Interactions between constituent single symmetries in multiple symmetry

Matthias Sebastian Treder · Gert van der Vloed ·
Peter A. van der Helm

Abstract As a rule, the discriminability of multiple symmetries from random patterns increases with the number of symmetry axes, but this number does not seem to be the only determinant. In particular, multiple symmetries with orthogonal axes seem better discriminable than multiple symmetries with nonorthogonal axes. In six experiments on imperfect two-fold symmetry, we investigated whether this is due to extra structure in the form of so-called *correlation rectangles*, which arise only in the case of orthogonal axes, or to the relative orientation of the axes as such. The results suggest that correlation rectangles are not perceptually relevant and that the percept of a multiple symmetry results from an orientation-dependent interaction between the constituent single symmetries. The results can be accounted for by a model involving the analysis of symmetry at all orientations, smoothing (averaging over neighboring orientations), and extraction of peaks.

M. S. Treder (✉) · P. A. van der Helm
Donders Institute for Brain, Cognition and Behaviour,
Radboud University Nijmegen,
Nijmegen, The Netherlands
e-mail: matthias.treder@tu-berlin.de

M. S. Treder
Machine Learning Laboratory, Berlin Institute of Technology,
Berlin, Germany

G. van der Vloed
Department of Human Technology Interaction,
Eindhoven University of Technology,
Eindhoven, The Netherlands

Detection of mirror symmetry (henceforth, simply *symmetry*) is believed to be an integral part of the perceptual organization process that is applied to any visual input. A large body of empirical studies have shown that human symmetry perception is quick, versatile, and resistant to noise and spatial jitter (for overviews, see Treder, 2010; Tyler, 1996; van der Helm & Leeuwenberg, 1996, 2004; Wagemans, 1995). Furthermore, symmetry processing was shown to interact with other factors in perceptual organization (e.g., stereo processing; Treder & van der Helm, 2007). Most empirical studies on symmetry perception have focused on one-fold symmetry alone; some have included repetition and Glass patterns; but only a few have focused on multiple symmetry. Yet, insight into all of these visual regularities is necessary to build general theories of perceptual organization. The aim of this study, therefore, is to provide more insight into the mechanisms underlying multiple-symmetry perception.

To this end, we investigated how the single symmetries in two-fold symmetry interact, and we modeled the interactions we found in a neurophysiologically plausible way. Although this study focuses on two-fold symmetry, its scope extends to multiple symmetry with any number of symmetry axes. This is illustrated by indicating how our research question was triggered by the few existing studies on multiple symmetry.

Most empirical research on multiple symmetry has involved two- and four-fold symmetry—that is, curiously, not three-fold symmetry. For one-fold, two-fold, and four-fold symmetry, data have consistently shown that

salience increases with the number of symmetry axes (e.g., Nucci & Wagemans, 2007; Palmer & Hemenway, 1978; Wagemans, Gool, & d'Ydewalle, 1991; Wenderoth, 1997). To explain this, one needs hardly more than the idea that the probability of detecting a symmetry axis increases with the number of symmetry axes (cf. Corbalis & Roldan, 1974; Palmer & Hemenway, 1978). But what about three-fold symmetry, which, in contrast to two- and four-fold symmetry, has only nonorthogonal symmetry axes?

The only three empirical studies that we know of on three-fold symmetry have been indecisive on this point. Hamada and Ishihara's (1988) data suggest that three-fold symmetry behaves like three-fold rotational symmetry (i.e., with three local mirror symmetries but without global mirror symmetries). Furthermore, for *n*-fold symmetry where *n* = 1, 2, 3, or 4, van der Vloed (2005) found that salience increases linearly with *n*, but Wenderoth and Welsh (1998) found that three-fold symmetry is not more salient than two-fold symmetry, but rather tends to be less salient. Models of one-fold symmetry perception that have been claimed to extend to multiple symmetry have also disagreed on this point, as we discuss next.

First, the so-called *transformational approach* (Garner, 1974; Palmer, 1983) models the salience of stimuli by the amount of regularity in a stimulus—namely, by the number of rigid transformations under which a stimulus is invariant. For instance, a one-fold symmetry is invariant under a 180° three-dimensional rotation about the symmetry axis. Likewise, an *n*-fold symmetry is invariant under *n* such three-dimensional rotations plus *n* two-dimensional rotations of $360°/n$ in the image plane. In other words, according to the transformational approach, the structural redundancy in a multiple symmetry increases with the number of symmetry axes. Accordingly, the transformational approach suggests that the salience of multiple symmetry, including three-fold symmetry, simply increases with the number of symmetry axes.

Second, the so-called *holographic approach* (van der Helm & Leeuwenberg, 1991, 1996, 1999, 2004) builds on the idea that symmetry allows for efficient stimulus representations (see, e.g., Attneave, 1954, 1955; Barlow, 1961, 2001; Leeuwenberg, 1969, 1971), and it models the salience of stimuli by the amount of regularity in such stimulus representations. In terms of symbol strings, this may be illustrated by the compression of the symmetrical string abcdeffedcba into the symmetry code S[(a)(b)(c)(d)(e)(f)]. Likewise, the double symmetry in the string abccbaabccba can be represented by a hierarchical combination of two symmetries. That is, the string can be encoded into the code S[(a)(b)(c)(c)(b)(a)], which can be encoded further to yield the code S[S[((a))((b))((c))]]. Hence, this code captures all symmetry by combining two separate one-fold symmetries, which suggests that both

symmetries add to the total salience. For the triple symmetry in the string abbaabbaabba, however, this coding approach implies that, after having captured the global symmetry in the code S[(a)(b)(b)(a)(a)(b)], only one local symmetry can be captured, yielding either S[S[((a))((b))] (a)(b)] or S[(a)(b) S[((b))((a))]].

Translated to 2-D multiple symmetries, the foregoing implies that the two extra symmetries in three-fold symmetry are predicted to add less to the total salience than does the single extra symmetry in two-fold symmetry (van der Helm & Leeuwenberg, 1996). The applied coding principles suggest that this is due to the relative orientation of the symmetry axes. That is, in case of three-fold symmetry (with nonorthogonal symmetries), the applied coding principles imply that capturing one of the two extra symmetries disables capturing the other or, in other words, that nonorthogonal symmetries impede each other. Here, this impediment is to be understood in representational coding terms, but one of the questions in this article is whether it is paralleled by an impediment in neurophysiological terms.

Third, in contrast to the holographic approach, the so-called *bootstrap approach* (Wagemans et al., 1991; Wagemans, Gool, Swinnen, & Horebeek, 1993) suggests mutual facilitation in the case of orthogonal symmetry axes (i.e., not mutual impediment in the case of nonorthogonal axes). This may be explicated as follows. In one-fold symmetry, two symmetrically corresponding points can be connected by a virtual line, and such virtual lines can in turn be joined to form virtual trapezoids. Distortion of these so-called lower- and higher-order structures has been shown to impair symmetry detection (Jenkins, 1983; Sawada & Pizlo, 2008; van der Vloed, Csathó, & van der Helm, 2005; Wagemans et al., 1991; Wagemans et al., 1993). Notice that the normals to the parallel virtual lines in a virtual trapezoid indicate the two directions in which the propagation process may continue—that is, the directions in which subsequent propagation steps should search for additional symmetry pairs (see Fig. 1a). In two-fold symmetry, however, some trapezoids are rectangles, indicating four propagation directions and implying that twice as many symmetry pairs can be added during each subsequent propagation step (see Fig. 1b). The correlation rectangles are therefore proposed to facilitate two-fold symmetry perception. Such correlation rectangles do not occur in three-fold symmetry, which therefore cannot benefit from their facilitatory effect.

Notice that, in the terms of Marr (1982), the three just-discussed approaches to multiple-symmetry perception are at the computational and algorithmic levels of description. In the General Discussion, we complement this type of approach by including the implementational level of description—that is, by casting the just-discussed facilita-

Fig. 1 Symmetry detection by way of bootstrapping in one-fold and two-fold symmetry. (**a**) In one-fold symmetry, symmetry pairs yield virtual lines (dashed lines) that join to form a correlation trapezoid (solid lines). These pairwise correlations suggest one propagation axis (up/down) along which additional symmetry pairs should be searched for during subsequent propagation steps. (**b**) In two-fold symmetry, two virtual lines (dashed lines) join to form a correlation rectangle (solid lines). These quartet-wise correlations suggest two propagation axes (up/down and right/left). This allows for the addition of twice as many symmetry pairs in each subsequent propagation step
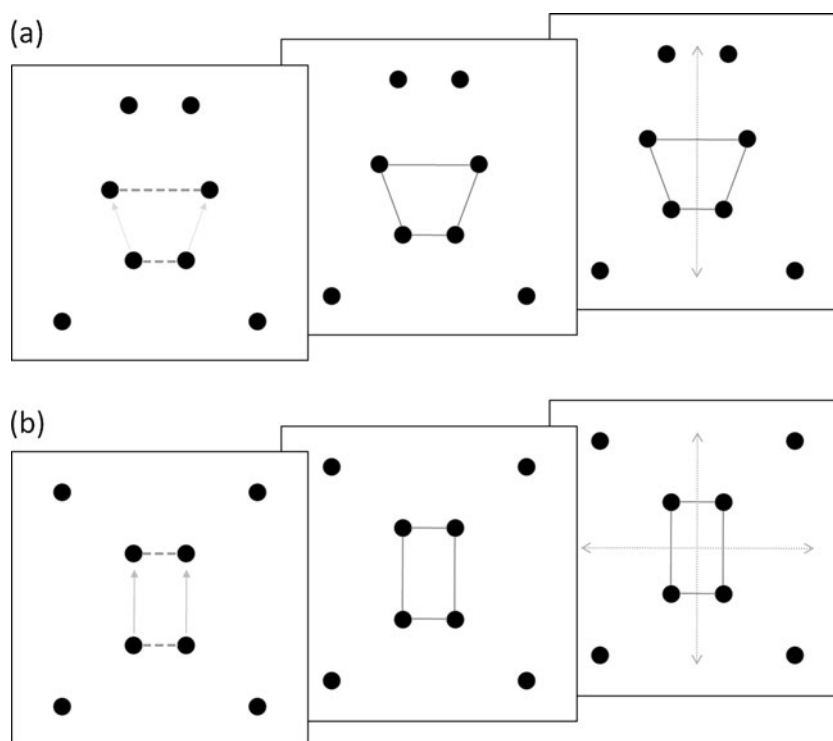
tion and impediment in terms of neural mechanisms. In any case, the foregoing discussion shows that the existing literature suggests that the number of single symmetries in a multiple symmetry might not be the only factor to be reckoned with. To investigate the interaction between the constituent single symmetries in more detail, we conducted two triples of experiments on imperfect two-fold symmetry.

Imperfect multiple symmetries abound in both nature and art (e.g., in flowers and band patterns), and our study does seem relevant to biology and art science (van der Helm, 2011), but the main reason we considered imperfect multiple symmetry is that we aimed to investigate properties of the human visual system. One of the methods to this end is to probe the detectability of symmetry in the presence of noise. Importantly, in the case of multiple symmetry, this allows for a decoupling of the number of symmetry axes, the presence of correlation rectangles, and the relative axis orientation—factors that cannot be decoupled in perfect multiple symmetry because, in the case of perfect symmetry, one factor dictates the others.

In the first triple of experiments, we focused on the idea that correlation rectangles, which occur only with orthogonal symmetry axes, facilitate two-fold symmetry perception. This idea is plausible enough to be taken seriously, but it had never been tested. In the second triple of experiments, we used only symmetries without correlation rectangles to investigate the more general characteristic of the relative orientation of symmetry axes.

Correlation rectangles

In Experiments 1a, 1b, and 1c, we examined whether correlation rectangles indeed have the facilitating effect on two-fold symmetry perception proposed in the bootstrap approach. To prevent the effect of correlation rectangles being contaminated with the effects of the number of symmetry axes or of relative axis orientation, we kept the latter factors constant. To this end, we focused on imperfect two-fold symmetries that, by means of various noise manipulations, had 50% noise about each symmetry axis. Humans are well able to detect such imperfect symmetries, and if correlation rectangles play a genuine perceptual role, they should also be effective in imperfect symmetries. Because the experimental setup was largely identical across the three experiments, we outline the general method below.

General method

Participants

A total of 21 undergraduate students participated in Experiment 1a, 19 other undergraduate students participated in Experiment 1b, and 37 undergraduate students participated in Experiment 1c. Two of the latter students had participated in one of the previous experiments more than 1 month before. All participants had normal or corrected-to-normal vision and received course credits or money for their participation.
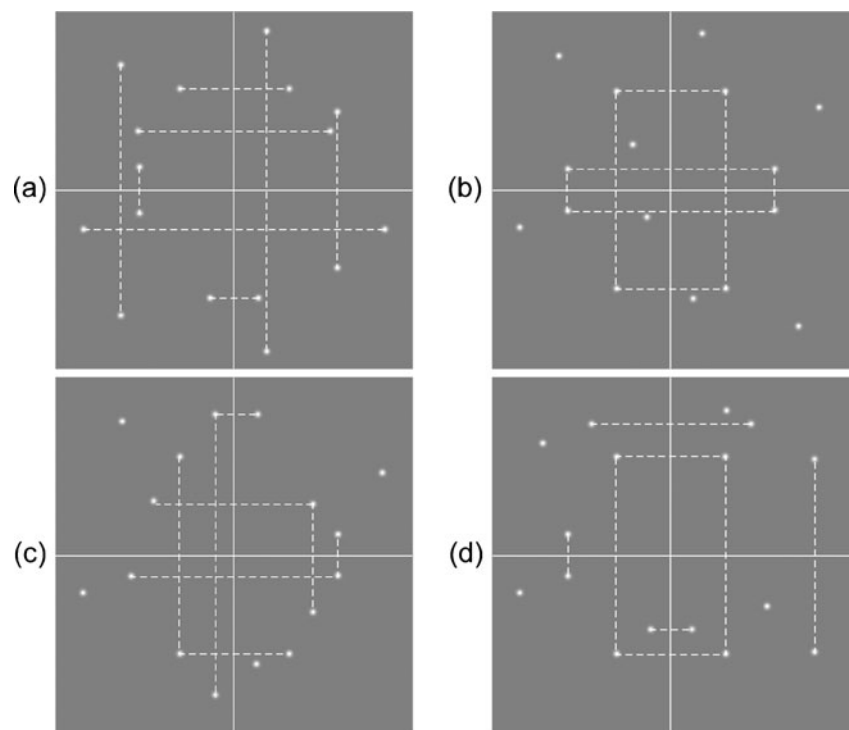
Stimuli

Each stimulus comprised 80 nonoverlapping low-pass elements. The luminance profile of an element was given by the radial Gaussian function $l(d) = 0.5 + 0.5 \cdot e^{-d2/(2\sigma2)}$, where $d \in [0; 16]$ is the distance in pixels from the center of the window and $\sigma = \sqrt{10}$. The stimuli were constructed such that they formed various two-fold symmetries featuring a vertical and a horizontal symmetry axis. Additionally, in all three experiments, we used one-fold symmetries as baseline trials and random patterns as catch trials.

More specifically, in Experiment 1a, we tested the effect of correlation rectangles by comparing the discriminability (from random patterns) of two-fold symmetries in the *no-rectangles* and *rectangles conditions* (see Fig. 2a and b). In the no-rectangles condition, the symmetries did not feature correlation rectangles; to this end, 50% of the elements were placed symmetrically about the horizontal axis alone, and the other 50% were placed symmetrically about the vertical axis alone. In the rectangles condition, the symmetries featured correlation rectangles; to this end, 50% of the elements were placed as quartets (i.e., elements were symmetrical about both axes), and the other 50% were randomly distributed so that they constituted global noise (i.e., elements that were symmetrical about neither axis). Hence, to each symmetry axis in each condition, 50% of the elements constituted symmetry pairs, and the other 50% of the elements constituted noise.

In Experiment 1b, to test the effect that global noise may have had in the rectangles condition, we replaced the rectangles condition by the *L condition* (see Fig. 2c). In the L condition, 75% of the elements were part of L-shaped triplets produced by first placing a symmetrical dot quartet (e.g., four dots symmetrical about both axes) and then removing one dot from the quartet so that the remaining three dots formed a virtual L instead of a rectangle. The other 25% of the elements constituted global noise. Consequently, neither the no-rectangles condition nor the L condition featured correlation rectangles. Furthermore, both kinds of two-fold symmetries contained 50% noise about either symmetry axis, but they differed in the proportion of global noise—namely, 0% in the no-rectangles condition and 25% in the L-condition.

In Experiment 1c, we further tested the above-mentioned factors by comparing the L condition and the *25%-rectangles condition* (see Fig. 2d). In the 25%-rectangles condition, 25% of the elements were placed in symmetrical dot quartets (i.e., elements were symmetrical about both axes), 25% were symmetrical about the horizontal axis only, 25% were symmetrical about the vertical axis only, and 25% constituted global noise. Hence, these two conditions were matched in terms of the amount of global noise (25%) and the amount of noise for each symmetry axis (50%). However, the L condition featured no correlation rectangles, whereas in the 25%-rectangles condition, 25% of the elements were part of correlation rectangles.

Fig. 2 Schematic overview of the two-fold symmetry stimulus manipulations in Experiments 1a, 1b, and 1c. For simplicity, each display shows only 16 elements (the actual stimuli contained 80 elements). The solid lines indicate the symmetry axes, and the dashed lines indicate symmetry pairs (these lines were not present in the actual stimuli). For details on the stimuli, see the General Method section. (**a**) No-rectangles condition. (**b**) Rectangles condition. (**c**) L condition. (**d**) 25%-rectangles condition

Example stimuli for all of the two-fold symmetry conditions are given in Fig. 2. Furthermore, in the baseline condition, we used one-fold symmetries in which 40 elements were positioned symmetrically about an axis (vertical or horizontal) while the remaining 40 elements were distributed randomly. In the catch trial condition, we used random stimuli in which all elements were randomly distributed. All stimuli were circular and subtended about 16° of visual angle. The absolute axis orientations were horizontal and vertical. The background of the display was set to a mean luminance of 16.64 cd/m^2. Using MATLAB, a unique set of stimuli was generated for each participant.

Procedure

Participants were seated 60 cm in front of a 19-in. monitor with a 100-Hz refresh rate and a resolution of 1,280 × 1,024 pixels. A chinrest was used to restrict head movements. Participants performed a two-alternatives forced choice (2AFC) task in which a symmetrical stimulus (either a one-fold or a two-fold symmetry) and a random stimulus were presented in subsequent intervals in a random order. Participants had to indicate which of the two intervals featured symmetry. The responses were recorded using a button box.

Prior to each experiment, each participant completed a practice block of 60 trials with feedback given after each trial. Each interval was preceded by a central fixation dot presented for 500 ms. During the intervals, stimuli were flashed for 120 ms. The number of experimental trials totaled 4 (symmetry conditions) x 120 (measurements) = 480. The experiments were self-paced.

Experiment 1a

In this experiment, we contrasted two-fold symmetries in which correlation rectangles were present (the rectangles condition) with two-fold symmetries in which correlation rectangles were absent (the no-rectangles condition). The rationale was that, if correlation rectangles facilitate two-fold symmetry perception (as proposed in the bootstrap approach), then the rectangles condition should be more salient than the no-rectangles condition.

Results

For each participant and subcondition, the discriminability measure d' was calculated on the basis of the correct scores on trials wherein the symmetry was presented in one interval, and of the error scores on trials wherein the symmetry was presented in the other interval (e.g., Swets, 1964; Wickens, 2002). Subsequently, d' values were compared using paired-samples t tests. Vertical one-fold

symmetry was more salient than horizontal one-fold symmetry, $t(20) = 3.540$, $p < .01$, and was less salient than two-fold symmetry—both with and without correlation rectangles, $t(20) = 8.066$, $p < .001$, and $t(20) = 10.613$, $p < .001$, respectively. Furthermore, d' on two-fold symmetry tended to be lower in the rectangles condition than in the no-rectangles condition, but this effect was not significant ($p = .084$). The results are depicted in Fig. 3a.

Discussion

In contrast to the predictions of the bootstrap approach, we did not find that two-fold symmetries with correlation rectangles were better detectable than symmetries without. This suggests that correlation rectangles do not play a role in symmetry perception. Note, however, that in the no-rectangles condition, every element was symmetrical about one of the axes, whereas in the rectangles condition, 50% of the elements constituted global noise (i.e., elements that were symmetrical about neither symmetry axis). Because, in principle, a negative effect of global noise could have overshadowed a positive effect of correlation rectangles, we conducted two control experiments to estimate the magnitude of both effects separately.

Experiment 1b

In this experiment, we probed the effect of global noise when correlation rectangles were absent. To this end, we contrasted the no-rectangles condition with the L condition. Neither condition featured correlation rectangles, but the no-rectangles condition does not feature global noise, whereas the L condition did (i.e., 25% of the elements formed global noise).

Results

We again found that vertical one-fold symmetry was more salient than horizontal one-fold symmetry, $t(18) = 3.470$, $p < .01$, and less salient than two-fold symmetry—both in the no-rectangles condition and in the L condition, $t(18) = 7.731$, $p < .001$, and $t(18) = 2.128$, $p < .05$, respectively. Furthermore, d' on two-fold symmetry was significantly lower in the L condition than in the no-rectangles condition, $t(18) = 5.218$, $p < .001$. These results are depicted in Fig. 3b.

Discussion

We found that performance was lower in the L condition than in the no-rectangles condition. Since these conditions
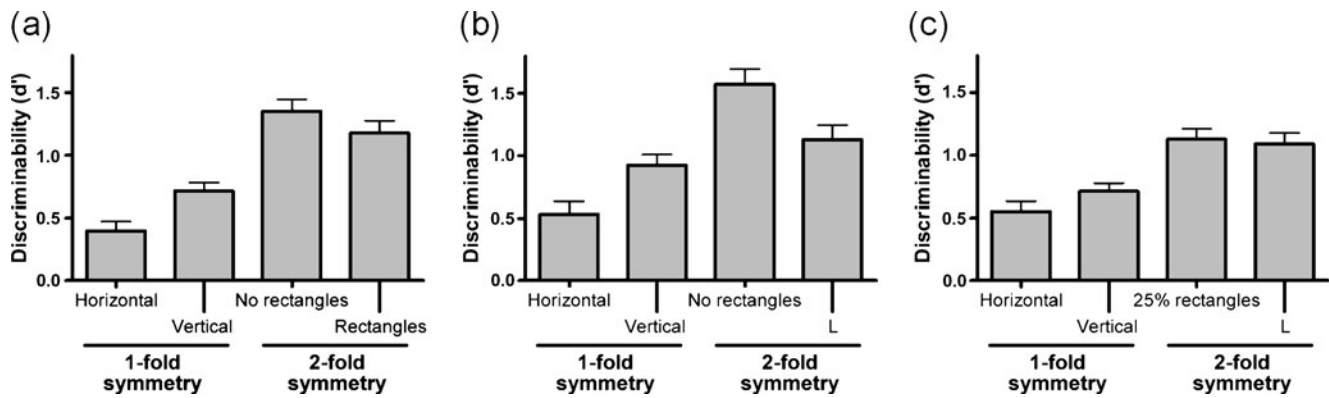
Fig. 3 Results of Experiments 1a, 1b, and 1c. Error bars represent 1 *SEM*. In all experiments, vertical symmetry is more salient than horizontal symmetry, and two-fold symmetry is more salient than one-fold symmetry. (a) Results of Experiment 1a. Performance in the rectangles condition is worse than in the no-rectangles condition. (b) Results of Experiment 1b. Performance in the L condition is worse than in the no-rectangles condition. Because there are no correlation rectangles in both conditions, this difference can be attributed to the global noise that is present in the L condition but not in the no-rectangles condition. (c) Results of Experiment 1c. For a fixed amount of global noise, performance in the L condition does not differ from performance in the 25%-rectangles condition

differed only in the amount of global noise, we may conclude that global noise hampers two-fold symmetry perception more than noise about the axes separately (the question of why this might be the case falls outside the scope of the article, since it would require a whole different series of experiments). This implies that, in Experiment 1a, a negative effect of global noise may indeed have overshadowed a positive effect of correlation rectangles. To control for the effect of global noise, we conducted the following experiment.

Experiment 1c

In this experiment, we probed the effect of correlation rectangles, if any, when the amount of global noise was fixed. To this end, we contrasted the L condition with the 25%-rectangles condition. Both conditions featured 25% global noise, but the L condition did not feature correlation rectangles, whereas the 25%-rectangles condition did (i.e., 25% of the elements formed correlation rectangles). To uncover an effect of correlation rectangles, even if it was small, the number of participants was considerably larger than in the previous experiments.

Results

One participant was excluded from the analysis because of an overall negative *d'*. This time, we did not find that vertical one-fold symmetry was more salient than horizontal one-fold symmetry ($p = .063$), but this was mainly due to 2 participants who apparently focused their attention on the horizontal axis: They performed better for horizontal symmetry than for all other symmetry conditions (removal

of these 2 participants would have yielded $p = .004$, with little effect on the other statistical comparisons). Furthermore, vertical one-fold symmetry was again less salient than two-fold symmetry—in both the 25%-rectangles condition and the L condition, $t(35) = 5.993$, $p < .001$, and $t(35) = 5.478$, $p < .001$, respectively. Crucially, we did not find a difference between the 25%-rectangles condition and the L condition ($p = .502$). These results are depicted in Fig. 3c.

Discussion

In spite of the large number of participants, we did not find a difference between two-fold symmetries with correlation rectangles and two-fold symmetries without. This suggests that the negative tendency in Experiment 1a was solely due to global noise, and it strengthens our conclusion that correlation rectangles do not play a role in multiple-symmetry perception.

Apart from the empirical evidence, there are also theoretical grounds to assume that correlation rectangles add little. Relatively local correlation trapezoids may occur anywhere along either axis, but relatively local correlation rectangles can occur only at the intersection of both axes. The detection process probably uses such relatively local structures as anchors to propagate from, so that the limited presence of relatively local correlation rectangles automatically limits the potential impact of correlation rectangles in general.

All in all, the foregoing suggests that the allegedly special status of orthogonal symmetry axes is not due to correlation rectangles; that is, correlation rectangles do not seem to facilitate two-fold symmetry perception. The question, then, is whether orthogonal symmetry axes have a special status at all. This was examined further in the

triplet of experiments reported next, in which we turned to a more general characteristic of multiple symmetry—namely, the relative orientation of the symmetry axes.

Relative orientations of symmetry axes

In Experiments 2a, 2b, and 2c, we investigated whether the relative orientation of symmetry axes affects symmetry processing. To get a clear picture, the effect of relative orientation would have to be isolated from any effect of the number of symmetry axes and, to be sure, also from a potential effect of correlation rectangles. To this end, we combined two different perfect one-fold symmetries in various relative orientations—in particular, orthogonal (90° relative orientation) and nonorthogonal (45°/135° relative orientation; i.e., one angle between the axes was 45°, and the other was 135°). A cross-correlation analysis of the stimuli showed that, going from coinciding to orthogonal axes, the physical signal-to-noise ratio for each axis drops quickly from 100% at 0° relative orientation to 50% at 10° relative orientation and stays there (see the supplemental materials). Hence, the physical signal-to-noise ratio was the same for the relative orientations we considered in these experiments. Furthermore, because the salience of single and multiple symmetry is known to vary with absolute orientation (Wenderoth & Welsh, 1998), we used absolute axis orientations of +22.5° and −22.5° with respect to the vertical and the horizontal, yielding four roughly equisalient symmetry axis orientations. As we report next, the three experiments differed regarding stimulus type and regarding the way the one-fold symmetries were combined.

Experiment 2a

In this experiment, we used dot stimuli and we composed two-fold symmetries by superimposing two different perfect one-fold symmetries in different relative orientations. This yields two-fold symmetries without correlation rectangles, and with 50% noise about each axis.

Method

Participants A total of 22 undergraduate students participated in the experiment. None of them had participated in any of the previous experiments. All had normal or corrected-to-normal vision and received course credits or money for their participation.

Stimuli To create two-fold symmetries, we superimposed two perfect but different one-fold symmetries of 40 elements each in either orthogonal (90°) or nonorthogonal (45°/135°) relative orientations. To have approximately equisalient symmetry axes across all conditions, we used absolute axis orientations of +22.5° and −22.5° with respect to the vertical and the horizontal.

By way of a baseline, we created one-fold symmetries, in which 40 elements were positioned symmetrically about an axis at one of the four designated absolute orientations, while 40 further elements were distributed randomly. By way of catch trials, we constructed random stimuli in which all elements were distributed randomly. Using MATLAB, a unique set of stimuli was generated for each participant. Examples are given in Fig. 4.

Procedure As before, and using the same apparatus and settings, participants performed a 2AFC task in which a symmetrical stimulus and a random stimulus were presented in subsequent intervals in a random order. Participants had to indicate which of the two intervals featured symmetry. Responses were recorded using a button box. Prior to the experiment, each participant completed a practice block of 60 trials, with feedback given after each trial. Each interval was preceded by a central fixation dot presented for 500 ms. During the intervals, stimuli were flashed for 200 ms. This was longer than in the previous experiments, but increasing the presentation time was necessary to maintain a comparable level of performance, since symmetries with symmetry axes off the four cardinal orientations (vertical, horizontal, left/right oblique) are more difficult to detect. The number of experimental trials totaled 3 (symmetry conditions) x 4 (axis orientations) x 50 (measurements) = 600. The experiment was self-paced.

Results

For the one-fold symmetries, we did not find a difference in salience between the absolute axis orientations. Therefore, for each condition, we pooled across the absolute axis orientations. Subsequently, the d' values obtained for one-fold, nonorthogonal, and orthogonal symmetry were compared using paired-samples t-tests. Nonorthogonal symmetry tended to be more salient than one-fold symmetry, but this was not significant ($p = .065$). Orthogonal symmetry, however, was significantly more salient than one-fold symmetry, $t(21) = 3.694$, $p < .01$. Crucially, orthogonal symmetry was also significantly more salient than nonorthogonal symmetry, $t(21) = 2.525$, $p < .05$. The results are depicted in Fig. 5a.
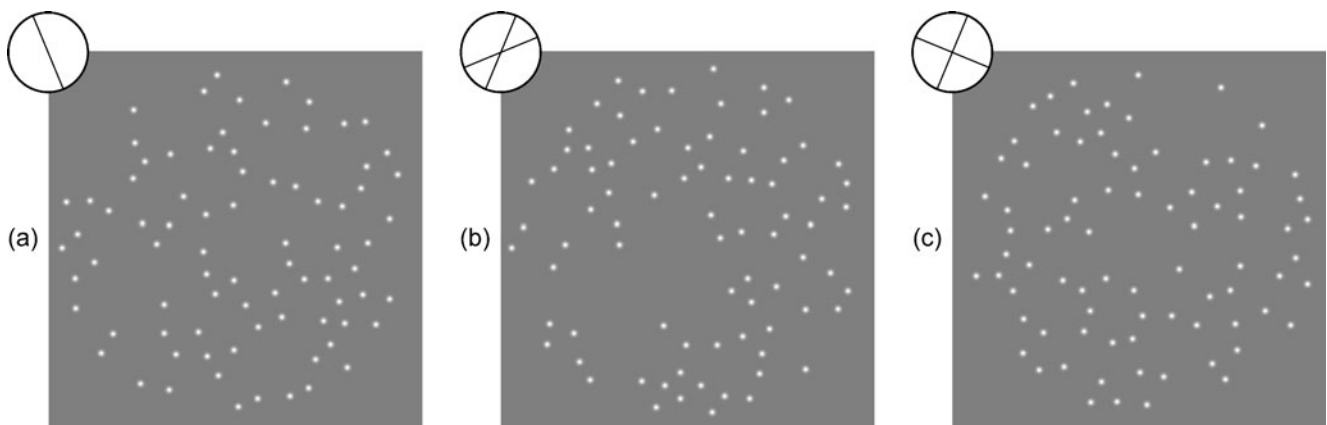
Fig. 4 Examples of the stimuli used in Experiment 2a. In the insets, the absolute orientations of the symmetry axes are indicated by straight lines. (a) One-fold symmetry. (b) Nonorthogonal two-fold symmetry with relative axis orientations of 45°/135°. (c) Orthogonal two-fold symmetry with a relative axis orientation of 90°

Discussion

The results show that two-fold symmetry is more salient when the axes are orthogonal than when the axes are nonorthogonal. Notice that this effect goes against the bootstrap approach, which attributes a positive effect of orthogonality solely to correlation rectangles, which here were absent. Apparently, the relative orientation of the symmetry axes is an independent factor affecting multiple-symmetry perception. To investigate whether this effect of relative orientation generalizes across stimulus types, we conducted the next experiment.

Experiment 2b

In this experiment, we used dense displays, composing two-fold symmetries by superimposing, again in different relative

orientations, two perfect one-fold symmetries that were confined to different spatial-frequency bands. This yields noisy two-fold symmetries, and the two different luminance distributions ensure the absence of correlation rectangles.

Method

Participants A total of 16 undergraduate students participated in this experiment. None of them had participated in any of the previous experiments. All had normal or corrected-to-normal vision and received course credits or money for their participation.

Stimuli All stimuli were constructed by superimposing two images filtered in the spatial-frequency (SF) domain
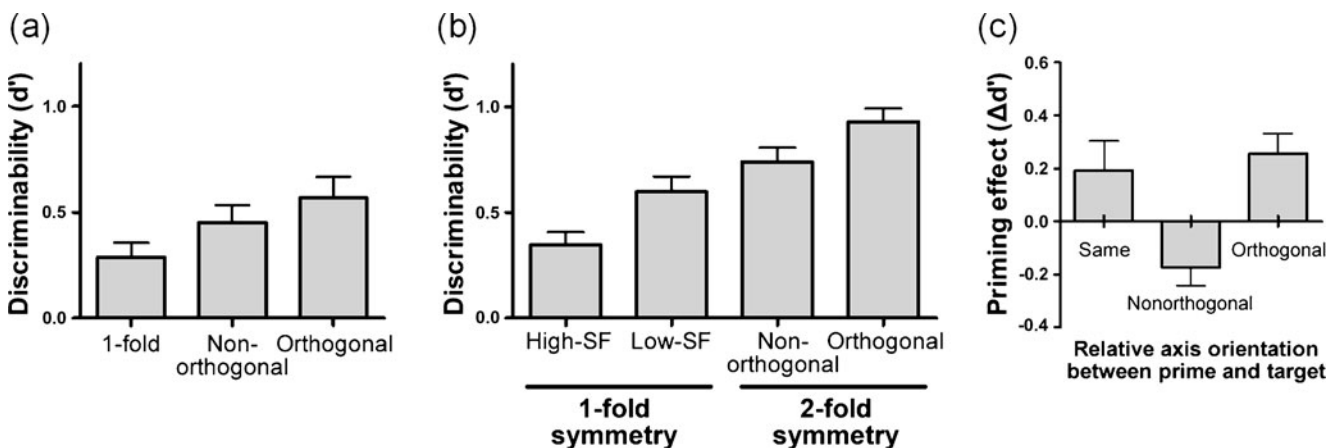

Fig. 5 Results of Experiments 2a, 2b, and 2c. Error bars represent 1 SEM. (a) Results of Experiment 2a. Nonorthogonal two-fold symmetry is more salient than one-fold symmetry but, crucially, less salient than orthogonal two-fold symmetry. (b) Results of Experiment 2b. Low-spatial-frequency (SF) one-fold symmetry is more salient than high-SF one-fold symmetry and less salient than both nonorthogonal and orthogonal two-fold symmetry. Most importantly, orthogonal two-

fold symmetry is more salient than nonorthogonal two-fold symmetry. (c) Priming effects obtained in Experiment 2c. The priming effect was defined by $\Delta d'$, calculated by subtracting d' on random primes from d' on one-fold symmetrical primes. The priming effects were not significant, negative, and positive when the symmetry axes in the prime and target had the same orientation, nonorthogonal orientations, and orthogonal orientations, respectively

(similar filtering techniques had been employed by, e.g., Dakin & Herbert, 1998; Julesz & Chang, 1979; Rainville & Kingdom, 1999, 2000). Each image was generated starting from randomly distributed Gaussian luminance ($\mu = 0$, $\sigma = 0.125$). To obtain a one-fold symmetrical image, an image half was filled with random luminance and reflected about the vertical. Using Fourier transformation, each image was transformed from the spatial domain into the SF domain. Each Fourier-transformed image was then filtered using an idealized isotropic bandpass filter with a bandwidth of one octave. The filter (H) was defined in the SF domain by

$$H = 1 \text{ if } f_{\text{low}} \leq f \leq f_{\text{high}}$$
$$ 0 \text{ otherwise}$$

where f_{low} and f_{high} were the lower and upper cutoff frequencies, respectively. To generate a relatively low-SF image, we selected a one-octave band with lower and higher cutoff frequencies of 0.31 cpd (cycles per degree) and 0.62 cpd, respectively. To generate a relatively high-SF image, we selected a one-octave band with lower and higher cutoff frequencies of 1.26 cpd and 2.52 cpd, respectively. This implies that the one-octave bands were separated by a gap of also one octave. To create a stimulus, low-SF and high-SF images were transformed back into the spatial domain by means of inverse Fourier transformation. The resulting images were rotated to the desired orientations and then superimposed by averaging their luminance at each pixel position. The resulting intensity distribution was shifted to mean luminance and normalized to 25% root-mean-square contrast.

As in Experiment 2a, we used absolute axis orientations of +22.5° and −22.5° with respect to the vertical and the horizontal. The two-fold symmetries were created by combining a near-vertical symmetry in one SF band with a near-horizontal symmetry in the other SF band, with a relative orientation of 90° in the orthogonal condition and orientations of 45°/135° in the nonorthogonal condition. By way of two baseline conditions, we constructed stimuli with one-fold symmetry in either the high SF band or the low SF band; in both cases, the complementary SF band contained noise. Catch trials featured stimuli in which both SF bands contained noise. All stimuli were circular with a diameter of about 10.6° of visual angle against a black background. Using MATLAB, a unique set of stimuli was generated for each participant. Figure 6 gives examples of the stimuli.

Procedure We used the same apparatus as in the previous experiments. Participants performed a signal detection (yes/no) task. In each trial, a symmetrical or random pattern was presented and participants had to indicate whether the pattern was symmetrical (irrespective of the number of symmetry axes that might be present) or random. Responses were recorded using a standard keyboard.

The experiment was preceded by 120 practice trials with feedback given after each trial. Each trial started with a central fixation cross presented for 500 ms. Subsequently, the stimulus was presented for 250 ms. Following the response and an intertrial interval of 100 ms, the next trial commenced automatically. The number of trials totaled 2 (symmetrical/random) x 4 (symmetry conditions) x 4 (axis orientations) x 20 (measurements) = 640. Breaks were given after every 128 trials.

Results

Based on the correct scores for symmetrical targets (hits) and error scores for random targets (false alarms), d' was calculated and used as a dependent variable in a series of paired-samples t-tests. Note that, although a yes/no paradigm was used here and a 2AFC paradigm was used in the previous experiments, signal detection theory shows a linear relationship between d' values obtained in these two ways (Wickens, 2002). This allows for a qualitative comparison of effects across experiments, irrespective of the method used.

For the one-fold symmetries, we did not find a difference in salience between the absolute axis orientations. Therefore, for each condition, we pooled the absolute axis orientations. In line with Julesz and Chang (1979), we found that the salience of low-SF one-fold symmetry was higher than the salience of high-SF one-fold symmetry, $t(15) = 2.826$, $p < .05$. Furthermore, d' for the two-fold symmetry with nonorthogonal axes was higher than the d's for both high-SF one-fold symmetry, $t(15) = 4.925$, $p < .001$, and low-SF one-fold symmetry, $t(15) = 2.467$, $p < .05$. Similarly, d' for the two-fold symmetry with orthogonal axes was higher than d' for both high-SF one-fold symmetry, $t(15) = 9.542$, $p < .001$, and low-SF one-fold symmetry, $t(15) = 5.85$, $p < .001$. Crucially, orthogonal symmetry was again more salient than nonorthogonal symmetry, $t(15) = 5.034$, $p < .001$. The results are depicted in Fig. 5b.

Discussion

Also for this stimulus type, we found that two-fold symmetry is more salient when the axes are orthogonal than when they are nonorthogonal. This demonstrates the robustness of this effect and suggests that it is not a stimulus-dependent artifact but, rather, a general characteristic of symmetry perception. An implication of these
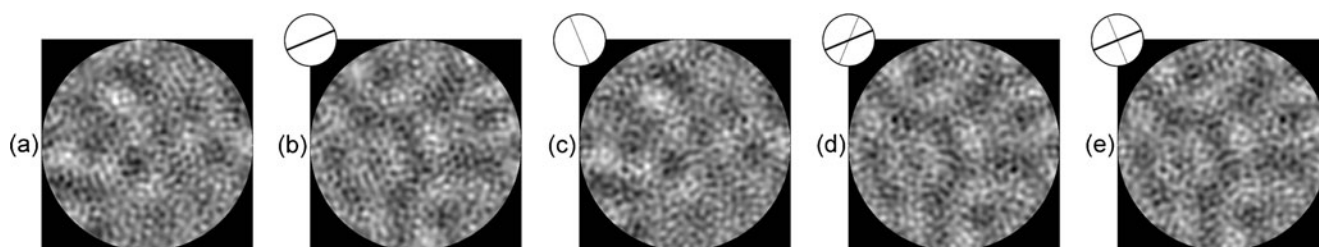
Fig. 6 Examples of the stimuli used in Experiment 2b. In the insets, the absolute orientations of the symmetry axes are indicated by straight lines; thick lines refer to the low spatial-frequency (SF) band, and thin lines refer to the high SF band. (a) Random pattern. (b) One-fold symmetry in the low SF band. (c) One-fold symmetry in the high SF band. (d) Two-fold nonorthogonal symmetry. (e) Two-fold orthogonal symmetry. For illustrative purposes, the low-SF and high-SF symmetries in panels (b) and (c) are the same as the ones that were combined in panels (d) and (e), respectively. Thus, the only difference between the nonorthogonal symmetry in panel (d) and the orthogonal symmetry in panel (e) is the relative orientation between the component symmetry axes

results is that the single symmetries constituting multiple symmetry apparently engage in an orientation-dependent interaction. To shed more light on the underlying mechanism, we conducted the following priming experiment in which we investigated the interaction between one-fold symmetries separated in the time domain (i.e., rather than in the SF domain).

Experiment 2c

The effect revealed in the previous two experiments suggests the involvement of orientation-selective processes in multiple-symmetry perception. It is unclear, however, whether this effect stems from a mutual facilitation of orthogonal symmetry axes, a mutual impediment of nonorthogonal symmetry axes, or both. To investigate this question, we used a paradigm wherein two symmetries were presented in succession, the first one serving as prime and the second one serving as target. The rationale was that, if orthogonal axes facilitate each other, the detection of a symmetry should be enhanced when it is preceded by an orthogonal prime (i.e., a prime whose symmetry axis is orthogonal to the symmetry axis of the target). Likewise, if nonorthogonal axes impede each other, the detection of a symmetry should be hampered when it is preceded by a nonorthogonal prime.

Method

Participants A total of 19 undergraduate students participated in the experiment. None of them had participated in the other experiments. All had normal or corrected-to-normal vision and received course credits or money for their participation.

Stimuli The stimuli consisted of spatially filtered Gaussian noise. They resembled the stimuli used in the previous experiment, but they consisted of only a single SF band instead of two. First, the Gaussian noise was transformed into the SF domain using Fourier transformation. An idealized isotropic bandpass filter, defined as in Experiment 2b, with lower and higher cutoff frequencies of 1.29 and 2.59 cpd, respectively, was used to retain only one octave of spatial frequencies. After inverse Fourier transformation into the spatial domain, the resulting intensity distribution was shifted to mean luminance and normalized to 25% root-mean-square contrast. We again used absolute axis orientations of +22.5° and −22.5° with respect to the vertical and the horizontal.

We created one-fold symmetrical primes and targets, as well as primes and targets containing randomly distributed Gaussian noise. The random targets were used in catch trials, and the random primes were included to establish a baseline against which the priming effect of the symmetrical primes could be measured. All stimuli were circular with a diameter of about 10° of visual angle on a black background. Using MATLAB, a unique set of stimuli was generated for each participant.

Procedure We used the same apparatus and parameters as before, except for the following. In each trial, a prime (one-fold symmetry or random) was presented, followed by a target (one-fold symmetry or random). Participants performed a signal detection (yes/no) task; that is, they had to indicate whether the target was symmetrical or random.

Apart from the random prime condition, there were three symmetrical prime conditions. The crucial conditions were the nonorthogonal condition, in which the relative orientation of the symmetry axes in prime and target was 45°/135°, and the orthogonal condition, in which it was 90°. We also included the same condition, in which the symmetry axes in prime and target had the same orientation. In all conditions, both the prime and the target were masked using a checkerboard pattern. Masking was necessary to remove apparent-motion artifacts, which, during pilot

experiments, had been found to particularly arise when both prime and target were symmetrical.

The experiment was preceded by 50 practice trials. Each trial started with a central fixation cross presented for 500 ms. Subsequently, the prime was presented for 500 ms. We chose a long presentation time to ensure that the participant really detected the symmetry in the prime. The prime was immediately followed by a checkerboard pattern, presented for 50 ms, and a blank screen lasting for 150 ms. Then, the target was presented for 200 ms, followed by another mask presented for 50 ms. Following the response and an intertrial interval of 100 ms, the next trial commenced automatically. The time course of a trial is depicted in Fig. 7. The number of trials totaled 2 (symmetrical/random prime) x 4 (prime orientations) x 2 (symmetrical/random target) x 4 (target orientations) x 10 (measurements) = 640. Breaks were given after every 128 trials.

Results

Again, based on the correct scores for symmetrical targets (hits) and error scores for random targets (false alarms), d' was calculated. As a measure of the priming effect, we used $\Delta d'$, calculated by subtracting d' on random primes from d' on symmetrical primes. This measure indicates how much a symmetrical prime facilitates (positive $\Delta d'$) or impedes (negative $\Delta d'$) symmetry detection relative to a random prime. One-sample t-tests were used to analyze $\Delta d'$. In the *same* condition, $\Delta d'$ tended to be positive, but this was not significant ($p = .093$). In the orthogonal condition, there was a significant positive priming effect, $t(18) = 3.326$, $p < .01$. These two conditions did not differ significantly from each other ($p = .529$). Finally, there was significant negative priming effect in the nonorthogonal condition, $t(18) = 2.631$, $p < .05$. The results are depicted in Fig. 5c.

Discussion

The results reveal a positive priming effect in the case where the prime and target have orthogonal symmetry axes, as well as a negative priming effect in the case where they have nonorthogonal symmetry axes. We also ran a control experiment using a thin (1-pixel) line as a prime instead of a symmetrical pattern. In this control experiment, we found no priming effects in the orthogonal and nonorthogonal conditions, which suggests that the above-mentioned priming effects are indeed inherent to symmetry processing. This, in turn, suggests that the results obtained in the previous two experiments are due to an orientation-dependent interaction between symmetry axes, with both a mutual facilitation of orthogonal symmetry axes and a mutual impediment of nonorthogonal symmetry axes. As we elaborate in the next section, this difference between orthogonal and nonorthogonal symmetry axes might be caused by a simple neural mechanism that accounts for our data in Experiment 2b.

General discussion

In this study, we investigated whether or not the number of single symmetries in a multiple symmetry is the only factor determining its salience. In Experiments 1a, 1b, and 1c, we found no evidence that two-fold symmetry with correlation rectangles is better detectable than two-fold symmetry without. Regarding the bootstrap model (Wagemans et al.,
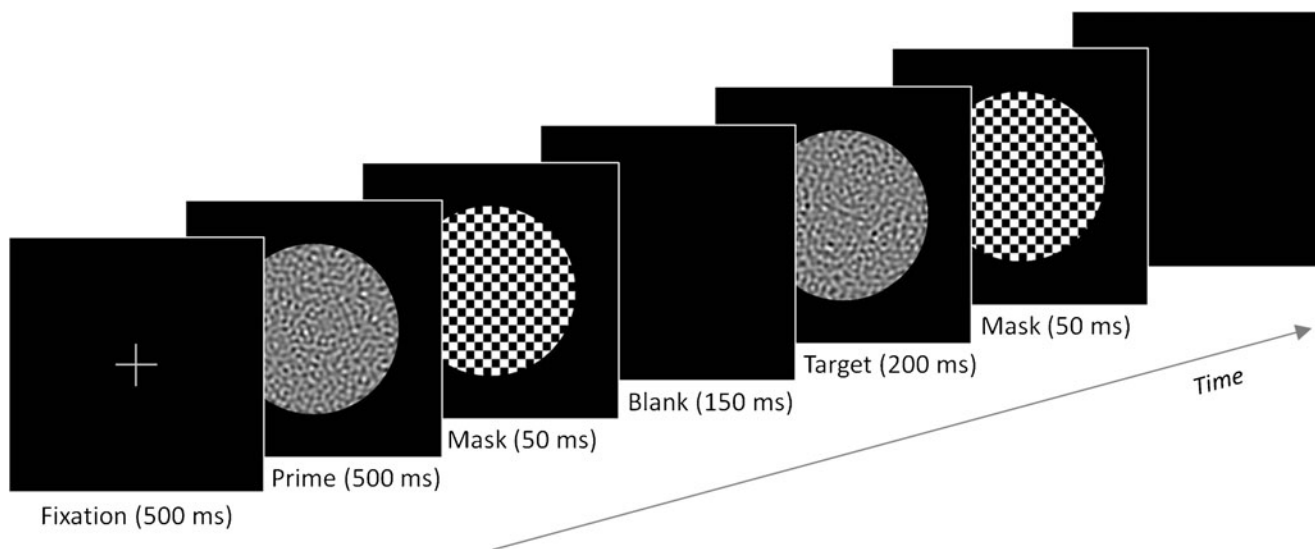

Fig. 7 The time course of a trial in priming Experiment 2c. See the procedure of Experiment 2c for details

1991; Wagemans et al., 1993), this does not challenge its primary assumption that correlation trapezoids are the anchors for the detection of symmetry about an axis, but it casts doubts on its secondary assumption that correlation rectangles facilitate the perception of multiple symmetries containing orthogonal symmetry axes.

In Experiments 2a, 2b, and 2c, we investigated whether the relative orientation between symmetry axes influences multiple-symmetry perception. We found that orthogonal symmetry is consistently more salient than nonorthogonal symmetry, even when correlation rectangles are absent. In Experiment 2c, we also found that detection of a target symmetry is both facilitated by an orthogonal symmetry prime and impeded by a nonorthogonal symmetry prime. As for the computational and algorithmic models outlined in the introduction, these results seem most consistent with the holographic approach, which predicts a mutually impeding effect of the symmetries in the case of non-orthogonal axes. That is, the transformational approach predicts no effect of the relative orientations between axes, whereas the bootstrap approach only predicts facilitation in the case of orthogonal axes—a facilitation solely due, however, to correlation rectangles, which we found to be effectless and which were even absent in Experiments 2a, 2b, and 2c. As we discuss next at the implementational level of description, a simple neural mechanism might be responsible for our empirical data.

Neurally, orientation-dependent interactions between orientation-selective neurons are ubiquitous. It is true that visual neurons are usually involved in the processing of luminance edges rather than symmetry axes, but Lee, Mumford, Romero, and Lamme (1998) found a late enhancement of responses in V1 cells of rhesus monkeys when the receptive fields were centered at a symmetry axis. Furthermore, symmetrical dot patterns and luminance contours were shown to elicit similar visually evoked potentials if the luminance contours and the symmetry axes were co-oriented (Beh & Latimer, 1997). Moreover, symmetry axes in dot patterns elicit tilt aftereffects comparable to those of lines (Joung & Latimer, 2003; Joung, van der Zwan, & Latimer, 2000; van der Zwan, Leo, Joung, Latimer, & Wenderoth, 1998). In line with this evidence, Gurnsey, Herbert, and Kenemy (1998) suggested that, within the visual system, symmetry axes form subjective contours that are processed in a way similar to luminance contours (see also Sally & Gurnsey, 2001).

The exact link between symmetry processing and orientation processing is still unclear, but our data call for a hopefully fruitful speculation about how known properties of the neural processing of orientation might be related to the interaction between single symmetries in multiple symmetry. To demonstrate this, we next present a model of symmetry detection involving smoothing—that is, averaging over neighbouring orientations.

There is consensus that symmetry is analyzed in parallel at multiple locations, orientations, and spatial frequencies (Dakin & Watt, 1994; Rainville & Kingdom, 1999; Wagemans, 1995; Wagemans et al., 1991). Despite the abundance of spatial-filtering models (e.g., Dakin & Watt, 1994; Gurnsey et al., 1998; Osorio, 1996; Rainville & Kingdom, 1999, 2002), there has been no account of how the information distributed across orientations and spatial frequencies is integrated. To account for our findings, we developed a four-stage model of symmetry detection, as follows.

First, symmetry is analyzed at all orientations and different spatial frequencies. For this stage, we adopted Dakin and Watt's (1994) symmetry detection model. We used this model because of its simplicity and computational ease, but other models, such as those of Rainville and Kingdom (1999) and Gurnsey et al. (1998), would yield comparable outputs. In the second stage, for each orientation, the output of the symmetry detection model is averaged across spatial-frequency bands. Examples of the resulting distributions are depicted in Fig. 8a. In the third stage, the output at each orientation is replaced by a weighted average of its neighborhood, mimicking the cross-orientation interactions ubiquitous in visual cortex. Hence, this smoothing operation is central to the model, because it removes most of the sharp noise peaks stemming from "spurious" symmetry (see Fig. 8a) while, at the same time, the peaks corresponding to the true axes of symmetry are preserved (Fig. 8b). In the fourth stage, the peaks of the smoothed distribution are extracted. The sum of the values of the peaks is taken as a measure for the total amount of symmetry in the stimulus (for details, see Appx. A).

To validate the model, we conducted a simulation using the same image type that was employed in Experiment 2b. After the application of the model to the images, we calculated d' (for details, see Appx. B). The results of the simulation, without and with smoothing, are shown in Fig. 9. Only with the smoothing operation are the empirical results (Fig. **5**b) replicated (see Fig. 9b). Hence, smoothing is crucial. Notice that smoothing affects nonorthogonal axes more than orthogonal axes. In other words, implicitly, smoothing implies an impediment in the case of nonorthogonal axes (which occur, e.g., in three-fold symmetry) rather than a facilitation, as in the case of orthogonal axes (which occur, e.g., in two-fold and four-fold symmetry).

Conclusion

Future research may expand our findings to more symmetry axes in more relative orientations, but our findings justify a
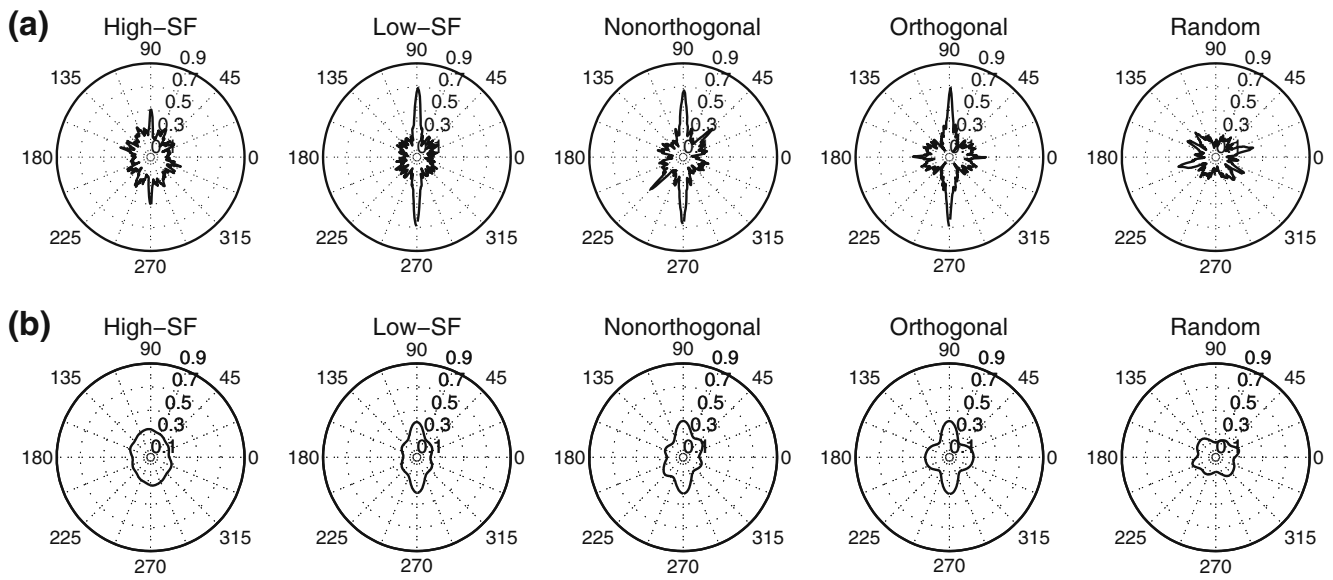
Fig. 8 Illustration of the application of the symmetry detection model to the five stimulus types used in Experiment 2b. Each plot is based on the analysis of 10 stimuli. (**a**) For each orientation, symmetry was estimated using Dakin and Watt's (1994) symmetry detection model. The results are shown as polar plots, whereby the magnitude represents the output of the symmetry detection model for each orientation. The first two plots give the results for one-fold symmetry in the high and low spatial-frequency (SF) bands, respectively, with a symmetry axis at 90 and noise in the other frequency band. For nonorthogonal symmetry, the symmetry axis in the low SF band was oriented at 90° and the symmetry in the high SF band was oriented at 45°. For orthogonal symmetry, the symmetry axis in the low SF band was oriented at 90° and the symmetry in the high SF band was oriented at 0°. These plots correspond to the output of the second stage of the symmetry detection model, where symmetry values have been averaged across two spatial-frequency bands. (**b**) After convolution with a Gaussian kernel, the output at each orientation is a weighted average of its neighboring orientations. In the smoothed representation, many of the sharp peaks in panel (**a**) representing spurious symmetry are removed, but there are still peaks coinciding with the true axes of symmetry

basic and novel conclusion about the phenomenal difference between orthogonal and nonorthogonal symmetry axes. Our study shows that the number of symmetry axes is not the only relevant factor in multiple-symmetry perception. Contrary to ideas in the literature, we found no evidence that correlation rectangles facilitate multiple-symmetry perception, but we did find evidence that, even in the absence of correlation rectangles, the salience of multiple symmetry is affected by the relative orientations of the symmetry axes. These findings argue against a special status of correlation rectangles in multiple-symmetry perception. Rather, they suggest an orientation-
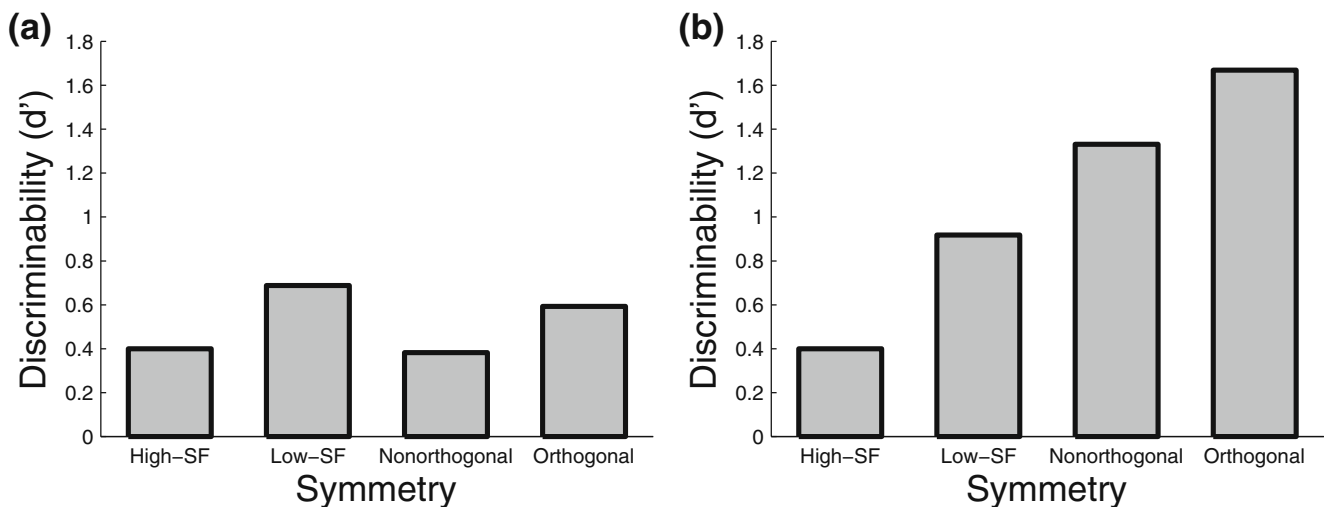

Fig. 9 Simulation results without smoothing (**a**) and with smoothing (**b**). Using smoothing, the simulation results are in line with the empirical data (Fig. 5b)

dependent interaction between the constituent one-fold symmetries. By means of a simulation, we demonstrated that a symmetry detection model involving smoothing across orientations can account for our empirical results.

Author Note Supplemental materials may be downloaded along with this article from www.springerlink.com.

Appendix A: symmetry detection model

The present symmetry detection model extends previous spatial-filtering models in that it does not focus on single symmetries only, but also gives an account of how multiple symmetry axes are integrated. The model operates in four successive stages. In the first stage, one-fold symmetry is analyzed in parallel at all orientations in different spatial-frequency bands. In the second stage, the symmetry outputs for each orientation are averaged across spatial-frequency bands. In the third stage, a smoothing operation is performed wherein the symmetry output at each orientation is replaced by a weighted average of neighboring orientations. In the fourth stage, the peaks of the smoothed representation are extracted. The sum of all peaks is taken as the total evidence for the presence of symmetry.

Analysis of symmetry at all orientations

In the first stage of the symmetry detection model, symmetry is analyzed at all orientations and for multiple spatial frequencies in parallel. To this end, we adopted a slightly modified version of Dakin and Watt's (1994) blob-alignment model. In the blob-alignment model, a stimulus is processed using three successive operations: spatial filtering, thresholding, and blob-alignment measurement.

First, the image is convolved with an elongated difference-of-Gaussian (DoG) filter oriented orthogonally to the putative symmetry axis. Dakin and Watt (1994) used an elongated DoG filter defined by

$$f_s(x,y) = \left(e^{-y^2/(2s^2)} - \frac{1}{2.23} \cdot e^{-y^2/(2(2.23s)^2)} \right) \cdot e^{-x^2/2(3s)^2}$$

where the space constant s determines the spatial-frequency (SF) sensitivity of the filter. Because our stimuli were confined to two SF bands, we used filters at two different spatial scales. Appropriate values for s were obtained by adjusting the filters' peak SF to roughly match the geometric mean frequencies of the two SF bands.

Second, the convolved image is thresholded to a binary image by setting all luminances in the range $(\mu - \sigma, \mu + \sigma)$ to 0, and all others to 1, where μ is the mean luminance and

σ is the standard deviation of the luminance distribution. This yields a number of blobs.

Third, the co-alignment of these blobs along a putative symmetry axis is taken as a measure for the degree of symmetry. This measure is normalized for blob size and numerosity, and it ranges from 0 (totally random) to 1 (perfectly symmetric). Assuming that the symmetry axis is located at $x = 0$, the alignment measure boils down to

$$A = \frac{1}{M} \cdot \sum_i exp\left(-\frac{c_i^2}{l_i} \right) \cdot \mu_i$$

where i ranges over the indices of blobs intersecting the symmetry axis, M is their total mass, c_i is the abscissa of the centroid of blob i, l_i is its length, and μ_i is its mass.

This measure differs slightly from Dakin and Watt's (1994) measure. As the denominator of the argument of the exponential, we chose l_i instead of $2l_i^2$. This gives less weight to the length of the blob, which proved to extract symmetry more reliably in our narrow-band stimuli (Dakin and Watt considered broad-band stimuli). A second modification was that we did not take into account blobs in the upper and lower end strips of 30 pixels. The reason is that the filter responds strongly to those parts of the stimulus contour that are approximately tangential to the orientation of the filter (i.e., the parts near both ends of the symmetry axis). Including those parts in our circular stimuli would have yielded comparatively large blob areas that overemphasized the amount of symmetry in every image, independently of the actually present symmetry (Dakin and Watt considered square-shaped stimuli, for which this artifact does not arise). In every other respect, we adhered to the model outlined by Dakin and Watt.

Averaging, smoothing, and extraction of peaks

This section expands on Stages 2, 3, and 4 of the symmetry detection model. In the second stage, the symmetry outputs for each orientation are averaged across spatial frequencies. This provides an estimate of the total amount of symmetry at each orientation. In the third stage, the symmetry output at each orientation is replaced by a weighted average of neighboring orientations. The smoothing operation removes many of the peaks stemming from spurious symmetry while preserving the global symmetry axes. Computationally, this operation is realized as convolution with a Gaussian kernel G, defined as

$$G = \frac{1}{\sigma\sqrt{2\pi}} \cdot e^{\frac{-x^2}{2\sigma^2}}$$

where $\sigma = 11$ is the standard deviation in degrees (the value 11 was empirically determined to give the best fit to our data).

In the fourth stage, the peaks are extracted from the smoothed representation. Usually, each peak will represent a global axis of symmetry. Hence, the sum of the values of the peaks represents the total amount of symmetry in the stimulus.

Appendix B: simulation

Using the symmetry detection model, we conducted a simulation. To this end, we created stimuli consisting of two superimposed Gaussian noise patterns, identical to the stimuli used in Experiment 2b. A total of 1,000 images were generated for each of five different stimulus types—namely, one-fold symmetry in the high SF band (with noise in the low SF band), one-fold symmetry in the low SF band, nonorthogonal symmetry, orthogonal symmetry, and random noise in both bands.

Each stimulus was processed through the four stages of the symmetry detection model. This yielded an estimate of the total amount of symmetry for each stimulus. Based on the distribution of symmetry estimates, d' was determined as specified in Geisler (2004):

$$d' = \frac{|E(s|sym_i) - E(s|ran)|}{\sqrt{\frac{Var(s|sym) + Var(s|ran)}{2}}}$$

where $E(s \mid sym_i)$ refers to the expectation value of the total amount of symmetry in stimulus s given one of the four classes of symmetry, $E(s \mid ran)$ refers to the expectation value of the total amount of symmetry in the stimulus given a random pattern, and Var is the corresponding variance (for more details, see Geisler, 2004).

References

Attneave, F. (1954). Some informational aspects of visual perception. *Psychological Review, 61*, 183–193.

Attneave, F. (1955). Symmetry, information, and memory for patterns. *The American Journal of Psychology, 68*, 209–222.

Barlow, H. B. (1961). The coding of sensory messages. In W. H. Thorpe & O. L. Zangwill (Eds.), *Current problems in animal behavior* (pp. 331–360). Cambridge: Cambridge University Press.

Barlow, H. (2001). The exploitation of regularities in the environment by the brain. *The Behavioral and Brain Sciences, 24*, 602–607.

Beh, H. C., & Latimer, C. R. (1997). Symmetry detection and orientation perception: Electrocortical responses to stimuli with real and implicit axes of orientation. *Australian Journal of Psychology, 49*, 128–133.

Corbalis, M. C., & Roldan, C. E. (1974). On the perception of symmetrical and repeated patterns. *Perception & Psychophysics, 16*, 136–142.

Dakin, S. C., & Herbert, A. M. (1998). The spatial region of integration for visual symmetry detection. *Proceedings of the Royal Society B, 265*, 659–664. doi:10.1098/rspb.1998.0344

Dakin, S. C., & Watt, R. J. (1994). Detection of bilateral symmetry using spatial filters. *Spatial Vision, 8*, 393–413.

Garner, W. R. (1974). *The processing of information and structure*. Potomac, MD: Erlbaum.

Geisler, W. S. (2004). Ideal observer analysis. In L. M. Chalupa & J. S. Werner (Eds.), *The visual neurosciences* (pp. 825–837). Cambridge, MA: MIT Press.

Gurnsey, R., Herbert, A. M., & Kenemy, J. (1998). Bilateral symmetry embedded in noise is detected accurately only at fixation. *Vision Research, 38*, 3795–3803.

Hamada, J., & Ishihara, T. (1988). Complexity and goodness of dot patterns varying in symmetry. *Psychological Research, 50*, 155–161.

Jenkins, B. (1983). Component processes in the perception of bilaterally symmetric dot textures. *Perception & Psychophysics, 34*, 433–440.

Joung, W., & Latimer, C. (2003). Tilt aftereffects generated by symmetrical dot patterns with two or four axes of symmetry. *Spatial Vision, 16*, 155–182.

Joung, W., van der Zwan, R., & Latimer, C. R. (2000). Tilt aftereffects generated by bilaterally symmetrical patterns. *Spatial Vision, 13*, 107–128.

Julesz, B., & Chang, J. J. (1979). Symmetry perception and spatial-frequency channels. *Perception, 8*, 711–718.

Lee, T. S., Mumford, D., Romero, R., & Lamme, V. A. (1998). The role of the primary visual cortex in higher level vision. *Vision Research, 38*, 2429–2454.

Leeuwenberg, E. L. (1969). Quantitative specification of information in sequential patterns. *Psychological Review, 76*, 216–220.

Leeuwenberg, E. L. (1971). A perceptual coding language for visual and auditory patterns. *The American Journal of Psychology, 84*, 307–349.

Marr, D. (1982). *Vision: A computational approach*. San Francisco: Freeman.

Nucci, M., & Wagemans, J. (2007). Goodness of regularity in dot patterns: Global symmetry, local symmetry, and their interactions. *Perception, 36*, 1305–1319.

Osorio, D. (1996). Symmetry detection by categorization of spatial phase, a model. *Proceedings of the Royal Society B, 263*, 105–110.

Palmer, S. E. (1983). The psychology of perceptual organization: A transformational approach. In J. Beck, B. Hope, A. Rosenfeld, & National Science Foundation (Eds.), *Human and machine vision* (pp. 269–339). New York: Academic Press.

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.

Rainville, S. J. M., & Kingdom, F. A. A. (1999). Spatial-scale contribution to the detection of mirror symmetry in fractal noise. *Journal of the Optical Society of America B, 16*, 2112–2123.

Rainville, S. J. M., & Kingdom, F. A. A. (2000). The functional role of oriented spatial filters in the perception of mirror symmetry—Psychophysics and modeling. *Vision Research, 40*, 2621–2644.

Rainville, S. J. M., & Kingdom, F. A. A. (2002). Scale invariance is driven by stimulus density. *Vision Research, 42*, 351–367.

Sally, S., & Gurnsey, R. (2001). Symmetry detection across the visual field. *Spatial Vision, 14*, 217–234.

Sawada, T., & Pizlo, Z. (2008). Detection of skewed symmetry. *Journal of Vision, 8(5)*, 14, 1–18. doi:10.1167/8.5.14

Swets, J. A. (1964). *Signal detection and recognition by human observers: Contemporary readings*. New York: Wiley.

Treder, M. S. (2010). Behind the looking-glass: A review on human symmetry perception. *Symmetry, 2*, 1510–1543. Available at www.mdpi.com/2073-8994/2/3/1510/pdf

Treder, M. S., & van der Helm, P. A. (2007). Symmetry versus repetition in cyclopean vision: A microgenetic analysis. *Vision Research, 47*, 2956–2967. doi:10.1016/j.visres.2007.07.018

Tyler, C. W. (1996). *Human symmetry perception and its computational analysis*. Zeist, The Netherlands: VSP.

van der Helm, P. A. (2011). The influence of perception on the distribution of multiple symmetries in nature and art. *Symmetry, 3*, 54–71.

van der Helm, P. A., & Leeuwenberg, E. L. J. (1991). Accessibility, a criterion for regularity and hierarchy in visual pattern codes. *Journal of Mathematical Psychology, 35*, 151–213.

van der Helm, P. A., & Leeuwenberg, E. L. (1996). Goodness of visual regularities: A nontransformational approach. *Psychological Review, 103*, 429–456.

van der Helm, P. A., & Leeuwenberg, E. L. J. (1999). A better approach to goodness: Reply to Wagemans (1999). *Psychological Review, 106*, 622–630.

van der Helm, P. A., & Leeuwenberg, E. L. J. (2004). Holographic goodness is not that bad: Reply to Olivers, Chater, and Watson (2004). *Psychological Review, 111*, 261–273.

van der Vloed, G. (2005). *The structure of visual regularities*. Ph. D. thesis. Nijmeegs Instituut voor Cognitie en Informatie: Nijmegen, the Netherlands.

van der Vloed, G., Csathó, Á., & van der Helm, P. A. (2005). Symmetry and repetition in perspective. *Acta Psychologica, 120*, 74–92.

van der Zwan, R., Leo, E., Joung, W., Latimer, C., & Wenderoth, P. (1998). Evidence that both area V1 and extrastriate visual cortex contribute to symmetry perception. *Current Biology, 8*, 889–892.

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

Wagemans, J., Gool, L. V., & d'Ydewalle, G. (1991). Detection of symmetry in tachistoscopically presented dot patterns: Effects of multiple axes and skewing. *Perception & Psychophysics, 50*, 413–427.

Wagemans, J., Gool, L. V., Swinnen, V., & Horebeek, J. V. (1993). Higher-order structure in regularity detection. *Vision Research, 33*, 1067–1088.

Wenderoth, P. (1997). The effects on bilateral-symmetry detection of multiple symmetry, near symmetry, and axis orientation. *Perception, 26*, 891–904.

Wenderoth, P., & Welsh, S. (1998). Effects of pattern orientation and number of symmetry axes on the detection of mirror symmetry in dot and solid patterns. *Perception, 27*, 965–976.

Wickens, T. D. (2002). *Elementary signal detection theory*. New York: Oxford University Press.