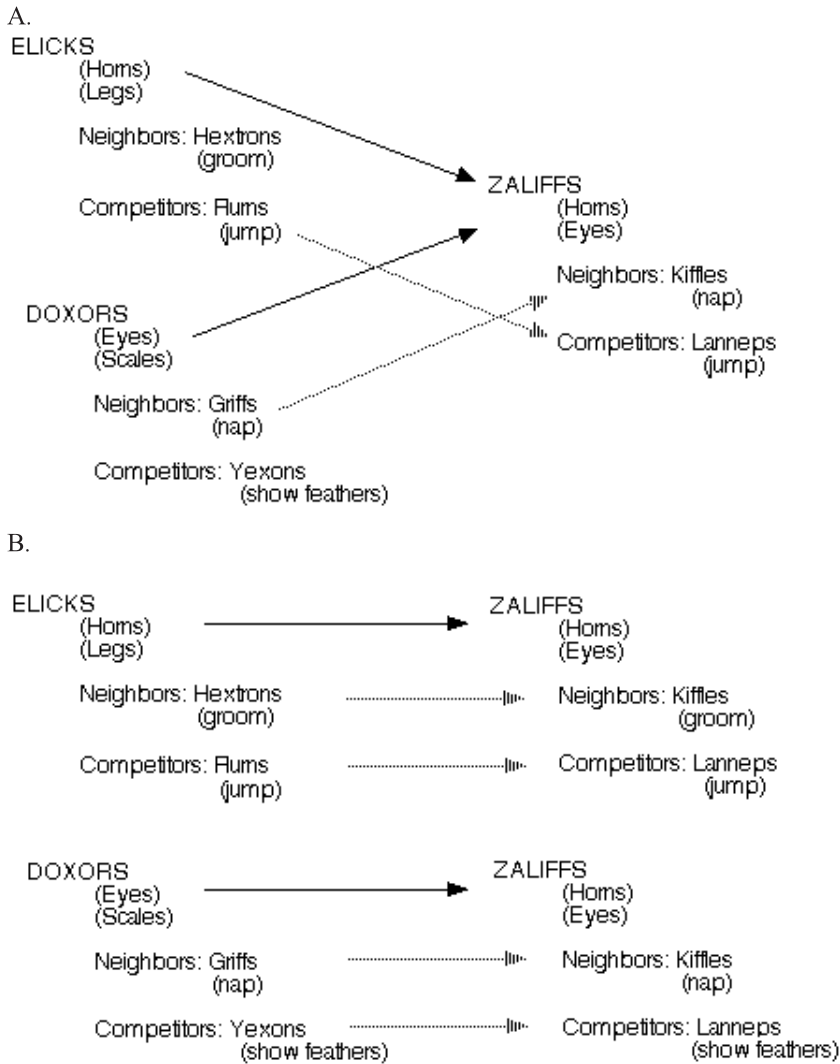


Fig. 1. A schematic illustration of the experiment. A. Representation of a homomorphic mapping (based on physical traits), along with related facts from the text of the experiment that show a mixed inference (the reactions from the neighbor and competitor of the dual mapped element, Zaliffs, come from different sources). B. Representation of two separate isomorphic mappings (based on physical traits). Note that in each isomorph the reactions of the neighbor and competitor of the repeated element, Zaliffs, come from the same source in this case.

Two other primary organisms from the target planet were described to allow for a one-to-one mapping; these shared only one trait with a primary organism from the source planet, whereas their other trait was new and unrelated to previously described traits. Following this description, the notes stated the names of a neighbor and a competitor. Beneath this text on each page of the mapping task there were headings for the source planet and target planet. Beneath the source heading was the name of one primary organism from the source planet. Beneath the target heading appeared five blank lines for the names of corresponding organisms from the target planet. Participants were instructed that there may be more than one corresponding organism, or that there may be no corresponding organisms, and that they should leave any unneeded lines blank. Participants were allowed to look back at the original information about the source planet in completing this task. The use of multiple response lines was intended to encourage the reporting of multiple mappings if the participant wished to provide them.

An inference task was included to further determine the types of representations participants used in the analogy task. The inference task was begun after participants read about a second plan Dr. Krick had to extend knowledge about the target planet. Krick guessed that the organisms from the target planet that corresponded to those on the source planet were likely to interact with their neighbors and competitors in ways that were similar to the interactions among the organism trios read about on the source planet. Participants were instructed to make inferences about reactions by the neighbor and competitor on the target planet in response to an action by the primary organism on that planet. Participants were also instructed to fill in as many lines for reactions as they believed might occur, but only to write down those actions that they felt were possible. The instructions stressed in boldface type that the participant should make *as many predictions as possible* and to note that *all* possible reactions should be inferred. These four inference collection sheets were the final pages in the booklet. Following the instructions, the next four pages included the name of the primary organism at the top, followed by the names of the neighbor and the competitor, and a reiteration of the primary organism traits (the same description from the mapping task). Further down the page the participant was presented with the statement that the primary organism obtained food. Following this statement were five fill-in-the-blank style answer spaces; for example, "While the ___ are grooming themselves, what activity or activities might the ___ be doing?" The instructions stated that only one pair of reactions should be provided on each line. Each of the four pages in the inference task followed this same format with four inference spaces, one for each reaction the participant wished to formulate.

The materials were counterbalanced to yield 32 distinct versions; each participant in the experiment received a unique version. The stories stated that four of the primary organisms, along with their particular neighbor and competitor, lived on the planet Corflew, and four other primary organisms and their neighbors and competitors lived on the planet Zorntel. Fig. 1 shows the sets of names from each planet. Materials were constructed so that Corflew was the source and Zorntel was the target planet in 16 versions, whereas Zorntel was the source planet and Corflew the target planet in the other 16 versions. The organisms described as residents of each planet remained with that planet's title for the counterbalancing switch. This procedure was intended to control for any differences in memorability or distinctiveness of certain planet or organism names. Within each set of 16 versions for which a planet was the source, there were four rotations of the ordering for each organism, so each primary organism and its partic-

ular neighbor and competitor appeared at each position in the four-paragraph ordering for initial trait descriptions and interaction descriptions on the pages describing the source planet. For each serial-position ordering there were four separate versions, allowing each neighbor reaction to be paired with each competitor reaction. Because the eight paragraphs describing the traits and interactions on the source planet were each distinctive, the order of these paragraphs was rotated so that each paragraph appeared at each position in the pages an equal number of times. These measures were carried out to control for any differences in memorability or distinctiveness among any particular combinations of organism names or reactions, or any particular paragraph ordering. In addition to these rotations, the ordering of organism names from the target planet was rotated in a similar manner so that each of the names appeared at each sequential position an equal number of times across all of the versions. These rotations occurred for the mapping order as well as the inference task order so that each organism name occupied the same position in both tasks in each particular version.

2. Results

2.1. Mapping task

Participants produced more mappings for the two-map organisms ($M = 1.90$) than for the one-map organisms ($M = 1.18$). A dependent measures t test revealed that these means were significantly different, $t(63) = 9.53, p < .01$. Participants generally produced the single correct mapping for the one-map organisms. Twenty-nine of the 32 participants mapped both one-map organisms to their appropriate matches, and 3 participants mapped a single one-map organism to its appropriate match. Six participants mapped one or both of the one-map organisms to both their correct match and one or more additional incorrect matches.

Most participants matched both of the two-map organisms to both of their appropriate matches only. Twenty-one participants mapped both of the two-map organisms to their two correct matches. Five participants mapped the two-map organisms to $\frac{3}{4}$ of their appropriate matches, 4 participants mapped the two-map organisms to $\frac{1}{2}$ of their appropriate matches, and 2 participants mapped the two-map organisms to $\frac{1}{4}$ of the correct matches. Four participants mapped one or both of the two-map organisms to one or more additional incorrect matches.

2.2. Inference task

Inferences generated during this task were scored as “consistent,” “mixed,” “unmapped-other,” or “irrelevant.” Consistent inferences followed directly from a possible mapping and had been paired with the appropriate reaction that had appeared in the description of the source planet. For the example depicted in Fig. 1, an inference that “Kiffles *would groom* while Lanneps *would show off their jumping ability*” would be consistent, as grooming and jumping were paired in the description of the source planet that followed from the mapping of the primary organisms Elick and Zaliff. Inferences were scored as mixed if they followed directly from a possible mapping, but were paired with the reaction from the alternative mapping. For example, “Kiffles *would groom* while Lanneps *would show off their feathers*” would represent a mixed inference, as grooming is a reaction derived from the main organism mapping of Elick (source) to

Table 1
Mean numbers of inferences of each type

Organism Condition	Consistent	Mixed	Unmapped-Other
One-map	.77	—	.10
Two-map	.66	.01	.04

Zaliff (target), whereas showing off feathers is a reaction derived from the main organism mapping of Doxor (source) to Zaliff (target). Mixed inferences could only be produced for the two-map primary organisms, as two possible mappings are necessary for their production. The mixed inference should only appear if participants are using a homomorphic representation to generate inferences (Fig. 1A), as this type combines information from one mapping with information that appeared in a separate mapping in the description of the source planet.

Inferences that included reactions that had appeared in nonmapped organism trios were scored as unmapped-other inferences. An example of this type would be the inference that “Kiffles *are sunbathing* while the Lanneps *are forming herds*.” Although these are reactions from the materials, they came from the organism trio involving the main organism Grimzel, which was not a possible mapping for the Zaliffs on the target planet (based on common physical traits). Inferences scored as “irrelevant” involved reactions that either did not appear in the materials, or contained facts that had appeared, but were unrelated to the organism interactions.

Table 1 summarizes the mean numbers of consistent, mixed, and unmapped-other inferences for the one-map and two-map organisms. The means in all inference tasks represent the mean number of inferences generated as a proportion of all inferences of that type that were possible. For one-map organisms, a one-way within-subjects analysis of variance revealed a significant difference between consistent ($M = .77$) and unmapped-other inferences ($M = .10$), $F(1, 31) = 53.33, p < .01$. A one-way within-subjects analysis of variance indicated a significant main effect of inference type for two-map organisms, $F(2, 62) = 81.02, p < .01$. A Newman-Keuls test indicated that for two-map organisms, significantly more consistent inferences ($M = .66$) were made than either mixed inferences ($M = .01$) or unmapped-other inferences ($M = .04$). Irrelevant inferences were not included in these analyses because they did not represent cases where participants were using analogy and were not reported often (a total of 15 were reported among all participants). The number of consistent inferences was significantly greater for the one-map organisms ($M = .77$) compared to the two-map organisms ($M = .66$), $F(1, 31) = 53.33, p < .01$.

3. Discussion

The results demonstrated that even under conditions in which participants often produced two-to-one mappings on an explicit mapping task, and exhaustive inference production was encouraged, participants almost never generated mixed inferences. This finding provides support for the interpretation that participants tend to form multiple isomorphic representations (see Fig. 1B) for objects that have two-to-one correspondences (Markman, 1997; Spellman & Holyoak,

1996). As consistent inferences were produced at overwhelmingly high levels compared to mixed and unmapped-other inferences, it appears that multiple isomorphic representations were formed in the mapping stage and then used to produce consistent inferences that kept information from each mapping separate. In several follow-up experiments that allowed for the possibility of mixed inferences based on homomorphic representations, we also found no evidence that people produce mixed inferences. These findings are consistent with a strongly implemented one-to-one constraint of the type enforced by the SME model (Falkenhainer et al., 1989).

These results establish a dissociation between overt production of many-to-one mappings in explicit mapping tasks, which are common, and production of mixed inferences based on multiple inconsistent mappings, which are virtually never made. This dissociation is consistent with the results of previous studies of mapping and inference, using nonisomorphic analogs (Markman, 1997; Spellman & Holyoak, 1992, 1996), but this study used stronger controls for confounding variables and a broader variety of analogical structures. These findings illustrate how people can use analogies to reason coherently, even when the mapping between the two analogs is fraught with ambiguity and competing alternatives. The power to extract coherent inferences in the face of such complexity is basic to the role of analogy in everyday cognition.

Acknowledgments

This project was supported by National Science Foundation grant SBR-9729023.

We thank Wei-Chung Ooi, Carla Webster, Christine Chow, Miguel Valenzuela, and Noemi Briano for help running participants and for their useful comments. We thank Art Markman and Dedre Gentner for helpful discussions and comments on an earlier draft.

References

- Falkenhainer, B., Forbus, K. D., & Gentner, D. (1989). The structure mapping engine: Algorithm and examples. *Artificial Intelligence*, *41*, 1–63.
- Gentner, D. (1982). Are scientific analogies metaphors? In D. S. Miall (Ed.), *Metaphor: Problems and perspectives* (pp. 106–132). Brighton, England: Harvester.
- Gentner, D. (1983). Structure mapping: A theoretical framework for analogy. *Cognitive Science*, *7*, 155–170.
- 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. Barden (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. (2003). A symbolic-connectionist theory of relational inference and generalization. *Psychological Review*, *110*, 220–264.
- Krantz, D. H., Luce, R. D., Suppes, P., & Tversky, A. (1971). *Foundations of measurement* (Vol. 1). New York: Academic.
- Markman, A. B. (1997). Constraints on analogical inference. *Cognitive Science*, *21*, 373–418.
- 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.