In Proceedings of ICCM - 2007- Eighth International Conference on Cognitive Modeling. 211 - 216. Oxford, UK: Taylor & Francis/Psychology Press.

The First Second of Symmetry: Towards a Model of Visual Search during Symmetry Verification

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

Abstract

This paper introduces a newly discovered pattern of eye-movements during symmetry judgments and a corresponding model of visual search. We analyze a corpus of eye movements containing over 19,000 symmetry judgments in two experiments. These eye movements show a "treble-clef-shaped" search strategy that is based on proximate feature inspection and cross-axis comparisons. We simulate this treble clef strategy and show that it accounts for the experimental data better than alternative models.

Introduction: Symmetry Verification

Symmetry judgment is a central process in perception, supporting perceptual organization and object-centered reference frames. Symmetry detection takes place during the earlier phase of the visual process and therefore has a significant impact on object recognition and perceptual organization. It is detected quickly and accurately, even after display times as short as 50 ms.

Although symmetry judgment seems effortless and instantaneous, there is considerable evidence that it is a two-stage process. As proposed by Palmer & Hemenway (1978), symmetry judgment begins with a quick but rough symmetry estimation during an *axis detection phase* that is 50-200 ms long. Then there is a 2000-4000 ms *verification phase* that checks the axis by closely examining the stimulus.

The verification phase of symmetry judgment is important for two reasons. First, because it uses information from the axis detection phase, the pattern of verification may help us implicitly understand what axis detection initially omits. Second, the verification phase provides general clues on how visual regularity guides visual search.

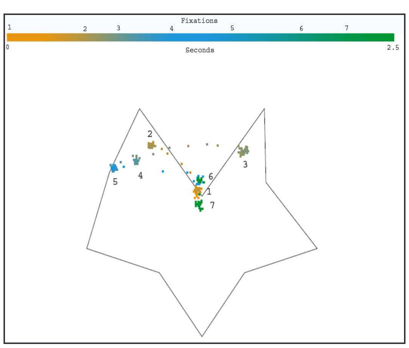

Figure 2: Although humans can detect a symmetry axis very quickly, verifying symmetry takes much longer. Here a participant uses seven fixations over 2500 ms to determine that the figure is asymmetric. The shading and color (color bar, top) indicates temporal order of samples.

Eye-tracking allows us to examine the verification phase directly to determine the pattern of fixations and saccades (Figure 2). (In contrast, the axis detection phase happens too quickly for saccades, leaving accuracy and response time as the typical dependent variables in symmetry studies.) Current eye-trackers measure eye movements with high precision and fast sample rates.

However, while eye trackers are increasingly precise, visual search paths remain highly variable, even for the same participant performing the same task on the same stimulus. Some researchers in motor control suggest that visual search uses stochastic processes to free the mind from the burden of computing the fine details of ocular control (see Mitchell 2003 for a discussion).

The variability of eye movements makes it difficult to model visual search during verification. Minimally, a large dataset is essential for analyzing underlying patterns despite this variability. In addition, general search strategies must be separated from those specific to symmetry judgment.

We address this challenge by using two large data sets of eye-tracking recordings from two previous experiments. These two experiments tracked the eye movements of 71 participants for over 19,000 symmetry judgments. We first analyze this data to determine a new model of visual search. We then use a computer simulation to test this model against alternative search strategies for eye movement. This new model of symmetry verification should allow more precise measurements of interactions between the axis detection and verification phases.

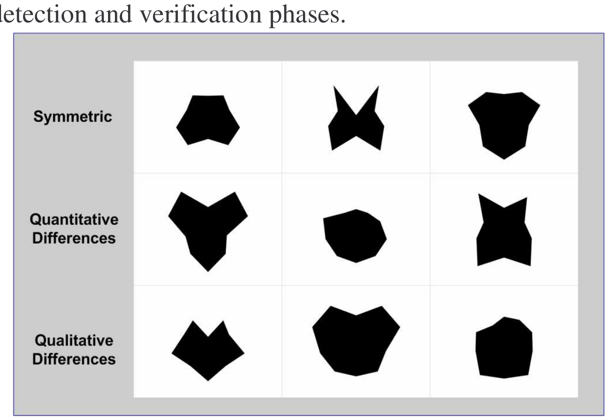

Figure 1: Example stimuli from Experiment 1, with three examples of each symmetry type.

	Experiment 1 (144 stimuli)	Experiment 2 (480 stimuli)
Design Conditions	*Symmetry Types* (3): symmetry, near-symmetry with qualitative difference, near-symmetry with quantitative difference *Fill Types* (2): filled / unfilled *Sizes* (3): small, medium, and large	*Symmetry types* (6): symmetry, near-symmetry with qualitative difference (small and large), near-symmetry with quantitative difference (small and large), asymmetric. *Complexity* (3): 10, 18, & 26 sides
Polygon radius	Small (50 ± 25 pixels), Medium (150 ± 50 pixels), Large (200 ± 100 pixels).	150 ± 70 pixels

Table 1: Summary of experimental designs in (Ferguson *et al.*, in preparation; Mappus *et al.*, 2005).

Review of Corpus Experiments and Data

The eye movement data analyzed here is from two recent experiments exploring how symmetry judgment is influenced by two types of near-symmetry (Ferguson et al., in preparation; Mappus et al., 2005). One type of near-symmetry is symmetric except for a *qualitative difference* (Figure 1), where pair of corresponding parts differs relationally (e.g., a concave vertex on one side of a polygon opposite a convex one). In contrast, the other type of near-symmetry involves a *quantitative difference*, which is one of degree only (e.g., opposing concave vertices where one indentation is larger).

Previous experiments (Ferguson *et al.*, 1996) showed that participants more accurately judge near-symmetries when differences are qualitative rather than quantitative. These two experiments retest this result, but also use eye tracking to see if qualitative differences influence the number and location of visual fixations.

The experimental designs are summarized in Table 1. Experiment 1 used 10-sided polygons (Figure 1) and measured the accuracy and eye-movements of each symmetry judgment. Experiment 2 used a wider range of symmetry types (Figure 3). Nearly all stimuli were symmetric or near-symmetric (except for Experiment 2, which also had totally asymmetric polygons).

The two experiments used different display times and conditions. Experiment 1 examined symmetry judgments of unlimited duration while testing for effects of stimulus size and fill. Experiment 2, in contrast, used a fixed presentation time (1000 ms) and tested for effects of stimulus complexity and difference size.

Both experiments tracked eye movements during the symmetry judgment task using a corneal reflection eye-tracker with a temporal resolution of 8.3 ms (120 Hz). In Experiment 1, participants responded, on average, within 3000 ms using 6 fixations. Experiment 2 reduced this to 1000 ms and 3 fixations.

Overall, 96 university students with normal or adjusted-to-normal vision participated in the studies for course credit. Experiment 1 used 55 participants, but 9 were dropped from the analysis due to high error rates or eye-tracker calibration errors. Similarly, Experiment 2 used 41 participants with 16 participants dropped.

The polygonal stimuli were randomly generated by connecting points along a set of evenly-spaced radii, as in (Palmer & Hemenway, 1978). Near-symmetric polygons were generated by changing a random vertex of a symmetric polygon by a random amount (as given in Table 1). Stimuli were shown in black on a white background subtending approximately 2 visual degrees.

The results showed clear processing differences for near-symmetric figures where differences were qualitative rather than quantitative. Participants made more judgment errors when differences were quantitative (Figure 4), and also used more fixations (an average of 6.3 fixations for quantitative and 5.3 for qualitative). Experiment 2 showed similar results (Figure 5), although the effect was greatly attenuated when quantitative differences were large or when the figure was extremely complex (e.g., for 26-gons).

These experiments showed that qualitative differences influence the number and placement of fixations, but did not predict the search path. The goal on this analysis is to determine if there is an underlying search strategy that participants used.

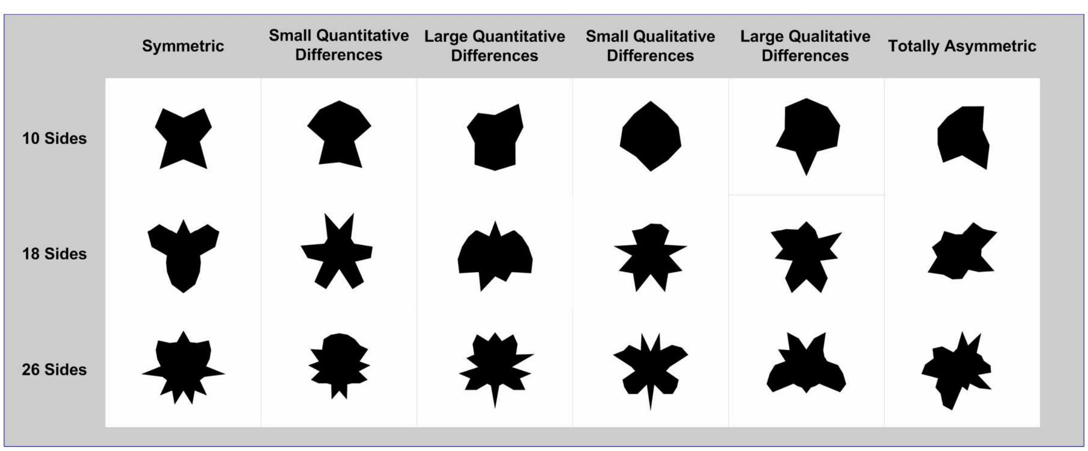

Figure 3: Example stimuli from Experiment 2. This chart shows one stimuli for each of the 18 different conditions.

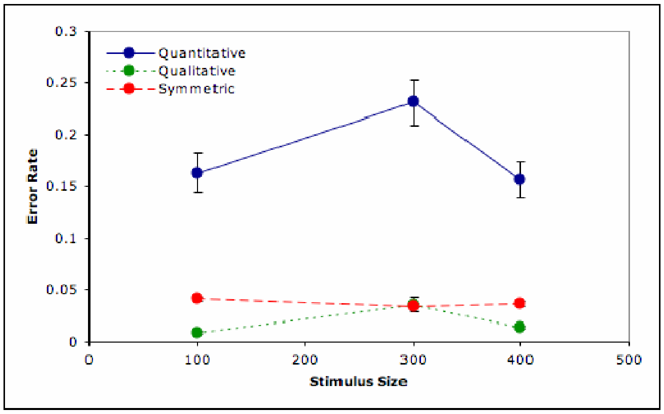

Figure 4: Error rates from Experiment 1 for symmetric figures, and near-symmetric figures with qualitative and quantitative differences.

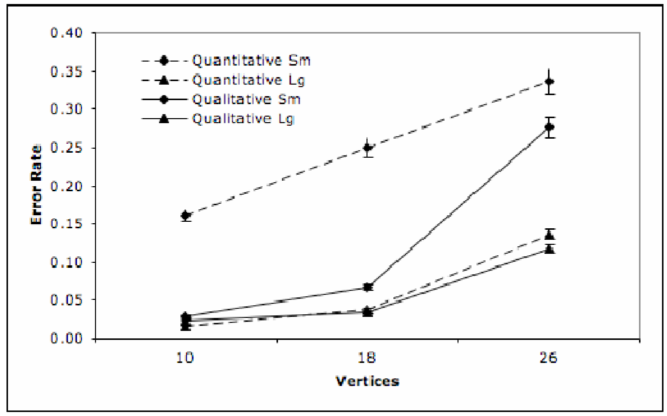

Figure 5: Error rates from Experiment 2 for near-symmetric figures with small and large qualitative and quantitative differences.

Analysis of visual search patterns

In general, the use of near-symmetric stimuli emphasized symmetry verification. Each near-symmetric stimulus contained a single differing vertex, and so verification was difficult and required visual search.

Interestingly, our analysis shows that while individual movement paths were highly variable, the aggregate paths show a predictable pattern that we have dubbed the "Treble Clef" strategy (Figure 6). In this strategy, participants initially fixate at the center of the figure, and then move up. They then descend while moving from the left to the right.

The Treble Clef strategy has several components, which we will cover in turn, using graphs of the vertical and horizontal movement components found in Figure 7.

Initial 200-250 ms. fixation. Most participants began in the center of the stimulus and remain there for the first 200-250 ms (Figure 7). This is expected because experiment protocol requires each subject to fixate on the center of the screen prior to stimulus onset. The 200-250 ms duration corresponds to the time needed for the axis detection phase.

Upward movement to tip of vertical axis. Then for the rest of the initial 500 ms the eye movements begin to spread out and migrate to the top of the figure.

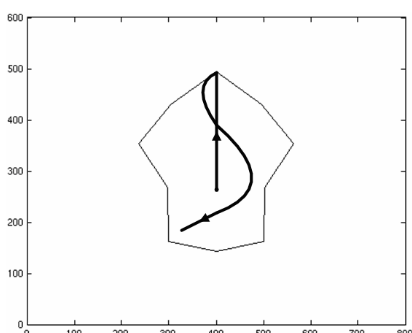

Figure 6: Diagram of Treble Clef search strategy. The diagram illustrates the general pattern of eye movements during the first 1,000 ms of a symmetry judgment task.

Oscillating downward movement. During the last 500 ms, the eye-tracking samples appear to spread out perpendicular to the center axis and slowly move down the figure. It is not until the end of the 1,000 ms that there are any samples at the bottom of the figure (Figure 8). This shows a clear top to bottom pattern; a pattern that has been noticed in both search tasks (Salvucci, 2000) and comparative search (Pomplun, 1998).

Other characteristics. There are several other characteristics of the visual search data. First, the data shows a clear undershoot bias; rather than overshooting the saccade endpoint and landing outside the figure, the vast majority of the sample points remain inside the figure.

Second, the treble clef pattern remains whether the display time is constrained or unlimited. Similarities between the mean vertical/horizontal positions for the first 1,000 ms of Experiment 1 and the mean positions from Experiment 2 suggests that search patterns are very similar for the first 1,000 ms despite the time constraints of Experiment 2.

This is also supported by the relatively similar accuracy levels for 10-sided polygons in Experiments 1 and 2 (Figures 4-5), which suggest that participants are able to maintain accurate verification even when display times are reduced from 3000 to 1000 ms. Participants are either able to optimize their search strategy for the shorter time or gain little from the visual fixations after the first 1000 ms.

Also note a slight lag in the peaks of Experiment 1's mean positions (Figure 7-A,B) relative to Experiment 2's (Figure 7-C,D). The apparent expanding of the waves in the graphs suggests that Experiment 2's time constraints may have improved the efficiency of the eye movements.

In summary, the treble clef pattern is very consistent. On average, participants start at the center, quickly move up, and then slowly move down, swerving from the left to the right. Separating the samples by stimulus type, we see a striking consistency across all symmetry categories despite other performance differences.

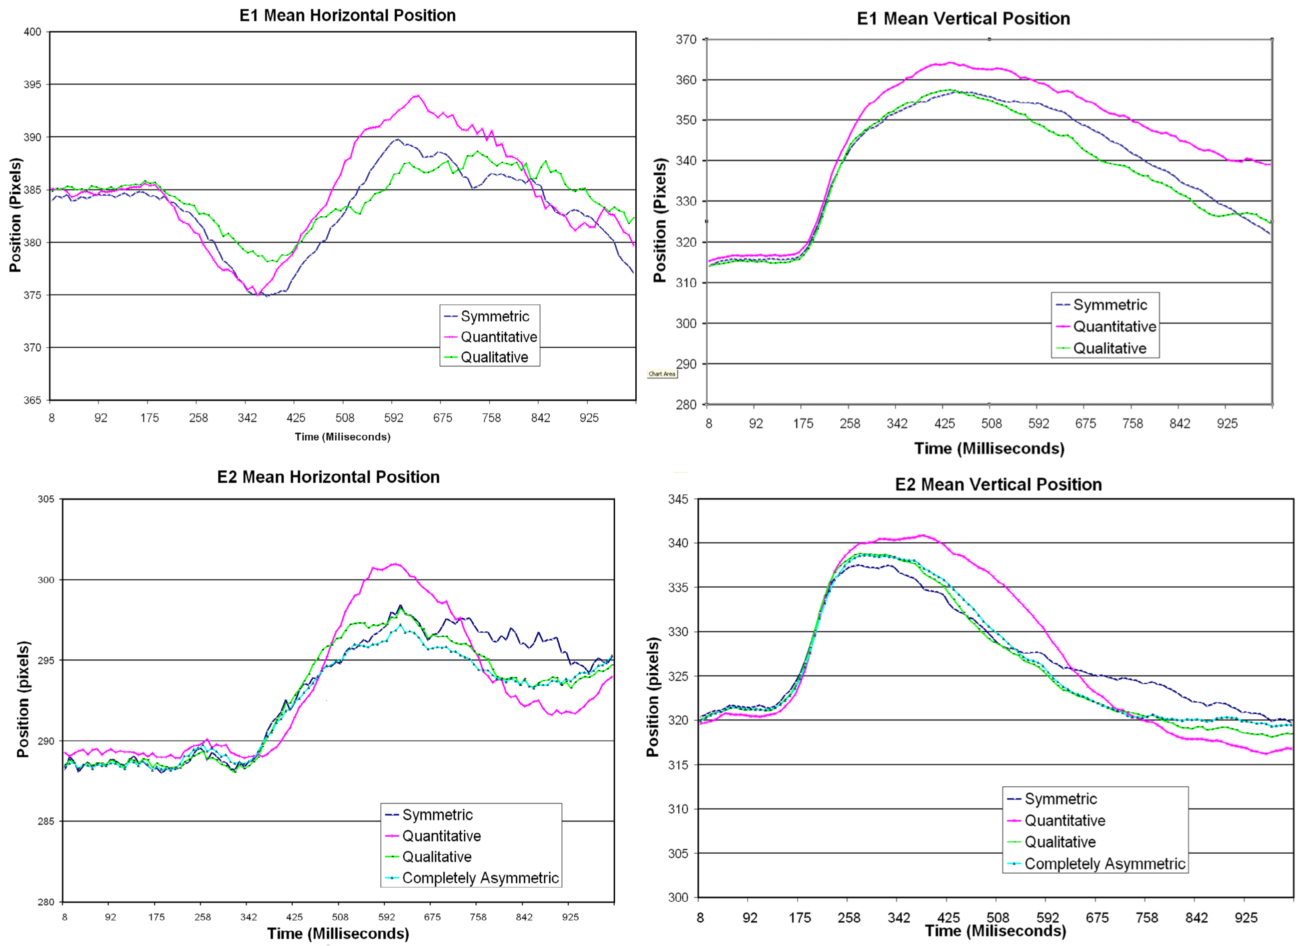

Figure 7: An overview of aggregate eye movements during Experiments 1 and 2, showing mean horizontal and vertical positions over time for all symmetry types. The Horizontal center is at 400 and the Vertical center is at 300.

Simulation

To further test the treble clef strategy, we compared it against three alternative search strategies: Random Search, Greedy Search, and the Area Activation Model (AAM).

Eye-movement model. To do this, we first built an eye movement model, based on known psychological results, to serve as the basis for all four search strategies. The model groups movement into two categories: fixations and saccades. A fixation is a period of stable movements closely clustered around a point. Saccades are fast ballistic movements that propel the eye across a visual scene from one fixation to another. The model first generates a fixation at the stimulus center. At the end of each fixation the search strategy selects a new fixation location and plots a saccade. This continues until 1,000 ms has expired.

To capture the stochastic nature of the fixations and saccades, variables such as the fixation duration and saccade duration are drawn from probabilistic distributions. In this model, the lengths of the fixation duration are sampled from a gamma distribution with a mean of 200 ms and a standard deviation of (1/3)*(200ms) (Epelboim & Suppes, 2001; Salvucci, 2000). During a fixation, locations for each time step are drawn from a 2D Gaussian distribution centered at the target with a standard deviation equal to one visual degree. To move from one fixation to another, a straight line saccade is charted with samples evenly distributed along the line. The saccade duration is 20 ms + .2 ms * visual angle. Saccade landings are not precise and instead of traveling the

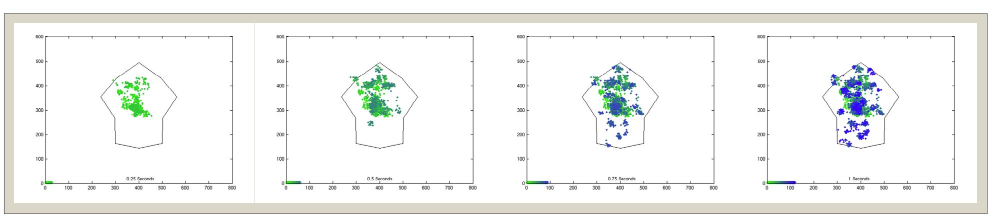

Figure 8: Time slices of eye-tracking samples for all subjects in Experiment 2. The diagram shows samples from a stimulus divided into four time slices: 0-250 ms, 0-500 ms, 0-750 ms, and 0-1000 ms.

full distance d from one fixation to the next, saccade distance is drawn from a Gaussian distribution with a mean of d and standard deviation equal to 0.1d (Salvucci 2000). Following the undershoot bias, there is a 90% chance that the actual saccade distance will be less than d and a 10% chance that the actual saccade distance is greater than or equal to d.

The eye movement model then serves as the basis for the following four search strategies.

Random Search Strategy. The Random search strategy assumes that there is no underlying motivation for selecting fixations. Each vertex has an equal chance of being selected, with the constraint is that no vertex is selected twice.

Greedy Search Strategy. The Greedy search strategy assumes that the visual system attempts to maximize information by reducing time spent on saccades, and always selects fixation locations on the closest unvisited vertex.

The Area Activation Model *(AAM)*. The Area Activation Model (Pomplun *et al.*, 2000) is a generalized model of eye-movements during search tasks. Like the greedy search strategy, AAM tries to optimize information. However, instead of selecting nearby vertices to minimize saccade time, AAM finds clusters of vertices where it can maximize information in a single fixation. One particularly interesting result of this is that fixations can occur at the center of a cluster of vertices rather than directly on a single vertex.

According to the model, a 3D activation mesh is created preattentively. The mesh consists of a mixture (summation) of 2D Gaussians centered on the stimuli's items (vertices in our case). The peaks of the mesh become the candidate fixation locations and are weighted by their relative heights. The first candidate is selected using a weighted probability. Subsequent candidates are selected based on their proximity to the current fixation.

Treble Clef. The Treble Clef search strategy mimics the pattern discovered in the two experiments. After the initial fixation, the next fixation is at the top of the figure. The remaining fixation locations are on vertices that are lower than the current fixation. This search strategy also implements the swinging motion by alternating between vertices on the left and right of the symmetric axis.

Comparing the Strategies

The goal of the simulation is to evaluate each search strategy's fit to the empirical data. Therefore, search strategies are tested by comparing the samples generated by the simulation (simulation samples) with the samples generated during Experiment 2 (experiment samples). The variability of the data requires the comparison of estimated sample distributions rather than a sample-to-sample comparison.

To quantitatively measure the accuracy of the simulations for a particular stimulus, we estimate the distributions of both the experimental samples and the simulation samples then calculate the divergence of the two distributions. The distributions are estimated using a mixture of Gaussians. One 2D Gaussian distribution, with a standard deviation equal to the radius of the foveal region, is placed at the coordinates of each sample point. The Gaussians are summed across the 2D space and scaled relative to the number of samples used to calculate the estimation. The percent of overlap, which we use as the measure of accuracy, is the summation of the minimum of the corresponding masses:

$$\sum_X \sum_Y \min[f(x,y), g(x,y)]$$

Where f(x,y) is the estimated probability distribution of the experimental samples and g(x,y) is the estimated probability distribution of the simulation samples. Since this calculation is specific to each stimulus, accuracy results are averaged across a random subset of Experiment 2 stimuli.

Figure 10 compares the mean accuracies of the different search strategies. Of the strategies tested, Treble Clef performs the best. Both Greedy and AAM perform with similar accuracy; this is expected because after selecting the first fixation the AAM strategy uses a greedy algorithm for selecting subsequent fixations. However, the deterministic nature of Greedy and the AMM strategies limits their ability to account for the variety of scan paths recorded from our experiment participants. The Random Strategy performs

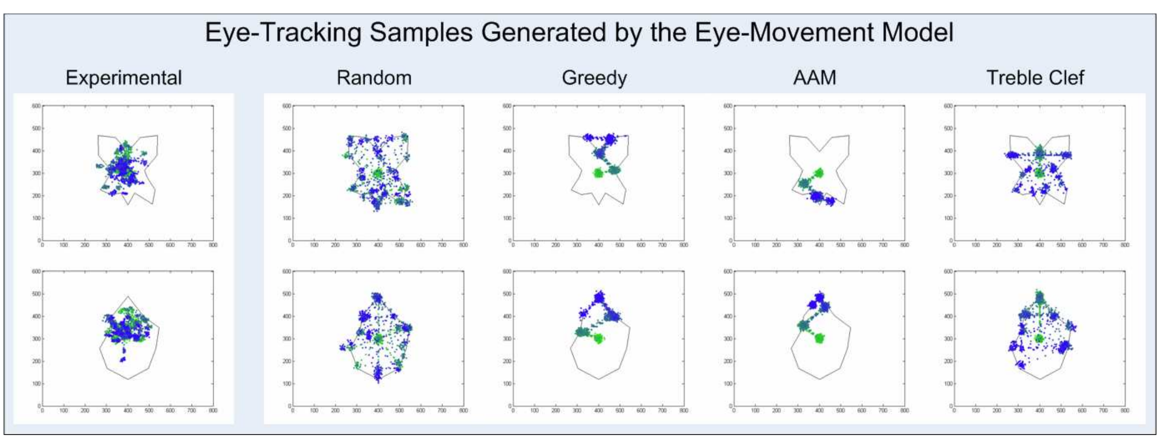

Figure 9: Eye-tracking samples generated by the simulation. The diagram shows how the different search strategies generate different sets of samples. The experimental samples from Experiment 2 are included on the left for comparison purposes.

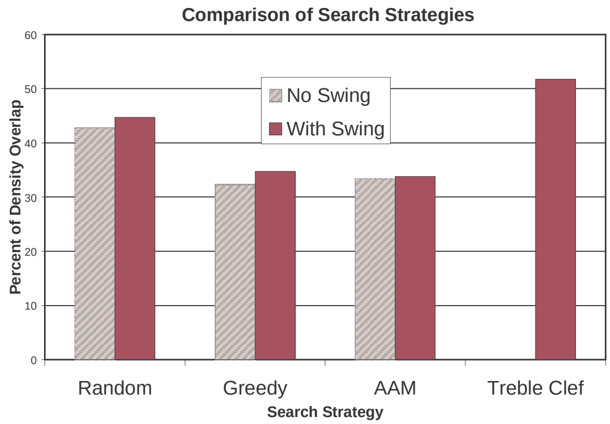

Figure 10: Mean accuracy of each search strategy. In this graph, the values are the percent of overlap with the estimated empirical density averaged across a randomly selected subset of Experiment 2 stimuli.

better than Greedy and AAM, but does not account for the vertex preferences that appear in the experimental data.

To simulate the comparative nature of the symmetry task we also tested a swing variant of Random, Greedy, and AAM. In all cases, adding the swing constraint (a constraint that successive fixations alternate between the left and right side of the symmetric axis, as is expected for symmetry judgment) improved performance.

Discussion and Future Work

The eye-tracking data and simulation results demonstrate the existence of the Treble Clef pattern. While the variability of eye-movements disguises the underlying strategy for selecting fixation locations, the use of a large corpus has enabled us to identify elements of a common strategy. Furthermore, the use of an eye-movement simulator has shown this strategy better represents the experimental data than several alternative strategies.

The discovery of this strategy provides valuable insight into symmetry verification and symmetry perception in general. Previous work shows that symmetry type influences detection through differences in accuracy, number of fixations, response time, and scan paths (Ferguson et al 1996; Mappus et al 2005). Similarly, our eye-tracking data shows that symmetry types differ with respect to their mean vertical and horizontal positions, especially after 500 ms (see Figures 7-10). These search patterns, like accuracy and response time, are affected by the symmetry type. This suggests that studying perturbations of the Treble Clef search pattern may indicate processing differences between symmetry types. It also suggests that one could determine how a symmetry type influences the visual system by determining when the eye-movement patterns diverge.

The Treble Clef pattern also suggests that the ocular system influences search strategies. Instead of quick straight-line movements and sharp angles between fixations, the pattern shows wavy movements. This can be explained by momentum-like forces impacting the ocular muscles, which suggests that search strategies are optimized to work within the physical constraints of the ocular system.

In the future, we will refine the Treble Clef search strategy and the eye-movement model to better account for the experimental data. An experiment designed to test a wider variety of symmetry types could improve the parameters of the eye-movement model and thus simulation accuracy. This is one goal for follow-on experiments.

We have also begun work on an experiment that uses electroencephalographic (EEG) recordings to find brain activation patterns during symmetry judgment. By evaluating eye-tracking along with EEG we may gain valuable insight into the cognitive processes involved in symmetry judgment.

Acknowledgments

This work is supported by a National Science Foundation grant under the Artificial Intelligence and Cognitive Science program and a Presidential Undergraduate Research Award to the first author. Special thanks to Paul Corballis for help with the experimental design and to the Georgia Tech Psychology Visual Attention Lab for use of their facilities. We also wish to thank reviewers' comments and suggestions.

References

Epelboim, J., & Suppes, P. (2001). A model of eye movements and visual working memory during problem solving in geometry. *Vision Research*, 41, 1561-1574.

Ferguson, R. W., Aminoff, A., & Gentner, D. (1996). Modeling qualitative differences in symmetry judgments. Cognitive Science Conference. Hillsdale, NJ: Erlbaum.

Ferguson, R. W., Mappus, R. L., Czechowski, K., & Corballis, P. M. (in preparation). Spotting differences: Effects of geometric relations on visual search patterns. Manuscript in preparation.

Mappus, R. L., Ferguson, R. W., Czechowski, K., & Corballis, P. M. (2005). Spotting differences: How qualitative asymmetries influence visual search. Cognitive Science Conference. Hillsdale, NJ: Erlbaum.

Mitchell, J., Zipser, D. (2003) Sequential memory-guided saccades and target selection: a neural model of the frontal eye fields. *Vision Research*, 43, 2669-2695.

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.

Pomplun, M. (1998). Analysis and models of eye movement in comparative visual search. Göttingen: Cuvillier.

Pomplun, M., Reingold, E. M., Shen, J., & Williams, D. E. (2000). The area activation model of saccadic selectivity in visual search. Cognitive Science Conference. Mahwah, NJ: Erlbaum.

Salvucci, D. D. (2000). A model of eye movements and visual attention. International Conference on Cognitive Modeling. Veenendaal, Netherlands: Universal Press.