

Modeling Qualitative Differences in Symmetry Judgments

Ronald W. Ferguson

Institute for the Learning Sciences
Northwestern University
1890 Maple Avenue
Evanston, IL 60201
ferguson@ils.nwu.edu

Alexander Aminoff

Harvard University
Cambridge, MA 02139
aminoff@wjh.harvard.edu

Dedre Gentner

Department of Psychology
Northwestern University
2029 Sheridan Road
Evanston, IL 60208
gentner@ils.nwu.edu

Abstract

Symmetry perception is an important cognitive process across many areas of cognition. This research explores symmetry as a special case of similarity—self-similarity—and proposes that qualitative relationships play a role in the early perception of symmetry. To support this claim, we present evidence from two psychological studies where subjects performed symmetry judgments for randomly constructed polygons. Subjects were faster and/or more accurate at detecting asymmetry for stimuli with *qualitative* asymmetries than for stimuli with equivalent *quantitative* asymmetries. Aspects of this effect are replicated using the MAGI computational model, which detects symmetry using a method of structural alignment. The results of this study suggest that qualitative information influences early perception of symmetry, and provides further support for the MAGI model.

Introduction

Symmetry serves as an organizing principle in several different areas of perception and cognition, including the Gestalt notion of figural goodness (Garner, 1974; Palmer, 1991), the visual reconstruction of 3D shape (McBeath, Schiano, & Tversky, 1994), and the computation of object-centered reference frames (Palmer, 1989). The breadth of these phenomena suggests that symmetry perception is an important and fundamental cognitive process.

Our research makes two distinctive claims about the perception of symmetry. First, we propose that early symmetry detection is a process of self-comparison that can be modeled as an alignment of maximally similar subsets of perceived structural relations in a figure. This assertion is supported by recent evidence suggesting that perceptual similarity can be modeled using the same kinds of structure-mapping processes that are used to model analogy (Falkenhainer, Forbus & Gentner, 1989; Goldstone, Medin & Gentner, 1991; Markman & Gentner, 1993; Medin, Goldstone & Gentner, 1993). With this in mind, we have implemented a computational model of symmetry detection called MAGI (Ferguson, 1994), which uses structure-mapping to detect symmetry in a way that has many of the characteristics of analogy, including robustness over incomplete or inexact descriptions. MAGI also has the ability to detect multiple axes of symmetry and repetition,

and to spontaneously make inferences from one half of a figure to another.

The second claim of this research is that early symmetry processes act over representations that include *qualitative relations*. Qualitative relations have been theorized to provide a foundation for our initial partitioning of the physical world (Forbus, 1984). Qualitative spatial relations have been shown to be important in human processing of spatial scenes (Glenberg & McDaniel, 1992; Palmer, 1989, 1991). Qualitative differences are important in visual similarity comparisons (Goldmeier, 1936/1972).

In this paper, we summarize recent work (Aminoff, Ferguson & Gentner, in preparation) indicating that humans utilize qualitative relationships in symmetry judgments. We then describe a replication of these psychological results using the MAGI computer model. Finally, the implications of this proposal are discussed.

Detecting the Effect of Qualitative Differences in Symmetry Judgments

If we assume that symmetry involves a structural alignment of perceived qualitative relations, we can test this hypothesis by observing how the misalignment of qualitative relations

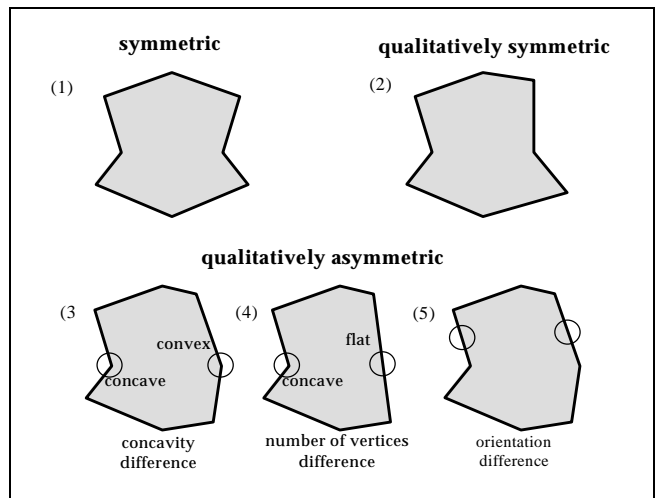


Figure 1: Symmetric, quantitatively asymmetric, and qualitatively asymmetric polygons

affects symmetry judgment. Qualitative differences—mismatched or misaligned relationships between sides of a figure—should therefore affect symmetry judgment more than would be predicted by the degree of quantitative difference between the sides of a figure.

Previous research gained insight into how humans process symmetry by examining the conditions under which symmetry is more easily perceived. For example, a large body of research shows that humans detect vertical symmetry more easily (i.e., more quickly or more accurately) than either horizontal or oblique symmetry (Corballis & Roldan, 1975; Palmer & Hemenway, 1978; Rock, 1983). Thus symmetry detection is not orientation-invariant, but depends on a frame of reference.

We are approaching this issue by asking when asymmetry is easy to perceive. Specifically, we hypothesize that figures containing qualitative differences should be easier to judge as asymmetric than figures without such differences, independent of any quantitative metric that we might use to measure asymmetry in a figure.

Although it is difficult to enumerate the full set of perceived qualitative differences, it is straightforward to choose viable candidates, as demonstrated by the polygons in Figure 1. This figure also illustrates a terminological distinction between qualitative and quantitative symmetry. The first polygon is exactly symmetric, with the left and right sides sharing equal dimensions. Thus it is *quantitatively symmetric*. In the second polygon, the left side is structurally similar to the right, but the lengths of corresponding lines differ. Because the sides differ quantitatively but not qualitatively, we call such objects both *quantitatively asymmetric* and *quantitatively symmetric*.

Finally, the bottom three polygons in Figure 1 are *qualitatively asymmetric*, containing three different types of qualitative difference. The first of these polygons contains a *concavity difference*. The polygon's left and right sides align somewhat, but there is a clear difference between the circled vertices—one is concave while the other is convex. The next polygon has a *number-of-vertices difference*, where a concavity on one side is missing on the other. The last of the three polygons has an *orientation difference*, where one line segment leans into the polygon and the other leans away. (Note that while orientation differences often co-occur with concavity differences, they are not equivalent.)

All of these qualitative differences could cause a misalignment or mismatch between opposing sides of the shapes. Just as vertical symmetry is easier to detect than horizontal symmetry, figures containing any of these qualitative differences should be easier to judge asymmetric than figures without them. We now summarize recent experiments that test this assumption.

Psychological Evidence

We tested these predictions in two experiments (Aminoff, Ferguson, & Gentner, in preparation). We presented polygonal stimuli to human subjects, and asked them to quickly judge whether each figure was symmetric. The crucial independent variable was the type of asymmetry—both quantitatively and qualitatively asymmetric objects were included in the stimulus set.

Of course, care must be taken to ensure that if the greater perceived asymmetry is found for qualitative differences that it is not the result of a correlated increase in quantitative difference. In order to control for quantitative differences, the stimuli were selected so that the most important quantitative metric, the sum of squared differences of radii, was equal across conditions. Because some variations in quantitative parameters was unavoidable, we also computed the correlation of our results with 26 other quantitative (and qualitative) measures of asymmetry. Over the two experiments, three types of qualitative difference were used: concavity differences, orientation differences, and number-of-vertices differences. The key measure is subjects' speed and accuracy at detecting asymmetry. (See Aminoff, Ferguson, & Gentner (in preparation) for more details on these experiments.)

Experiment 1

In experiment 1, sixteen subjects were sequentially shown forty 16-sided polygons from a stimulus set of eighty. After a very brief masked presentation (50 ms.) subjects indicated if the polygon was symmetric by pressing one of two computer keys. The stimulus set was evenly divided between symmetric and asymmetric stimuli, with the latter evenly divided between qualitatively and quantitatively asymmetric polygons. Qualitatively asymmetric polygons were further subdivided by qualitative difference type. The forms of qualitative difference used were concavity differences and number-of-vertices differences.

Although experiment 1 showed no significant effect of qualitative difference on reaction time, it did show a significant effect for accuracy (Figure 2). Subjects were

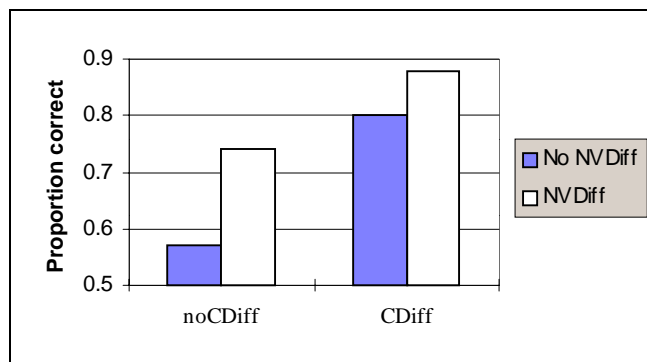


Figure 2: Human accuracy results from experiment 1 (asymmetric figures only).

much more accurate for polygons that contained either concavity or number-of-vertices differences. This effect was roughly additive: subjects were most accurate at stimuli that had both concavity and number-of-vertices differences. Subjects were also more accurate at correctly classifying polygons with number-of-vertices differences than those with concavity differences, and at accurately classifying symmetric polygons than asymmetric polygons.

Experiment 2

Experiment 2 used the same method as Experiment 1 with a slightly easier perceptual task. Simpler 12-sided polygons were displayed at brighter contrast levels in one of two conditions—a fast condition in which the polygon was displayed for 50 ms, and a slow condition in which the stimulus remained on the screen until the subject pressed a key. The number of stimuli was doubled to 160. Of 89 subjects in this study, 54 were assigned to the fast condition and 35 to the slow condition. The qualitative differences used were concavity differences and orientation differences.

Again, subjects were significantly better at stimuli containing qualitative differences. Subjects' accuracy for asymmetric stimuli was uniformly high in the slow condition, but in the fast condition showed a significant effect for the presence of either concavity or orientation differences (Figure 3). Subjects were not significantly more accurate for figures with concavity differences over those with orientation differences, or for symmetric over asymmetric figures.

The reaction time data from experiment 2 showed an effect for figures with concavity differences, but no significant effect for figures with orientations differences. The effect was most significant in the slow condition (Figure 4), but was also marginally significant in the fast condition.

Results

Across two experiments, as predicted by the MAGI model, significant effects were found for qualitative differences in symmetry judgment. For asymmetric objects, subjects were faster and/or more accurate when the asymmetry was manifested in a qualitative difference between the halves of the figure. Subjects responded slower and/or less accurately for asymmetric objects without such qualitative differences.

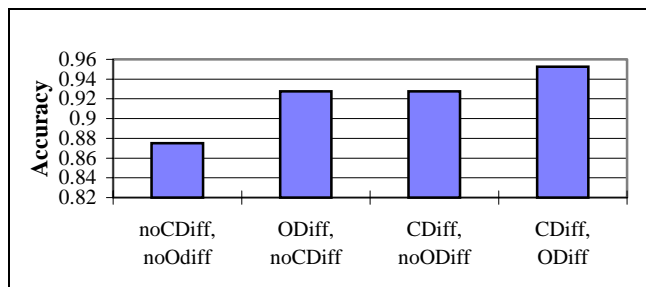


Figure 3: Experiment 2, fast condition. Human accuracy for asymmetric figures

(We have no principled reason to predict whether the early advantage of qualitative asymmetry over quantitative asymmetry should show up in greater accuracy or in faster processing.) Along with supporting the effect of qualitative differences in symmetry detection, these experiments also suggest that some kinds of differences are more important to symmetry detection than others.

In experiment 1, number-of-vertices differences had a greater effect on accuracy than concavity differences, while experiment 2 in turn showed a greater or equal effect for concavity differences over orientation differences in reaction time measurements. Symmetric figures were classified more accurately and/or more quickly than asymmetric figures in both experiments.

Testing the Results Using the MAGI Model

If MAGI is an accurate model of symmetry detection, it should be able to replicate the results of these experiments, not only in terms of higher accuracy for asymmetrical figures with qualitative differences, but also in terms of which qualitative differences are most important.

The MAGI Model of Symmetry Detection

MAGI (Ferguson, 1994) models symmetry detection as a relational self-similarity mapping that aligns a qualitative representation of a figure with itself. MAGI has been implemented as a computational model using an extension of the Incremental Structure Mapping Engine (I-SME; Falkenhainer, Forbus & Gentner, 1989; Forbus, Ferguson & Gentner, 1994). In essence, MAGI computes a structural alignment between two sides of a figure.

In constructing a mapping, MAGI follows the constraints of I-SME's analogical mapping. Matches must be one-to-one, and arguments of matched expressions must match as well. Only expressions with identical predicates (or non-identical functions that are arguments of other matched expressions) can match. A scoring mechanism encourages

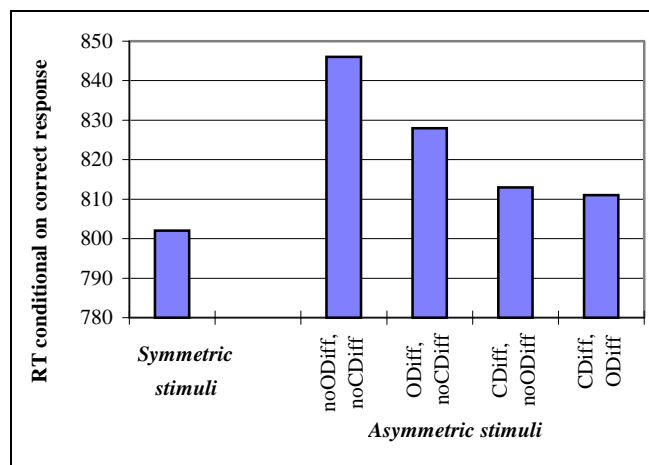


Figure 4: Experiment 2, slow condition. Human reaction time for symmetric and asymmetric stimuli.

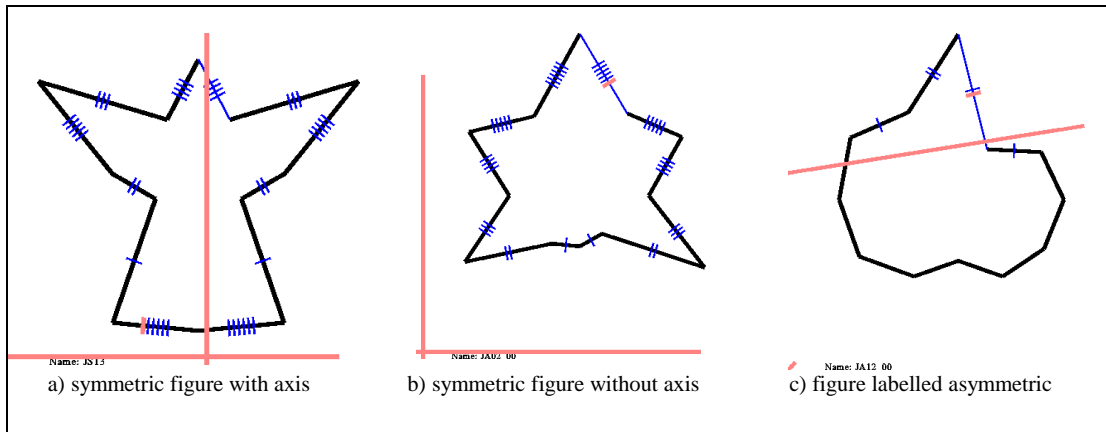


Figure 5: Representative output MAGI for experiment 2. Mapped parts of figures are indicated by an equal number of hash marks. Gray lines indicate the axis and reference frame suggested by MAGI.

relationally deep interconnected systems of matched expressions. In addition, because MAGI maps a description to itself, it blocks expressions from matching to themselves, allowing it only when the self-match is an argument of two different matching parents. MAGI then analyzes the mapping to determine if the mapping merely found a repeated pattern, or found a core of symmetrical matches that makes the whole mapping symmetric.

By mapping qualitative relationships, MAGI can greatly constrain the quantitative calculations and comparisons that it performs. If the mapping is symmetrical, MAGI can compute an axis by using a Hough transform over all the bisecting lines between mapped lines in the figure. Since MAGI's axis-detection only considers a potential axis between symmetrically mapped lines, it is both extremely efficient and robust in the presence of distracters.

To test whether MAGI fits the human results, we ran MAGI on the same stimulus sets used in experiments 1 and 2. The version of MAGI used the same constraints described in Ferguson (1994), but also contained an extension allowing mapping of commutative relationships (such as corner relations and line groups).

Representations Used

Any relational model of perception must make assumptions about types of visual relations that are perceived (Pinker, 1984). The representations given to MAGI are generated using a geometric representation system called GeoRep. GeoRep is not strictly a model of the perception process, but is designed to produce plausible visual representations given simple vector drawings. From the original stimulus data files used in the two experiments, which give each polygon as a set of line segments, GeoRep generates the following polygonal relations: corners, corner concavity or convexity, the presence of perpendicular, obtuse, or acute corners, the presence of protrusions or indentations in the figure (defined, respectively, as adjacent sets of convex or concave corners), and the relative position of protrusions relative to the gravitational reference frame. Relationships are computed only between proximate objects, using a simple proximity metric based on object size and distance. Since some relationships (such as corners or protrusions) can also be the arguments of other geometric relations, the representations tend to be hierarchic.

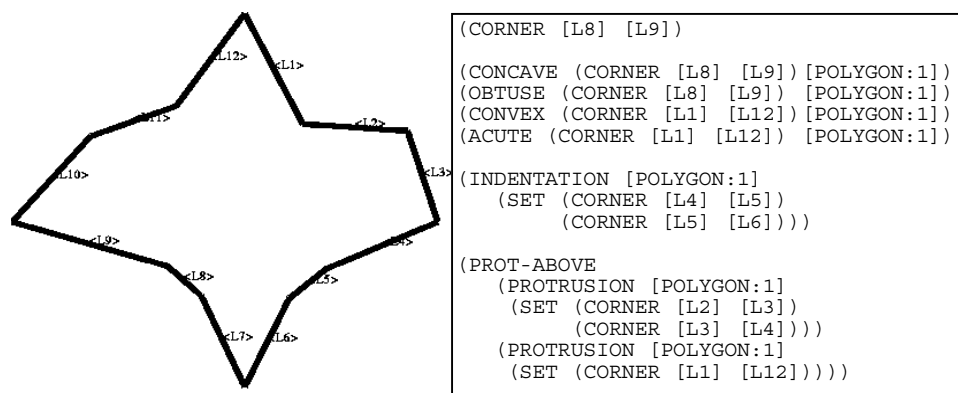


Figure 6: Representative relationships from description of experiment 2 stimulus

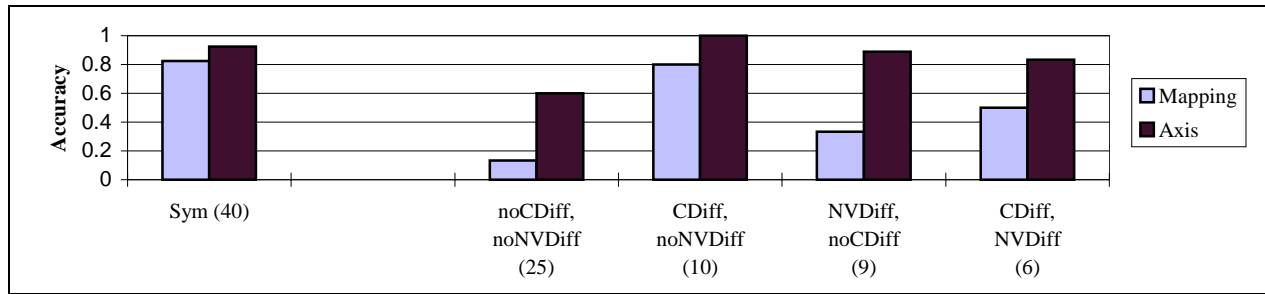


Figure 7: Accuracy rates of MAGI on experiment 1 stimuli, by qualitative difference and criterion type. The number of stimuli in each condition is given in parentheses.

Figure 6 contains a prototypical subset of an actual representation. Note that although the ABOVE relationships for protrusions in the figure imply a gravitational frame of reference, MAGI does not assume a vertical axis in the figure, although it does encourage vertical over horizontal symmetry. However, this is a preference it shares with humans, and so it carries some cognitive validity. However, MAGI can still find horizontal symmetry when a figure has a good intrinsic horizontal axis.

Qualitative differences clearly have an effect upon mapping in representation produced by GeoRep, and so affect the mapping done by MAGI. Number-of-vertices and orientation differences can cause changes in the alignment of the two sides, and affect the perceived protrusions in the figure. Since the representation directly represents concavity, a concavity difference removes an incentive to match corresponding corners from the two sides.

Because MAGI uses a Hough transform to compute an axis, it has a possibly unique characteristic among symmetry-detection algorithms, which is that it can find an object qualitatively symmetric, but then fail to find a straight axis. For this reason, MAGI is equipped with two criteria for judging if a figure is symmetric. A presented stimulus passes the *mapping* criterion for symmetry if more than half the lines in the figure can be mapped symmetrically and if only a small subset of the mapping (less than 20%, in terms of its structural score) is mapped non-symmetrically. A figure can pass the *axis-detection* criterion for symmetry if the corresponding lines actually produce a vertical axis computed using a Hough transform.

Figure 5 shows the output by MAGI for three of the figures from the study. Figure 5(a) maps symmetrically and

produces an axis, thus passing both the mapping and axis-detection criteria. In contrast, Figure 5(b) passes only the mapping criteria, because MAGI finds a symmetric alignment of the parts of the figure (as indicated by the hash marks), but cannot find a straight axis based on that alignment (although it does find a reference orientation, as indicated by the gray lines to the bottom and left in the figure). Figure 5(c) passes neither criterion, and is judged asymmetric by MAGI.

Results

The results for running MAGI on the stimuli sets from experiments 1 and 2 are shown in Figure 7 and Figure 8. For the 80 figures used in experiment 1, the results were suggestive, but not conclusive. MAGI performed extremely well on the symmetric stimuli, classifying over 90% of them correctly. Also, as expected, it performed significantly better on asymmetric figures with qualitative differences (judging them asymmetric in 56% of all instances based on the mapping criterion, and 92% of all instances based on the axis criterion) than on asymmetric figures with only quantitative differences (judging them asymmetric via mapping 13% of the time, and asymmetric via the axis criterion 60% of the time). However, while human subjects clearly were able to use some combinations of qualitative differences better than others, MAGI was unable to replicate that result from experiment 1. We suspect that the lack of strong congruence with the human results might be due to the variance resulting from having a very small number of stimuli. Experiment 2, with a larger number of stimuli, remedies this problem.

On experiment 2's stimuli (Figure 8), MAGI performed

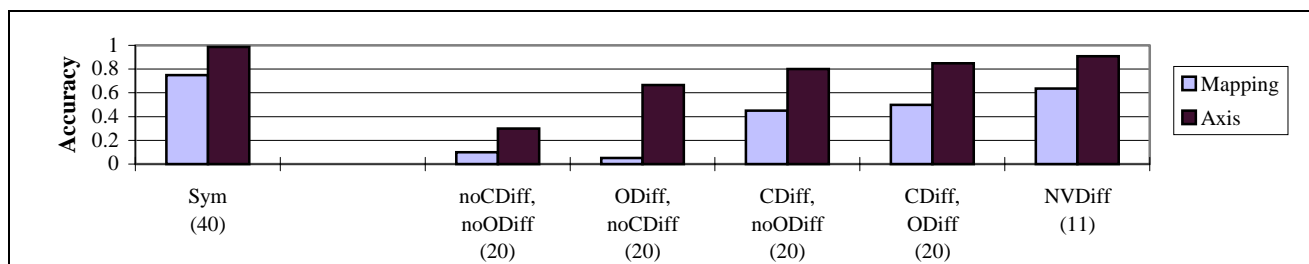


Figure 8: MAGI's accuracy on stimuli sets from experiment 2 for symmetrical figures, and figures with concavity, orientation, and number-of-vertices differences. Number of stimuli in each set is given in parentheses.

in a way that was more congruent with the human data, producing the same ordering among figures' concavity differences and orientation differences that was found in the two psychological experiments. Symmetric figures were classified more accurately than asymmetric figures, and concavity differences affected accuracy more than orientation differences did. MAGI was also tested on a subset of experiment 2's stimuli that had marginal number-of-vertices differences, showing that such differences had a more significant effect than either concavity differences or orientation differences, which matches the results from experiment 1.

These results suggest that the human data can be accounted for within the MAGI model. We see these results as a promising lead for future research. For example, the relatively large effect for number-of-vertices differences in MAGI's replication of experiment 2 (Figure 8), leads us to conjecture that figures with number-of-vertices differences may be easier to detect as asymmetric than figures with either concavity or orientation differences.

Conclusion

Qualitative relations are central to human symmetry perception. Just as vertical symmetry is easier to detect, asymmetric figures with qualitative differences are more easily judged asymmetric than figures with quantitative differences. Further, some types of qualitative differences are easier to detect than others.

This preference for qualitative differences in symmetry judgments implies a model that utilizes qualitative perceptual relationships. MAGI currently can model many aspects of this preference, including the distinctions between different kinds of qualitative relations. More research is needed to understand the limitations of the model (which does not yet include perceptual grouping), and to validate GeoRep's assumptions about perceptual representation. MAGI's ability to run on moderately complex line drawings identical to those given human subjects suggest that it is not only a viable psychological model, but also a useful tool for conducting further research into symmetry's fundamental role in cognition.

Acknowledgments

We would like to thank Doug Medin, Ken Forbus, and the analogy and similarity group at Northwestern for many contributions to this research. This research was supported by the Office of Naval Research and a Northwestern University Fellowship to the second author.

References

Aminoff, A., Ferguson, R.W., & Gentner, D. (in preparation). Early detection of qualitative symmetry. Manuscript being prepared for submission.

- Corballis, M.C., & Roldan, C.E. (1975). Detection of symmetry as a function of angular orientation. *Journal of Experimental Psychology: Human Perception and Performance*, 1, 221-230.
- Falkenhainer, B., Forbus, K.D., & Gentner, D. (1989). The Structure-Mapping Engine: Algorithms and examples. *Artificial Intelligence*, 41(1), 1-63.
- Ferguson, R.W. (1994). MAGI: Analogical encoding using symmetry and regularity. In *Proceedings of the Cognitive Science Society*. Atlanta, GA.
- Forbus, K.D. (1984). Qualitative Process Theory. *Artificial Intelligence*, 24, 85-168.
- Forbus, K.D., Ferguson, R.W., & Gentner, D. (1994). Incremental Structure Mapping. In *Proceedings of the Cognitive Science Society*. Atlanta, GA.
- Garner, W.R. (1974). *The processing of information and structure*. Potomac, MD: Erlbaum.
- Glenberg, A.M., & McDaniel, M.A. (1992). Mental models, pictures and text: Integration of spatial and verbal information. *Memory and Cognition*, 20(5), 458-460.
- Goldmeier, E. (1936/1972). Similarity in visually perceived forms. *Psychological Issues*, 8(1), 14-133.
- Goldstone, R.L., Medin, D.L., & Gentner, D. (1991). Relational similarity and the nonindependence of features in similarity judgments. *Cognitive Psychology*, 23, 222-264.
- McBeath, M., Schiano, D., & Tversky, B. (1994). Three-dimensional bilateral symmetry assumed in judgments of figural identity and orientation. Paper presented at *The 35th Annual Meeting of the Psychonomic Society*, St. Louis, MO, Nov. 11-13.
- Markman, A.B., & Gentner, D. (1993). Structural alignment during similarity comparisons. *Cognitive Psychology*, 23, 431-467.
- Medin, D.L., Goldstone, R.L., & Gentner, D. (1993). Respects for similarity. *Psychological Review*, 100(2), 254-278.
- Palmer, S.E. (1989). Reference frames in the perception of shape and orientation. In B. E. Shepp & S. Ballesteros (Eds.), *Object perception: Structure and process* (pp. 121-163). Hillsdale: Lawrence Erlbaum Associates.
- Palmer, S.E. (1991). Goodness, gestalt, groups, and Garner: local symmetry subgroups as a theory of figural goodness. In Lockhead & Pomerantz (Eds), *The Perception of Structure*.
- Palmer, S.E., & Hemenway, K. (1978). Orientation and symmetry: Effects of multiple, rotational, and near symmetries. *Journal of Experimental Psychology: Human Perception and Performance*, 4, 691-702.
- Pinker, S. (1984). Visual cognition: An introduction. In Pinker (Ed), *Visual Cognition*. Cambridge, MA: The MIT Press.
- Rock, I. (1990). *The Logic of Perception*. Cambridge, MA: The MIT Press.