INDIVIDUAL DIFFERENCES IN HOW STRUCTURE AFFECTS SIMILARITY

David H. Landy
dhlandy@gmail.com
Eric G. Taylor
etaylor4@uiuc.edu
Department of Psychology, 603 E. Daniel St.
Champaign, IL 61820 USA

ABSTRACT

A general disconnect exists between most models of structure-based similarity and the empirical literature they attempt to capture. Models of structure processing often purport to explain behavior at the level of the individual, but most studies have analyzed structure sensitivity at the level of mean aggregates across individuals. We begin to address this issue by presenting several re-analyses of similarity ratings among simple objects originally presented by Larkey and Markman (2005), specifically attempting to identify sources of individual variability. Individual variability was generally coherent across items of a given type, suggesting that sample noise is to some degree reflective of individual differences, and not mere response noise (Goldstone, 1994). We find three primary sources of variation among individuals: value placed on feature matches which compete with identical feature matches on other objects (a kind of MOP, see Goldstone, 1994), value placed on feature matches that support analogical matches (MIPs), and value placed on the holistic well-matchedness of each object in the scene. We suggest that the coherent variability of individuals across items provides a valuable constraint on the processes of structure sensitivity, and may shed light on the relationship between analogical reasoning and other aspects of human cognition. Furthermore, successful process models of structure processing should capture variation among individuals through systematic variation of parameters

INTRODUCTION

Structural information is critical throughout cognition. It plays a role in many, perhaps most, high-level cognitive acts (Gentner, 2003; Hofstadter, 2001), from recognizing objects (Biederman, 1987), to comparing scenes (Goldstone, 1994; Markman and Gentner, 1993), to making decisions (Petkov & Kokinov, 2006), to building explanations (Falkenhainer and Forbus, 1990; Hummel, Landy & Devnich, 2008). Computational models have been instrumental in understanding the effects of structure on performance. However, while the models typically characterize the cognitive processes of individuals, in practice, performance and model fits have always been presented at the level of groups (e.g., Falkenhainer et al, 1989; Hummel & Holyoak, 1997; Goldstone, 1994; Larkey and Love, 2003). If we assume that structure processing is an unvarying fact about cognition, that individuals do not differ systematically in structure processing, then analyses on group means are ideally revealing. However, to the extent that the mechanisms that underlie performance in structure sensitive tasks do vary, the standard research practice runs the risk of missing a potentially interesting aspect of human structure processing.

In this paper, we will analyze judgments of similarity between simple objects. Aspects of such judgments have frequently been attributed to structure-sensitive processes (e.g., Goldstone, 1994; Larkey & Markman, 2005; Taylor & Hummel, in press), and many models attempt to capture both analogical correspondence and the role of structure in similarity judgments (SME, Falkenhainer et al, 1989; CAB, Larkey & Love, 2003; SIAM, Goldstone, 1994; Taylor & Hummel, in press). In some of these models (e.g., SIAM), similarity and structural evaluation are tightly linked processes; in others similarity is computed on the basis of structural correspondence (e.g., LISA; see Taylor & Hummel, in press). In the latter case, our discussion applies primarily to the application of structure matches in similarity judgments, rather than to the evaluation of structural correspondences themselves.

BACKGROUND: CONSTRAINTS ON STRUCTURE PROCESSING

A major goal of analogy research has been to identify the constraints on structure processing based on relational and non-relational properties of potential analogs. Gentner (1983) originally argued for the importance of structural relations in analogy, distinguishing relational matches from object-based matches. Falkenhainer, Forbus, and Gentner (1990) developed the Structure Mapping Engine (SME) to evaluate analogies based on systems of relations, preferring analogies with greater structural consistency. SME accounts for people's general tendency to prefer relational matches to purely object-based matches (e.g., Gentner et al, 1993) and for people's similarity judgments, which show a greater impact of differences among objects that that play similar relational roles (alignable differences) than objects that play different roles (non-alignable differences; Markman & Gentner, 1993).

Emphasis on the role of structure and relational alignment is clearly crucial for understanding human behavior; however, in at least some domains in which structure plays a crucial role, relational alignability interacts with other factors in determining behavior. Goldstone (1994; Goldstone & Medin, 1994) demonstrated that while feature matches do contribute to similarity, even for structurally similar objects, these feature matches interact with structural properties. Early stages of structure processing rely crucially on "matches out of place" (MOPs), or feature matches across items that play different relational roles. After structural alignments become more obvious, "matches in place" (MIPs) come to dominate. Furthermore, feature similarities show a preference for 1-to-1 correspondences, so that MOPs that compete with a MIP contribute little if anything to similarities (a MOP competes with a MIP when they both relate the same feature of the same object, but one (the MIP) aligns with the higher level structure). Larkey and Markman (2005) replicated these finding for non-speeded judgments and showed that only one extant model of structure processing, SIAM (Goldstone, 1994), provided qualitative fits to MOP-sensitive similarity judgments. In particular, these data provide constraints on candidate models of analogical reasoning, selecting only those that incorporate a role for featural/semantic, as well as structural representations of objects under comparison (c.f. Taylor & Hummel, in press).

INDIVIDUAL CONSTRAINTS ON STRUCTURE PROCESSING

Although the prior literature has capitalized on cases where people tend to agree about what is a good analogy, it is readily apparent that substantial disagreement about the quality of particular alignments exists as well, perhaps especially among people interested in analogy. Indeed, in one lab meeting during which we discussed this topic, conversation descended into congenial name-calling, divided between those who found the pair of pairs of items in Figure 1A most appealing, and those who preferred the comparison illustrated by Figure 1B.

The distribution of traits that run in our lab are admittedly unlikely to generalize to the larger population, but fortunately, other empirical evidence also points toward individual differences in relational reasoning. In the domain of similarity judgment, Simmons and Estes (2008) found that people systematically differed in their ratings of the similarity of items that were thematically or analogically related. Furthermore, individuals that gave higher ratings for thematically similar items had a lower need for cognition (c.f., Cacioppo, Petty, & Kao, 1984). Another source of individual differences in relational reasoning stems from the work on culture and cognition. Nisbett and colleagues (e.g, Nisbett, Peng, Choi, & Norenzayan, 2001) argue from an extensive database of empirical results that members of East Asian culture show more sensitivity to the relations of objects to context than do members of Western culture.

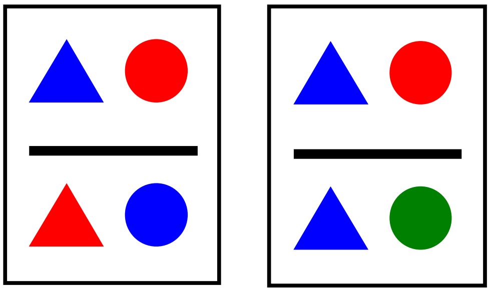

Figures 1a (left; pair AB/BA) and 1b (right; AB/AC). Two object pairs with disputable relative similarities.

These data argue compellingly that different individuals do, in fact, differ in the way they treat relational and associative information, but of course it is unlikely that they exhaust the sources of variation. Furthermore, it is not clear what the impact of these results is on theories of structure processing. For instance, does variability in the use of thematic (co-occurrence) information suggest variability in a structure processing, or variability in how people interpret the word "similarity"? Consonant with the latter interpretation, Simmons and Estes (2008) found that people who identified thematic proximity as a valid source of similarity were far more likely to use thematic proximity in making similarity judgments, suggesting that the primary differences lay how the participants construed similarity. Thus, it is not clear how relevant such differences are to theories of structure mapping. Nevertheless, there is some suggestion that how people invoke structure in making decisions may differ among individuals.

In this paper we report evidence of individual differences in structure processing mechanisms from a task designed to test models of similarity. With this evidence, we will then focus our discussion on implications for models and theories of structure-based similarity judgments more generally.

ANALYSES

Description of Data

We analyzed a set of similarity judgments over pairs of object pairs. This data was originally collected by Levi Larkey and Art Markman. Larkey and Markman (2005) presented a portion of this data set, along with the fits of a variety of models to the mean ratings of various types of items. Particulars of the experimental design and procedure are presented in detail in the original paper. A total of 116 people participated in two studies conducted at the University of Texas at Austin. The first study (n=58) manipulated the color and shape of simple object pairs; the second study (n=58) manipulated color and texture.
Participants rated 162 displays, each consisting of two pairs of objects. Participants were instructed to rate the similarity of the pairs on a scale from 1 to 6. The objects varied in either color or shape, or in color and texture pattern; there were four possible values along each dimension. One pair in any rated item (i.e., a single display) consisted of two objects with different values on each dimension. The other pair varied systematically.

Rated items (pairs of object pairs) are coded based on the relations between the object features (see Table 1). Following Larkey and Markman's coding scheme, object pairs are coded by their values on the relevant dimensions. The base pair is always coded AB/AB: that is, along one dimension, the value that one object had is named A, the value that the other object had is named B; the second dimension is named identically. The second pair is then named for the corresponding feature values in corresponding objects. To demonstrate, Figure 1b depicts an AB/AC pair. The shapes from the top two objects (triangle=A and circle=B) are repeated for the bottom two objects, hence, the first AB. The color from the top left object (color=A) is repeated for the bottom left object, but the color for the other object is new (color=C), hence AC. Using this scheme, Larkey and Markman created nine patterns of single-dimension feature shifting. These were presented in all possible pairwise combinations, across two different spatial relationships.

Goldstone (1994) identified two types of similarities between object pairs like those in Table 1—"matches in place" (MIPs) and matches out of place (MOPs). Take for example the item pairs with code AB/AA from Table 1. An example MIP is the match in the shapes of the two squares; the match is "in place", because the square on the left maps (corresponds) best to the square on the right. An example MOP is the matching colors on the square and the circle; this match is "out of place" because the square and circle do not map.

Following Goldstone (1994), we distinguish between MOPs that compete with a corresponding MIP in the same feature (C-MOPs), and those which do not (NC-MOPs). An example of a competing MOP is given by the AB/AA item in Table 1. The matching color blue between the square on the left and the circle on the right is a competing MOP, because the color blue also creates a MIP between the square on the left and the square on the right. That is, the MOP dimension "blue" also contributes to a MIP for the same pairs. The AB/CA item shows a similar item with an NC-MOP, rather than a C-MOP.

We re-analyzed the entire data set from Larkey and Markman, collapsing across dimension and spatial position, since the authors found little effect of spatial relations on the judgments. After simplification across notational symmetries, this analysis separates items into 16 distinct types (see Table 1 for various relevant properties of the transformations).

Correlations among judgments

Our first goal was to verify whether there are, in fact, sources of noise based in individual preferences (other than stochastic noise).

Table 1: Types of stimuli from Larkey and Markman (2005)

Code	Example	MIP	C-MOP	NC-MOP	Num. Good
AA/AA		2	2	0	3
AA/AC		2	1	0	3
AA/BC		2	1	0	4
AA/CD		1	1	0	4
AB/AA		3	1	0	4
AB/AB		4	0	0	4
AB/AC		3	0	0	4
AB/BA	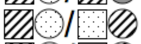	2	0	2	4
AB/CA		2	0	1	4
AB/CD		2	0	0	4
AC/AC		2	0	0	2
AC/BC		2	0	1	3
AC/CA		1	0	1	3
AC/CB		2	0	0	4
AC/CD		1	0	0	2
CD/CD		0	0	0	0

First consider a null model, in which each individual really is identical. To reiterate: this is not a view anyone holds explicitly, but it is the implicit working assumption of formal modeling efforts. If the impact of environmental structure on similarity calculations really were fixed, one might imagine that each participant in a study arrives with identical preference judgments, and that both inter- and intra- participant noise comes from noise on the response to each item. Goldstone (1994) followed this approach by modeling variability in responses by adding Gaussian noise to the output of the SIAM model. Taken as a theoretical commitment (which is certainly not what Goldstone intended), this assumption predicts that ratings made on different items by particular subjects would be independent. The alternative approach—assuming that noise results from genuine (possibly parametric) differences between individuals in the processes that generate or use structural alignments, predicts that item preferences should be systematic across participants. We therefore explored the correlations, across participants, for each of the 128 item types. Table 2 presents all of the pairings that were significant after Bonferroni correction ($p<0.00046$) or marginally significant results ($p<0.00094$).

Table 2. Correlations

1st pair	2nd pair	r
AB/AB	AB/AC	.32*
AB/AB	AA/BC	-.44
AB/BA	AB/CA	.39
AB/BA	AA/AC	-.41
AB/AA	AB/AC	.40
AB/AA	AC/CD	-.5
AB/AC	AA/BC	-.56
AB/AC	AA/AA	-.32
AB/AC	AA/CD	-.61
AB/AC	AC/AC	.33
AB/AC	AB/CA	-.33
AB/CA	AB/CD	.54
AB/CA	AA/AA	-.46
AB/CA	AA/AC	-.56
AB/CD	AA/BC	-.34
AB/CD	AA/AA	-.45
AB/CD	AA/AC	-.56
AA/BC	AA/AA	.46
AA/BC	AA/CD	.47
AA/BC	AA/AC	.32*
AA/BC	AC/AC	-.34
AA/AA	AA/CD	.36
AA/AA	AA/AC	.49
AA/AA	CD/CD	-.33
AA/CD	AC/AC	-.33
AC/BC	AC/CD	.38
AC/CD	CD/CD	.50

As can be seen, even after correcting for the large number of comparisons performed, many substantial correlations were found. Participants in this study appear to vary systematically across multiple items. Furthermore, some patterns can already be discerned in these data. Correlations tend to be strongly positive among pairs that for both items had the same value along some dimension (i.e., items with an "AA"), and negative between "AA" pairs and other pairs.

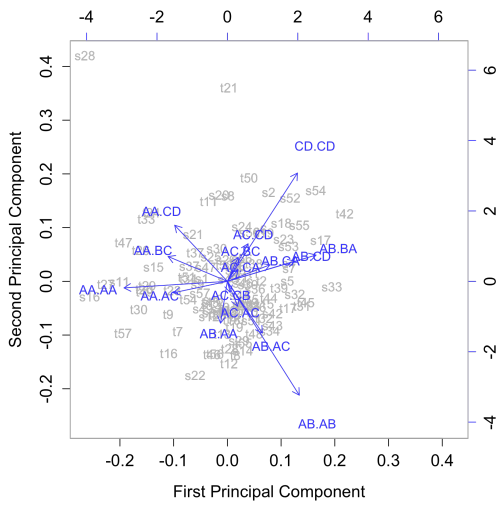

Figure 2: Relations between principle components 1 and 2.

Principal components of variation

Having established the existence of patterns among individual subjects, we used a Principal Components Analysis to attempt to clarify the sources of meaningful variation.
We analyzed the components of the space of participant responses, using each item as an input dimension. We independently analyzed Larkey and Markman's Texture and Shape experiments, and found very similar, and generally independently significant patterns in each group (with one noted exception). For ease and clarity of exposition, we present data here from an analysis of the combined data set. The first four dimensions of this analysis accounted for significant proportions of the variance among subjects by the broken stick test (c.f. Jolliffe, 2001); we were only able to satisfactorily interpret the first three. Between them, these three dimensions accounted for 57% of the subject-wise variance. Biplots of the first and second, and first and third dimensions are shown in figures 2 and 3.

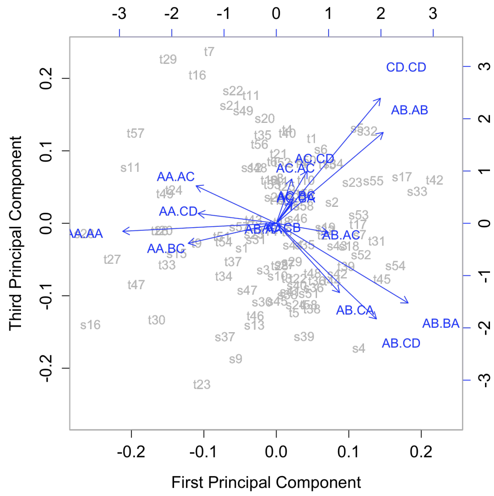

Figure 3: Relations between principle components 1 and 3.

The first principle component accounted for 24% of the total variance. Loadings on this dimension are shown in Figure 4a. We interpret this dimension as consistent primarily of a sensitivity to the presence and number of C-MOPs in a particular stimulus type (shown in Figure 4b). The correlation between number of C-MOPs and the first principal component was r=.86 (p<.001). Participants at one extreme of this dimension preferred a C-MOP to a simple mismatch (for instance, they preferred AB/AA to AB/AC); participants at the other end gave higher ratings to items with simple mismatches.

The second principal component, which accounted for 18% of the variance (Figure 5a) is readily interpreted as the number of MIPs a stimulus has (r=.92; p<.001). MIPs were generally positive—for 115 of the 116 participants, number of MIPs correlated positively with ratings; however, the magnitude of that positive influence varied substantially across individuals.

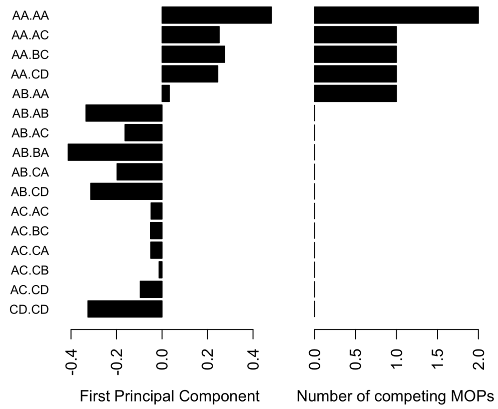

Figures 4a (left) and 4b (right): Loading of principle component 1 and number of competing MOPs for each item type.

Both of the first two components corresponded quite closely to factors that have previously been identified as relevant to similarity judgments. The third principal component has no such ready interpretation. However, this dimension does correspond reasonably well (r=.69, p<.01) to a rather intuitive idea: the number of objects in a scene that have some shared features, ignoring correspondences.

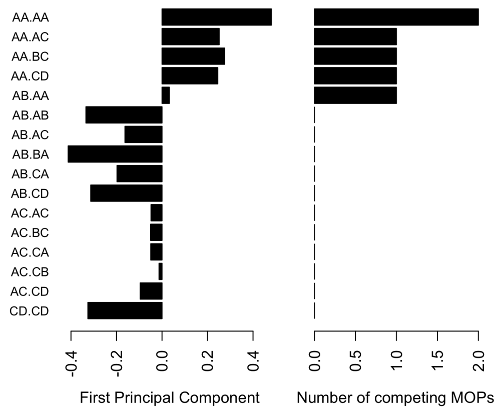
Figures 4a (left) and 4b (right): Loading of principle component 1 and number of competing MOPs for each item type.

Consider the pair AC/BD. In this item, three of the objects have some featural overlap with an item in the comparison pair. The fourth object (the one with features "C" and "D") has no featural overlap with anything. This item therefore gets coded with the value three.

The correspondence between the number of good feature matches and the third principal component was significant, as mentioned above. However, in this case, the correspondence was only significant in the texture experiment (r=.81, p<.0001). In the shape experiment, the correlation was less robust (r=.49, p=.054). Although the dimension correlates quite well with the number of objects with a match, it substantially mis-predicts the value of the third principal component on identical items (see Figures 6a and 6b), and fails to capture qualitatively important aspects of the third component, such as the difference between the three largest positive values, and the other values. This source of variation requires further study before it can be confidently interpreted.

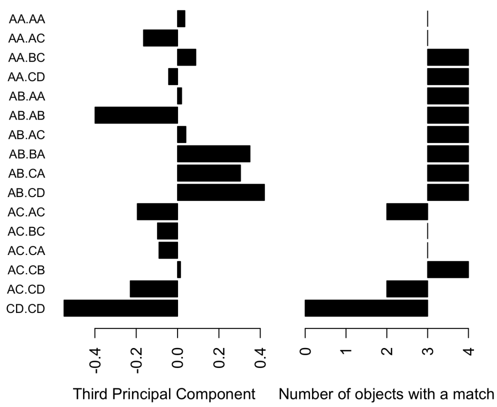

Figures 6a (left) and 6b (right): Loadings of principle component 3 and number of objects with a match for each item type.

Magnitude of the variation

Generally, different groups of individuals identified by the PCA analysis showed different but related ordinal rating patterns. As an illustration, Table X shows the order among mean ratings for the first and fourth quartile of dimension 1. As expected, people in one extreme rated AB/AA systematically higher than people inhabiting the opposite extreme. Indeed, neither group's average matches the overall average patterns reported by Larkey & Markman (2005). In the first quartile, AB/AA is stronger than AB/BA, while in the fourth, it is weaker than AB/AC (for this group, having an extra MOP actually appears to weaken the mapping strength!). Moreover, the shift generalizes quite systematically. Fourth quartile subjects systematically dis-prefer patterns with C-MOPs (items with code --/AA or AA/--) relative to those in the first quartile by 2 to 4 positions.

Table 3: Ordinal mean ratings among different quartiles of dimension 1

Rating	1st Quartile	4th Quartile
1 (highest)	AB/AB	AB/AB
2	AB/AA	AB/BA
3	AB/BA	AB/AC
4	AB/AC	AB/AA
5	AA/AA	AB/CA
6	AA/AC	AB/CD
7	AB/CA	AC/AC
8	AA/BC	AA/AA
9	AB/CD	AA/AC
10	AA/CD	AC/CB
11	AC/AC	AA/BC
12	AC/CB	AC/CA
13	AC/CA	AC/BC
14	AC/BC	AA/CD
15	AC/CD	AC/CD
16 (lowest)	CD/CD	CD/CD

On this analysis, the ordering presented by Larkey and Markman appears to be an average across disparate individuals rather than a uniform absolute. Nonetheless, it is also the case that many individuals did conform to the average ordering. About 40% of subjects fit the ordinal relationships Larkey and Markman report as significant. Thus, a successful model of similarity will be compatible with the mean behavior pattern—but it will have to be compatible with other orderings as well.

DISCUSSION

We examined a collection of similarity ratings for pairs of object pairs differing along two simple features. Our analysis revealed significant and interpretable patterns of individual variation in endorsement of various kinds of comparisons, corresponding to the first three components of variation in the principal components analysis. The first principle component revealed that participants varied in how they weighed repeated matches of the same value along a particular dimension (C-MOPs). This is particularly interesting: while such matches have been noted to have little positive impact on match quality (Goldstone, 1994), previous analyses have not considered whether this fact resulted from a consistently small positive weight accorded to such features, or to differences among participants. Here, we see that in fact most subjects do weight these repeated features, but ‘disagree’ about whether they positively impact scene similarity.

We interpreted the second principal component as variation in how much the presence of matches in place affected participants’ judgments of similarity. All but one participant positively weighted feature matches among corresponding objects; however there was a substantial range in just how much these MIPs mattered to similarity.

Finally, and somewhat more speculatively, the third principle component suggests that some participants highly weight scenes in which each object had at least one “good match”, regardless of the higher-order structural correspondences. Again, for all participants more matches contributed to similarity, but there was substantial range in the value placed on this kind of holistic integration.

Although we found substantial variability across participants, our analyses also show robust commonalities. For essentially every subject, matches in place increased match quality—though such matches mattered substantially more to some raters than to others. This finding is consistent with the hypothesis that most participants were using structured correspondences to evaluate similarity but that this process had different outcomes for different individuals.

Put this way, a tempting account of the individual differences may be to assume that all participants used a general-purpose structure-processing-and-similarity algorithm, but some participants used it more than others. However, an account based purely on degrees of a single structure processing mechanism seems unable to account for the results. Both MIPs and C-MOPs are only defined in the context of structural matches, and variation in their affect on similarity was largely orthogonal. Since both sources of variation require structural analysis, these findings are more readily interpreted by assuming that different people are using structural information in different ways. In the case of a model such as SIAM, this suggests variation in the mappings themselves, since similarity and mapping are tightly integrated. Other models which propose a sequential process in which the current best mapping is evaluated independently of similarity, which is then computed from the mapping (e.g., LISA) might posit differences in how a fixed mapping translates into similarities. Either way, it seems unlikely that our results can be accommodated by a simple assumption of more or less use of structure.

A promising approach to accounting for individual variability, in our view, is to incorporate into computational models theories and commitments about how parameters within the model vary across individual reasoners and reasoning instances (see Larkey and Love, 2003, for one example of this kind of reasoning).

In fact, we believe that research into individual differences in structure processing will have substantial consequences for the evaluation of models. Consider for example the model comparison study by Larkey and Markman (2005). The authors concluded that, of the models they tested, only SIAM could correctly predict the ordinal pattern of average judgments using a single set of parameters. Following standard procedures for the evaluation of models in the analogy literature, they are quite correct (although LISA has since been shown also to be able to match this aggregate pattern, Taylor & Hummel, in press). However, to the extent that individual differences play a substantial role in structure processing, future studies following this approach to model comparison will be incomplete, and the conclusions potentially incorrect. Such studies should also consider whether the models capture the distribution of similarity ratings using a principled set of parameter ranges, rather than a single set of parameters. If alternative models such as CAB (Larkey & Love, 2003) and SME (Falkenhainer et al, 1989) are able to successfully describe this distribution, then their accounts remain quite viable.

Another more empirically-based direction for the study of individual differences is to explore personality, cultural, or situational factors that influence structure-sensitivity. For example, Kim and Markman (2006) found with both reasoning and memory tasks that people with higher "fear of isolation" (FOI) show increased sensitivity to relations between objects and contexts. This was true in populations with chronic FOI and in experimentally induced FOI participants. Kim, Narvaez, and Markman (2007) also found that individuals with an independent, as opposed to interdependent self-construal, showed heightened sensitivity to contextual relations. Taken together, modeling and empirical approaches that incorporate a role for individual differences promise to shed substantial light on theories of structure processing.

A final, more speculative direction to consider is to evaluate the role of other structure sensitive processes in leading to the individual differences in similarity judgment. It could be that people with different similarity profiles bring different goals to bear on the similarity task. If so, then differences in structure processing may point to the influence of other processes on the fringe of analogy and comparison. Future work could integrate similarity judgment with other tasks to test for these relationships.

CONCLUSION

Our extended analyses of the data collected by Larkey and Markman (2005) show that individuals' similarity judgments varied systematically as a function of (a) number of competing MOPs, or MOPs with the same feature value as a separate MIP, (b) number of MIPs, and (c) number of object with at least one feature match. The weightings of these three dimensions varied orthogonally across individuals and each explained a substantial portion of variance. These findings pose a challenge to extant theories and models of structure-based similarity judgment, which characterize structure processing at the individual's behavior but are traditionally fit to aggregate data.

ACKNOWLEDGMENTS

We thank Levi Larkey and Art Markman for providing access to their data. John Hummel, Robert Goldstone, and Wookyoung Jung provided valuable comments and suggestions. This research was supported by AFOSR Grant # FA9550-07-1-0147.

REFERENCES

Cacioppo, J.T., Petty. R.E., & Kao, C.F. (1984). The efficient assessment of need for cognition. Journal of Personality Assessment, 48, 306–307.

Forbus, K. D. & Falkenhainer, B. (1990). Self-Explanatory Simulations: An Integration of Qualitative and Quantitative Knowledge. AAAI 1990, 380-387.

Falkenhainer, B., Forbus, K. D., and Gentner, D., 1990, The Structure mapping engine: algorithm and examples, Artificial Intelligence 41:1-63.

Gentner, D. (2003). Why We're So Smart. In D. Gentner and S. Goldin-Meadow, Language in mind: Advances in the study of language and thought. Cambridge, MA, MIT Press: 195-235.

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

Gentner, D., & Rattermann, M.J., & Forbus, K. (1993). The roles of similarity in transfer: Separating retrievability and inferential soundness. Cognitive Psychology, 25, 524-575.

Goldstone, R. L. (1994). Similarity, Interactive Activation, and Mapping. Journal of Experimental Psychology: Learning, Memory, and Cognition 20(1), 3-28.

Goldstone, R. L. & Medin, D. L. (1994). Time Course of Comparison. Journal of Experimental Psychology: General 20(1), 29-50.

Hofstadter, D. (2001). Analogy as the Core of Cognition. In Dedre Gentner, Boicho Kokinov, and Keith Holyoak (eds.), The Analogical Mind: Perspectives from Cognitive Science. Cambridge, MA: MIT Press/Bradford Book, 499-538.

Hofstadter, D. R. and Mitchell, M. (1994). The Copycat project: A model of mental fluidity and analogy-making. In K. Holyoak and J. Barnden (editors), Advances in Connectionist and Neural Computation Theory, Volume 2: Analogical Connections. Norwood, NJ: Ablex Publishing Corporation.

Hummel, J., Holyoak, K. J. (1997). "Distributed Representation of Structure: A Theory of Analogical Access and Mapping." Psychological Review, 104(3), 427-466.

Hummel, J. E., Landy, D. H., & Devnich, D. (2008). Toward a process model of explanation with implications for the type-token problem. AAAI Fall Symposium on Naturally-Inspired Artificial Intelligence. Arlington, VA.

Jolliffe, I.T. (2002). Principal Components Analysis. Springer-Verlag: NY.

Kim, K. & Markman, A. B. (2006). Differences in Fear of Isolation as an Explanation of Cultural Differences: Evidence from Memory and Reasoning. Journal of Experimental Social Psychology, 42, 350-364.

Kim, K., Narvaez, L.R., & Markman, A.B. (2007). Self-construal and the processing of covariation information in causal reasoning. Memory and Cognition 35(6), 1337-1343.

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

Larkey, L. B. & Love, B. C. (2003). CAB: Connectionist Analogy Builder. Cognitive Science, 27, 781-794.

Markman, A. B. & Gentner, D. (1993). Structural Alignment during Similarity Comparison. Cognitive Psychology 25, 431-467.

Nisbett, R. E., Peng, K., Choi, I., & Norenzayan, A. (2001). Culture and systems of thought: Holistic vs. analytic cognition. Psychological Review, 108, 291-310.

Petkov, G. & Kokinov, B. (2006). JUDGEMAP—Integration of Analogy-Making, Judgment, and Choice. Proceedings of the 28th Annual Conference of the Cognitive Science Society. Erlbaum, Hillsdale, NJ.

Simmons, S. & Estes, Z. (2008). Individual differences in the perception of similarity and difference. Cognition, 108, 781-795.

Taylor, E. G., Hummel, J. E. (in press). Finding similarity in a model of relational reasoning. Cognitive Systems Research. Special Issue on Analogies: Integrating Cognitive Abilities.