Spotting Differences: How Qualitative Asymmetries Influence Visual Search

Rudolph L. Mappus IV (cmappus@cc.gatech.edu)
Ronald W. Ferguson (rwf@cc.gatech.edu)
Kenneth Czechowski (kentcz@cc.gatech.edu)
College of Computing, Georgia Institute of Technology
Atlanta, GA 30332 USA

Paul M. Corballis (paul.corballis@psych.gatech.edu)
Department of Psychology, Georgia Institute of Technology
Atlanta, GA 30332

Abstract

While our current understanding of symmetry perception is based on the perception of exact symmetry, there is increasing evidence that humans are sensitive to *qualitative* symmetry, which is based on a figure's pattern of similar alignable features rather than its geometric invariance about an axis. Previous research on alignment-based models of symmetry perception found evidence that *qualitative differences* (which break the pattern of alignment in otherwise symmetric figures) disproportionately improve the overall speed and accuracy of symmetry judgments. In this experiment, we examine whether qualitative differences affect the earliest stage of symmetry detection by examining their effect on visual search. There are two central results. First, qualitative differences reduce fixations in visual search. Participants spend less time and fewer fixations on qualitative differences than other differences. This suggests an early role for alignment in symmetry detection. Second, participants are significantly more accurate at judging symmetry of figures with qualitative differences than other differences. This result replicates Ferguson, Aminoff & Gentner (1996) while generalizing that result to stimuli with different fill characteristics displayed both foveally and parafoveally.

Introduction

Symmetry is a basic quality of many objects in the visual environment, playing a role in perceptual organization and figure reconstruction (Wagemans, 1995). The form of symmetry we perceive is usually understood to be exact or *quantitative* symmetry, where (for mirror symmetric figures) quantities such as angle and length are identical on both sides of an axis. Understanding symmetry as exact symmetry has lead to simple but useful models of symmetry detection based on the transformational invariance of a figure. Yet as useful as these models are, they fall short when applied to approximate or *qualitative* symmetry, which is problematic given that many real-world objects (such as human figures) display approximate symmetry.

The MAGI model of regularity detection (Ferguson, 1994, 2001) accounts for qualitative symmetry detection by modeling it as a mapping process that aligns similar qualitative relations and features (such as line intersections and boundary concavities) using a structure mapping process like that used to model similarity and analogical comparison (Gentner, 1983). While MAGI handles exact symmetry like transformational invariance, MAGI also readily detects qualitative symmetry, finding the axis and corresponding parts of near-symmetric figures in a way that appears to approximate human performance.

MAGI's performance on qualitative symmetry leads to a testable psychological prediction: that there are two different classes of asymmetry (Figure 1) caused by two difference types. Qualitative deviations from symmetry, which change the set of qualitative features, may block MAGI's alignment process, allowing quick classification of the figure as asymmetric. In contrast, quantitative deviations from symmetry, which break exact symmetry but preserve alignable qualitative features, may initially fool the alignment process, requiring additional scrutiny to detect the asymmetry. Thus, humans should judge figures with qualitative differences faster or more accurately than figures with quantitative differences.

We can make this prediction more concrete by considering polygons as our stimuli. If we consider the vertices of a non-uniform polygon, each vertex (feature) has a concavity characteristic (being concave, or convex). Corresponding features match if they match in their qualitative concavity and quantitative value. A polygon contains a *quantitative* difference when two corresponding

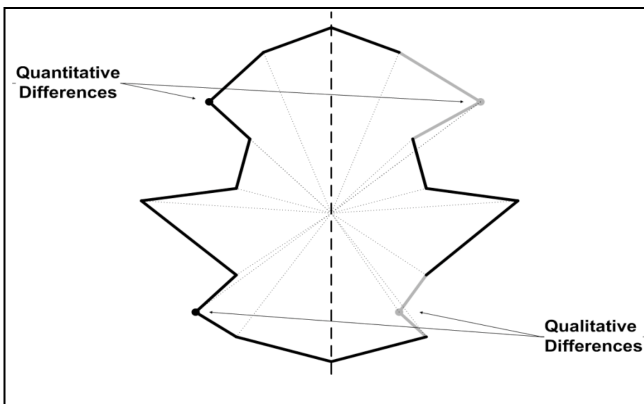

Figure 1: Polygon demonstrating qualitative and quantitative differences.

features have the same qualitative value but differ in their exact value (e.g., both are concave, but one is more concave than the other). A polygon contains a *qualitative* difference if a pair of corresponding features differ in their qualitative value (e.g. one is concave and one is convex).

Exactly this effect was shown for qualitative and quantitative differences in two experiments by Ferguson, Aminoff, & Gentner (1996). In these experiments, participants judged the symmetry of random 12- and 16-gons displayed for 50 msecs. The results showed that qualitative differences in a stimulus improved participant accuracy and response time. In both experiments, human participants were faster or more accurate at judging asymmetric figures with qualitative differences than with quantitative differences. This result supports use of an alignment process in human symmetry detection.

But if humans use an alignment-based process to detect symmetry, when is it performed? Palmer & Hemenway (1978) proposed a two-stage model of symmetry detection. In their framework, a first stage detects one or more potential axes of symmetry, while a second verification stage confirms the correct axis.

Human sensitivity to qualitative symmetry early in perception (after 50 msec display times) suggests that the alignment process would be the first stage. Symmetry recognition would then involve an interaction between an alignment process that finds the qualitative symmetry, and a subsequent verification stage that uses these correspondences to verify exact symmetry. We note that qualitative symmetry, though approximate, is adequate for guiding visual search during verification.

Additional evidence for a two-stage model of symmetry detection can be found in the *symmetry-based lateral bias effect*. Locher & Nodine (1973) found that visual search patterns for some tasks differ significantly for symmetric and asymmetric figures. They recorded participants' eye movements during a complexity judgment task for random polygons. These polygons were either symmetric about the vertical axis or completely asymmetric (symmetric along no axis). For symmetric figures, participants' fixations were heavily biased to one half of each figure, while fixations for asymmetric figures were unbiased. As noted by M. Corballis (1976), this indicates that some form of symmetry was detected before the first saccade. One interpretation based on the two-stage model is that first stage of processing occurs before visual search. If so, we can determine if qualitative symmetry is recognized in the first stage by examining second-stage visual search patterns.

We tested this hypothesis using a modification of Locher & Nodine's methodology. If the first stage of symmetry detection occurs before the first saccade, and this stage is sensitive to qualitative symmetry, then visual search patterns in the second stage should be different for qualitative and quantitative differences in the figure. This should not just affect the final accuracy (as in Ferguson et al., In preparation) but also the visual search pattern. By analyzing the visual search pattern for asymmetric stimuli

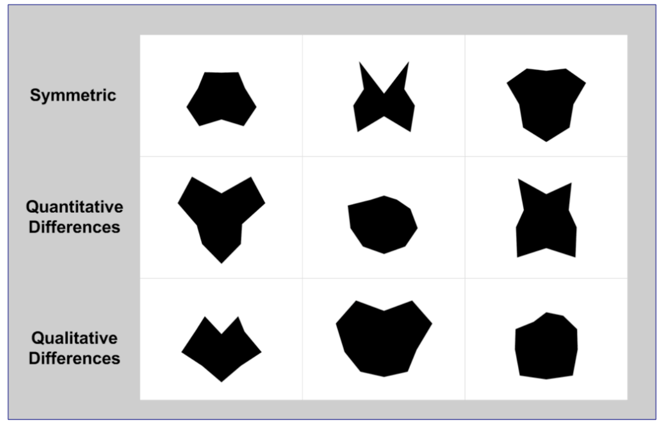

Figure 2: Representative stimuli from the experiment. The complete stimulus set was composed of filled and unfilled figures at three sizes for each symmetry type.

with qualitative and quantitative differences, it should be possible to isolate this effect, thus providing evidence of an alignment process in the first stage of symmetry detection.

To further generalize earlier results, we also looked at two critical factors that might influence symmetry detection and visual search. First, the stimulus size relative to the foveal area determines the amount of visual information that is available before the first saccade, and so could influence the pattern of visual search. An alignment-based model predicts that while added fixations may be required to capture the salient features of the stimulus, the accuracy of judgment should remain, even as size changes. Second, whether the polygon is filled or unfilled may affect the ability to determine the figure-ground information necessary to isolate particular concavities. A filled polygon may assist an alignment-based process by making concavity information more salient or more rapidly available.

Experiment

Method

Participants. 55 university students with normal or adjusted-to-normal vision participated in the study for course credit. Data from nine participants were dropped. Seven participants were omitted due to a high error rate (more than 8% of samples), while two others were omitted due to calibration errors with the eye-tracker.

Materials. A set of 144 randomly generated polygons was used as experimental stimuli, evenly divided between three symmetry types: symmetric polygons, near-symmetric polygons with qualitative differences, and near-symmetric polygons with quantitative differences (Figure 2). Stimuli were shown on a 19 in. monitor set to a resolution of 800x600 pixels and a refresh rate of 60 Hz. Participants were seated at a viewing distance of 81 cm. At this distance, a 30-pixel radius subtended 2 degrees of visual angle. All stimuli were displayed as black on a white background.

Stimuli were created using the method described in (Palmer & Hemenway, 1978), which was modified to

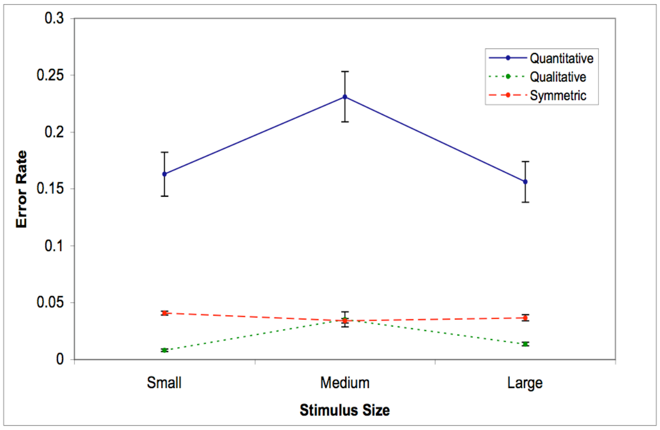

Figure 3: Error rate for symmetry judgments of the three symmetry types at the three stimulus sizes.

generate polygons that varied according to three independent variables: symmetry quality (symmetric, quantitative asymmetric, and qualitative asymmetric), fill quality, and size (with three approximate radii: 50 pixels, 150 pixels, 200 pixels). Qualitative and quantitative differences were generated by taking a generated symmetric shape and changing one randomly selected vertex by a random amount. The range of the amount differed for each size: ±25 pixels for small, ±50pixels for medium, ±100 pixels for large. Polygons were generated as line drawings that were either filled in or given a 3 pixel line thickness. Three additional stimuli were generated for practice trials.

Design. The design of the experiment was within-subjects with the three independent variables for stimuli: symmetry type (3), fill (2), and stimulus size (3). The dependent variables were accuracy and number of fixations.

Procedure. The experiment task was symmetry judgment. Participants were briefed on the experiment task and given three practice trials. Before displaying each stimulus, a fixation point was displayed at the center of the screen until the eye tracker detected the participant's fixation on the point. This centered participants' attention at stimulus onset, and also validated the eye tracking calibration. Stimuli were displayed until participants made a verbal response to the judgment task, at which point the experimenter advanced to the next trial.

Eye movements (Figure 4) were recorded using a corneal reflection eye-tracking device. Eye positions were sampled at a rate of 120 Hz. For analysis, a fixation was detected if a minimum of 200ms of samples were in the same location. A microphone recorded participants' responses. Participants were given 144 trials, where factors were interleaved. Because the stimulus order was fixed, we checked for order effects in the mean fixations between the first and second halves of the stimulus set but found no evidence of an order effect on accuracy ($F(1,142)=0$, $p>0.9$).

Results

Our analysis focused on participant accuracy and the length and number of fixations. We expected figures with qualitative differences to be judged more accurately and with fewer fixations than figures with quantitative differences. We also hoped to be able to see differences in the number and length of local fixations near the differences themselves. In our analysis, we first consider accuracy and fixations for the three symmetry types. We then examine effects of size and fill. Finally, we looked for a symmetry-based lateral bias.

Participants were indeed more accurate and spent fewer fixations judging symmetry of figures with qualitative differences than figures with quantitative differences. Further, results from an analysis of fixations show how visual search is affected by symmetry type. To characterize these effects, we calculated the general pattern of fixations using two different methods: as a proportion of fixations on left and right sides of each stimulus and as fixations occurring closest to qualitative or quantitative differences in the near-symmetric figures.

Effects of symmetry type.

Accuracy. As predicted, participants were significantly more accurate judging figures with qualitative differences (M=98.1%) than either figures with quantitative differences (M=81.6%) or symmetric (M=96.2%) types (Figure 3). These differences between symmetry types are significant ($F(2,43)=15.75$, $p<0.001$). As Figure 4 reveals, this pattern held across all three size conditions. There was no main effect of size ($F(2,45)=0.49$, *ns*) nor was the interaction between size and symmetry type significant ($F(4,41) = 0.32$, *ns*).

This result is consistent with the use of an alignment model for symmetry detection: qualitative differences give earlier feedback to the participant than quantitative differences, improving accuracy for qualitative differences. The analysis of fixations makes this clearer.

Fixations. Symmetry type also significantly influenced the pattern of fixations, but only for the medium and large figures (Figure 5). In general, participants spent more fixations on symmetric figures (M=7.19) than for figures with quantitative differences (M=6.44) or figures with

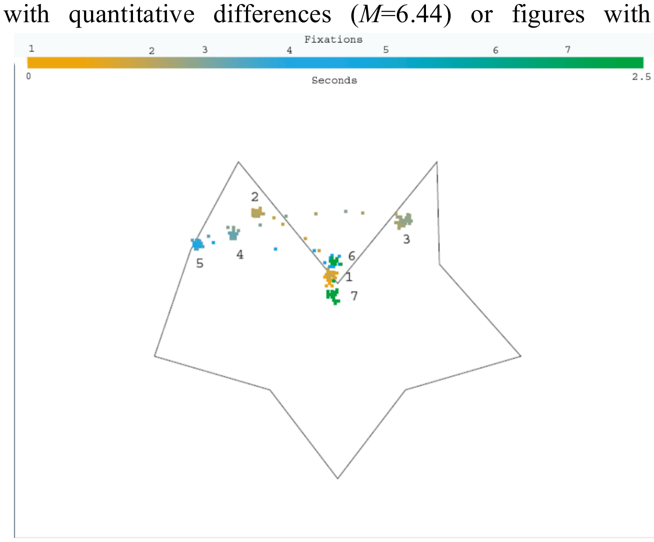

Figure 4: Visual search pattern for single participant. Color (color bar, top) indicates temporal order of samples.

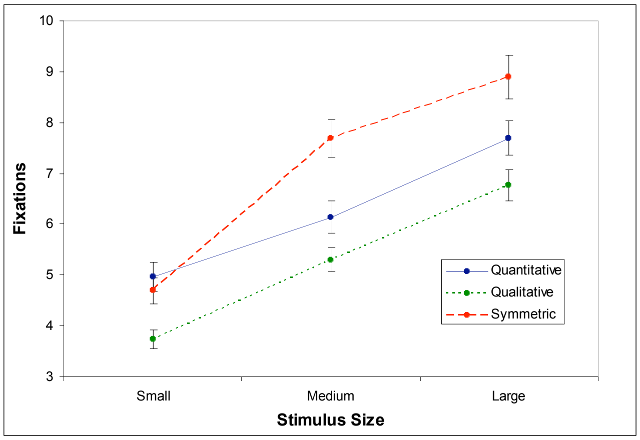

Figure 5: Interval plot of fixation distributions for symmetry type, size factors

qualitative differences (M=5.27, $F(2,45)=106$, $p<0.001$).

The ANOVA revealed significant main effects of symmetry type and size ($F(2,43)=102.45$, $p<0.001$ and $F(2,43)=342.20$, $p<0.001$ respectively) but not of fill ($F(1,44)=0.12$, *ns*). There was also a significant interaction between symmetry type and size ($F(4,40)=10.64$, $p<0.001$). The interaction between symmetry type and fill was marginally significant ($F(2,43)=2.35$, $p<0.1$). The interaction of stimulus size and fill was not significant ($F(2,43)=1.45$, *ns*). The significance of the main effects and the interaction of symmetry type and size both indicate that although fixations increase due to stimulus size (as might be expected), there is a difference in the number of fixations for the different symmetry types. As the stimulus boundary is increased farther from the foveal view, more fixations are required to navigate to the boundary, but fewer fixations are needed to assess figures with qualitative asymmetries than figures with quantitative asymmetries.

Fill. Whether a figure was filled or unfilled did not significantly affect either accuracy ($F(1,142)=0.12$, *ns*) or the number of fixations ($F(1,142)=0$, *ns*). The two way interactions involving fill also were not significant in both analyses. Three-way interactions for both response time and fixations were significant ($F(4,41)=6.00$, $p<0.001$ and $F(4,41)=4.03$, $p<0.01$ respectively). A plot of the distributions of fixations at the factor levels for symmetry type and fill as size increases showed that an interaction with fill was only noticeable in the large size condition.

Eye movement strategies. Capturing eye movements in the symmetry judgment task allows us to test whether qualitative differences guided specific fixations in visual search. To see if participants looked longer at quantitative differences than qualitative differences, we classified each vertex in each asymmetric stimulus as matching (being part of a symmetric feature), quantitative mismatch (being part of a quantitative difference), qualitative mismatch (being part of a qualitative difference), or on axis. We then assigned each sample to its closest vertex. Since there were four times as many symmetric vertexes as asymmetric vertexes, we scaled the symmetric sample counts accordingly.

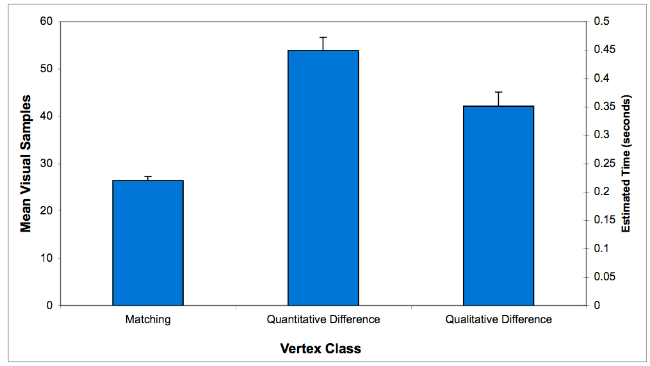

Figure 6: Differences of distributions of mean visual samples (and equivalent time in seconds) proximate to pairs of quantitative, qualitative, and matching differences.

The results (Figure 6) show that participants looked significantly longer at quantitative differences (M=0.44s) than qualitative differences (M=0.35s; $F(1,46)=26.30$, $p<0.001$). They also looked significantly longer at either difference type than at matching vertices (M=0.22s, $F(2,45)=212.5$, $p<0.001$). Again, this suggests that some form of symmetry was known before visual search began.

Symmetry-based Lateral Bias Effect. The experiment by Locher & Nodine (1973) asked participants to rate the complexity of presented stimuli that varied in symmetric quality as well as the number of sides (complexity). Using eye-tracking data, Locher & Nodine reported 11 out of 16 symmetric trials showed a bias of fixations of at least 70/30 to one side of the stimulus relative to the symmetric axis. For asymmetric figures, they reported 14 of 16 shapes showed a distribution of 50/50 or 60/40 between top and bottom axis (asymmetric figures were bisected in the horizontal axis). These results indicate a lateral bias effect for symmetric but not asymmetric figures. We calculated fixation bias over the three symmetry types to test whether the lateral bias effect extends to near-symmetric figures with qualitative or quantitative differences.

If we compare the mean bias ratios of stimuli in the three symmetry types for the different levels of size, we notice bias ratio values that correspond to bias values found in symmetric trials by Locher & Nodine. For all sizes and symmetry types, the bias ratio values are at least 0.70. In addition, figures with qualitative differences (M=0.7876) show significantly more bias than those with quantitative differences (M=0.7729), which in turn show more bias than symmetric figures (M=0.7256, $F(2,45)=77.81$, $p<0.001$) (Figure 7). It is possible this is an artifact of the number of fixations. Participants searched longer and spent more fixations on symmetric than near-symmetric figures, and similarly searched longer and spent more fixations on figures with quantitative rather than qualitative differences. If lateral bias tended to occur early in search, these longer search times would reduce lateral bias, resulting in the greatest bias for the quickest judgments (figures with

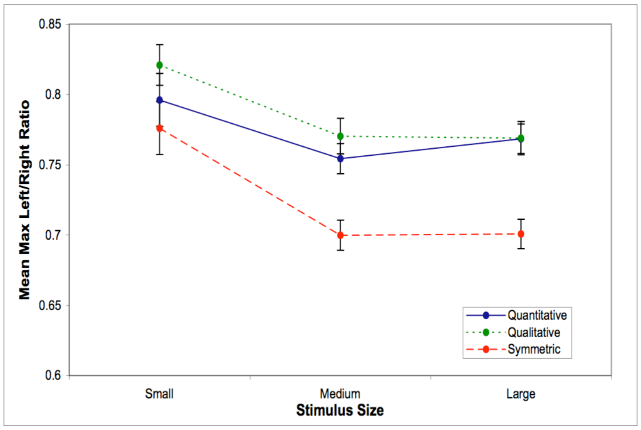

Figure 7: Mean max left/right bias ratio in symmetry axis

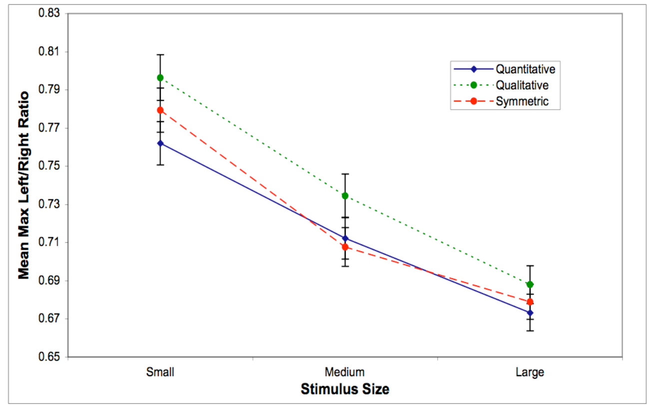

Figure 8: Left/right bias ratio using 1.67s window

qualitative differences) and the least bias for the slowest judgments (symmetric figures).

To compare our result with Locher & Nodine's, we analyzed fixations in asymmetric figures across the horizontal axis. In this case, mean bias ratios were at least 0.70 in all factor levels. This indicates that our asymmetric shapes are more similar to symmetric shapes than the Locher and Nodine asymmetric stimuli. To obtain a clearer understanding of the bias, we analyzed a 1.67s window of eye movements (Figure 8). The window analysis indicates that the bias exists early in processing, and decreases over time. Lateral bias also decreases as size increases.

Discussion

These results support the findings of symmetry processing found in (Ferguson et al., 1996): participants judged near-symmetric figures more accurately when they contained qualitative rather than quantitative differences. This replicates the result of the earlier experiment across two fill and three size conditions. This experiment also shows that qualitative and quantitative differences affect the pattern of visual search. In general, participants looked longer at figures with quantitative rather than qualitative differences, and also fixated on them more. In addition, participants were more likely to fixate on any individual vertex when it was part of a quantitative difference than when it was part of a qualitative difference. All of these factors support the assertion that the visual system is significantly more sensitive to visual differences in near-symmetric figures when those differences are qualitative and involve a relational difference, rather than a difference of degree. In addition, it provides some evidence that symmetry-based lateral bias occurs in near-symmetric as well as symmetric figures.

However, we must consider at least two possible alternative explanations for effects of symmetry type. First, we must consider whether participants were wholly better at classifying figures with qualitative differences, or were improving their accuracy by delaying their response. Second, we must consider whether the distinction between qualitative and quantitative differences is one of degree rather than type (i.e, do humans simply see qualitative differences as larger quantitative differences, and so are more accurate?). We consider each in turn.

Checking for speed/accuracy tradeoff. We tested for a speed/accuracy tradeoff by estimating response time. Although the experiment procedure used did not allow for a precise calculation of response time, we were able to perform a *post hoc* estimate of response time based on the number of samples collected by the eye tracker for each trial. During each trial, participants were allowed to take as much time as needed to judge symmetry and to allow visual search (M=356 samples, 2.97s). When the participant responded vocally, the experimenter pushed a button to stop sampling and present the next stimulus. Although this technique could theoretically introduce experimenter bias into the sample count, such bias might be limited by the speed of the trials, the fact that RT was not a factor of interest, and a tendency to focus on the voice response rather than the displayed stimulus. In fact, a greater problem could be latency added by the experimenter response, which would tend to increase the variance of the sample count.

With these limitations firmly in mind, we performed an analysis of mean response time for each symmetry type at each stimulus size (Figure 9). Participants judged figures with qualitative differences either equally fast (for small

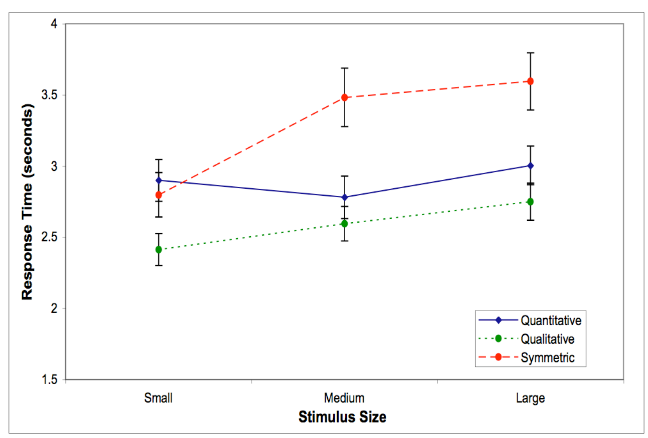

Figure 9: Interval plot of response time for symmetry type as size increases

stimuli) or significantly faster (for medium and large stimuli) than figures with quantitative differences, ruling out a speed/accuracy tradeoff. The two-way interaction between symmetry type and size was significant ($F(4,40)=8.17$, $p<0.001$). We note in passing that participants were also faster (fewer samples) for smaller than larger stimuli ($F(2,43)=63.07$ $p<0.001$, $F(2,43)=21.82$ $p<0.001$ respectively). Differences for fill were not significant ($F(1,44)=0.05$, *ns*).

Two-way interactions between symmetry type and size were significant ($F(4,40)=8.17$, $p<0.001$). Other two-way interactions were not significant (symmetry type, fill $F(2,43)=0.39$, *ns* and stimulus size, fill $F(2,43)=1.17$, *ns*).

Qualitative differences as larger quantitative differences. It is possible that our accuracy effect was due to the qualitative differences simply being larger than the quantitative differences, allowing them to be perceived more easily. To check this possibility, we turned to a psychologically-tested metric model of asymmetry, the Continuous Symmetry Measure (CSM) (Zabrodsky *et al.*, 1992). Using a weighted sum of squared radial differences, CSM measures a figure's difference to the figure with a symmetric exemplar shape. In our stimulus set, figures that are asymmetric (quantitative or qualitative differences) are asymmetric in one feature. The closest symmetric shape (minimum CSM) is the minimum CSM of two potential figures: one setting the asymmetric feature to match the corresponding feature or vice-versa. We calculated CSM for each of the asymmetric stimuli. A one-way ANOVA showed that the CSM of figures with qualitative differences was indeed larger ($F(1,94)=7.46$ $p<0.01$).

We then tested to see if the CSM predicted our accuracy results. We compared the calculated CSM with the observed mean accuracy for each asymmetric trial (quantitative and qualitative). A Pearson correlation test done to test whether CSM was correlated with accuracy in the asymmetric trials showed no correlation between CSM and accuracy in the stimulus set ($r=0.059$, *ns*). Correlation tests based on levels of size showed no significance: small-size, ($r=0.101$, *ns*) medium-size ($r=-0.076$, *ns*), large-size ($r=-0.156$, *ns*). CSM was marginally predictive of accuracy in two cases: one for unfilled figures ($r=0.289$, $p<0.049$) and for small figures with qualitative differences ($r=-0.497$, $p<0.05$) but otherwise not a significant predictor in these subconditions. These results suggest that our earlier results for qualitative differences are not due to CSM. As a result, the effect for qualitative differences is not one of degree, as measured by the most cognitively plausible metric.

Future Work

The results presented here provide further evidence that in a two-stage process model of symmetry perception, qualitative features are handled in the early stage consistent with an alignment-based process. In future work, we expect to conduct further experiments aimed at refining these results. An important follow-on to this experiment will add a control over stimulus presentation time, and counterbalance the amount of quantitative and qualitative difference in the asymmetric stimuli. By using a fixed stimulus presentation time, we hope to eliminate the variance created by experimenter-based advancement of stimuli, while still capturing the salient eye movements. Using the CSM to constrain stimulus generation, we also hope to obtain better comparison results between asymmetric stimuli.

Acknowledgments

This work is supported by a grant from the National Science Foundation under the Artificial Intelligence and Cognitive Science program. The third author is funded by a Presidential Undergraduate Research Award from the Georgia Institute of Technology. We wish to thank the Georgia Tech Psychology Visual Attention Lab for the use of their facilities and helpful suggestions. We also wish to thank anonymous reviewers for comments and suggestions.

References

Corballis, M. C., & Beale, I. L. (1976). *The psychology of right and left*. Hillsdale, NJ: Erlbaum Associates.

Ferguson, R. W. (1994). MAGI: Analogy-based encoding using symmetry and regularity. In A. Ram & K. Eiselt (Eds.), *Proceedings of the sixteenth annual conference of the cognitive science society* (pp. 283-288). Atlanta, GA: Lawrence Erlbaum Associates.

Ferguson, R. W. (2001). *Symmetry: An analysis of cognitive and diagrammatic characteristics.* Unpublished Ph.D., Northwestern University, Evanston, Illinois.

Ferguson, R. W., Aminoff, A., & Gentner, D. (1996). Modeling qualitative differences in symmetry judgments. In *Proceedings of the eighteenth annual conference of the cognitive science society* (pp. 12). Hillsdale, NJ: Lawrence Erlbaum Associates.

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

Locher, P., & Nodine, C. (1973). Influence of stimulus symmetry on visual scanning patterns. *Perception and Psychophysics, 13*(3), 408-412.

Palmer, S., & Hemenway, K. (1978). Orientation and symmetry: Effects of multiple, rotational, and near symmetries. *Journal of Experimental Psychology, 4*(4), 691-702.

Wagemans, J. (1995). Detection of visual symmetries. *Spatial Vision, 9*(1), 9-32.

Zabrodsky, H., Peleg, S., & Avnir, D. (1992). *A measure of symmetry based on shape similarity.* Paper presented at the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, Champaign, IL USA.