

Brief article

Cultural commonalities and differences in spatial problem-solving: A computational analysis

Andrew Lovett*, Kenneth Forbus

EECS Department, Northwestern University, 2133 Sheridan Rd., Evanston, IL 60208, United States

ARTICLE INFO

Article history:

Received 24 September 2010

Revised 21 June 2011

Accepted 23 June 2011

Available online 30 July 2011

Keywords:

Spatial cognition

Computational modeling

Structure-mapping

Cross-cultural

Qualitative representation

ABSTRACT

A fundamental question in human cognition is how people reason about space. We use a computational model to explore cross-cultural commonalities and differences in spatial cognition. Our model is based upon two hypotheses: (1) the structure-mapping model of analogy can explain the visual comparisons used in spatial reasoning; and (2) qualitative, structural representations are computed by people's visual systems and used in these comparisons. We apply our model to a visual *oddy task*, in which individuals are shown an array of two-dimensional images and asked to pick the one that does not belong. This task was previously used to evaluate understanding of geometric concepts in two disparate populations: North Americans, and the Mundurukú, a South American indigenous group. Our model automatically generates representations of each hand-segmented image and compares them to solve the task. The model achieves human-level performance on this task, and problems that are hard for the model are also difficult for people in both cultures. Furthermore, ablation studies on the model suggest explanations for cross-cultural differences in terms of differences in spatial representations.

© 2011 Elsevier B.V. All rights reserved.

1. Introduction

A key question in cognition is how we represent and reason about space. Dehaene, Izard, Pica, and Spelke (2006) demonstrated that aspects of human reasoning about two-dimensional space are universal, rather than culturally-specific. They used an *oddy task* (Fig. 1), in which participants were shown an array of six images and asked to pick the image that did not belong. Solving this task requires sensitivity to two-dimensional geometric concepts such as parallel lines, right angles, and axial symmetry. Dehaene et al. compared the performance of two population groups: North Americans and the Mundurukú, a South American indigenous group. They found that while the Americans performed better than the Mundurukú overall, the Mundurukú performed above chance on nearly

all problems. Furthermore, the error patterns correlated across cultures. Thus, the Mundurukú appeared to utilize the same geometric concepts as the Americans, despite having few words for such concepts in their language and no formal schooling in geometry.

These results have been taken as evidence for *core knowledge* of geometry (Spelke & Kinzler, 2007), an innate, universal cognitive module that deals specifically with geometric concepts. However, the performance of the two groups was not identical: Mundurukú of all ages performed at about the same level, while Americans tended to improve over time (Newcombe & Uttal, 2006), suggesting that culture-specific learning is a factor. Further research is needed to determine which aspects of spatial cognition are universal and which are learned. Here we use a computational model to study this task. Our model performs similarly to both groups on the *oddy task*. By ablating different aspects of the model, we can cause it to perform more like one group or the other, thereby producing novel, testable hypotheses about how spatial cognition varies across cultures.

* Corresponding author. Address: 807 Church St., Apt. 210, Evanston, IL 60201, United States. Tel.: +1 865 776 1430.

E-mail address: Andrew@cs.northwestern.edu (A. Lovett).

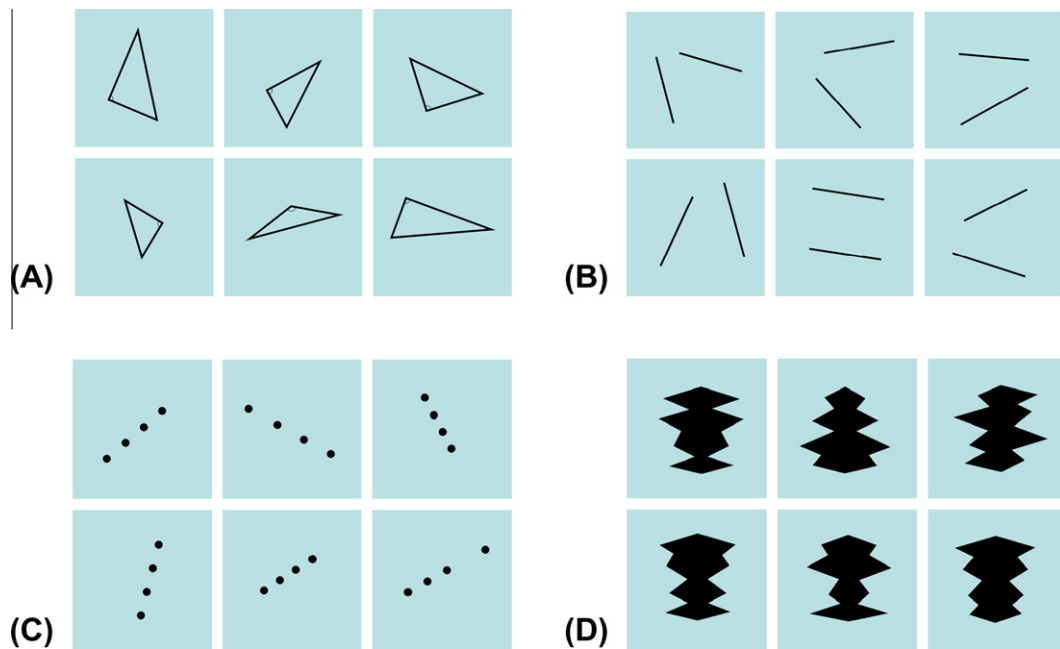


Fig. 1. Four oddity task problems from (Dehaene et al., 2006). Participants must pick the image that does not belong.

Our model is based on three claims about human spatial cognition. Firstly, we believe people use similar processes for concrete visual comparisons and abstract analogical comparisons. We model comparison as structure-mapping (Gentner, 1983), a domain-general process of aligning common relational structure. Although structure-mapping was originally proposed to explain analogy, there is evidence it may also be used in people's visual comparisons (Lovett, Gentner, Forbus, & Sagi, 2009; Lovett, Tomai, Forbus, & Usher, 2009; Markman & Gentner, 1996). The oddity task presents a further test of structure-mapping's generality.

Secondly, we believe that when possible, people use *qualitative* or *categorical* representations of space (e.g., Biederman, 1987; Forbus, Nielsen, & Faltings, 1991; Kosslyn et al., 1989), rather than quantitative or coordinate representations, to capture how objects relate to each other in a visual scene. These spatial representations are hierarchical (Hochstein & Ahissar, 2002; Palmer, 1977), capturing both larger-scale relations between groups of objects, and smaller-scale relations between parts of an individual object. Solving the oddity task requires determining a granularity at which there is a salient qualitative difference between one image and the others.

Finally, we propose that differential performance across cultures may be linked to differences in encoding of spatial representations. Research suggests spatial representations vary based on exposure to external artifacts such as language (e.g., Haun, Rapold, Call, Janzen, & Levinson, 2006; Hermer-Vazquez, Moffet, & Mukholm, 2001; Loewenstein & Gentner, 1995). While the oddity task study demonstrated the Mundurukú have many of the same geometric concepts as Americans, there may be differences in when and how they encode those concepts. Here, we make the

simplifying assumption of a single, broad set of concepts available to both groups. We then test whether the groups vary in how easily they can use certain concepts.

This paper describes a simulation of the oddity task that operates over the stimuli used by Dehaene et al. (2006). Our model automatically constructs qualitative spatial representations and compares them via structure-mapping to identify the odd image out. The model achieves results comparable to human performance on this task, matches the error patterns found in both Americans and Mundurukú, and suggests explanations for cultural performance differences based on different spatial encoding.

2. Method

Our model utilizes two pre-existing systems: the Structure-Mapping Engine (Falkenhainer, Forbus, & Gentner, 1989) for comparison, and CogSketch (Forbus, Usher, Lovett, Lockwood, & Wetzell, 2008) for generating qualitative spatial representations.

2.1. Model components

The Structure-Mapping Engine (SME) is a computational model of comparison based on Gentner's (1983) structure-mapping theory of analogy. Given two symbolic representations, it compares them by aligning their common relational structure. It operates on *structural descriptions*, symbolic representations consisting of *entities*, *attributes*, and *relations*. Given two structural descriptions, SME returns one or more mappings between them. Each mapping consists of: (1) a set of correspondences between elements in the two descriptions; (2) a *structural evaluation*

score which estimates similarity based on the breadth and depth of aligned structure; and (3) a set of *candidate inferences* representing information that could be projected from one description to the other, based on their aligned structure. SME is useful because it can identify corresponding elements in two representations, evaluate their similarity, and identify differences between them (i.e., the candidate inferences).

CogSketch (Forbus et al., 2008) is a sketch understanding system. Given a two-dimensional sketch containing objects drawn by a user or imported from PowerPoint, CogSketch constructs a qualitative representation of the attributes of each object, as well as spatial relationships between objects. Attributes include open vs. closed shapes, or straight vs. curved, while relations include relative location and containment. See Table 1 for the terms used in this simulation. Note that some terms are *orientation-specific*, meaning they would change if an image rotated in space. These include the orientation and relative location of objects. We return to this distinction later.

On demand, CogSketch can generate representations at three levels in a spatial hierarchy. The default is the *Object* level. The lower level of *Edges* describes relations between edges within an object. The higher level of *Groups* describes relations between groups of objects. The Edges level is computed by segmenting a shape into its component edges (see Lovett, Tomai, et al., 2009) and identifying qualitative spatial relations between the edges (see Table 2 for a full list, and Lovett & Forbus, 2010 for a discussion).

The Groups level is computed by grouping objects together based on similarity and proximity—that is, objects should be similar in size and shape and equidistant from each other. Objects must also be axis-aligned (Palmer, 1980)—that is, the line connecting two objects should align with their axes of elongation or symmetry. The Group qualitative vocabulary is similar to the Object vocabulary (Table 1), and if no groups are found, a Group-level representation will be identical to an Object-level one.

There are two Object level features that can only be computed by comparing Edge level representations. Firstly, the model uses SME to identify transformations between two objects' shapes, by comparing their edge-level representations to determine if there is a rotation or reflection between them. Thus we use structure-mapping to model mental rotation (Lovett, Tomai, et al., 2009; Shepard & Metzler, 1971).

Using a similar approach (Ferguson, 1994), the model compares an object's edge-level representation to itself to identify axes of symmetry, such as those found in Fig. 1D. The model distinguishes between four types of symmetry based on the number and orientation of any axes of symmetry found (Table 1). Additionally, when no axial symmetry is detected, an object may be classified as *Non-Elongated* if it lacks an axis of elongation (Sekuler, 1996). This feature may be seen as a crude measure of symmetry, since many regular, symmetric shapes lack an axis of elongation (circles, squares, equilateral triangles, etc).

2.2. Oddity task model

The operation of the model is as follows (Fig. 2):

- (1) Using CogSketch, encode qualitative spatial representations for each image at a given level in the spatial hierarchy. The model always begins at the highest level of Groups (Hochstein & Ahissar, 2002). Thus, the dots in Fig. 1C would be grouped together and treated as a single element.
- (2) Using the Structure-Mapping Engine, compare half the images to determine what is constant across them. This happens in two steps: (A) Compare the images looking for differences. Specifically, if there are differences in any *orientation-specific* features (see Tables 1 and 2), such as relative location, then filter out all orientation-specific features. This

Table 1

Qualitative vocabulary for Objects and Groups (Terms marked with an "O" are orientation-specific).

Basic Attributes	Edge-Based Attributes ^a	Symmetry Attributes	Location Attributes
<ul style="list-style-type: none"> • 2D-Shape-Generic • 2D-Shape-Open/Closed • 2D-Shape-Forked • 2D-Shape-Oblique^O • VerticalEdge/HorizontalEdge^O • Dot-Shape 	<ul style="list-style-type: none"> • 2D-Shape-Convex • 2D-Shape-Curved/Straight/Ellipse • 2D-Shape-Axis-Aligned^O • 2D-Shape-Perpendicular 	<ul style="list-style-type: none"> • Symmetric-Shape • Perpendicular-Symmetric-Shape • Multiply-Symmetric-Shape • Fully-Symmetric-Shape • Non-Elongated-Shape 	<ul style="list-style-type: none"> • Centered-Element • OnTop-Element/OnBottom-Element^O • OnRight-Element/OnLeft-Element^O
<i>Color/Texture Attributes</i>			
<ul style="list-style-type: none"> • (ObjectsColoredFn color) • (ObjectsBorderColoredFn color) • TexturedObject 			
<i>Spatial Relations</i>	<i>Alignment Relations</i>	<i>Transformation Relations</i>	
<ul style="list-style-type: none"> • rightOf/above^O • onRightHalfOf/onLeftHalfOf^O • onTopHalfOf/OnBottomHalfOf^O • elementsIntersect • elementsOverlap • elementContains 	<ul style="list-style-type: none"> • parallelElements • perpendicularElements • collinearElements • centeredOn 	<ul style="list-style-type: none"> • reflectedShapes-XAxis • reflectedShapes-YAxis • reflectedShapes • rotatedShapes-90 • rotatedShapes-180 • rotatedShapes 	

^a Edge-based attributes are encoded when a feature (such as being straight, or having perpendicular corners, see Table 2) holds for every edge in the object.

Table 2

Qualitative vocabulary for edges (Edge Cycle Relations are higher-order relations encoded between edges in a cycle).

Attributes	Simple Edge Relations	Edge Cycle Relations
<ul style="list-style-type: none"> • PerceptualEdge • StraightEdge/CurvedEdge/EllipseEdge • Length (Tiny/Short/Medium/Long) 	<ul style="list-style-type: none"> • edgesPerpendicular • edgesParallel • edgesCollinear • edgesCurveCompatible • edgeCurveCompatibleWith • elementsConnected • elementsIntersect • elementIntersects 	<ul style="list-style-type: none"> • convex/concaveAngleBetweenEdges • cycleAdjacentAngles • adjacentAcuteToObtuseAngles/adjacentObtuseToAcuteAngles • perpendicularCorner • parallelEdgeRelation • collinearEdgeRelation
<p><i>Orientation-Specific Terms</i></p> <ul style="list-style-type: none"> • VerticalEdge • HorizontalEdge • ObliqueEdge-Upward/Downward • CurvedEdge-Right/Left/Up/Down Bumped • axisAligned 	<ul style="list-style-type: none"> • rightOf/above 	<ul style="list-style-type: none"> • leftToRightCorner/rightToLeftCorner/topToBottomCorner/bottomToTopCorner • verticallyOrientedCorner/horizontallyOrientedCorner

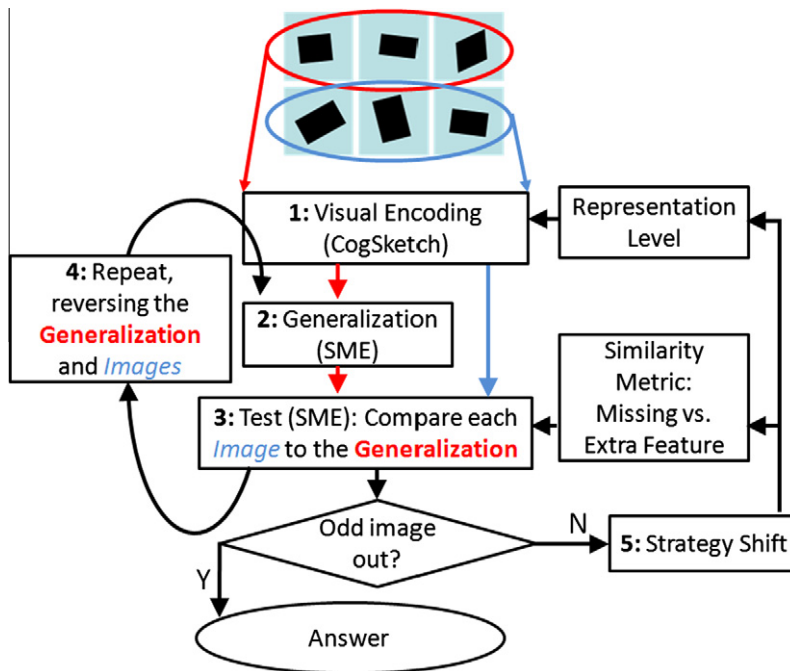


Fig. 2. Flowchart of the computational model.

simulates the decision to not worry about orientation on problems like Fig. 1B, where the relative location of the lines is irrelevant. (B) Construct an *analogical generalization* (Kuehne, Forbus, Gentner, & Quinn, 2000) over the images. An analogical generalization is built by comparing the images and abstracting out features that fail to align, leaving only the features found in all images.

- (3) Compare each remaining image to the generalization to find the least similar image. For example, in Fig. 1A, a generalization over the top row would indicate that every image contains a right angle between two edges. The lower middle image lacks this right angle, so it is less similar to the generalization.

- (4) Because the model does not know where the odd image out will be, it always performs steps (2 and 3) twice, first generalizing across the top row and then generalizing across the bottom row.
- (5) If no solution is found, make a strategy shift. Either drop the representation level from Groups down to Edges, or change the similarity metric. The default similarity metric seeks drops in similarity which indicate that an image lacks a feature of the generalization. The alternate metric seeks an image that has an extra feature beyond the generalization.

Fig. 1A requires switching from the Group level to the Edge level to notice that all the images but one contain a

right angle, whereas 1B requires switching similarity metrics to notice that only one image contains parallel elements.

2.3. Possible sources of cultural variation

Recall one of our claims is that cultural groups may vary in how they encode spatial concepts. Based on our model, we identified four factors that might be culturally dependent. Firstly, cultures might vary in their ability to join objects up into Groups or segment objects down into Edges. Secondly, cultures might have difficulties with two types of features at the Object level: Shape Transformations, and Shape Symmetry. In our model, these features require extra effort to compute, as the objects must be compared at the edge level. These four factors represent a hypothesis space for evaluating cultural differences. Of course, this is only a first pass; many other factors may play a role.

2.4. Materials

Our simulation used the stimuli from (Dehaene et al., 2006) as input. The original experimenters provided PowerPoint slides with 41 of the 45 stimuli, and we recreated the other four before importing the 45 slides into CogSketch. One problem was touched up in PowerPoint: separate parts of a shape were redrawn as a single shape. Additionally, five were altered in CogSketch: extra lines meant to draw the participant's attention to a particular feature were removed, since CogSketch cannot use this information. Aside from these minor changes, all stimuli were identical to those used in the original study. These stimuli were hand-segmented, in the sense that CogSketch treated each PowerPoint shape as a separate object, rather than automatically finding objects in the visual scene. However, CogSketch automatically identified edges and groups while generating representations.

2.5. Participants

We compared the model's performance to participants from (Dehaene et al., 2006), considering only participants for whom complete data was available. The four groups considered were 40 young children (Americans, aged 4–8 years, mean = 6.12 years, 24 female/16 male), 64 older children (Americans, aged 8–12 years, mean = 10.40 years, 33 female/31 male), 47 adults (Americans, aged 18–52, mean = 26.41 years, 28 female/19 male), and 44 Mundurukú (aged 5–83 years, mean = 37.73 years, 24 female/20 male). We collapsed across all ages for the Mundurukú because Dehaene et al. found no age-related differences in Mundurukú performance.

3. Results

Overall, the model correctly solved 39/45 problems, or 87%. This places it above the average performance for the human populations (see Table 3), demonstrating that the model was sufficient for performing the task. Table 4 shows the correlations between the model's accuracy and the accuracy of each population group across the 45 problems, as well as correlations between groups. The model's performance correlated significantly with all groups ($p < .05$), suggesting that the problems on which the model failed were also difficult for human participants. The model had the highest correlation with American adults ($r = .77$), and the lowest correlation with the Mundurukú ($r = .49$).

To explore the relationship between the model's performance and human error patterns, we ranked each of the problems according to how difficult it was for a given group, such that 1 indicated the easiest problem for that group and 45 indicated the most difficult. We then considered each group's rankings for the six problems on which the model failed (see Table 5). For all groups, the median ranking was 41 or above. For all three American groups, the minimum rank was above 30/45, meaning all six problems were difficult for those groups. For the Mundurukú,

Table 3
Accuracy of the model and average accuracy of each population group on the 45 oddity task problems.

	Accuracy
Model	0.87
Young children	0.55
Older children	0.75
Adults	0.83
Mundurukú	0.67

Table 5
Rankings of the six problems the model fails to solve (1 = easiest for a group, 45 = hardest).

	Median rank	Min rank	Max rank
Young children	41.5	34	45
Older children	42	33	45
Adults	42.5	37	45
Mundurukú	41	12	44

Table 4
Correlations in accuracy on each of the 45 problems (Pearson's r).

	Model	Young children	Older children	Adults	Mundurukú
Model	1	.58	.66	.77	.49
Young children	.58	1	.91	.84	.83
Older children	.66	.91	1	.94	.75
Adults	.77	.84	.94	1	.72
Mundurukú	.49	.83	.75	.72	1

Note: All correlations are significant ($p < .05$).




Fig. 3. One of six problems the model fails to solve. Average human performance: 68% (American adults), 86% (Mundurukú).

one problem ranked much lower (12/45). This problem (Fig. 3) is unique in that it was difficult for both the model and the Americans, but relatively easy for the Mundurukú.

What can the model tell us about the representations each group uses? As described above, we identified four factors that might contribute to difficulty: Groups, Edges, Shape Transformation, and Shape Symmetry. Restricting our analysis to the 39 problems the model solved, we scored each problem for whether it required each factor. This was done via ablation, selectively removing the model's ability to perform an operation and identifying the problems which it was no longer able to solve. When Groups were ablated, the model built representations at the Object level instead. When Symmetry was ablated, the model never computed axial symmetry, but it still identified Non-Elongated shapes.

We built linear models for each group, using the four factors as independent variables (Table 6). All four linear models have R^2 values of around .5, indicating they account for about half the variance in human performance. Importantly, there are differences in which factors contribute significantly to each group's model. In particular, the Edges

Table 6

Linear models for each group's accuracy (shaded cells are factors making statistically significant contributions, $p < .05$).

Model	R^2 (adj R^2)	Intersect	Groups	Edges	Shape Transformation	Shape Symmetry
Young children	.57 (.52)	0.71	−0.07	−0.21	−0.34	−0.41
Older children	.61 (.57)	0.91	−0.04	−0.20	−0.28	−0.39
Adults	.47 (.40)	0.94	−0.01	−0.08*	−0.18	−0.17
Mundurukú	.57 (.52)	0.81	−0.31	−0.07	−0.42	−0.14

Notes. All linear models have a significant correlation with their respective groups ($p < .05$).

* $p = .058$.

Table 7

Linear models for each group's reaction time (shaded cells are factors making statistically significant contributions, $p < .05$).

Model	R^2 (adj R^2)	Intersect	Groups	Edges	Shape Transformation	Symmetry
Young children	.24 (.15)	8494	1700	1179	5055	512
Older children	.69 (.65)	5866	2970	4959	4938	5071
Adults	.63 (.58)	5445	2888	3592	5374	4103

Note. All linear models have a significant correlation with their respective groups ($p < .05$) except young children ($p = .052$).

factor contributes significantly to two of the three American models (young children and older children), and is marginally significant ($p = .058$) in the third model (adults). In contrast, the Groups factor contributes significantly in the Mundurukú model. Thus, Americans appeared to have greater difficulty on those problems which required focusing on the individual edges of objects, rather than the objects as a whole (e.g., Fig. 1A), whereas the Mundurukú had greater difficulty on problems that required looking at overall configurations of objects (e.g., Fig. 1C).

The other difference may be related: whereas all groups had greater difficulty on problems involving Shape Transformation, only the Americans had difficulty on problems involving Shape Symmetry. These problems (e.g., Fig. 1D), appear to require a close consideration of the edges.

Finally, we built linear models for the average reaction times of each group. This analysis was restricted to the three American groups, as precise timing data was unavailable for the Mundurukú. It was also restricted to 37 of the 39 problems, since the first two were used in (Dehaene et al., 2006) as training problems. We considered only correct responses.

The results (Table 7) show a clear split between the young children and the other groups. The model for the young children, while a marginally significant predictor ($p = .052$), accounts for substantially less of the variance than the models for the older children and adults. This may be because the younger children were less systematic or more easily distracted while performing the task.

The timing models for the older children and adults both account for over half the variance. Every factor is a significant contributor. This suggests that *objects* are the most basic level for American representations of two-dimensional space. Participants required more time on problems that involved working with a different level of representation (Groups or Edges). They also required more time on problems involving Shape Transformation or Shape Symmetry, confirming our model's prediction that these features require additional operations.

4. Discussion

Our simulation shows that structure-mapping combined with a hierarchical, qualitative representation of space is sufficient for achieving human-level performance on the oddity task, given hand-segmented images. The correlation between human and model error patterns further supports the generality of our model, which has been used to simulate other visual comparison tasks (Lovett, Gentner, et al., 2009; Lovett, Tomai, et al., 2009).

While our model demonstrates that some spatial reasoning can be accomplished with qualitative representations, people also utilize quantitative information. It is possible that the Mundurukú and young American children, whose performance correlated least with the model, relied more on quantitative information, perhaps because of less familiarity with the words for various geometric concepts. Exploring the interaction between qualitative and quantitative representations is one important avenue for future research.

Our simulations suggest cultural differences in spatial cognition may manifest as differences in representational focus (Nisbett & Masuda, 2003; Nisbett & Miyamoto, 2005). Americans are biased to focus first on individual objects, leading to greater difficulty reasoning about the edges within an object. The Mundurukú, in contrast, do well with edges or objects, but have more difficulty considering entire groups of objects. Indeed, the impressive performance of the Mundurukú on one problem that was difficult for both the Americans and our model (Fig. 3) suggests they may divide up space to consider even quadrants within a single edge.

One possible reason for this cultural difference is that formal schooling in geometry, and the many shape names in the English language, might focus Americans on shapes when they consider visual scenes. The Mundurukú, lacking such exposure to formal geometric concepts, may be more inclined to focus on the parts that make up a shape. Understanding these cultural differences and how they arise, using a combination of psychological and computational studies, will help shed light on what is universal versus not in human spatial cognition.

Acknowledgements

This work was supported by NSF SLC Grant SBE-0541957, the Spatial Intelligence and Learning Center (SILC).

We thank Elizabeth Spelke and Veronique Izard for generously providing stimuli and data from their original study, and Dedre Gentner for many productive discussions and feedback.

Preliminary results for this work were presented at the 30th Annual Meeting of the Cognitive Science Society, and at Spatial Cognition 2008.

References

- Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, *94*, 115–147.
- Dehaene, S., Izard, V., Pica, P., & Spelke, E. (2006). Core knowledge of geometry in an Amazonian indigene group. *Science*, *311*, 381–384.
- Falkenhainer, B., Forbus, K., & Gentner, D. (1989). The structure mapping engine: Algorithm and examples. *Artificial Intelligence*, *41*, 1–63.
- Ferguson, R.W. (1994). MAGI: Analogy-based encoding using regularity and symmetry. In *Proceedings of the 16th Annual Conference of the Cognitive Science Society*.
- Forbus, K., Usher, J., Lovett, A., Lockwood, K., & Wetzel, J. (2008). CogSketch: Open-domain sketch understanding for Cognitive Science research and for education. In *Proceedings of the Eurographics Workshop on Sketch-Based Interfaces and Modeling 2008*.
- Forbus, K., Nielsen, P., & Faltings, B. (1991). Qualitative spatial reasoning: The CLOCK project. *Artificial Intelligence*, *51*, 417–471.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, *7*, 155–170.
- Haun, D. B. M., Rapold, C. J., Call, J., Janzen, G., & Levinson, S. C. (2006). Cognitive cladistics and cultural override in Hominid spatial cognition. *PNAS*, *103*, 17568–17573.
- Hermer-Vazquez, L., Moffet, A., & Mukholm, P. (2001). Language, space, and the development of cognitive flexibility in humans: The case of two spatial memory tasks. *Cognition*, *79*, 263–299.
- Hochstein, S., & Ahissar, M. (2002). View from the top: Hierarchies and reverse hierarchies in the visual system. *Neuron*, *36*, 791–804.
- Kosslyn, S. M., Koenig, O., Berrett, A., Cave, C. B., Tang, J., & Gabrieli, J. D. E. (1989). Evidence for two types of spatial representations: Hemispheric specialization for categorical and coordinate relations. *Journal of Experimental Psychology: Human Perception and Performance*, *15*, 723–735.
- Kuehne, S., Forbus, K., Gentner, D., & Quinn, B. (2000). SQL: Category learning as progressive abstraction using structure mapping. In *Proceedings of the 22nd Annual Conference of the Cognitive Science Society*.
- Loewenstein, J., & Gentner, D. (1995). Relational language and the development of relational mapping. *Cognitive Psychology*, *50*, 315–353.
- Lovett, A., & Forbus, K. (2010). Shape is like space: Modeling shape representation as a set of qualitative spatial relations. In *Proceedings of the AAAI Spring Symposium*. Palo Alto, CA.
- Lovett, A., Gentner, D., Forbus, K., & Sagi, E. (2009). Using analogical mapping to simulate time-course phenomena in perceptual similarity. *Cognitive Systems Research*, *10*(3), 216–228 (Special Issue on Analogies – Integrating Cognitive Abilities).
- Lovett, A., Tomai, E., Forbus, K., & Usher, J. (2009). Solving geometric analogy problems through two-stage analogical mapping. *Cognitive Science*, *33*, 1192–1231.
- Markman, A. B., & Gentner, D. (1996). Commonalities and differences in similarity comparisons. *Memory & Cognition*, *24*, 235–249.
- Newcombe, N. S., & Uttal, D. H. (2006). Whorf versus Socrates, round 10. *Trends in Cognitive Science*, *10*, 394–396.
- Nisbett, R. E., & Masuda, T. (2003). Culture and point of view. *Proceedings of the National Academy of Sciences*, *100*(19), 11163–11170.
- Nisbett, R. E., & Miyamoto, Y. (2005). The influence of culture: Holistic versus analytic perception. *Trends in Cognitive Science*, *9*(10), 467–473.
- Palmer, S. (1977). Hierarchical structure in perceptual representation. *Cognitive Psychology*, *9*, 441–474.
- Palmer, S. E. (1980). What makes triangles point: Local and global effects in configurations of ambiguous triangles. *Cognitive Psychology*, *12*, 285–305.
- Sekuler, A. B. (1996). Axis of elongation can determine reference frames for object perception. *Canadian Journal of Experimental Psychology*, *50*, 270–278.
- Shepard, R. N., & Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science*, *171*, 701–703.
- Spelke, E., & Kinzler, K. D. (2007). Core knowledge. *Developmental Science*, *10*, 89–96.