



Perception of Symmetries in Drawings of Graphs

Felice De Luca¹(✉), Stephen Kobourov¹, and Helen Purchase²

¹ University of Arizona, Tucson, USA

felicedeluca@email.arizona.edu, kobourov@cs.arizona.edu

² University of Glasgow, Glasgow, UK

helen.purchase@glasgow.ac.uk

Abstract. Symmetry is an important factor in human perception in general, as well as in the visualization of graphs in particular. There are three main types of symmetry: reflective, translational, and rotational. We report the results of a human subjects experiment to determine what types of symmetries are more salient in drawings of graphs. We found statistically significant evidence that vertical reflective symmetry is the most dominant (when selecting among vertical reflective, horizontal reflective, and translational). We also found statistically significant evidence that rotational symmetry is affected by the number of radial axes (the more, the better), with a notable exception at four axes.

1 Introduction

Many objects in nature, from plants and animals to crystals and snowflakes, have symmetric patterns. Humans and other animals have a nearly perfect reflective symmetry along a single axis; sea stars and snowflakes have repetitive patterns along two or more radial axes; leaves and flowers often have translational patterns of symmetry; see Fig. 1.

The perception of symmetry is one of the key concepts in Gestalt theory which studies how humans perceive different types of objects. Symmetry has also been considered an important feature of well-drawn graphs, on the basis that depicting symmetries will reveal a graph's structure and properties [8]. A natural question that arises is: which types of symmetry are easier to perceive and how does this affect drawings of graphs? In this study we investigate this question, focusing on the reflective (also called “mirror”), translational and rotational (also called “radial”) types of symmetries:

- **Vertical:** A pattern is reflected across a vertical axis (*reflective symmetry*)
- **Horizontal:** A pattern is reflected across an horizontal axis (*reflective symmetry with a 90° rotation*)
- **Translational:** A pattern is repeated and shifted in the space
- **Rotational:** A pattern is repeated across radial axes with a given angle

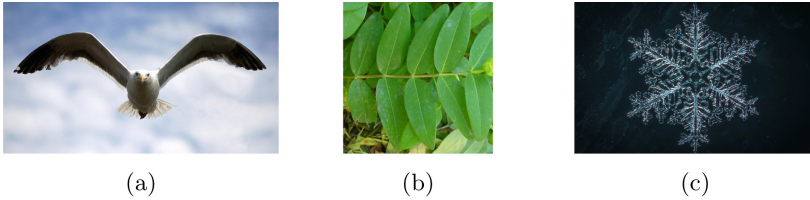


Fig. 1. Symmetry in nature: (a) reflective, (b) translational, (c) rotational.

We used synthetically generated graph drawings in a human subjects experiment to answer several questions related to the perception of symmetry. Specifically, we created graph layouts that exhibit different types of symmetries and asked our participants to select the more symmetric ones. We found statistically significant evidence that vertical reflective symmetry is the most dominant (when selecting among vertical, horizontal and translational). We also found statistically significant evidence that rotational symmetry is affected by the number of radial axes (the more, the better) with a notable exception at four axes.

2 Related Work

Gestalt theorists studied how objects are perceived, their view being that the perception of a whole object cannot be reduced to the sum of the perception of its parts [17]. They focused on describing the patterns we see in visual stimuli [2, 30], and in how we distinguish background from foreground [23], devising a set of laws describing fundamental perceptual phenomena. Well-known Gestalt principles include *proximity* (things that are close together are perceived as being in a group), *similarity* (things that look “similar” are perceived as being in a group), *closure* (closed shapes are preferred to open shapes), and *symmetry*, which we discuss in more detail next.

Giannouli [9] defines symmetries on the plane as “transformations that preserve equal geometric distance,” and identifies the three standard types: translational, rotational and reflective symmetries. Each of these have variations: the distance and direction of the transposition for translational symmetry, the number of times around a circle the object is repeated (known as its “order”) for rotational symmetry, and the angle of the axis (with horizontal and vertical being the most common) for reflective symmetry.

The human perception of symmetry has been studied in different contexts. For example, there is evidence that people remember figures as more symmetric and closer than they really are [28] and that symmetry aids in the recall of abstract patterns [12, 26].

Several studies compare the effectiveness of different types of symmetry. An early experiment by Corbalis and Roldan [7] compared vertical reflective symmetry with translation symmetry (limited to only horizontal translation). They compared these two conditions in two forms – where the two components were

touching each other (i.e., creating a single holistic visual object) or where there was a horizontal distance between the two components (i.e., they were perceived as two separate objects). They concluded that vertical symmetry was more salient in the holistic case, but translation more salient when the two components are disjoint. Other early work [3] concluded that vertical symmetry is easier to detect than translational symmetry. Royer [24] prioritizes horizontal and vertical reflective symmetry over diagonal reflection, with “centric” (loosely comparable to rotational, despite Royer’s stimuli being square in form) performing worse. The work by Palmer and Hemenway [18] confirms the ordering: vertical reflective, horizontal reflective, diagonal reflective.

Cattaneo *et al.* [6] investigate the neurological basis for the perception of vertical and horizontal reflective symmetry and conclude that there is a “partial” difference between the regions of the brain used in detecting these two symmetries. Giannouli’s [9] review of research on the visual perception of symmetry finds that vertical reflective symmetry is more readily perceived than any other type. Similar findings are reported in an earlier review by Wageman [29], who ranks reflective symmetries as follows: vertical produces better recognition performance (faster or more accurate) than horizontal, which in turn performs better than diagonal.

Jennings and Kingdom [10] conducted experiments to compare the perception of different orders of rotational symmetry (3, 5 and 7), together with a vertical reflection condition. They conclude that it is easier to detect rotational symmetry as the order increases (as measured by response time), and that, in comparison with rotational symmetry, the vertical reflective symmetry condition performs better than 3rd order, the same as 5th order, and worse than 7th order.

Note that none of this work involved graphs or drawings of graph.

While some researchers have considered the application of Gestalt principles to graph drawing, such work is rather fragmented. Wong and Sun [27] create key criteria for the depiction of UML class diagrams based on the Gestalt theories, and then evaluate three UML diagram tools based on these criteria. Bennet *et al.* [1] review the literature on graph drawing aesthetics with reference to Norman’s [16] stages of perception (visceral, behavioral and reflective), including the Gestalt theories in the visceral stage. They conclude that more work needs to be done to validate the common graph drawing aesthetic criteria with respect to perceptual theories. Nesbitt and Freidrich [15] discuss some of the Gestalt theories in relation to graphs that evolve over time, although symmetry is not explicitly considered. Lemon *et al.*’s experiments [13] show that the principles of similarity, proximity and continuity affect the comprehension of complex software diagrams. Rusu *et al.* [25] focus on the principle of continuity, and proposed a method for reducing visual clutter created by edge crossings by creating gaps in the edges. This is embodied in the partial edge drawing algorithms of Bruckdorfer *et al.* [4] and Burch *et al.* [5]. Marriott *et al.* [14] conduct an experiment looking at what features of small graph drawings made them most memorable. Their experimental conditions explicitly relate to the Gestalt principles of symmetry,

continuity, orientation and proximity and their finding indicate that drawings that exhibit symmetry and continuity are amongst those most readily recalled.

Eades advocates the use of algorithms that aim to draw graphs with “as much symmetry as possible” [8]. Early experiments investigating the relative importance of different graph drawing aesthetics find support for the depiction of symmetry in terms of performance on graph-reading tasks [20,21]. Two computational methods have been proposed for measuring the extent of symmetry in a graph drawing, a non-trivial task in and of itself. The method proposed by Purchase [19] considers only reflective symmetry. It generates potential axes of symmetry between all pairs of vertices, and determines the existence of symmetric sub-graphs (edges reflected around the axis, with a tolerance) for each axis. Klapaukh’s method [11] uses an edge-based metric that includes rotational and translation symmetries in addition to vertical ones. Welch and Kobourov [31] studied which of these two algorithms best correlates with the human perception of symmetry, with results that suggest that a graph drawing with vertical symmetry is considered more symmetric than the identical drawing presented at a slightly different orientation, and that the greater the extent of symmetry in a drawing, the faster the participants’ response.

While there has been extensive experimental research in the perception literature comparing the different types of symmetry in a variety of artificial stimuli, no comparable work has been performed to investigate the perception of symmetry in graph drawings. We therefore extend the work done so far by conducting experiments that specifically considers which of the three types of symmetry are more salient in drawings of graphs, including several additional variations.

3 Research Questions

We investigate the perception of graph drawings that exhibit three types of symmetry: reflective (vertical and horizontal), translational and rotational. Rather than attempting to draw existing graphs that embody such symmetries (a difficult task), we create symmetric graph drawings by duplicating graph-substructures. Specifically, we draw a small graph, make a duplicate, place the duplicate(s) appropriately (according to the type of symmetry), and join the components together to create a graph drawing that exhibits the desired symmetry.

Since rotational symmetry is visually very different from reflective and translational symmetry, we address two separate research questions:

1. What is the relative ranking of reflective and translational symmetries for drawings of graphs?
2. What is the impact of the number of axes (order) for rotational symmetry?

3.1 The Symmetric Graph Drawings

Our experiment considers several different types of symmetry: horizontal (H), horizontal with rotation (Hr), vertical (V), vertical with rotation (Vr), translational (T), translational with rotation (Tr), rotational with fixed components

(RC) and rotational with fixed vertices (RV). For a baseline, we also have a non-symmetric version (NS); see Fig. 2. We consider two variants of rotational symmetry (RC and RV) in order to take into account the effect of the number of rotational axes and the effect of different graph sizes.

Ideally, all the stimuli should represent exactly the same graph, but since this would be impossible (especially for the rotational drawings), we attempt to impose some consistency by using the same “base graph” from which the larger graphs are derived.

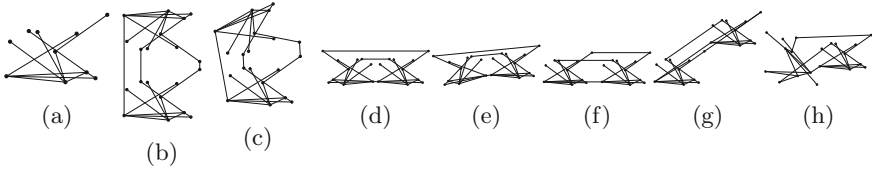


Fig. 2. Example layouts in the dataset: (a) base graph, and its (b) H, (c) Hr, (d) V, (e) Vr, (f) T, (g) Tr, and (h) NS.

We create the horizontal, vertical and translational drawings (HVT, for short) as follows. Let $G_b = (V_b, E_b)$ be a base graph drawn with a random layout such that each vertex $v_b \in V_b$ has positive coordinates; see Fig. 2a. Let $G_c = (V_c, E_c)$ be a copy of G_b with the same layout. Then $G_s = (V_b \cup V_c, E_b \cup E_c \cup E)$ is created from the two graphs together with edge set E connecting the vertices in V_b to their copied version in V_c . We fix $|E| = 3$ and the vertices that are chosen for the connection are chosen at random. We use G_s to create the layouts of the graphs in the HVT set by changing the coordinates of the vertices $v_c \in V_s$ as follows:

- **H**: If $v_b = (x, y)$ then $v_c = (x, -y)$; see Fig. 2b.
- **Hr**: H version with a rotation with angle in $[0, 45]$; see Fig. 2c.
- **V**: If $v_b = (x, y)$ then $v_c = (-x, y)$; see Fig. 2d.
- **Vr**: V version with a rotation with angle in $[0, 45]$; see Fig. 2e.
- **T**: If $v = (x, y)$ then $v_c = (x - \delta, y)$ where δ is a shifting factor such that the bounding boxes of V_b and V_s do not overlap; see Fig. 2f.
- **Tr**: T version with a rotation with angle in $[0, 45]$; see Fig. 2g.
- **NS**: Non symmetric (random) placement of the vertices in V_s ; see Fig. 2h.

We have two types of rotational drawings: maintaining the base graph component (the “fixed component” version), and limiting the maximum number of vertices (the “fixed vertices” version). We use both methods in order to control for the possible confounding factor of different graph sizes in the first variant.

The rotational fixed component (RC for short) versions are symmetric layouts that repeat the base graph drawing around each axis. The base graph $G_b = (V_b, E_b)$ is a component of the layout of the symmetric graph G_s and we create the different graphs as follows:

- **RC{X}**: The drawing of each component is replicated on each of the $X = [4, \dots, 10]$ axes of symmetry. By choosing two random vertices from a component, we connect each pair of rotationally consecutive components with edges to the corresponding vertices; see Fig. 3.

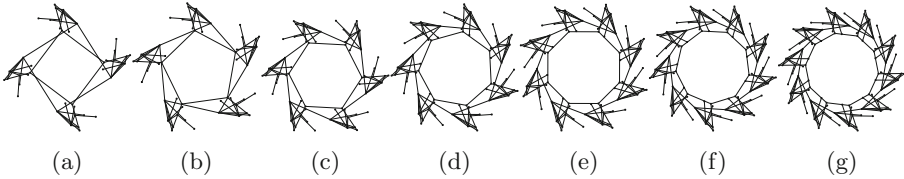


Fig. 3. Example of rotational layouts with fixed components: (a) RC4, (b) RC5, (c) RC6, (d) RC7, (e) RC8, (f) RC9, and (g) RC10.

The rotational fixed vertices (RV for short) versions are symmetric layouts with a limited maximum number of total vertices. They are created as follows:

- **RV{X}**: The base graph is reduced in size by removing as many vertices (at random) as needed so that when it is replicated on each of the $X = [4, \dots, 10]$ axes of symmetry, the total number of vertices does not exceed 50. As before, we connect each pair of rotationally consecutive components with two edges; see Fig. 4. Inevitably, the graphs in this set are not exactly of the same size, but they are within 20% of each other; Table 1 shows the number of vertices and average number of edges.

Table 1. RV sizes

	V	Avg E
RV4	40	52
RV5	50	65
RV6	48	60.3
RV7	49	60.9
RV8	48	59.6
RV9	45	55.8
RV10	50	62

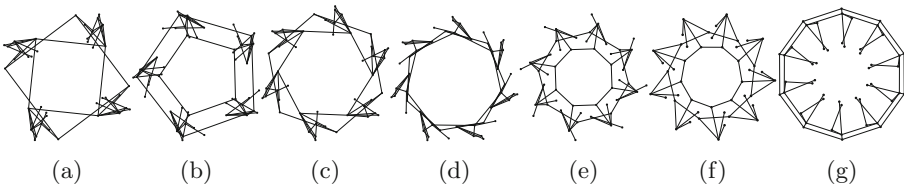


Fig. 4. Example of rotational layout with fixed number of vertices: (a) RV4, (b) RV5, (c) RV6, (d) RV7, (e) RV8, (f) RV9, and (g) RV10.

We use small and sparse base graphs ($|V| = 10$ and $|E| = 11$) as these would be copied and interconnected when creating the various experimental sets. Each base graph is drawn by placing the vertices at random, to avoid accidental symmetries within the base graph and in order to focus on the symmetries

created by the base graph replication. We use two (RC and RV) or three (HVT) edges to connect the replicated components for the same reasons.

By construction, all pairs of HVT graphs used as stimuli are the same, while pairs of RC and RV graphs are structurally different. We used the yFiles library [32] to generate and draw our stimuli.

4 Experimental Methodology

We conducted three separate experiments: Reflective and Translational (RT), Rotational Fixed Component (RFC) and Rotational Fixed Vertices (RFV). We used the same methodology but different sets of stimuli. The participants' task was to look at a series of pairs of graph drawings, and, for each pair, indicate which one of them they thought was "more symmetric." This task was designed to address the following experimental questions:

- **Q1:** Which type of symmetry among H, V, T, Hr, Vr, Tr is most recognizable as symmetry?
- **Q2:** How many rotations is most recognizable as rotational symmetry, using the fixed-component generation method?
- **Q3:** How many rotations is most recognizable as rotational symmetry, using the fixed-vertices generation method?

For Q1, we expected vertical symmetry (V) will be most prominent, with horizontal (H) being noticed more than translational (T), and that non-rotated versions will be easier to detect than the corresponding rotated versions [3, 24, 29].

For Q2 and Q3 we expected that the higher the degree, the greater the extent of symmetry recognition [10]. We were unsure about the effect of the additional visual clutter inherent in the RC drawings (as the size of graph increases with the degree) but anticipated that both should follow the same trend.

4.1 Stimuli

Each experiment uses five base graphs. Specifically, we generated 20 random simple graphs, called base graphs (each with 10 vertices and 11 edges), and drew them using a random layout. We randomly chose five of these base graphs as generators for the stimuli in each experiment.

Reflective and Translational (RT): There are 7 conditions in this experiment: horizontal (H), vertical (V), translational (T), horizontal rotated (Hr), vertical rotated (Vr), translational rotated (Tr), and non-symmetric (NS). In this experiment, all drawings are of the same graph; they are all based on the same base component. With 5 different versions of the base graph, we have $5 * 7 = 35$ stimuli for this experiment.

Rotational Fixed Component (RFC): There are 7 conditions in this experiment: rotational orders of 4–10 (designated as RC4, . . . , RC10). Since the number of incidences of the base component drawing is increased for each order, and two edges are added every time a new base component is added, the size of the graph increases with every order. With 5 different base graphs, we have $5 * 7 = 35$ stimuli for this experiment.

Rotational Fixed Vertices (RFV): There are 7 conditions in this experiment: rotational orders of 4–10 (designated as RV4, . . . , RV10). Here, the number of vertices in each graph varies from 40 to 50, as described in Table 1. With 5 different versions of the base graph, we have $5 * 7 = 35$ stimuli for this experiment.

5 Experimental Process

We use a “two-alternative forced choice” methodology, where a pair of stimuli are presented and participants must choose one of them. Specifically, the participants are asked to select the drawing that they considered “more symmetric.” In each experiment, we show all possible pairs twice (switching between left and right for the second presentation); with 7 stimuli we get $2 * (7 * 6) / 2 = 42$ pairs and this is done for each of the 5 different versions of the base graph for a total of $5 * 42 = 210$ trials. As a within-participants’ experiment, all participants see all 210 pairs. To mitigate against the learning effect, the experimental stimuli are preceded by 10 practice trials for which the data is not collected (using a sixth version of the base graph), and each participant is presented with the 210 trials in a different random order. We collected data on the choice made by the participant for each pair, and the time taken to make the choice.

The experiment is conducted online with the online system randomly assigning one of the three experiments, and randomly selecting 5 base graphs from the 20 available for that participant-experiment combination. By asking each participant to do only one of the three experiments and by choosing only five versions of the base graph, we anticipated that the experiment would not be too lengthy (therefore minimizing the chance of participants not finishing the experiment). We expected the experiment to take approximately 10 min.

Participants are required to give consent at the start, and a set of instructions followed before the practice trials began. A self-timed break is offered every 20 trials. At the end of the experiment, participants are asked to give demographic data: gender, age, educational background, familiarity with networks and with symmetries. As a reward for taking part of the study, statistics about the participant’s answers in relation to the answers of other participants are presented. These statistics show an example of the layout used in the task, the number of selections from the participants and from the specific participant grouped by each version, the number of clicks for any pair of versions, the number of left and right clicks, and the average answer time for the current and all participants.

Participants are not given details about the problem we are considering or the task that they would perform and we are intentionally vague: we just ask

them to “*choose the layout that looks more symmetric*” between the proposed pair of layouts. We do not provide information about the concept of symmetry or our interpretations thereof, so as to not accidentally influence the participants.

5.1 Pilot Study

We use a simple online system made of four main parts: an introduction page, the main experiment, a demographics page (gender, age, graph expertise, symmetry expertise), and the statistics page. We conducted a pilot study to test our setup. In its early stages our system had an introductory screen with a brief description about the concept of “*graphs*” and “*graph drawings*” but several participants suggested to use “*network*” instead, as it is more common and easier to understand. We also added information about the expected duration of the experiment. Initially we showed pairs of graphs starting from 10 base graphs for a total of 420 pairs, but the pilot study showed that this resulted in a test that was too long. We decided to reduce the number of base graphs to 5 and to introduce a break every 20 pairs. The demographics page was augmented to allow for feedback from the participants. The end of the study was modified from a “Thank you” page to a page that provides statistics about the participants results in comparison to the average previous participants. Finally, following feedback from the pilot study, we changed the background color from white to light-gray and added a black border to highlight the selected image.

6 Data Analysis

We used Reddit [22] and personal communications to crowdsource our study. We collected data from a total of 97 participants.

6.1 Experimental Conduct

We removed data from all participants who did not complete the full experiment ($n = 39$). It’s disappointing that so many people did not finish experiment (which was intended to take less than ten minutes); informal feedback from participants are that they found the task boring and, in some cases, difficult. One participant in the RT experiment gave an exceptionally high number of votes to the non-symmetric drawing; we removed this participant from our analysis. Two response times were particularly higher than others (17s and 100s) - these data points were replaced by the mean of the other response times for the respective participants. One participant gave almost exactly the same votes to all conditions in the RFV experiment: we removed this participant’s data. This left us with 19 participants for RT, 19 for RFC and 18 for RFV. The demographic information of the 56 participants (13 female and 43 male) is summarized in Fig. 5.

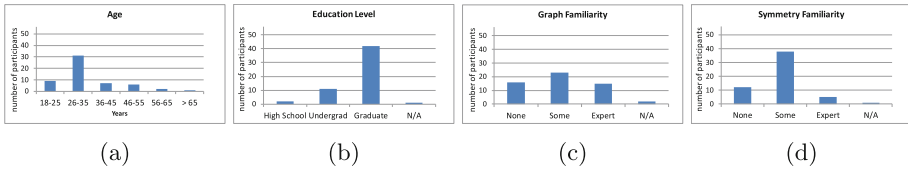


Fig. 5. Participants demographic information (a) age, (b) education level, (c) graph familiarity, and (d) symmetry familiarity.

6.2 Analysis Process

We (conveniently) have seven conditions for each of the three experiments, so the form of the data is the same. Since we include only those participants who completed the entire experiment, each participant “voted” 210 times, each vote being associated with one of the seven conditions. We also have an average response time associated with each vote. For both response times and votes, we use ANOVA and adjusted planned comparison pair-wise tests between the pairs of conditions of interests to determine which conditions are (a) favored over the others, and (b) responded to most quickly. We use a significance level of 0.05 throughout, adjusted as appropriate for the number of planned comparisons made. We do not compare all pairs of conditions, only those of interest to our research question as doing this reduces the extent of required adjustments. The mean vote and mean response time (over all participants) charts are depicted in Fig. 6, while the exact values are shown in Table 2.

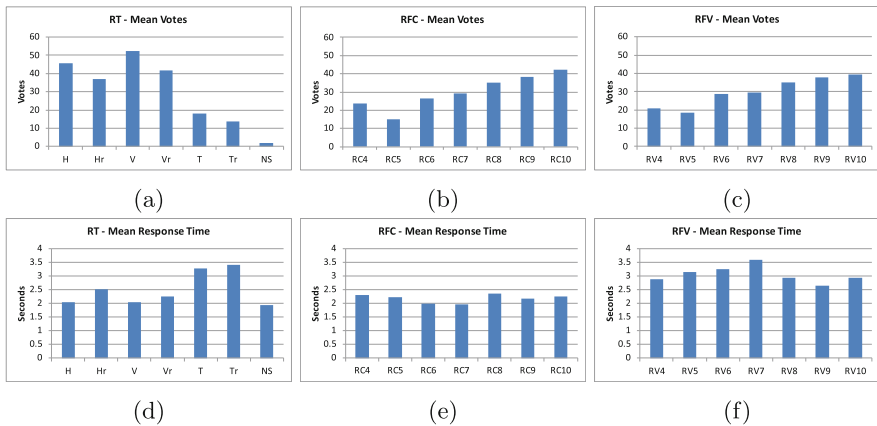


Fig. 6. Mean votes charts for (a) RT (b) RFC, (c) RFV and mean response time charts for (d) RT (e) RFC, and (f) RFV.

Voting Responses

Reflective and Translational (RT): A repeated-measures ANOVA reveals a significant difference between the average votes ($F = 240.5, df = 6, p < 0.001$). Five pairwise comparisons (at adjusted $p = 0.01$), reveal the following results;

- H obtained significantly more votes than Hr ($p < 0.001$)
- V obtained significantly more votes than Vr ($p < 0.001$)
- T obtained significantly more votes than Tr ($p < 0.001$)
- V obtained significantly more votes than H ($p < 0.001$)
- H obtained significantly more votes than T ($p < 0.001$)

Rotational Fixed Component (RFC): A repeated-measures ANOVA reveals a significant difference between the average votes ($F = 12.2, df = 6, p < 0.001$). Two pairwise comparisons (at adjusted $p = 0.025$), reveal the following results;

- RC6 obtained significantly more votes than RC5 ($p < 0.001$)
- No significant difference is between the votes for RC4 and RC5 ($p = 0.050$)

Rotational Fixed Vertices (RFV): A repeated-measures ANOVA reveals a significant difference between the average votes ($F = 10.9, df = 6, p < 0.001$). Two pairwise comparisons (at adjusted $p = 0.025$), reveal the following results;

- RV6 obtained significantly more votes than RV5 ($p = 0.021$)
- No significant difference is between the votes for RV4 and RV5 ($p = 0.63$)

Table 2. Mean vote and mean response time for each task.

	Number of participants	H	Hr	V	Vr	T	Tr	NS
Mean vote	19	45.63	36.89	52.16	41.74	18	13.63	1.95
Mean response time (s)	19	2.05	2.53	2.03	2.25	3.27	3.41	1.93
		RC4	RC5	RC6	RC7	RC8	RC9	RC10
Mean vote	19	23.63	15.11	26.42	29.05	35.21	38.32	42.26
Mean response time (s)	19	2.30	2.21	1.98	1.96	2.36	2.18	2.24
		RV4	RV5	RV6	RV7	RV8	RV9	RV10
Mean vote	18	20.78	18.67	28.78	29.44	35.22	37.72	39.39
Mean response time (s)	18	2.89	3.15	3.25	3.58	2.94	2.65	2.93

Response Time

Reflective and Translational (RT): A repeated-measures ANOVA reveals a significant difference between the average response time ($F = 4.84, df = 6, p < 0.001$). Five pairwise comparisons (at adjusted $p = 0.01$), reveal;

- The response time for H is significantly less than that for T ($p < 0.001$)
- No significant differences between the response times for the following pairings: H/Hr; V/Vr; T/Tr; V/H.

Rotational Fixed Component (RFC): A repeated-measures ANOVA reveals no significant difference between the average response time ($F = 1.29$, $df = 6$, $p = 0.267$).

Rotational Fixed Vertices (RFV): A repeated-measures ANOVA reveals no significant difference between the average response time ($F = 1.535$, $df = 6$, $p = 0.174$).

7 Discussion

With respect to question **Q1**, we found statistically significant effects confirming that vertical symmetry is more recognizable as symmetry, followed by horizontal and translational. In all three variants adding even slight rotation has a significant effect. Vertical symmetry is more recognizable than vertical with rotation, horizontal is more recognizable than horizontal with rotation, and translational is more recognizable than translational with rotation. Mean response time also follows this trend, although we have only one statistically significant finding that horizontal is faster than translational. An immediate implication of this is that vertical symmetry is the best perceived because it is frequently seen by people and that can be exploited in future graph layout algorithms.

The RFC and RFV experiments helped us answer questions **Q2** and **Q3**. Our experimental results provide evidence of a greater symmetry recognition for high number of rotation axes. We believe that this is due to the fact that layouts with high number of rotational axes tend to become more and more circular, and the circle is a very symmetric shape. We also find that the increased size of the graphs in the RFC experiments does not seem to affect the better perception for high number of rotational axes as the results for RFC and RFV are similar.

Of particular interest is one exception: RC4 is considered more symmetric than the RC5, which goes against the general trend of better perception for high number of rotation axes. We discuss possible explanations below.

8 Conclusions and Future Work

The conclusions from our experiments are limited by the specifics of our study; e.g., we only consider graphs of similar sizes, starting with a base graph of a fixed size, and connecting copies thereof in symmetric sub-structures using two or three edges. Despite such limitations, our experiment does provide potentially useful information about the relative effects of different types of symmetries in drawings of graphs. The results from our study suggest that humans recognize vertical reflective symmetry over all other types of symmetry, followed by horizontal and translational symmetries and that rotational symmetry is affected by the number of radial axes. These findings can help guide algorithms that identify features to be displayed using these types of symmetries. Vertical symmetry can be used to call attention to isomorphic subgraphs and cycles can be highlighted by laying them out as regular n -gons that have high rotational symmetry.

The results of our experiment, in particular for the RFC and RFV tasks, show that a rotationally symmetric layout with 4 axes is considered more symmetric than that with 5 axes, which goes against the general tendency of recognizability towards higher number of axes. This leads to an interesting question: what is it about these rotationally symmetric layouts that gives such different results? Is it because of something in particular about the layout with 5 axes or the one with 4? Is it because rotational 4 is not perceived as a rotational symmetry but as a combination of horizontal and vertical symmetries? If so, is the rotational symmetry with many axes perceived better than the reflective and translational ones? Answers to such questions can help guide the design of algorithms that visualize symmetries. It would also be interesting to repeat this study using different base graphs that were not drawn with random layout.

Acknowledgements. We thank Eric Welch for help with the experiment, Hang Chen for help with statistics, and yFiles for providing their software for research and education. This work is supported by NSF grants CCF-1423411 and CCF-1712119.

References

1. Bennett, C., Ryall, J., Spalteholz, L., Gooch, A.: The aesthetics of graph visualization. *Comput. Aesthet.* **2007**, 57–64 (2007)
2. Bruce, V., Green, P.R., Georgeson, M.A.: *Visual Perception: Physiology, Psychology, & Ecology*. Psychology Press, Hove (2003)
3. Bruce, V.G., Morgan, M.J.: Violations of symmetry and repetition in visual patterns. *Perception* **4**(3), 239–249 (1975)
4. Bruckdorfer, T., Kaufmann, M.: Mad at edge crossings? break the edges!. In: Kranakis, E., Krizanc, D., Luccio, F. (eds.) *FUN 2012*. LNCS, vol. 7288, pp. 40–50. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-30347-0_7
5. Burch, M., Vehlow, C., Konevtsova, N., Weiskopf, D.: Evaluating partially drawn links for directed graph edges. In: van Kreveld, M., Speckmann, B. (eds.) *GD 2011*. LNCS, vol. 7034, pp. 226–237. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-25878-7_22
6. Cattaneo, Z., Bona, S., Silvanto, J.: Not all visual symmetry is equal: partially distinct neural bases for vertical and horizontal symmetry. *Neuropsychologia* **104**, 126–132 (2017)
7. Corballis, M.C., Roldan, C.E.: On the perception of symmetrical and repeated patterns. *Percept. Psychophys.* **16**(1), 136–142 (1974)
8. Eades, P., Hong, S.H.: Symmetric graph drawing. In: Tamassia, R. (ed.) *Handbook on Graph Drawing and Visualization*, pp. 87–113. Chapman and Hall/CRC, Boca Raton (2013)
9. Giannouli, V.: Visual symmetry perception. *Encephalos* **50**, 31–42 (2013)
10. Jennings, B.J., Kingdom, F.A.A.: Searching for radial symmetry. *i-Perception* **8**(4), 2041669517725758 (2017)
11. Klapaukh, R.: An empirical evaluation of force-directed graph layout. Ph.D. thesis, Victoria University of Wellington (2014)
12. Lai, H.C., Chien, S.H.L., Kuo, W.Y.: Visual short-term memory for abstract patterns: effects of symmetry, element connectedness, and probe quadrant. *J. Vis.* **9**(8), 593 (2009)

13. Lemon, K., Allen, E.B., Carver, J.C., Bradshaw, G.L.: An empirical study of the effects of gestalt principles on diagram understandability. In: 2007 First International Symposium on Empirical Software Engineering and Measurement. ESEM 2007, pp. 156–165. IEEE (2007)
14. Marriott, K., Purchase, H., Wybrow, M., Goncu, C.: Memorability of visual features in network diagrams. *IEEE Trans. Visual. Comput. Graph.* **18**(12), 2477–2485 (2012)
15. Nesbitt, K.V., Friedrich, C.: Applying gestalt principles to animated visualizations of network data. In: 2002 6th International Conference on Information Visualisation, pp. 737–743. IEEE (2002)
16. Norman, D.A.: *Emotion Design: Why We Love (or Hate) Everyday Things* (2004)
17. Palmer, S.E.: Theoretical approaches to vision. In: *Vision Science: Photons to Phenomenology*, pp. 45–92 (1999)
18. Palmer, S.E., Hemenway, K.: Orientation and symmetry: effects of multiple, rotational, and near symmetries. *J. Exp. Psychol. Hum. Percept. Perform.* **4**(4), 691–702 (1978)
19. Purchase, H.: Metrics for graph drawing aesthetics. *J. Vis. Lang. Comput.* **13**(5), 501–516 (2002)
20. Purchase, H.: Which aesthetic has the greatest effect on human understanding? In: DiBattista, G. (ed.) *GD 1997. LNCS*, vol. 1353, pp. 248–261. Springer, Heidelberg (1997). <https://doi.org/10.1007/3-540-63938-1.67>
21. Purchase, H.C., Cohen, R.F., James, M.: Validating graph drawing aesthetics. In: Brandenburg, F.J. (ed.) *GD 1995. LNCS*, vol. 1027, pp. 435–446. Springer, Heidelberg (1996). <https://doi.org/10.1007/BFb0021827>
22. Reddit: The front page of the internet. <https://www.reddit.com>
23. Rock, I., Palmer, S.: The legacy of gestalt psychology. *Sci. Am.* **263**(6), 84–91 (1990)
24. Royer, F.L.: Detection of symmetry. *J. Exp. Psychol. Hum. Percept. Perform.* **7**(6), 1186–1210 (1981)
25. Rusu, A., Fabian, A.J., Jianu, R., Rusu, A.: Using the gestalt principle of closure to alleviate the edge crossing problem in graph drawings. In: 2011 15th International Conference on Information Visualisation, pp. 488–493. IEEE (2011)
26. Schnore, M., Partington, J.: Immediate memory for visual patterns: symmetry and amount of information. *Psychon. Sci.* **8**(10), 421–422 (1967)
27. Sun, D., Wong, K.: On evaluating the layout of UML class diagrams for program comprehension. In: 2005 Proceedings of 13th International Workshop on Program Comprehension. IWPC 2005, pp. 317–326. IEEE (2005)
28. Tversky, B.: Visuospatial reasoning. In: *The Cambridge Handbook of Thinking and Reasoning*, pp. 209–240 (2005)
29. Wagemans, J.: Detection of visual symmetries. *Spat. Vis.* **9**(1), 9–32 (1995)
30. Ware, C.: *Information Visualization: Perception for Design*. Elsevier, Amsterdam (2004)
31. Welch, E., Kobourov, S.: *Measuring symmetry in drawings of graphs*, vol. 36, pp. 341–351. Wiley (2017)
32. yWorks: The diagramming company. <https://www.yworks.com>