Published in final edited form as:

Psychol Sci. 2014 June ; 25(6): 1206–1217. doi:10.1177/0956797614530725.

A perceptually completed whole is less than the sum of its parts

Jason M. Gold

Department of Psychological and Brain Sciences, Indiana University, 1101 East 10th St.,
Bloomington, Indiana 47405, USA

Abstract

How efficiently do we integrate across the disconnected image fragments that fall upon our eyes when we view partly occluded objects? We used a psychophysical summation-at-threshold technique to address this question by measuring discrimination performance with both isolated and combined features of physically fragmented but perceptually complete objects. If visual completion promotes superior integration efficiency, we would expect performance with a visually completed object to exceed what would simply be expected from performance with the individual object parts shown in isolation. Contrary to this prediction, we find that discrimination performance with both static and moving versions of physically fragmented but perceptually complete objects is significantly worse than would be expected from performance with their constituent parts. These results present a challenge for future theories of visual completion.

Keywords

visual completion; perceptual organization; information integration

Introduction

A fundamental challenge faced by any real-world visual pattern recognition system is the ability to accurately identify objects that are partially hidden from view (Pessoa, Thompson & Noe, 1998). There is considerable evidence that the human visual system solves this problem by interpolating missing image fragments in order to construct unified object representations (Gold, Murray, Bennett & Sekuler, 2000; Kellman, Yin & Shipley, 1998; Ringach & Shapley, 1996; Sekuler & Palmer, 1992; von der Heydt, Peterhans & Baumgartner, 1984). Traditional Gestalt views of perceptual grouping and organization (Koffka, 1935) maintain that the formation of such a representation brings with it certain emergent properties, such that a visually completed whole is perceived to be something different than just the sum of its individual parts. Although this view may offer a compelling description of our phenomenological experience when we encounter partially occluded objects, it is less clear how these properties are manifested at the level of behavioral performance in object recognition tasks.

Previous experiments have demonstrated that presenting features in configurations or contexts where visual completion is thought to take place can promote faster reaction times and higher accuracies in discrimination tasks. For example, Sekuler and Palmer (1992) found that when observers were primed with a partly occluded version of an object, they were faster to make subsequent decisions about filled-in versions of the object than incomplete versions matching the stimulus they had seen previously. Similarly, Pomerantz and colleagues (Eidels, Townsend & Pomerantz, 2008; Pomerantz, 2003; Pomerantz & Portillo, 2011; Pomerantz & Pristach, 1989; Pomerantz, Sager & Stoeve, 1977) have found that reaction times for discriminating amongst sets of items can be much faster when presented within the context of features that are entirely redundant but that elicit the percept of a perceptually complete figure (what they call the ‘configural superiority effect’). Ringach and Shapley (1996) have found that orienting the elements of a Kanisza figure so that observers no longer perceive illusory or occluded contours can result in dramatic decreases in discrimination accuracy.

The results of such experiments suggest that an observer’s ability to integrate information across the disparate image fragments of a partly occluded object is aided by visual completion. Thus, we might expect observers’ performance with perceptually complete wholes to exceed what one would simply expect from their ability to use each image fragment shown in isolation. One approach that has been used successfully to test this kind of prediction in other domains is the ‘summation-at-threshold’ technique (Gold, Mundy & Tjan, 2012; Graham, Robson & Nachmias, 1978; Nandy & Tjan, 2008). This approach involves measuring observers’ contrast sensitivities (i.e., the reciprocal of their contrast thresholds) when discriminating the isolated elements of a set of stimuli and using them to quantitatively predict what their sensitivity should be when discriminating the combined versions of the stimuli. Under certain conditions, if an observer simply uses the information from each element in a constant fashion, regardless of whether they appear in isolation or in combination, it can be shown (Nandy & Tjan, 2008) that the sum of their squared sensitivities to each of the individual elements should equal their squared sensitivity to the combined stimulus (referred to as ‘optimal’ integration). If presenting the elements in combination allows an observer to make better use of information than would be predicted by their performance with the individual elements, their squared sensitivity to the combination should exceed the sum of their individual squared sensitivities (‘super-optimal’ integration). Alternatively, if there is a significant cost to processing all of the elements in combination, their squared sensitivity to the combination should fall below the sum of their individual squared sensitivities (‘sub-optimal’ integration). These predictions can be conveniently expressed by computing an integration index Φ :

$$\Phi = S_{combined}^2 / \sum_{i=1}^n S_{part_i}^2 \quad (1)$$

Here, S denotes sensitivity (i.e., $1/\text{contrast threshold}$) and n equals the number of individual parts that make up a combined stimulus. $\Phi = 1$ indicates optimal integration, $\Phi > 1$ indicates super-optimal integration, and $\Phi < 1$ indicates sub-optimal integration (Gold et al., 2012).

Given the above, we applied this summation-at-threshold technique to a series of tasks in which the stimuli were physically fragmented but perceptually complete. We reasoned that, if the process of visual completion enhances an observer's ability to make use of information carried by the individual elements of an object, we would expect to find super-optimal integration for perceptually complete but not perceptually fragmented figures.

Method

Participants

Three volunteers between the ages of 19 and 42 as well as the author participated in the study. All had normal or corrected to normal visual acuity, and all provided consent within a protocol approved by the Indiana University Internal Review Board.

Stimuli

All four observers performed a series of discrimination tasks in which objects appeared as either perceptually complete or as collections of disconnected fragments (Figure 1). One of the requirements for making the prediction described in Equation 1 is that the isolated fragments be orthogonal to each other (i.e., their dot-product must be equal to zero; Nandy & Tjan, 2008). Thus, we designed three different tasks whose features were spatially non-overlapping, and thus met this requirement.

In the 'Bent Bar' task (Figure 1a), two pac-man-like circles, each with a rectangular section missing from one side, were displayed with either their mouths facing each other or facing in the same direction. Each pac-man was slightly rotated clockwise or counter-clockwise. When the mouths of the pac-men were facing each other, it created the percept of the presence of an occluding bar that was slightly bent either to the left or right ('Complete' stimuli; Figure 1a, bottom left). When the mouths of the pac-men were facing the same direction, no percept of an illusory bar was perceived, and both pac-men simply appeared to be slightly oriented to the left or right ('Fragmented' stimuli, Figure 1a, bottom right). For each of these pairs of stimuli, we varied the contrast of the images across trials and asked observers to classify them as either oriented towards the 'left' or 'right', in order to obtain 71% correct contrast discrimination thresholds. To make the task more difficult, we also embedded the stimuli on each trial in a random sample of Gaussian white pixel noise (see Supplementary Figure S1 for a demonstration of Figure 1a with added noise at near-threshold signal contrast). In addition to these two main Bent Bar conditions, we tested each observer's ability to classify the individual pac-man elements that made up the stimuli ('Top Only' and 'Bottom Only', Figure 1a, top and middle rows). Thus, there were three conditions tested for each kind of stimulus: Combined, Top Only and Bottom Only.

This same approach was applied to two other tasks that were based upon the occluded rotating square tasks developed by Lorenceau and Shiffrar (1992) and Murray, Sekuler & Bennett (2001). In the 'Rotating Square' task (Figure 1b), a set of four white line segments rotated either clockwise or counter-clockwise. When a set of four solid black squares appeared at the corners of the stimulus, the line segments appeared to unite into a single rotating square that was partially occluded by four black corner elements ('Complete'

condition; Figure 1b, bottom left). When these corner elements were painted the same color as the background (mid-gray), the square appeared as a set of four disconnected rotating fragments ('Fragmented' condition; Figure 1b, bottom right). Much like in the Bent Bar task, observers were asked to classify the stimulus as rotating either 'clockwise' or counter-clockwise', with either all of the elements present ('Combined') or just a single element present ('Bottom Only', 'Top Only', 'Left Only', 'Right Only').

The 'Shrinking/Expanding Square' task (Figure 1c) was similar to the Rotating Square task, except that the line segments moved in towards the center of the figure or out towards the edge of the figure. In the presence of the four black corner elements, this created the percept of an occluded shrinking or expanding square ('Complete' condition; Figure 1c, bottom left); in the absence of the corner elements, it appeared as a disconnected set of line segments moving either in or out from the center of the figure ('Fragmented' condition; Figure 1c, bottom right; also see Supplementary Figures S2–S5 for dynamic versions of the Rotating and Shrinking/Expanding Square stimuli, with and without added noise). The specific properties of each kind of stimulus used in the experiments were as follows:

Bent Bars—Each 'pac-man' element of the Bent Bar stimuli was 46 pixels (0.75° of visual angle) in diameter, with a mouth that was 15 pixels (0.25°) wide and 24 pixels (0.39°) deep. Depending upon the condition, each element was rotated by $\pm 5^\circ$ from facing either straight upwards or downwards. In the combined conditions, the distance between the edges of the two elements was 45 pixels (0.74°). In the top and bottom only conditions, the elements would appear in the same locations that they appeared in the corresponding combined condition. In all conditions, the elements were embedded in a background of average luminance that was 147 pixels (2.40°) in height x 64 pixels (1.05°) in width. The background region was surrounded by a dark outline that was 2 pixels (0.03°) in width. The noise that was added to the stimulus covered the entire stimulus region. The stimulus duration was 43 frames (505 ms).

Rotating Squares—Each rotating square was constructed by placing a 76 x 76 pixel ($1.25^\circ \times 1.25^\circ$) outline of rotating square within a 138 x 138 pixel ($2.26^\circ \times 2.26^\circ$) region of average luminance. The corners of the rotating square were removed by replacing the each corner of the 138 x 138 pixel region with a 52 x 52 pixel ($0.85^\circ \times 0.85^\circ$) square that was either darker than the background (for the completed condition) or the same as the background (for the fragmented condition). This produced 4 separate rotating line segments that were 5 pixels (0.08°) thick, and that varied between 32 (0.52°) and 35 (0.57°) pixels in length across the different angles of rotation. Depending upon the location of the segment and the direction of rotation, each swept through a series of 21 angular rotations that covered a 10° range, yielding a rotation rate of $40.5^\circ/\text{sec}$. The stimulus duration was 21 frames (247 milliseconds). As with the Bent Bar stimuli, the line segments in the top, bottom, left and right only conditions appeared in the same locations that they appeared in the combined conditions. A unique sample of noise was added to the entire stimulus region on each frame.

Shrinking/Expanding Squares—Each shrinking or expanding square was constructed by placing the outline of a square within a 216 x 216 ($3.53^\circ \times 3.53^\circ$) pixel region of average luminance. The corners of the square were removed by replacing the each corner of the 216

x 216 pixel region with a 96 x 96 pixel ($1.57^\circ \times 1.57^\circ$) square that was either darker than the background (for the completed condition) or the same as the background (for the fragmented condition). This produced 4 separate line segments that were 5 pixels (0.08°) thick and 24 pixels (0.39°) in length. Depending upon the location of the segment and whether the square was shrinking or expanding, each swept between a distance of 30 and 60 pixels (0.49° and 0.98°) from the center of the display (sampled every 2 pixels (0.03°)), yielding a velocity of 165 pixels/sec (2.70° /sec). The stimulus duration was 31 frames (365 ms). As with Rotating Square stimuli, the line segments in the top, bottom, left and right only conditions appeared in the same locations that they appeared in the combined conditions. A unique sample of noise was added to the entire stimulus region on each frame.

All stimuli were shown on a CRT display with MATLAB, using in-house software and the extensions provided by the psychophysics toolbox (Brainard, 1997). The resolution of the display was 1024 x 768 pixels, which subtended $73.6^\circ \times 77.6^\circ$ from a viewing distance of 130 centimeters. The frame rate of the display was 85 Hz. Luminance on the display ranged from 5.5 to 120.5 cd/m², with an average (background) luminance of 38 cd/m². The contrast noise added to the stimuli on each trial was white and Gaussian, and had a standard deviation of 0.032 (noise spectral density = $2.66\text{e-}07$).

Threshold and Sensitivity Measurement

Thresholds in each condition were measured by varying the RMS contrast of the stimuli across trials using a 2-down, 1-up adaptive staircase procedure. For the bent-bar task, 200 trials were measured per observer in each condition (i.e., top only, bottom only, combined) and stimulus type (i.e., complete, fragmented). For the rotating and shrinking/expanding squares, 150 trials were measured in each condition and stimulus type (only 150 trials were measured in each condition due to the additional time it took to generate dynamic rather than static noise on each trial). Weibull psychometric functions were fit to the staircase data for each condition, and threshold was defined as the contrast that yielded 71% correct performance (chosen for its location at approximately the center of the psychometric function). The contrast of an isolated stimulus feature was defined as the contrast of the combined stimulus from which the feature was extracted (called 'nominal contrast'). That is, the combined stimulus was first set to the specified level of contrast and then individual features were extracted from this image according to the specific condition. Sensitivity was then computed as $1/\text{contrast threshold}$.

Experimental Procedure and Design

On a given trial, a fixation point appeared at the center of the screen, followed by the stimulus display for the specified duration and a selection window. In the case of the bent-bar task, high contrast, noise-free versions of the two possible choices were shown (e.g., bottom feature rotated counter-clockwise, bottom feature rotated clockwise) and corresponding response keys were displayed beneath the choices. For the rotating and expanding square tasks, text that described each possible responses (i.e., 'clockwise' vs. 'counter-clockwise'; 'in' vs. 'out') were shown, along with a corresponding response key. Observers were given unlimited time to make a key press response, and were given auditory accuracy feedback. For a given stimulus type in a given task, trials for the combined and

individual feature conditions were randomly intermixed within a given session. A random task order was assigned to each observer.

Results & Discussion

Squared sensitivities and integration indexes for each observer as well as the means across observers in each condition are plotted in Figures 2 and 3. Surprisingly, these data show that none of the observers exhibited super-optimal integration in any of the conditions where the stimuli were perceived as being complete. Instead, mean integration efficiency for the complete stimuli was significantly *less* than the prediction of optimal summation for all three tasks (one-sample t-test (2-tailed) for Bent Bars: $t(3) = 5.76$, $p = .01$; Rotating Squares: $t(3) = 19.79$, $p < .001$; and Shrinking/Expanding Squares: $t(3) = 7.98$, $p < .005$). Perhaps even more surprising is the fact that integration efficiency for the fragmented stimuli was generally *higher* than for the corresponding completed versions of each stimulus type. A 2 (completeness) $\times$ 3 (stimulus type) repeated measures ANOVA confirmed that there was a significant effect of completeness ($F(1,3) = 14.49$, $p < .05$), as well as an effect of stimulus type ($F(2,2) = 226.75$, $p < .01$), with no significant interaction between these two factors ($F(2,2) = .98$, $p = .51$).

So how might we explain these surprising results? One possibility is that the presence of externally added noise coupled with threshold signal contrast may have led observers to adopt a strategy where they only relied upon the single feature to which they were most sensitive, i.e.:

$$\Phi_{best\ feature} = S_{combined}^2 / \arg \max_{i \in [1,n]} (S_{part_i}^2) \quad (1)$$

The predictions of such a ‘best-feature’ model are plotted as small triangles on top of each panel in Figure 3. The best-feature model actually does a fairly good job at predicting integration efficiency when the combined stimuli are perceived as being complete. However, this was not the case for the fragmented stimuli, where the best feature model generally *under* predicted integration efficiency.

Although we cannot rule out the possibility that observers were using a best-feature strategy when discriminating amongst the complete but not the fragmented stimuli, it does force us to draw the awkward conclusion that observers used only a single feature when engaging in visual completion, yet relied on multiple features when a stimulus was perceived as being composed of disconnected fragments. Additionally, there are good reasons to believe that adding external noise and presenting signals at low contrast generally have little impact upon observers’ strategies. First, previous experiments using reverse correlation have shown that observers engage in completion under nearly identical conditions as those used in the current experiment (Gold et al., 2000). Second, many experiments have demonstrated the observers’ contrast energy thresholds for detecting, discriminating and identifying a wide variety of visual patterns (i.e., gratings, objects, faces) are linearly related to the noise spectral density of an externally added noise (Pelli & Farrell, 1999). This pattern of results is

exactly what one would predict if observers were employing a strategy that was independent of stimulus contrast (Pelli, 1981).

Another possible explanation for why we found such poor integration efficiency with our perceptually complete stimuli is that the process of interpolation itself may have paradoxically served to counteract any potential benefits gained from feature binding. The finding that visual completion can serve to hinder performance under certain circumstances is not without precedence. In particular, decrements in performance have been reported when perceptual completion is placed into conflict with other cues, such as stereoscopic depth (Hou, Lu, Zhou & Liu Liu, 2006; Jacobs & Basri, 1999). Given the strong behavioral (Gold et al., 2000; Keane, Lu & Kellman, 2007; Ringach & Shapley, 1996; Sekuler & Palmer, 1992), physiological (Pillow & Rubin, 2002; von der Heydt et al., 1984) and theoretical (Grossberg & Mingolla, 1985; Kellman & Shipley, 1991) support for the idea that illusory and occluded contours are added into an observer's internal representation of a stimulus, we would predict that interpolated edges should have corresponding direct effects on task performance. In the context of our current set of experiments, recall that all three tasks that we asked observers to perform involved adding random samples of visual noise to each pixel in the display. This noise fell in regions of the stimulus that both carried information (the physically present edges) and carried no information (the regions in between the physically present edges). As a result, if an observer were to compare a completed representation to such a stimulus, the regions of interpolation would only serve to contribute noise to the ultimate response of the visual system, resulting in a reduction in integration efficiency.

We decided to test this possibility by having three new observers and the author perform a slight variant of our Illusory Bent-Bar discrimination task (Figure 4a), in which a 3 pixel (0.05°) wide ring was drawn around the perimeter of each pac-man element. This gave the impression of an occluded rather than illusory right or left pointing bar when the top and bottom elements were shown in combination (e.g., Gold et al., 2000; Ringach & Shapley, 1996). Because the interpolated edges were perceived as continuing behind rather than in front of an intervening background, this new stimulus allowed us to add a 'Local Noise' to just the circular regions where the two pac-man elements could appear (a 'Local Noise') without running the risk of interrupting the interpolation process. We also chose to increase the contrast of the added noise ($\sigma = 0.14$, noise spectral density = 2.68×10^{-4}), in order to insure that it would be limiting observers' performance rather than any internally generated noise (Burgess, Wagner, Jennings & Barlow, 1981; Pelli, 1981). All other aspects of the experiment were the same as in the previous Illusory Bent Bars condition.

Given the above conditions, we reasoned that if observers' use of noisy regions falling between the stimulus elements was responsible for their sub-optimal integration efficiency in our original experiments, these effects should be ameliorated by restricting the presentation of noise to just the regions where the inducing elements appeared. Contrary to this prediction, we found that integration efficiency was sub-optimal for all four observers in the presence of Local Noise, as well as in the presence of a Global Noise that covered the entire rectangular region within which the stimulus appeared (Figure 4b-c). One-sample t-tests (2-tailed) confirmed that the mean index in the presence of both kinds of noise was

significantly less than the prediction of optimal summation (Local Noise: $t(3) = 10.74$, $p < .002$; Global Noise: $t(3) = 27.65$, $p < .0001$) and a correlated two-sample t-test (2-tailed) showed that the difference in integration efficiency between Local and Global Noise was not significant ($t(3) = 2.49$, $p = .09$). Thus, it appears that the sub-optimal integration efficiency associated with visual completion stems from a source other than observers' reliance on stimulus regions that correspond to perceptually interpolated (but physically uninformative) features.

Of course, it could also be that our summation-at-threshold technique suffers from some unforeseen methodological flaw. Although there is no particular reason to suspect this to be the case, only a handful of previous studies have used the technique (Gold et al., 2012; Graham et al., 1978; Nandy & Tjan, 2008), and none have reported evidence of super-optimal integration. To address this possible concern, we carried out an additional proof-of-concept experiment with a stimulus and task that should unquestionably lead to super-optimal integration. We used a Vernier acuity task, which has been traditionally employed to evaluate the resolving ability of the visual system (Westheimer, 1965). A Vernier acuity task requires observers to determine whether two abutting line segments are either perfectly aligned or slightly misaligned relative to one another (bottom row of Figure 5a). In such a task, observers will undoubtedly rely upon the relative rather than absolute positions of the two lines, due to the limiting effects of intrinsic spatial uncertainty in the human visual system (Tjan, Lestou & Kourtzi, 2006; Zeevi & Mangoubi, 1984). However, this relational cue is not available when only one of the two line segments is presented, as was the case in our 'Top Only' and 'Bottom Only' isolated line segment conditions (top and middle rows of Figure 5a)¹. Thus, we would expect observers to benefit much more from the combined presentation of the two lines segments than would be predicted from their performance with the individual isolated line segments. In fact, this is exactly what we found (Figure 5b–c): integration efficiency was super-optimal for all observers, with a mean integration index of 1.5. A one-sample t-test (2-tailed) confirmed that the mean index was significantly greater than the prediction of optimal summation ($t(3) = 3.32$, $p < .05$). In fact, these data are consistent with recent results reported by Pomerantz and Portillo (2011) and Pomerantz and Cragin (2013) who found that co-linearity of line segments in a manner akin to our Vernier stimuli yielded stronger configural superiority effects than Kanisza-square stimuli defined by illusory contours.

Conclusions

Taken together, our results offer compelling new evidence that there are unanticipated costs associated with the process of visual completion. Admittedly, the tasks and stimuli we used in our experiments were artificially generated to allow us to quantitatively measure the efficiency of feature integration with both perceptually complete and incomplete objects.

¹Each line segment of the Vernier stimulus was 2 pixels (0.03°) thick and 36 pixels (0.59°) in height. For the Aligned stimulus, the top and bottom lines were abutting, such that they appeared as a single continuous line when shown in combination. For the Misaligned stimulus, the bottom line was shifted to the right by 5 pixels (0.08°) and the top line was shifted to the left by 5 pixels. As in the other tasks, the top and bottom only lines would appear in the same locations that they appeared in the corresponding combined condition. In all conditions, the elements were embedded in a background of average luminance that was 128×128 pixels ($2.09^\circ \times 2.09^\circ$) in size. The background region was surrounded by a dark outline that was 2 pixels in width. The noise that was added to the stimulus covered the entire stimulus region. The stimulus duration was 43 frames (505 ms). There were 300 trials per observer in each condition.

However, the history of experimental psychology is filled with similar examples of experiments and demonstrations in which the underlying mechanics of psychological phenomena have been elucidated using tasks and stimuli that only exist in the laboratory. Unlike our relatively artificial laboratory conditions, object recognition in the natural world typically involves dealing with a tremendous degree of uncertainty about various properties of objects, such as their viewpoint, size, reflectance and shape (Tjan et al., 2006). Because the true properties of an object are generally underspecified by the raw data that our visual system receives, strong assumptions must often be made in order to reliably recover these properties and accomplish tasks such as discriminating one object from another (Ramachandran, 1988). It is perhaps under these kinds of conditions that the benefits of visual completion may outweigh the costs, allowing it to play a more constructive role in promoting efficient information integration.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We would like to thank Michael Simmons, the IU Vision Lab, James Pomerantz and two anonymous reviewers for helpful input and suggestions. This research was supported by National Institute of Health grant R01-EY019265.

References

- Brainard DH. The psychophysics toolbox. *Spatial Vision*. 1997; 10(4):433–436. [PubMed: 9176952]
- Burgess AE, Wagner RF, Jennings RJ, Barlow HB. Efficiency of human visual signal discrimination. *Science*. 1981; 214(4516):93–4. [PubMed: 7280685]
- Efron, B.; Tibshirani, R. An introduction to the bootstrap. New York: Chapman & Hall; 1993.
- Eidels A, Townsend JT, Pomerantz JR. Where similarity beats redundancy: The importance of context, higher order similarity, and response assignment. *J Exp Psychol Hum Percept Perform*. 2008; 34(6): 1441–63. [PubMed: 19045985]
- Gold JM, Mundy PJ, Tjan BS. The perception of a face is no more than the sum of its parts. *Psychol Sci*. 2012; 23(4):427–34. [PubMed: 22395131]
- Gold JM, Murray RF, Bennett PJ, Sekuler AB. Deriving behavioural receptive fields for visually completed contours. *Curr Biol*. 2000; 10(11):663–6. [PubMed: 10837252]
- Graham N, Robson JG, Nachmias J. Grating summation in fovea and periphery. *Vision Res*. 1978; 18(7):815–25. [PubMed: 676089]
- Grossberg S, Mingolla E. Neural dynamics of form perception: Boundary completion, illusory figures, and neon color spreading. *Psychol Rev*. 1985; 92(2):173–211. [PubMed: 3887450]
- Hou F, Lu H, Zhou Y, Liu Z. Amodal completion impairs stereo acuity discrimination. *Vision Research*. 2006; 46:2061–2068. [PubMed: 16472836]
- Keane BP, Lu H, Kellman PJ. Classification images reveal spatiotemporal contour interpolation. *Vision Research*. 2007; 47(28):3460–75. [PubMed: 18053850]
- Kellman PJ, Shipley TF. A theory of visual interpolation in object perception. *Cogn Psychol*. 1991; 23(2):141–221. [PubMed: 2055000]
- Kellman PJ, Yin C, Shipley TF. A common mechanism for illusory and occluded object completion. *J Exp Psychol Hum Percept Perform*. 1998; 24(3):859–69. [PubMed: 9627421]
- Koffka, K. Principles of gestalt psychology. New York: Harcourt, Brace and Company; 1935.
- Liu Z, Jacobs DW, Basri R. The role of convexity in perceptual completion: beyond good continuation. *Vision research*. 1999; 39(25):4244–57. [PubMed: 10755161]

- Lorenceanu J, Shiffrar M. The influence of terminators on motion integration across space. *Vision Research*. 1992; 32:263–273. [PubMed: 1574843]
- Murray RF, Sekuler AB, Bennett PJ. Time course of visual completion revealed by a shape discrimination task. *Psychonomic Bulletin & Review*. 2001; 8(4):713–720. [PubMed: 11848590]
- Nandy AS, Tjan BS. Efficient integration across spatial frequencies for letter identification in foveal and peripheral vision. *J Vis*. 2008; 8(13):3, 1–20. [PubMed: 19146333]
- Pelli, DG. Thesis. Cambridge: University of Cambridge; 1981. Effects of visual noise.
- Pelli DG, Farrell B. Why use noise? *Journal of the Optical Society of America A*. 1999; 16(3):647–653.
- Pessoa L, Thompson E, Noe A. Finding out about filling-in: A guide to perceptual completion for visual science and the philosophy of perception. *Behav Brain Sci*. 1998; 21(6):723–48. discussion 748–802. [PubMed: 10191878]
- Pillow J, Rubin N. Perceptual completion across the vertical meridian and the role of early visual cortex. *Neuron*. 2002; 33(5):805–13. [PubMed: 11879656]
- Pomerantz JR. Wholes, holes, and basic features in vision. *Trends Cogn Sci*. 2003; 7(11):471–3. [PubMed: 14585437]
- Pomerantz, JR.; Cragin, AI. Emergent features and feature combination. In: Wagemans, J., editor. *Oxford Handbook of Perceptual Organization*. Oxford University Press; (in press)
- Pomerantz JR, Portillo MC. Grouping and emergent features in vision: Toward a theory of basic gestalts. *J Exp Psychol Hum Percept Perform*. 2011; 37(5):1331–49. [PubMed: 21728463]
- Pomerantz JR, Pristach EA. Emergent features, attention, and perceptual glue in visual form perception. *J Exp Psychol Hum Percept Perform*. 1989; 15(4):635–49. [PubMed: 2531201]
- Pomerantz JR, Sager LC, Stoeber RJ. Perception of wholes and of their component parts: Some configural superiority effects. *J Exp Psychol Hum Percept Perform*. 1977; 3(3):422–35. [PubMed: 886276]
- Ramachandran VS. Perception of shape from shading. *Nature*. 1988; 331(6152):163–6. [PubMed: 3340162]
- Ringach DL, Shapley R. Spatial and temporal properties of illusory contours and amodal boundary completion. *Vision Res*. 1996; 36(19):3037–50. [PubMed: 8917767]
- Sekuler AB, Palmer SE. Perception of partly occluded objects: A microgenetic analysis. *Journal of Experimental Psychology: General*. 1992; 121(1):95–111.
- Sorkin RD, Hays CJ, West R. Signal-detection analysis of group decision making. *Psychol Rev*. 2001; 108(1):183–203. [PubMed: 11212627]
- Tjan BS, Lestou V, Kourtzi Z. Uncertainty and invariance in the human visual cortex. *J Neurophysiol*. 2006; 96(3):1556–68. [PubMed: 16723410]
- von der Heydt R, Peterhans E, Baumgartner G. Illusory contours and cortical neuron responses. *Science*. 1984; 224(4654):1260–2. [PubMed: 6539501]
- Westheimer G. Visual acuity. *Annu Rev Psychol*. 1965; 16:359–80. [PubMed: 14268891]
- Zeevi YY, Mangoubi SS. Vernier acuity with noisy lines: Estimation of relative position uncertainty. *Biol Cybern*. 1984; 50(5):371–6. [PubMed: 6477979]

Figure 1.

Stimuli and conditions used to test the efficiency of feature integration during visual completion. (a) Bent Bars; (b) Rotating Squares; and (c) Shrinking/Expanding Squares. Figures (b) and (c) are single frames from the entire dynamic stimulus sequences that observers actually saw. See main text for details about each stimulus type. See supplementary Figures S2 and S3 for full dynamic versions of the Rotating and Shrinking/Expanding Square stimuli, respectively.

Figure 2.

Squared sensitivities for individual observers and the mean across observers in each condition that was used to evaluate the efficiency of feature integration during visual completion. The top row corresponds to the conditions under which the combined stimulus was perceived as being a complete object; the bottom row corresponds to conditions under which the combined stimulus was perceived as collection of fragments. Each column corresponds to a stimulus type described in Figure 1. Error bars on individual observer sensitivities correspond to ± 1 sd, and were estimated through bootstrap simulations (Efron & Tibshirani, 1993). Error bars for mean sensitivities correspond to ± 1 s.e.m.

Figure 3.

Integration indexes for individual observers and the mean across observers in for each stimulus type described in Figure 1. Error bars were calculated as described in Figure 2.

Figure 4.

(a) Stimuli used in the Occluded Bent Bar task. (b) Sensitivities and (c) integration indexes for individual observers and the mean across observers in the Occluded Bent Bar task, in the presence of either Local Noise (top panels) or Global Noise (bottom panels). Error bars were calculated as described in Figures 2–3.

Figure 5.

(a) Stimuli used in the Vernier acuity task. (b) Sensitivities and (c) integration indexes for individual observers and the mean across observers in the Vernier acuity task. Error bars were calculated as described in Figures 2–3.