Perspectives on Similarity from the LISA Model

Eric G. Taylor (etaylor4@uiuc.edu)
Department of Psychology, 603 East Daniel St.
Champaign, IL 61820

John E. Hummel (jehummel@cyrus.psych.uiuc.edu)
Department of Psychology, 603 East Daniel St.
Champaign, IL 61820

Keywords: similarity; analogy; neural networks; modeling

Abstract

Similarity is a central construct in perceptual and cognitive science, with implications for everything from basic color and pattern perception, to object and phoneme recognition, memory retrieval, analogical reasoning and problem solving. Previous studies have examined people's overt similarity judgments to understand the roles of similarity in cognition, and several model have been proposed to account for the data. This paper presents our simulations of these data with a model originally designed to simulate analogical reasoning (Hummel & Holyoak's 1997, 2003a, LISA model), not similarity judgments. We show that the same mechanisms that LISA uses to simulate analogy also provide a natural account of disparate data on similarity judgments. These successes speak to the utility of LISA's symbolic-connectionist knowledge representations as an account of human mental representations.

Similarity in Perception and Cognition

The concept of "similarity" plays a central role in perception and cognition, and numerous models have been proposed to account for it. Most of these efforts have met with mixed success.

Shepard's (1962) Multi-Dimensional Scaling (MDS) account of similarity and Tversky's (1977) feature-based *contrast model* capture many "featural" aspects of similarity, but fail to capture structural aspects (e.g., Gentner & Markman, 1997; Goldstone, 1994; Goldstone, Medin & Gentner, 1991). Gentner and colleagues have modeled "structural" aspects of similarity with the SME model of analogy (e.g., Gentner, Ratterman, & Forbus, 1993; Markman & Gentner, 1993), but their approach does not capture basic "featural" aspects of similarity (e.g., Goldstone's "Matches Out of Place", or MOPs; Larkey and Markman, 2005).

These complementary successes and failures suggest that what is needed is an account of mental representation that captures both the semantic/featural content of percepts and concepts (like MDS or the contrast model, but unlike SME) and also captures their relational/propositional organization (like SME, but unlike MDS and the contrast model).

Hummel and Holyoak (1997, 2003) have made similar observations about the complementary featural/structural requirements on mental representation, not in the context of similarity judgment, but in the context of relational reasoning, especially reasoning based on analogies, schemas and rules. Their LISA model incorporates both featural and structural knowledge and has accounted for numerous results in analogical reasoning, including analogical reminding and mapping (Hummel & Holyoak, 1997), analogical inference and schema induction (Hummel & Holyoak, 2003a), the development of these abilities (Morrison, Doumas, & Richland, 2006) and their decline in normal aging as a result of varieties of fronto-temporal dementia (Morrison, et al., 2004; Viskontas et al., 2004), among other phenomena (see Hummel & Holyoak, 2003b, for a review). We hypothesize that overt similarity judgments rely on the same representations that coordinate such high-level cognitive functioning. If this is the case, then LISA may provide a powerful platform for simulating such judgments. Moreover, its ability or inability to do so stands as an important test of its account of knowledge representation.

Knowledge Representation in LISA

LISA's knowledge representations ("LISAese") are connectionist representations made explicitly relational (i.e., symbolic) by virtue of their ability to dynamically bind relational roles to their fillers. LISA represents propositions such as "the machine repairs the car", or *repair* (machine, car), using a hierarchy of distributed and progressively more localist codes (Figure 1). At the bottom of the hierarchy, "semantic units" represent objects (e.g., machine and car) and relational roles (e.g., the *repairer* and *repairee* roles of *repair* (x, y)) in a distributed fashion. For example, "car" might be represented in terms of features such as "artifact", "mechanical", "transport", etc., and "machine" as "artifact", "mechanical", "labor", etc. These distributed semantic representations serve to make explicit what objects (and roles) have in common and how they differ.

At the next level of the hierarchy, *object* and *role* units (large circles and triangles, in Figure 1) represent objects and relational roles in a localist fashion, sharing bi-directional excitatory connections with the corresponding semantic units. For storage in LTM (and predication in working memory [WM]; see Hummel and Holyoak, 2003a), objects are bound to relational roles (e.g., "machine" to *repairer* and "car" to *repairee*) by means of localist *sub-proposition* units (SPs; rectangles in Figure 1). Collections of role-filler bindings are bound into complete propositions by means of localist *proposition* (P) units (ovals in Figure 1).

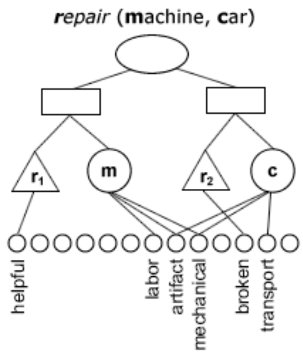

Figure 1. LISAese representation of *repair* (machine, car). The labels r_1, r_2, m and c refer to *repairer*, *repairee*, machine and car, respectively.

When a proposition enters WM, roles fire in synchrony with the fillers to which they are bound and out of synchrony with other role-filler bindings, resulting in a systematic pattern of firing on the semantic, object, role and SP units. For example, to represent *repair* (machine, car), the units representing car would fire in synchrony with those representing *repairee*, those representing machine would fire in synchrony with *repairer*, and the car+*repairee* units would fire out of synchrony with the machine+*repairer* units. The resulting representations are simultaneously symbolic, by virtue of dynamically binding roles to their fillers, and distributed, by virtue of the semantic units. As such, they suggest a potential basis for simultaneously capturing both the featural/semantic aspects and the structural aspects of similarity.

Computing Similarity Judgments with LISA

Although LISA was not designed to simulate explicit similarity judgments, it is nonetheless possible to extract similarity judgments from the model by exploiting two functions the model currently uses for analogical inference (Hummel & Holyoak, 2003a) and relational match-to-sample (Kroger, Holyoak & Hummel, 2004).

People use the quality of the mapping between a target analog and a potential source to decide whether to make inferences from the source to the target (Lassaline, 1996). LISA implements this constraint by computing how well a target analog, T, maps to a potential source, S (Hummel & Holyoak, 2003, Eq. A15):

$$q(T,S) = \frac{\sum_{t,s \in T,S} i_t [m(t,s_{\max}) - m(t,s_{\max 2})]}{\sum_{t \in T} i_t}, \tag{1}$$

where $q(T, S)$ is the quality of the mapping of T onto S, i_t is the "importance" (i.e., pragmatic centrality; Holyoak & Thagard, 1989) of unit t in T, $m(x, y) \in 0\ldots1$ is an index of mapping strength from unit x to unit y, $s_{\max}$ is the unit in S with the largest such mapping to t and $s_{\max2}$ is the unit with the second largest. Eq. 1 expresses the (importance-weighted) proportion of the structures in T that map uniquely to structures in S. In all the simulations reported here, $i = 1$ for all s and all t, allowing us to replace the sum in the denominator with $n(T)$, the number of token units in T.

Eq. 1 can be generalized to provide a measure of mapping-based *similarity* by adding a measure of the semantic similarity of t to $s_{\max}$. We chose $\cos(t, s_{\max})$, the cosine of the angle between the vector of semantic weights leading into unit t and the vector leading into $s_{\max}$:

$$\sigma(T,S) = \frac{\sum_{t,s \in T,S} [m(t,s_{\max}) - m(t,s_{\max 2})]\cos(t,s_{\max})}{n(T)}. \tag{2}$$

Note that $\sigma(S, T)$ will not necessarily equal $\sigma(T, S)$, especially if $n(T) \neq n(S)$.

Finally, we generalize Eq. 2 by adding a Weber constant to the denominator:

$$\sigma_{\mathrm{MIP}}(T,S) = \frac{\sum_{t,s \in T,S} [m(t,s_{\max}) - m(t,s_{\max 2})]\cos(t,s_{\max})}{1 + n(T)}. \tag{3}$$

(See Hummel & Holyoak, 1997, for a discussion of the utility of the Weber law.)

Eq. 3 represents LISA's mapping-based estimate of the similarity of T to S. It is at a maximum when each t maps uniquely to one s (i.e., $m(t, s_{max}) = 1$ and $m(t, s_{max2}) = 0$) and is maximally similar to that s (i.e., $\cos(t, s_{max}) = 1$); in this case, $\sigma_{\mathrm{MIP}}(T, S) = n(T)/(1+n(T))$. In the parlance of Goldstone (1994), it is the contribution to similarity of "matches in place", or MIPs. However, Goldstone and others (e.g., Larkey & Markman, 2005) have shown that "matches out of place" (i.e., shared features on non-corresponding elements; MOPs) also contribute to judgments of similarity. To obtain a measure of the total similarity, $\sigma_{\mathrm{total}}(T, S)$, we therefore add a measure of the similarity, $\sigma_{\mathrm{MOP}}(T, S)$, of all objects and roles across T and S, whether they map or not:

$$\sigma_{\mathrm{MOP}}(T,S) = \frac{\sum_{\text{objects } t,s} \cos(t,s) + \sum_{\text{roles } t,s} \cos(t,s)}{1 + n_o(T)\, n_o(S) + n_r(S)\, n_r(T)}, \tag{4}$$

$$\sigma_{\text{total}}(T,S) = \sigma_{\text{MIP}}(T,S) + \sigma_{\text{MOP}}(T,S), \quad (5)$$

where $n_o(T)$, $n_r(T)$ and $n_o(S)$, $n_o(S)$ are the numbers of object and role units in T and S. Similarity in LISA (Eq. 5) is a simple sum of MIP and MOP similarities.

The components of Eq. 5—the assessment of mapping quality (Eq. 1), the use of the Weber fraction (Eq. 2) and the cosine measure of similarity (Eqs. 3 and 4)—were all part of LISA prior to their use in these equations. That is, Eq. 5 represents our best interpretation of LISA's estimate of similarity as embodied in the 2003 version of the model. With this formulation, we can now run "standard" LISA on several tasks that have been reported in the literature to test whether LISAese produces the same kinds of similarity judgments as human subjects.

Simulations

To assess LISA's adequacy as a model of similarity, we simulated both similarity effects previously captured by feature-based models and effects captured by structural models.

Violations of the Metric Axioms (Tversky, 1977)

A classic view of similarity states that concepts are represented as points in a "mental space" and that the similarity of two concepts is the inverse of distance between those points (e.g., Shepard, 1962a, 1962b). If concepts are represented as points in mental space, then the similarity of any pair of concepts must obey the metric axioms. However, Tversky (1977) presented evidence that human similarity judgments violate the metric axioms. As a first test of LISA's account of similarity, we used the model to simulate these violations of the metric axioms.

Symmetry: Symmetry states that the distance, $d(x, y)$, from point x to point y is equal to $d(y, x)$, the distance from y to x (and thus the similarity $s(x, y)$ must equal $s(y, x)$). Following Tversky (1977), we modeled a violation of symmetry by computing the similarity of China to North Korea and the similarity of North Korea to China. Motivated by Tversky's observation that his subjects likely knew more about China than North Korea, LISA's representation of China contained more propositions than its representation of North Korea (9 vs. 3; see Table 1). To simulate the process of making a similarity judgment, we allowed LISA to map China onto Korea (or vice-versa) and then applied Eq. 5, above.

Like Tversky's subjects, LISA rated the similarity of North Korea to China (σ_{total} = 0.877) greater than that of China to North Korea (σ_{total} = 0.385). This difference comes entirely from the σ_{MIP} term (Eq. 3; the σ_{MOP} term was the identical for the two ratings). When a sparse analog (North Korea) is mapped to a more complex analog (China), all the structures in the sparser analog can map uniquely to structures in the complex one, resulting in a proportion close to one. For the reverse judgment, many aspects of China did not map to anything Korea, yielding a smaller proportion of matches. Thus, LISA explains the violation of symmetry in terms of an asymmetry of mapping between analogs with different numbers of propositions.

Minimality: Minimality states that the distance between any point and itself is always zero, and thus equal for all points. Minimality implies that all concepts are equally self-similar. However, Tversky argued that people may view complex objects (e.g., a cubist painting) as more self-similar than simpler objects (e.g., a square).

We simulated this effect by computing the similarity of China (relatively complex) to itself and North Korea (relatively simple) to itself. Interestingly, the σ_{MIP} and σ_{MOP} measures of similarity yielded opposite patterns. For σ_{MIP}, China was more self-similar than Korea (0.851 vs. 0.662, respectively), whereas for MOPs, Korea was more self-similar than China (0.677 vs. 0.122). Goldstone and Medin (1994) showed that the effect of MOPs decreases during the time-course of comparison. LISA thus predicts that the relative self-similarity of simple and complex objects may reverse over time, with simple objects being more self-similar early in processing and complex objects being more self-similar later in processing. To our knowledge, this remains an untested prediction of the LISA model.

SYMMETRY and MINIMALITY	
Propositions	**CHINA** *people* (Chinese-people) *food* (Chinese-food) *government* (Chinese-government) *climate* (Chinese-climate) *history* (Chinese-history) *geography* (Chinese-geography) *customs* (Chinese-customs) *clothing* (Chinese-clothing) *economy* (Chinese-economy) **KOREA** *people* (Korean-people) *government* (Korean-government) *economy* (Korean-economy)
TRIANGLE INEQUALITY	
Propositions	**RUSSIA** *government* (Russian-government) **CUBA** *government* (Cuban-government) *climate* (Cuban-climate) **JAMAICA** *climate* (Jamaican-climate)

Table 1. Propositions used in the simulations of the violations of the metric axioms. Shared predicates across countries (e.g., *government*, across China and North Korea, or Russia and Cuba) are semantically identical, although their arguments (e.g., Chinese-government, Korean-government, etc.) are not.

Triangle inequality: The triangle inequality states that $d(x, z)$ can be no greater than $d(x,y)$ plus $d(y,z)$. Tversky

(1977) translated this inequality into similarities by noting that s(x, y) and s(y, z) together set a "lower limit" on the similarity s(x, z). An example that appears to violate this rule, taken from William James, is that Jamaica and Cuba are similar (geographically), and Cuba and Russia are similar (politically), but Jamaica and Russia are not similar at all (cf. Tversky & Gati, 1982).

LISA produces a similar violation given representations of Jamaica, Cuba, and Russia including only "geographical" and "political" propositions (see Table 1). LISA gave Jamaica and Cuba a total similarity of 1.006, and Cuba and Russia a similarity of 0.735, but Jamaica and Russia a similarity of zero. Clearly, whatever lower limit the triangle inequality should impose based on the first two similarities was violated.

Alignable Differences (Markman & Gentner, 1996)

Markman and Gentner (1996) argued that there are two psychologically distinct kinds of differences: *alignable* and *non-alignable*. A difference is alignable when the different elements form a correspondence (e.g., the number of wheels on a car vs. motorcycle). Conversely, a seat belt is a non-alignable difference between a car and a motorcycle, because nothing from a motorcycle corresponds to the seatbelt of a car. Markman and Gentner demonstrated that alignable differences have a greater impact on similarity than non-alignable differences.

Markman & Gentner (1996)	
Propositions	**A** *fixing* (machine, car) **B** *fixing* (machine, truck) **C** *fixing* (machine, robot) **D** *fixing* (machine, car); truck **E** *fixing* (machine, car); robot
Semantics	car: s1, s2, s3, s4, s5 machine: s1, s6, s7, s8, s9 truck: s1, s2, s3, s10, s11 robot: s1, s6, s13, s14, s15

Table 2. Propositions: The propositions and isolated objects (truck and robot in D and E) in the five situations depicted in Markman and Gentner (1996), Experiment 2 and simulated using LISA. Semantics: The semantic features representing the objects in the LISA simulations.

We modeled this effect using stimuli from their second experiment. In this study, participants rated the similarity of a base item, A, to four other items (B–E), which varied in their alignable and non-alignable differences with A. The base item was a machine fixing a car, represented by *fixing* (machine, car), B was *fixing* (machine, truck), C was *fixing* (machine, robot), D was *fixing* (machine, car) with a truck in the background, and E was *fixing* (machine, car) with a robot in the background. The truck and robot in B and C formed an alignable difference with the car in A, whereas in D and E they formed non-alignable differences with A. In our simulations, the truck shared three of five features with the car, whereas the robot shared one (see Table 2).

Comparison		Alignable advantage
M&G 1996	A-B minus A-C	1.450
	A-D minus A-E	0.360
LISA	A-B minus A-C	0.089
	A-D minus A-E	0.018

Table 3. Alignable advantage for human subjects and LISA. A-B, for example, refers to the judged similarity of item A to item B.

As shown in Table 3, for both people and LISA, the effect of the truck vs. the robot is roughly four times greater when they are MIPs (A-B minus A-C) than when they are MOPs (A-D minus A-E). Both the σ_{MIP} and the σ_{MOP} terms lead to this effect. The A-B minus A-C difference is large because *car* maps to either *truck* (A-B) or *robot* (A-C), and *car* shares more semantic units with *truck* than with *robot*, resulting in a higher σ_{MIP} term for A-B relative to A-C. Additionally, the greater similarity of *car* to *truck* yields a higher σ_{MOP} term for A-B than for A-C. The A-D minus A-E effect is smaller because the semantic similarities between car & truck and car & robot are only captured by the σ_{MOP} term; the σ_{MIP} term is dominated by the alignable difference of *car* and *truck*, and thus, differs little from A-D to A-E.

MIPs and MOPs

Goldstone demonstrated that MIPs and MOPs both affect similarity judgments (Goldstone, 1994; Goldstone & Medin, 1994; Larkey & Markman, 2005). An example MIP is matching colors on the t-shirts of two persons; a MOP would be matching colors on different garments.

To show how MOPs affect similarity, Larkey and Markman (2005) created pairs of objects that varied in shape and color. Their participants rated the similarity of one pair of objects to another pair, where in each comparison, the second pair was a systematic transformation of the first. A transformation was either a change in shape or a change in color. Pairs compared to an "AB" stimulus are represented in Figure 2 along the abscissa, along with the human data and LISA's fits. Matching values (of shape or color) are represented as matching letters.

LISA represented each pair with two propositions per object: *color* (obj1), *shape* (obj1), *color* (obj2), and *shape* (obj2). Even with these very simple representations, LISA captures the vast majority of the variance in the human data (r^2 = 0.91) and fits all ordinal differences in the similarity ratings (see Larkey and Markman, 2005). To highlight the major results, LISA's ratings were sensitive to both MIPs and MOPs and were sensitive to the interactions between MIPs and MOPs in the same way as human ratings.

Larkey and Markman (2005) further explored their results by fitting three similarity models to their data: SME (Falkenhainer, et al., 1989), CAB (Larkey and Love, 2003), and SIAM (Goldstone, 1994). Only SIAM fit the full pattern of results shown in Table 2. One major difference between SIAM and LISA is that similarity in SIAM is always mediated by mapping connections, or correspondences, between representation elements. In contrast, LISA's σ_{MOP} term factors MOPs into similarity independent of mapping connections. In future work, we are interested in comparing the assumptions of these two mechanisms empirically.

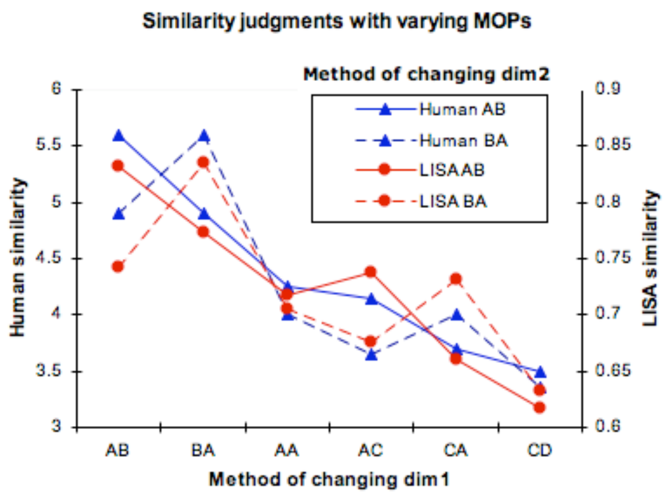

Figure 2. Larkey and Markman (2005) data and LISA fits.

Discussion

LISA was designed to simulate analogical reasoning, not similarity judgments. However, LISA simulated numerous findings from the similarity literature using only algorithmic ingredients already in place for other functions. These successes show that LISA's knowledge representations may accurately simulate those underlying peoples' explicit similarity judgments, in particular, and relational thought more generally.

References

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

Gentner, D., & Markman, A.B. (1997). Structure mapping in analogy and similarity. *American Psychologist*, *52*, 45-56.

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

Goldstone, R. L., Medin, D. L., & Gentner, D. (1991). Relational similarity and the non-independence of features in similarity judgments. *Cognitive Psychology, 23,* 222-262.

Goldstone, R. L., & Medin, D. L. (1994). The time course of comparison. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 29-50.

Goldstone, R. L. (1994). Similarity, interactive activation, and mapping. *Journal of Experimental Psychology: Learning, Memory,* & *Cognition*, *20*, 3-28.

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*, 220-264.

Kroger J. K., Holyoak K. J., Hummel J. E. (2004). Varieties of sameness: the impact of relational complexity on perceptual comparisons. *Cognitive Science*, *28*, 335-358.

Lassaline, M. E. (1996). Structural alignment in induction and similarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22,* 754 -770.

Larkey, L. B., & Love, B. C. (2003). CAB: Connectionist analogy builder. *Cognitive Science*, 27, 781–794.

Larkey, L. B., & Markman, A. B. (2005). Processes of similarity judgment. *Cognitive Science*, *29*, 1061-1076.

Markman, A. B., & Gentner, D. (1996). Commonalities and differences in similarity comparisons. *Memory and Cognition*, *24*, 235–249.

Nosofsky, R. M. (1992). Exemplar-based approach to relating categorization, identification, and recognition. In *Multidimensional models of perception and cognition. Scientific psychology series*, F. G. Ashby, Editor. Lawrence Erlbaum Associates, Inc: Hillsdale, NJ, England.

Shepard, R. N. (1962). The analysis of proximities: Multidimensional scaling with an unknown distance function. Part 1. *Psychometrika*, *1*, 125-140.

Tversky, A. (1977). Features of similarity. *Psychological Review*, *84*, 327–352.

Viskontas, I., Morrison, R., Holyoak, K. J., Hummel, J. E., & Knowlton, B. J. (2004) Relational integration, inhibition, and analogical reasoning in older adults. *Psychology and Aging*, *19*, 581 – 591.