What do we see behind an occluder?

Amodal completion of statistical properties in complex objects

Thomas Cherian & SP Arun

Centre for Neuroscience, Indian Institute of Science, Bangalore 560012

#Correspondence to: sparun@iisc.ac.in

ABSTRACT

When a spiky object is occluded, we have the compelling percept that its spiky features continue behind the occluder. Although many real-world objects contain complex features, it is unclear how their features are amodally completed and whether this process is automatic. To this end, we asked participants to search for oddball targets among distractors and asked whether similarity relations in visual search between occluded displays match better with global or local completions of these displays. In Experiment 1, when objects with curved/straight corners were occluded, they were perceived as continuing with the same features than with them exchanged. In Experiment 2, we obtained similar results for objects with irregular/symmetric features. Analogous investigations on deep neural networks revealed similar results for curved/straight contours but not for irregular/symmetric features. Thus, amodal completion involves extending not only simple contours, but also their more global statistical properties.

INTRODUCTION

If you want to become whole,

let yourself be partial.

- Tao Te Ching (Mitchell, 1988)

When we see a partially occluded object, we have the vivid impression that its edges continue behind the occluder *in specific ways*. What do we really see behind the occluder? Broadly there are three possible interpretations: (1) a mosaic where two objects are juxtaposed and not occluding; (2) a local completion where the occluding edges are extrapolated based on local extrapolation of straight contours; and (3) a global completion where the occluded edges are extrapolated consistent with global properties such as symmetry. There have been extensive studies to distinguish these interpretations in behavior and neural recordings (Nanay, 2018; Thielen et al., 2019).

Evidence in favor of local/global completions comes largely from two paradigms: priming and visual search. In the priming paradigm, an occluded display is shown as prime, after which participants were required to make same/different response to a stimulus pair. In these studies, participants made faster matching responses to global/local completions compared to mosaic displays after being primed by the occluded display, suggesting that viewing an occluded display gives rise to the global and local interpretations but not the mosaic (Sekuler and Palmer, 1992; Sekuler, 1994; van Lier et al., 1994; Tse, 1998, 1999). There is even a graded priming advantage across a variety of possible completions (van Lier et al., 1995; De Wit and Van Lier, 2002). These interpretations are complicated by the fact that even unprimed completed displays elicit

faster responses than mosaic displays (Sekuler and Palmer, 1992), leaving open the possibility that the priming effect is at least partially driven by the presence of prominent common contours in mosaic displays. Moreover, the primes in these studies were typically shown for 500-1000 ms, with participants usually instructed to attend to the prime, leaving open the possibility that their responses are driven by higher level concepts or reasoning as opposed to automatic visual processing.

In the visual search paradigm, participants had to search for a target among distractors, revealing their automatic visual processing. In these studies, search for an occluded display among global/local completions is easier than a mosaic display among the same distractors, suggesting that the occluded display is similar to the global/local completions but not the mosaic (Rensink and Enns, 1998; Chen et al., 2018). However, this interpretation is complicated by the fact that these searches have strong asymmetries: finding the mosaic among completed is easy, but finding the completed among mosaic is hard (Rensink and Enns, 1998; Rauschenberger and Yantis, 2001; Rauschenberger et al., 2004), presumably because the mosaic contains prominent new features. These findings are illustrated in Figure 1A.

Investigations of amodal completion using other paradigms largely support these insights. Objects are easier to recognize when their contours are occluded than when cut (Johnson and Olshausen, 2005). Occluded objects are represented akin to single whole objects, as evidenced by easier feature processing for occluded compared to separate objects with equivalent features (Murray et al., 2001; Johnson and Olshausen, 2005; Hadad and Kimchi, 2018). However, the occluded edges in these studies are usually relatively simple and lend themselves to straight line local completions. Finally, there is

recent evidence that humans show robust recognition even with extensive occlusion, in contrast to deep networks trained (Zhu et al., 2019; Kang and Druckmann, 2020; Kortylewski et al., 2020).

Although the above studies have yielded many important insights, they have two fundamental limitations. First, most studies of amodal completion have studied only one aspect of global completion, namely symmetry. As a result, we do not fully understand how complex object features are amodally completed. For instance, when a spiky object is occluded, we have the vivid impression that its spiky features continue behind the occluder. Since we do not have any specific expectation about the exact position or size of the spiky features, we consider this as an expectation that the statistics of the visible features will continue behind the occluder. Whether such complex and sophisticated forms of amodal completion occur automatically or only through higher level reasoning is unclear.

Second, nearly all studies involve comparisons with the mosaic display, which contains prominent new features, namely the cut edges now present on both objects (Figure 1B). We lack a clear theoretical framework to understand how such new features affect the underlying representation. For instance, the mosaic display might be so different from both occluded and local/global completions that *any* completion might yield similar results. Likewise, the occluded and completed displays differ not only in the occlusion of contours but also their spatial arrangement, whose effects are unclear. Thus, comparing across display types can be problematic. A potential solution to these confounds is to compare within display types – for instance, perceptual relations between

occluded objects can be compared with the perceptual relations between objects with various types of completion.

Figure 1. Amodal completion in visual search

(A) Schematic of classic results on amodal completion (Rensink and Enns, 1998). Finding an occluded target among separated displays is hard (*top row*) but finding a mosaic target among separated displays is easy (*middle row*). These two observations suggest that the occluded display representation is more similar to the separated displays than the mosaic display. However, this interpretation is complicated by the fact that finding the occluded display among mosaic displays is hard (*bottom row*), suggesting that even these two displays are similar.

(B) Comparing object representations across display types is problematic because occluded, mosaic and separated display differ not only in terms of the occlusion properties but also in terms of other features, such as the additional cut in the mosaic display, or the change in position of the occluded object in the separated display. We overcome this problem in our study by comparing search representations within display types, as depicted in Figure 2.

Overview of this study

We sought to address the above limitations in our study. First, to investigate global completion based on statistical properties, we created pairs of displays with identical edges at the point of occlusion but with occluded portions exchanged. Figure 2A shows an example pair of two occluded shapes. When the visible portion of the occluded display has sharp corners (Figure 2A, top row), the likely interpretation is a shape with sharp

94 corners consistent with the visible features and the unlikely interpretation is a shape with
95 round corners. When the visible portion of the occluded display has round corners (Figure
96 2A, bottom row), the likely interpretation is a shape with round corners consistent with the
97 visible statistics and the unlikely interpretation is a shape with sharp corners.

98 We denoted these completions as likely and unlikely because they are both global
99 completions in the sense used in the literature. Note that the top-right quadrant of the
100 likely and unlikely completions in each row are made of exactly the same contour. As a
101 result, any representational difference between likely and unlikely shapes can only arise
102 due to the congruence between the visible and occluded portions but not due to
103 differences in shape. Figure 2B shows another example pair of shapes demonstrating a
104 more complex global completion.

105 Second, to compare perceptual relations within a given display type, we measured
106 perceptual dissimilarities between pairs of occluded displays, and compared these with
107 pairwise dissimilarities between likely completions and between unlikely completions
108 (Figure 2C). To measure perceptual dissimilarities between a given pair of displays, we
109 created search arrays with one of the images as target and many instances of the other
110 as distractors. Example searches of each type are shown in Figure 3: search for an
111 occluded target among occluded distractors (Figure 3A), for a likely completion target
112 among another such distractor (Figure 3B), and an unlikely completion target among
113 another such distractor (Figure 3C).

Figure 2. Amodal completion of statistical properties in complex objects.

(A) *Left:* Example pair of occluded displays demonstrating simple statistical completion. Both shapes have identical local edges at the point of occlusion. *Middle:* Likely completions of these shapes, with statistics consistent with the visible features. *Right:* Unlikely completions with statistics inconsistent with the visible features, created by swapping the contours in the likely completions. We investigated amodal completion of these shapes in Experiment 1.

(B) Same as (A) but with objects demonstrating complex statistical completion. We investigated amodal completion of these shapes in Experiment 2.

(C) Approach for comparing occluded, likely, and unlikely shape representations. Comparing occluded displays with likely or unlikely displays is confounded by feature differences (see text). We therefore hypothesized that the pairwise dissimilarities between occluded displays will match better with likely or unlikely shape representations.

130 Since search time increases with target-distractor similarity, we took the reciprocal
131 of mean search time ($1/RT$) as a measure of perceptual dissimilarity. Intuitively, the
132 reciprocal of search time can be thought of as the underlying feature dissimilarity signal
133 that drives visual search (Sunder and Arun, 2016). It behaves as a mathematical distance
134 metric (Arun, 2012), and a variety of feature differences combine linearly in terms of this
135 distance ($1/RT$) but not using search time (Vighneshvel and Arun, 2013; Pramod and
136 Arun, 2014, 2016; Sunder and Arun, 2016; Jacob and Arun, 2020). Thus $1/RT$ is an
137 insightful measure to characterize dissimilarity in visual search.

138 Our main hypothesis is that perceptual dissimilarities in occluded displays will
139 match more closely with likely interpretations compared to unlikely interpretations. We
140 performed two experiments. In Experiment 1, we tested this hypothesis on simpler shapes
141 with rounded or curved corners. In Experiment 2, we tested it on a more diverse set of
142 shapes with varying protrusions. Finally, in Experiment 3, we used this approach to query
143 how deep networks represent occluded displays.

Figure 3. Example search displays for occluded, likely, and unlikely conditions.

- (A) Example search array for an occluded target among identical other distractors of the same type.
- (B) Example search array for a target with likely completion among identical other distractors of the same type.
- (C) Example search array for a target with unlikely completion among identical other distractors of the same type.

EXPERIMENT 1: AMODAL STATISTICAL COMPLETION OF SIMPLE SHAPES

In Experiment 1, we created simple shapes containing straight or curved edges, and asked if their occluded representation matched better with completions consistent with the visible contours (likely completions) as against completions inconsistent with their visible contours (unlikely completions).

METHOD

All experiments were performed according to a protocol approved by the Institutional Human Ethics Committee at the Indian Institute of Science, Bangalore.

Participants. A total of 16 participants (6 female, age, mean $\pm$ s.d.: 27 ± 6 years) were recruited for the experiment. We selected this sample size because previous studies of visual search in our lab yielded highly consistent search response times. Participants were naïve to the purposes of the experiments, had normal or corrected-to-normal vision and were monetarily compensated for their time and efforts.

Stimuli. We created 8 base shapes, each with roughly 2 edges per quadrant intersecting in sharp corners. We ensured that only two edges of each shape were present in the occluded quadrant and the angle of intersection of these occluded edges varied across shapes. For each of the 8 base shapes, we created additional versions in which the edges intersected in curved instead of sharp corners. This yielded a total of 16 shapes containing either all sharp or all curved corners, which were designated as likely completions. Unlikely versions were created by switching the top-left sharp/curved portions of each

sharp/curved shape pair. This procedure ensured that the unlikely shapes had no sharp discontinuities. Example stimuli from this experiment are shown in Figure 2A & Figure 4A, and the complete stimulus set is shown in Figure S1A.

We had three display conditions for each shape: occluded, likely, and unlikely. To create the occluded displays, we overlaid a brick texture onto the upper left quadrant of each likely shape (by design, an identical image would be produced by overlaying the brick texture on the corresponding unlikely shape). Likely and unlikely displays were likewise created by placing the same brick texture on the left quadrant of each shape but by adding a constant position offset such that there was no contact between the brick texture and the shape. In total there were 16 occluded, likely, and unlikely displays used in this experiment.

All shapes were created in grayscale color space using Inkscape (Inkscape Project, <https://inkscape.org>). These (along with a brick texture) were imported into MATLAB (R2019b, The Mathworks Inc) for creating the occluded, likely, and unlikely displays. All displays used a solid black background and were exported at a resolution of 200 pixels by 147 pixels.

Procedure. Each participant completed a visual search task online. The task was designed and built using PsychoPy v2020.1.3 (Peirce et al., 2019) and was exported to PsychoJS (a browser compatible JavaScript library available within PsychoPy). The experiment codes were then ported to the online behavioral platform Pavlovia (<https://pavlovia.org>). Participants performed the experiment through the internet using their own personal computers. To counter the possible change in image size due to

differences in computer monitor, we ran a custom scaling routine (code adapted from <https://gitlab.pavlovlab.org/Wake/screenscale>, courtesy of Wakefield Morys-Carter) at the beginning of each session, in which participants adjusted an on-screen rectangle to match a standard credit card. The resulting dimensions were used to scale images such that they measured a fixed physical size regardless of monitor dimensions. To calculate the size of stimuli in degrees, visual angle we assumed participants to be sitting ~57 cm from the screen.

Despite controlling for image size in this manner, it is possible that data collected from online participants is unreliable due to variations in the participants displays, computers or surrounding environment. To verify this, we ran an independent visual search experiment both online and in controlled lab conditions and confirmed that visual search times were highly correlated across these conditions. Similar findings have been reported by other studies (de Leeuw and Motz, 2016; Bridges et al., 2020; Grootswagers, 2020).

Participants completed three blocks of visual search: a practice block, a block containing only occluded displays, and a block containing likely and unlikely completions. The occlusion block was always presented first to avoid biasing participants towards likely or unlikely completions. The practice block consisted of 20 oddball searches based on five unrelated shapes so that participants learned to perform the task correctly. In the occlusion block, the 16 stimuli were used to form all possible $^{16}C_2 = 120$ possible pairs. We created two search arrays for each pair with either item as target and the other as distractor, with target position chosen randomly as left or right of screen center. This resulted in a total of 240 searches in the occlusion block. The likely/unlikely completion

block consisted of 240 search arrays made from the 16 likely completions, and 240 from the 16 unlikely completions. Incorrect trials were repeated after a random number of other trials, to ensure that there were 240 correct responses for each display type. The entire session was completed in 45-60 minutes.

In the visual search blocks, each trial started with the presentation of a yellow fixation cross ($0.2^\circ \times 0.2^\circ$) presented for 500 ms. This was followed by the visual search array, consisting of 12 items arranged in a 3 x 4 grid, comprising a single oddball target among 11 identical distractors. Each item in the array measured $4.5^\circ \times 3.3^\circ$, with a center-to-center item spacing of 6.4° along the horizontal direction and 5.5° along the vertical direction. Each item was jittered in position (± 0.5 dva sampled from a uniform distribution, along both horizontal and vertical directions) to prevent alignment cues from driving search. A thin vertical red line (0.1° wide) bisected the screen vertically.

Participants were asked to indicate the position of the oddball target stimulus relative to the central line with a key press ("Z" for left, "M" for right). This search display remained on the screen till participants responded or 10 seconds, whichever occurred earlier. Inter-trial interval was set at 500 ms. Participants could pause the task and resume at will but were instructed to avoid taking too many breaks.

Data analyses. We analyzed response times from correct trials. For each search pair, reaction times were averaged across AB and BA searches and the inverse of the averaged reaction time was taken as a measure of perceptual dissimilarity (Arun, 2012).

244 *Multidimensional scaling.* To visualize the arrangement of shapes in perceptual space,
245 we used multidimensional scaling ('mdscale' function in MATLAB). This algorithm finds
246 the best-fitting 2D coordinates of each shape such that their pairwise distances match the
247 observed pairwise distances. Since two representations could be identical short of
248 shift/scale/rotation operations, we aligned them using a Procrustes transformation on the
249 likely/unlikely/mosaic display coordinates to rotate/shift them to match with the occluded
250 display arrangement ('procrustes' function in MATLAB; Figure 4F-H, 5 G-J).

RESULT

Example shapes used in this experiment are shown in Figure 4A. Participants were highly accurate across all display types (accuracy, $97\% \pm 1\%$ across subjects: $97\% \pm 1\%$ for occlusion block, $98\% \pm 1\%$ for likely completions; $97\% \pm 1\%$ for unlikely completions). They were also highly consistent in their response times across searches (split-half correlation between average search times of odd- & even-numbered participants: $r = 0.84$ for occlusion; $r = 0.86$ for likely; $r = 0.84$ for unlikely).

Our goal was to investigate the underlying representation of occluded displays using visual search. For each pair of occluded displays, we calculated the average search time across trials and participants and took its inverse as a measure of perceptual dissimilarity. We observed a striking correlation between the pairwise dissimilarities between occluded displays and the pairwise dissimilarities between the likely completions ($r = 0.89$, $p < 0.00005$; Figure 4B). By contrast, there was a slightly lower correlation between occluded dissimilarities and unlikely completions ($r = 0.79$, $p < 0.00005$; Figure 4C). These high correlations are expected since three of the four quadrants of the shapes used in the occluded, likely, and unlikely conditions are identical.

To investigate whether the match between occluded and likely completion is indeed significantly better than the match between occluded and unlikely completions, we performed two analyses. First, we repeatedly sampled search pairs with replacement and calculated the correlations each time between occluded and likely or unlikely dissimilarities. To calculate statistical significance, we calculated the fraction of times a given trend was violated out of 1,000 such bootstrap samples. Indeed, the correlation between occluded-likely was significantly larger than occluded-unlikely ($p < 0.005$; Figure

4D). Second, we performed an error analysis to compare the deviations between occluded dissimilarities and the best possible prediction from the likely or unlikely dissimilarities. To this end, we scaled and shifted the likely dissimilarities to match the occluded dissimilarities and calculated the squared error between the two and did likewise for the unlikely dissimilarities (Figure 4E). This comparison revealed a smaller error for the occluded-likely match compared to the occluded-unlikely match (average error, mean $\pm$ s.e.m. across search pairs: $0.0079 \pm 0.0010 \text{ s}^{-2}$ for occluded-likely; $0.0144 \pm 0.00416 \text{ s}^{-2}$ for occluded-unlikely, $p = 1.19 \times 10^{-4}$, sign-rank test across 120 pairs; Figure 4E).

To visualize the shape representations for the three display types, we performed a multidimensional scaling analysis. This analysis finds the 2-dimensional coordinates of a set of stimuli such that their distances match the observed pairwise dissimilarities. In the resulting plot, nearby stimuli represent hard searches. The resulting visualizations are shown for occluded displays (Figure 4F), likely completions (Figure 4G) and unlikely completions (Figure 4H). We observed that the occluded display representation is more similar to the likely compared to unlikely completions.

The above results could be driven strongly by some image property common to most shapes. One such possible property is symmetry, since many shapes used in this experiment were symmetric. To investigate this possibility, we repeated the above analysis for symmetric and asymmetric shapes separately. This revealed qualitatively similar results for both types of shapes (Section S2).

We conclude that, while all three shape representations are highly correlated, the occluded representation has a slight but significantly better match with likely completions compared to unlikely completions.

Figure 4: Amodal statistical completion of simple shapes (Experiment 1)

- (A) Example pair of shapes showing occluded displays and their likely and unlikely completions.
- (B) Pairwise dissimilarity between occluded displays plotted against the corresponding dissimilarities between likely completions. Asterisks indicate statistical significance (* is $p < 0.05$, ** is $p < 0.005$, etc).
- (C) Same as (B) but for occluded displays and unlikely completions.
- (D) Mean and s.d. of bootstrapped correlation to occluded dissimilarities across 1000 samples with replacement for both likely and unlikely dissimilarities. Asterisks indicate statistical significance (number of times observed trend is violated across the 1000 bootstrap samples, ** equals $p < 0.005$)
- (E) Mean and s.e.m. of the squared residual error when predicting occluded dissimilarities using likely and unlikely dissimilarities. Asterisks indicate statistical significance calculated using a sign-rank test comparing the two groups.

313 (F) Visualization of occluded display dissimilarities using multidimensional scaling.
314 Numbers in the plot correspond to the 16 shapes used, and nearby numbers indicate
315 hard searches. The correlation coefficient on the top right indicates the match between
316 the 2D coordinates in the plot and the observed dissimilarities in visual search.
317 (G) Same as (F) for shapes with likely completion. For ease of comparison, the 2D
318 arrangement is rotated and shifted to match the arrangement of occluded shapes.
319 (H) Same as (G) but for shapes with unlikely completions.

320

EXPERIMENT 2. AMODAL COMPLETION OF COMPLEX SHAPES

In Experiment 1, we showed that when simple shapes are occluded, the occluded display representations match better with likely completions whose statistics match the visible portions. In Experiment 2, we sought to establish the generality of this effect by testing more complex shapes. We also included additional searches corresponding to mosaic completions in this experiment, to investigate their underlying representation.

METHOD

All methods were identical to Experiment 1, and only the details specific to this experiment are summarized below.

Participants. A total of 16 individuals (6 female, age, mean $\pm$ s.d.: 25 ± 3 years) participated in this experiment.

Stimuli. We created 8 shape pairs with protrusions such that the two items in each pair differed in the number and frequency of protrusions. Example pairs are shown in Figure 2B & Figure 5A, and the full set of stimuli are shown in Section S1. The resulting 16 shapes were designated as likely completions. Their unlikely counterparts were created by simply exchanging the top left quadrants of the likely shapes in each shape pair. This ensured no discontinuities in the unlikely shapes. We further created mosaic counterparts where from the likely completion the occluded quadrant pixels were removed.

343 **Search pairs.** We included 16 displays corresponding to mosaic completion in this
344 experiment, which resulted in an additional $^{16}C_2 = 120$ search pairs. We created two trials
345 for each search pair, with either item as target, resulting in a total of 240 trials for this
346 display type. Together with 240 trials each for occluded-occluded, likely-likely, and
347 unlikely-unlikely searches, there were $240 \times 4 = 960$ trials in the experiment.

RESULT

Example shapes used in this experiment are shown in Figure 5A. Participants were highly accurate across all display types (accuracy, $97\% \pm 3\%$ across subjects: $97\% \pm 2\%$ for occlusion block, $97\% \pm 3\%$ for likely completions; $96\% \pm 4\%$ for unlikely completions; $96\% \pm 3\%$ for mosaic completions). They were also highly consistent in their response times across searches (split-half correlation between average search times of odd- & even-numbered participants: $r = 0.89$ for occlusion; $r = 0.88$ for likely; $r = 0.88$ for unlikely; $r = 0.89$ for mosaic).

As before we observed a striking correlation between the occluded and likely dissimilarities ($r = 0.91$, $p < 0.00005$; Figure 5B). We obtained a similar correlation between occluded and mosaic dissimilarities ($r = 0.92$, $p < 0.00005$; Figure 5D). By contrast, there was a lower correlation between the occluded and unlikely dissimilarities ($r = 0.67$, $p < 0.00005$; Figure 5C). Note that a high correlation is expected between all four display types because three of four quadrants are identical for any search pair.

To establish whether the occluded representation matched better with likely or unlikely, we compared the correlation as well as squared residual error as before. This yielded no difference in the quality of match between occluded-likely and occluded-mosaic correlations, but these two matched better than occluded-unlikely completions ($r = 0.91$ for occluded-likely; $r = 0.92$ for occluded-mosaic; $r = 0.67$ for occluded-unlikely; $p > 0.5$ for likely vs mosaic; $p < 0.005$ for likely vs unlikely; $p < 0.005$ for mosaic vs unlikely; across 1,000 bootstrap-resampled correlations; Figure 5E). Similarly, we obtained similar errors for likely and mosaic predicting occluded dissimilarities, but a lower error for unlikely completions (squared residual error, mean $\pm$ s.e.m.: $0.0068 \pm 0.0007 \text{ s}^{-2}$ for likely;

371 $0.0220 \pm 0.0025 \text{ s}^{-2}$ for unlikely; $0.0059 \pm 0.0011 \text{ s}^{-2}$ for mosaic; $p = 0.1137$ for likely vs
372 mosaic; $p = 7.056 \times 10^{-10}$ for likely vs unlikely; $p = 1.687 \times 10^{-10}$ for mosaic vs unlikely,
373 sign-rank test across 120 pairs; Figure 5F). We obtained qualitatively similar results for
374 both symmetric and asymmetric shapes (Section S1).

375 Next, we visualized the representation of occluded displays and likely, unlikely and
376 mosaic completions using multidimensional scaling. The resulting plots show that the
377 occluded display representation (Figure 5G) matched qualitatively with the likely
378 completions (Figure 5H) as well as mosaic completions (Figure 5J) but not as well with
379 unlikely completions (Figure 5I).

380 We conclude that, while all four shape representations are highly correlated for
381 complex shapes, the occluded representation has a slight but significantly better match
382 with likely and mosaic completions compared to unlikely completions.

Figure 5: Amodal statistical completion of complex shapes (Experiment 2)

(A) Example shapes showing occluded displays and their likely/unlikely completions.

(B) Pairwise dissimilarity between occluded displays plotted against the corresponding dissimilarities between likely completions. Asterisks indicate statistical significance (* is $p < 0.05$, ** is $p < 0.005$, etc).

(C) Same as (B) but for occluded displays and unlikely completions.

(D) Same as (B) but for occluded displays and mosaic completions.

(E) Mean and s.d. of bootstrapped correlation to occluded dissimilarities across 1000 samples with replacement for both likely, unlikely and mosaic dissimilarities. Asterisks indicate statistical significance (number of times observed trend is violated across the 1000 bootstrap samples, ** equals $p < 0.005$)

(F) Mean and s.e.m. of the squared residual error when predicting occluded dissimilarities using likely, unlikely and mosaic dissimilarities. Asterisks indicate statistical significance calculated sign-rank test comparing the two groups.

(G) Visualization of occluded display dissimilarities using multidimensional scaling. Numbers in the plot correspond to the 16 shapes used, and nearby numbers indicate hard searches. The correlation coefficient on the top right indicates the match between the 2D coordinates in the plot and the observed dissimilarities in visual search.

- 402 (H) Same as (G) for shapes with likely completion. For ease of comparison, the 2D
403 arrangement is rotated and shifted to match the arrangement of occluded shapes.
404 (I) Same as (H) but for shapes with unlikely completions.
405 (J) Same as (H) but for shapes with mosaic completions.

EXPERIMENT 3. OCCLUSION IN DEEP NETWORKS

Deep networks have revolutionized computer vision, and their representations closely match those observed in biological vision (Serre, 2019). Yet, they do not generalize well to occluded objects (Rajaei et al., 2019; Zhu et al., 2019). Their representations of occluded displays have recently found not to match with likely completions (Jacob et al., 2021). However, none of these studies have compared dissimilarities within a given display type as done in our study.

We therefore set out to investigate how the representation of occluded displays in deep neural networks would compare to their representation of likely, unlikely and mosaic completions. To this end, we analyzed a deep neural network pre-trained for object classification (Simonyan and Zisserman, 2014) on a large image dataset containing 1,000 categories (Deng et al., 2010).

METHOD

Stimuli: We used stimuli from Experiments 1 and 2 (Figure S1), including their occluded, likely, unlikely and mosaic completions.

Deep neural network. We selected a VGG-16 architecture pre-trained on the ImageNet database for 1000-way object categorization (*Deep Learning Toolbox, MATLAB 2020b, Deep Learning Toolbox Model for VGG-16 Network*). All stimuli were resized to 224 x 224 x 3 pixels while maintaining aspect ratio and given as input to the network. We extracted activation vectors from the first layer to the final fully connected layer (layer 37, final fully

connected layer, 'fc8', excluding two 'dropout' layers – layer 35 and layer 38), and took them as the representation of the image in each layer. Distances between pairs of images were measured by calculating the 'city-block' distance between the feature vectors in each layer. We obtained qualitatively similar results with the Euclidean distance.

Amodal Completion Index. To track the representations of occluded, likely, and unlikely completions across layers, we defined an Amodal Completion Index (ACI) as follows:

$$ACI = r(d_{\text{Occluded}}, d_{\text{Likely}}) - r(d_{\text{Occluded}}, d_{\text{Unlikely}})$$

where r is the Pearson's correlation, d_{Occluded} is the set of pairwise occluded display dissimilarities, d_{Likely} is the set of likely completion dissimilarities and d_{Unlikely} is the set of unlikely completion dissimilarities. Because occluded-likely correlation is larger than occluded-unlikely correlation in humans, the ACI for humans will be a positive number. A positive value for the deep network in each layer would thus suggest a representation that is similar to humans.

Data Analysis. Custom MATLAB codes for plotting shaded error bars in Figure 5C and Figure 6F were provided by Kelly Kearney (<https://github.com/kakearney/boundedline-pkg>).

RESULT

Amodal completion of simple shapes by deep networks

For the simple shapes in Experiment 1, we first compared the representation of occluded, likely, unlikely and mosaic displays in the final output layer of the VGG-16 neural network (layer 37, final fully connected layer, 'fc8'). The occluded dissimilarities matched best with the mosaic dissimilarities, matched moderately with the likely dissimilarities and least with the unlikely dissimilarities (Figure 6A). All comparisons were statistically significant using the same bootstrap resampling as before. We obtained analogous results using squared residual error (Figure 6B). Thus, even in the deep network VGG-16, occluded displays activate likely completions more than unlikely completions, which resembles the finding observed in humans (Figure 4).

Next, we investigated the evolution of amodal completion across layers by calculating the Amodal Completion Index (ACI). The ACI remained close to zero across layers and became positive only in the fully connected layers (Figure 6C).

Amodal completion for complex shapes by deep networks

For the complex shapes in Experiment 1, we first compared the representation of occluded displays with likely, unlikely and mosaic displays. The occluded display representation matched best with mosaic completions as before, and relatively weakly with both likely and unlikely completions (Figure 6D). Here again, we obtained analogous results using squared residual error (Figure 6E). This pattern is qualitatively different compared to humans in Experiment 2 (Figure 5).

469 Next, we characterized the evolution of the amodal completion by plotting the
470 ACI across layers for these shapes (Figure 6F). Surprisingly, the ACI was initially above
471 zero but decreased across layers and converged to zero by the fully connected layers.
472 This means that occluded displays do not activate likely completions any stronger than
473 unlikely completions in the deep network VGG-16, unlike humans.

Figure 6: Occlusion in deep neural networks

- (A) Correlation with occluded display dissimilarities for likely (*blue*), unlikely (*red*) and mosaic (*purple*) dissimilarities, calculated based on distances between feature activations in the final fully convolutional layer of the deep network for Experiment 1 stimuli (simple shapes). Bars show mean $\pm$ s.d. based on resampling image pairs with replacement 1,000 times, and asterisks represent the fraction of times the observed trend was violated (** indicates $p < 0.005$).
- (B) Squared residual error in predicting occluded dissimilarities using likely (*blue*), unlikely (*red*) and mosaic (*purple*) dissimilarities. Asterisks indicate statistical significance as indicated by a sign-rank test across $16C_2 = 120$ feature activation dissimilarities (* is $p < 0.05$, ** is $p < 0.005$ etc).
- (C) Amodal completion index (ACI) for Simple shapes (Experiment 1) stimuli across the 37 layers of deep network. Human perception is a positive number (*shaded green region*) and the value computed from Experiment 1 is shown (*dashed line*). Shaded regions around each line is the standard deviation of the ACI based on resampling

491 stimulus pairs 1,000 times with replacement from each layer. Symbols represent layer
492 type in the deep network (open circles: convolutional layer; filled circles: ReLU layer;
493 diamonds: max pooling layer; squares: fully connected layer).
494 (D-F) Same as panels A-C for complex shapes in Experiment 2.

SEARCH COMPARISONS ACROSS DISPLAY TYPES

The above results are based on comparing visual search within display types to avoid interpretational difficulties due to search asymmetries. To validate this approach, we performed an additional experiment involving searches between one display type to another (e.g., occluded vs mosaic or mosaic vs occluded). This revealed strong asymmetries that render the results difficult to interpret, as we had originally predicted (Section S3). By contrast the within-display type searches reported in Experiments 1-2 had far fewer asymmetries, allowing for much simpler interpretations.

GENERAL DISCUSSION

We set out to examine the fidelity of amodal completion using a series of visual search experiments on complex shape. Our main finding is that occluded display representations matched better with likely completions whose statistics match the visible portions compared to unlikely completions with inconsistent statistics. These findings suggest that amodal completion may be more complex than previously thought. Below we discuss these findings in context of the existing literature.

Our main finding that amodal completion involves extrapolating the statistics of the visible portions is consistent with previous findings that amodal completion favors global completions driven by symmetry (Sekuler and Palmer, 1992; Sekuler, 1994; van Lier et al., 1994; Tse, 1998, 1999). Our results go even further to show that amodal completion is driven not just by symmetry, but a more general process of extrapolating the statistics of the visible features. Indeed, we have found that occluded display representations match better with likely compared to unlikely completions for a variety of simple and complex shapes.

We have also found that occluded representations match equally well with likely as well as mosaic completions, suggesting that both representations are activated equally strongly. However, the match between occluded and mosaic is unsurprising because any pair of occluded displays differ in exactly the same features as the equivalent pair of mosaic displays. By contrast the match between occluded and likely completion is surprising because the likely completion contains many additional features. More generally, this finding is consistent with suggestions that amodal completion activates

multiple parallel representations including the mosaic and likely completions (van Lier, 1999; Rauschenberger and Yantis, 2001; De Wit and Van Lier, 2002).

We have used a novel approach to understand amodal completion. Previous studies have compared occluded and unoccluded displays that differ not only in occlusion but also in the spatial separation of items, or in the presence of prominent new features whose effect on the shape representation is unclear, and which can result in strong search asymmetries (Rensink and Enns, 1998; Rauschenberger and Yantis, 2001; Rauschenberger et al., 2004). Our approach overcomes this confound by comparing within occluded displays or within unoccluded (likely and unlikely) displays. This allows us to ask whether the perceived arrangement of occluded displays match better with that of likely, unlikely or mosaic completions.

One major limitation with our approach is that it can be used for a collection of shapes but cannot elucidate the amodal completion of a particular shape. One way to overcome this limitation by characterizing how a particular shape is represented relative to a variety of other fixed shapes with known completions. Nonetheless the fact remains that our approach rests on measuring dissimilarity, which inevitably involves more than one shape. Perhaps a deeper understanding of amodal completion might come from recordings of single neurons in the early visual cortex, where amodally completed contours can be detected by placing their small receptive fields at various locations on the occluder. Even then, the possibility remains that there is no single amodal completion but an entire distribution of completions that are activated in a graded fashion according to their consistency with the visible portions.

548 Finally, our results suggest that pre-trained feedforward deep networks trained for
549 object categorization do not amodally complete occluded displays in the same way that
550 humans do. It is possible that our results arise due to the stimuli differences (abstract vs
551 real images) or due to the lack of occlusion instances in the training images for this
552 network. But, even with real images, deep networks are heavily compromised by
553 occlusion (Zhu et al., 2019). There have been attempts to train deep networks to be
554 invariant to object occlusions (Tang et al., 2018; Kang and Druckmann, 2020; Kortylewski
555 et al., 2020) though it is not clear if these occlusion-invariant networks would match better
556 with likely more than unlikely completions. We propose that our comparison of occluded
557 with likely and unlikely completions can be added to the list of qualitative differences
558 between deep networks and brains as reported previously (Jacob et al., 2021).

559

REFERENCES

- Arun SP (2012) Turning visual search time on its head. *Vision Res* 74:86–92.
- Bridges D, Pitiot A, MacAskill MR, Peirce JW (2020) The timing mega-study: comparing a range of experiment generators, both lab-based and online. *PeerJ* 8:e9414.
- Chen S, Schnabl L, Müller HJ, Conci M (2018) Amodal Completion of a Target Template Enhances Attentional Guidance in Visual Search. *Iperception* 9.
- de Leeuw JR, Motz BA (2016) Psychophysics in a Web browser? Comparing response times collected with JavaScript and Psychophysics Toolbox in a visual search task. *Behav Res Methods* 48:1–12.
- De Wit TCJ, Van Lier RJ (2002) Global visual completion of quasi-regular shapes. *Perception* 31:969–984.
- Deng J, Dong W, Socher R, Li L-J, Kai Li, Li Fei-Fei (2010) ImageNet: A large-scale hierarchical image database. In: 2009 IEEE Conference on Computer Vision and Pattern Recognition, pp 248–255.
- Grootswagers T (2020) A primer on running human behavioural experiments online. *Behav Res Methods* 52:2283–2286.
- Hadad BS, Kimchi R (2018) Perceptual completion of partly occluded contours during childhood. *J Exp Child Psychol* 167:49–61.
- Jacob G, Arun SP (2020) How the forest interacts with the trees: Multiscale shape integration explains global and local processing. *J Vis* 20:20.
- Jacob G, Pramod RT, Katti H, Arun SP (2021) Qualitative similarities and differences in visual object representations between brains and deep networks. *Nat Commun* 12:1872.
- Johnson JS, Olshausen BA (2005) The recognition of partially visible natural objects in the presence and absence of their occluders. *Vision Res* 45:3262–3276.
- Kang B, Druckmann S (2020) Object recognition under occlusion revisited: Elucidating algorithmic advantages of recurrent computation. *bioRxiv*:44.
- Kortylewski A, He J, Liu Q, Yuille A (2020) Compositional convolutional neural networks: A deep architecture with innate robustness to partial occlusion. *Proc IEEE Comput Soc Conf Comput Vis Pattern Recognit*:8937–8946.
- Mitchell S (1988) *Tao te ching: A new English version*. Harper Collins.
- Murray RF, Sekuler AB, Bennett PJ (2001) Time course of amodal completion revealed by a shape discrimination task. *Psychon Bull Rev* 8:713–720.
- Nanay B (2018) The Importance of Amodal Completion in Everyday Perception. *Iperception* 9:2041669518788887.
- Peirce J, Gray JR, Simpson S, MacAskill M, Höchenberger R, Sogo H, Kastman E, Lindeløv JK (2019) PsychoPy2: Experiments in behavior made easy. *Behav Res Methods* 51:195–203.
- Pramod RT, Arun SP (2014) Features in visual search combine linearly. *J Vis* 14:1–20.
- Pramod RT, Arun SP (2016) Object attributes combine additively in visual search. *J Vis* 16:8.
- Rajaei K, Mohsenzadeh Y, Ebrahimpour R, Khaligh-Razavi SM (2019) Beyond core object recognition: Recurrent processes account for object recognition under occlusion. *PLoS Comput Biol* 15:1–30.
- Rauschenberger R, Peterson MA, Mosca F, Bruno N (2004) Amodal completion in visual search: preemption or context effects? *Psychol Sci* 15:351–355.

- Rauschenberger R, Yantis S (2001) Masking unveils pre-amodal completion representation in visual search. *Nature* 410:369–372.
- Rensink RA, Enns JT (1998) Early completion of occluded objects. *Vision Res* 38:2489–2505.
- Sekuler AB (1994) Local and global minima in visual completion: effects of symmetry and orientation. *Perception* 23:529–545.
- Sekuler AB, Palmer SE (1992) Perception of Partly Occluded Objects: A Microgenetic Analysis. *J Exp Psychol Gen* 121:95–111.
- Serre T (2019) Deep Learning: The Good, the Bad, and the Ugly. *Annu Rev Vis Sci* 5:399–426.
- Simonyan K, Zisserman A (2014) Very Deep Convolutional Networks for Large-Scale Image Recognition. *arXiv* 1409.1556.
- Sunder S, Arun SP (2016) Look before you seek: Preview adds a fixed benefit to all searches. *J Vis* 16:3.
- Tang H, Schrimpf M, Lotter W, Moerman C, Paredes A, Ortega Caro J, Hardesty W, Cox D, Kreiman G (2018) Recurrent computations for visual pattern completion. *Proc Natl Acad Sci U S A*.
- Thielen J, Bosch SE, van Leeuwen TM, van Gerven MAJ, van Lier R (2019) Neuroimaging Findings on Amodal Completion: A Review. *Iperception* 10:2041669519840047.
- Tse PU (1998) Illusory volumes from conformation. *Perception* 27:977–992.
- Tse PU (1999) Complete mergeability and amodal completion. *Acta Psychol (Amst)* 102:165–201.
- van Lier R (1999) Investigating global effects in visual occlusion: from a partly occluded square to the back of a tree-trunk. *Acta Psychol (Amst)* 102:203–220.
- van Lier R, van der Helm P, Leeuwenberg E (1994) Integrating global and local aspects of visual occlusion. *Perception* 23:883–903.
- van Lier RJ, Leeuwenberg EL, van der Helm PA (1995) Multiple completions primed by occlusion patterns. *Perception* 24:727–740.
- Vighneshvel T, Arun SP (2013) Does linear separability really matter? Complex visual search is explained by simple search. *J Vis* 13:1–24.
- Zhu H, Tang P, Park J, Park S, Yuille A (2019) Robustness of Object Recognition under Extreme Occlusion in Humans and Computational Models. *arXiv* 1905.04598.

ACKNOWLEDGEMENTS

This research was supported by the DBT/Wellcome Trust India Alliance Senior Fellowship awarded to SPA (Grant# IA/S/17/1/503081) and by a fellowship from the Indian Council of Medical Research (ICMR) to TC.

DATA AVAILABILITY

All data and code required to reproduce the results are publicly available at <https://osf.io/8e9m3/>

1
2
3
4
5
6

SUPPLEMENTARY MATERIAL

for

What do we see behind an occluder?

Amodal completion of statistical properties in complex objects

Thomas Cherian & SP Arun

Centre for Neuroscience, Indian Institute of Science, Bangalore

SECTION S1. STIMULI USED IN ALL EXPERIMENTS

Figure S1: Stimuli used in Experiments 1 & 2. Likely and Unlikely (and Mosaic in experiment 2) shapes were always flanked by the occluder in the actual experiments but are not shown here due to space constraints.

SECTION S2. ANALYSIS OF SYMMETRY

The finding that occlusion display dissimilarity matches likely completions better than unlikely completions could be driven strongly by some image property common to most shapes. One such possible property is symmetry since many shapes in both experiments were symmetric. To investigate this possibility, we repeated the above analysis for symmetric and asymmetric shapes separately. The results are summarized in Figure S2.

In Experiment 1, there were 8 symmetric and 8 asymmetric shapes. Occluded dissimilarities matched better with likely completion dissimilarities for both groups (Figure S2A) and yielded smaller squared residual error (Figure S2B). However, the difference was statistically significant only for symmetric shapes ($p < 0.0005$) but not for asymmetric shapes ($p = 0.21$).

In Experiment 2, there were 10 symmetric and 6 asymmetric shapes. Occluded dissimilarities matched better with likely and mosaic, compared to unlikely dissimilarities in both groups (Figure S2C) and yielded smaller squared residual error (Figure S2D). The difference in residual error was statistically significant for both symmetric shapes (sign-rank test for difference between groups, $p < 0.005$ between likely and unlikely dissimilarities and $p < 0.005$ between mosaic and unlikely dissimilarities) and for asymmetric shapes (sign-rank test for difference between groups, $p < 0.05$ between likely and unlikely dissimilarities and $p < 0.005$ between mosaic and unlikely dissimilarities).

Figure S2. Amodal completion in symmetric and asymmetric objects

- (A) Mean and s.d. of bootstrapped correlation to occluded dissimilarities across 1000 samples with replacement for both likely (blue) and unlikely (red) dissimilarities in Experiment 1. Asterisks indicate statistical significance (number of times observed trend is violated across the 1000 bootstrap samples, ** equals $p < 0.005$)
- (B) Squared residual error in predicting occluded dissimilarities using likely (blue) and unlikely (red) dissimilarities. Asterisks indicate statistical significance as indicated by a sign-rank test across ${}^8C_2 = 28$ search pairs (* is $p < 0.05$, ** is $p < 0.005$ etc).
- (C) Same as (A) but for symmetric and asymmetric shapes in Experiment 2 (mosaic dissimilarities in purple).
- (D) Same as (B) but for symmetric and asymmetric shapes in Experiment 2. ${}^{10}C_2 = 45$ search pairs for symmetric shapes and ${}^6C_2 = 15$ search pairs for asymmetric shapes in Experiment 2. Asterisks as in (B)

SECTION S3. COMPARING SEARCHES ACROSS DISPLAY TYPES

Classic studies of amodal completion in visual search have shown that searching for an occluded display among completed displays is harder than searching for a mosaic display among the same distractors. However, this difference could potentially arise because the mosaic display has a prominent additional feature (the repeated edge present in both occluder and the occluded object), and not necessarily because the occluded display activates the completed display. Indeed, there are search asymmetries reported in these studies: finding a mosaic among completed is easy, but finding completed among mosaic is hard (Rensink and Enns, 1998; Rauschenberger and Yantis, 2001; Rauschenberger et al., 2004). These asymmetries complicate the interpretation of visual search exclusively as a measure of similarity.

To investigate these issues, we performed a visual search experiment with searches across display type. We selected the following target-distractor search pairs of interest: occluded-cut, occluded-likely, occluded-unlikely, cut-likely, cut-occluded, likely-occluded, unlikely-occluded and likely-cut.

METHOD

The general methodology was identical to Experiment 1, and only details different in this experiment are described below.

Participants. A total of 16 individuals (7 female, age, mean $\pm$ s.d.: 27.5 ± 3.7 years) participated in this experiment.

Stimuli. We used all stimuli from Experiments 1 and 2. Along with the Occluded, Likely and Unlikely displays, we also created mosaic displays from the likely displays by cutting out the top left quadrant of the shapes (the exact portion covered by the occluder in occluded display).

Search Pairs. For each of the 32 shapes, we created all possible ($n = {}^4C_2 = 6$) search pairs among the 4 conditions (Occluded, Cut, Likely and Unlikely). This resulted in a total of $32 \times 6 = 192$ search pairs. We created two search arrays for each pair, with either item in each pair as target or distractor, resulting in a total of 384 trials. All searches were interleaved randomly and shown in a single block.

RESULT

Participants were highly accurate across all searches (accuracy, mean $\pm$ sd across subjects: $95\% \pm 4\%$). They were also highly consistent in their response times across searches (split-half correlation between average response times on odd- & even-numbered participants: $r = 0.87$).

Similarity relations in the presence of asymmetries

Example searches from this experiment are shown in Figure S3A-D, and the average search times for all pairs of search types is shown in Figure S3E. It can be seen that searching for an occluded target among likely completions (Figure S3A) is harder than finding a mosaic target among likely completions (Figure S3B). This finding

replicates the classic result reported by Rensink and Enns (1998), suggesting that the occluded display activates the likely completion more strongly than the mosaic completion.

However, the exact opposite interpretation is suggested by comparing search for an occluded target among likely completions (Figure S3A) with searching for an occluded target among mosaic completions (Figure S3C). Here, the latter search is harder, implying that the occluded display activates the mosaic interpretation more strongly. This confirms our initial premise that comparing searches across display type can be problematic.

Proceeding further, it can be seen that search for occluded among likely is harder than mosaic among likely, but reversing the identity of the target and distractor leads to a much weaker effect (Figure S3E).

Similarity relations averaging across asymmetries

Finally, if we simply ignore search asymmetries and compare occluded-likely, occluded-cut and occluded-unlikely searches, we find that all three conditions are equivalent in their difficulty (Figure S3F).

To visualize how these display types are related in visual search space (averaging across asymmetries), we performed a multidimensional scaling analysis on the inverse of the average search times for each pair of conditions. In the resulting plot, shown in Figure S3G, nearby items represent hard searches. It can be seen that the likely and unlikely completions are extremely close by, which is expected since they differ only in the upper left quadrant. Importantly, the occluded displays are equidistant from mosaic and likely/unlikely completions. These large intrinsic differences between occluded, mosaic and likely/unlikely groups make it difficult to reveal subtle representational differences.

Search asymmetries across-display types compared to within-display types

We wondered whether there was any difference in the frequency of search asymmetries in this experiment (across-display types) compared to the within-display type searches in Experiments 1 & 2. To this end, we compared the response times corresponding to each item in a search pair as target with the other as distractor, using a sign-rank test (with criterion = 0.05).

For the simple shapes used from Experiment 1, we obtained 8% (8/96) asymmetries when searches were performed across-display types in this experiment. By contrast, the within-display type searches in Experiment 1 using the same shapes had more asymmetries ($51/360 = 14\%$). However, this difference was not statistically significant using a chi-squared test comparing the observed number of asymmetries with the number expected from the incidence in Experiment 1 ($p = 0.14$).

For the complex shapes used from Experiment 2, we obtained 34% (33/96) asymmetries when searches were performed across-display types. By contrast, the within-display type searches in Experiment 2 had far fewer asymmetries ($84/480 = 17\%$). This difference was statistically significant using a chi-squared test comparing the observed number of asymmetries with the number expected from the incidence in Experiment 2 ($p < 0.00005$). A closer investigation showed that 83 of the 84 asymmetries involved pairs of mosaic shapes. Thus, there were very few asymmetries involving

searches within occluded/likely/unlikely displays ($1/360 = 0.2\%$), and more asymmetries across display type.

The above results suggest that there are strong search asymmetries in any search involving mosaic displays for complex shapes, and this renders any across-display type comparisons uninterpretable.

Figure S3. Visual search across display types (Experiment 3).

(A) Example search for an occluded target among likely completions.

(B) Example search for a mosaic completion among likely completions.

(C) Example search for an occluded target among mosaic completions.

(D) Example search for a mosaic completion among occluded displays.

(E) Average (mean $\pm$ s.e.m.) search times for all searches tested in this experiment. All statistical comparisons were performed using a sign-rank test. Asterisk indicates statistically significant differences between groups being compared (* is $p < 0.05$, ** is $p < 0.005$ etc.).

(F) Average search times across key searches of interest averaging across search asymmetries.

(G) Similarity relations between display types obtained using multidimensional scaling. In the resulting plot, nearby points represent hard searches. For instance, search for a likely completion among unlikely completions (or vice-versa), is hard compared to all other comparisons.

SUPPLEMENTARY REFERENCES

- Rauschenberger R, Peterson MA, Mosca F, Bruno N (2004) Amodal completion in visual search: preemption or context effects? *Psychol Sci* 15:351–355.
- Rauschenberger R, Yantis S (2001) Masking unveils pre-amodal completion representation in visual search. *Nature* 410:369–372.
- Rensink RA, Enns JT (1998) Early completion of occluded objects. *Vision Res* 38:2489–2505.