Dynamics of unconscious contextual effects in orientation processing

Isabelle Mareschal[1] and Colin W. G. Clifford

School of Psychology and Australian Centre of Excellence in Vision Science, University of Sydney, Sydney 2006, Australia

Edited by Barbara Anne Dosher, University of California, Irvine, CA, and approved April 2, 2012 (received for review January 17, 2012)

Contextual effects abound in the real world; how we perceive an object depends on what surrounds it. A classic example of this is the tilt illusion (TI) whereby the presence of a surround shifts the perceived orientation of a target. Surprisingly, the magnitude and direction of this shift depend on the orientation difference between the target and surround: when their orientations are similar, the perceived difference is amplified and the target appears repelled in orientation from the surround (i.e., the TI). However, when their orientations are close to perpendicular, the difference is decreased and the target appears attracted in orientation toward the surround (i.e., the indirect TI). These misperceptions of orientation have revealed much about the underlying detectors involved in visual processing and how they interact with each other. What remains at stake are the levels of processing involved. To examine this, we designed a reverse-correlation technique whereby observers are blind to the orientation of the surround. We find that the TI and indirect TI occur reliably and over a similar time course, supporting the role of a single mechanism underlying orientation biases that operates in the early stages of visual processing before the conscious extraction of the surround orientation.

visual illusions | classification images | awareness

Contextual effects arise when surrounding stimuli alter our perception of a target. Understanding how and when these illusions occur is critical for elucidating the mechanisms underlying visual processing. A common example is the tilt illusion (TI) whereby the perceived orientation of a target is repelled from the orientation of the surround when the angle between them is small (1–7). However, when the orientation difference between the target and surround is large (60°–80°), the perceived target orientation is attracted toward that of the surround, a phenomenon known as the indirect TI (8, 9). The indirect TI is smaller in magnitude and has been reported to differ in several fundamental ways from the TI (e.g., it is reduced by having a frame around the stimulus, it is immune to the addition of a gap between target and surround, and it is not dependent on spatial frequency) leading to the suggestion that it results from higher-level feedback that reflects the operation of global processing and involves a separate mechanism to the one responsible for the TI (8–10).

Early models account for the TI as a normalization process toward cardinal orientations (1), or as inhibitory lateral interactions between oriented neurons in early cortical areas (refs. 3, 4, 11, but see ref. 12), but many fail to address the indirect TI. More recent models account for the indirect TI but remain agnostic as to where it occurs and whether it involves a separate processing stream from the direct TI. Clifford et al. (13) present a model whereby responses from a population of neurons with tuning properties similar to those in early cortical areas are shifted and scaled in a manner that could arise from changes in their tuning curves. The shifting process produces only repulsion, but, when scaling is added, the model captures the TI and the indirect TI. In a more recent model, Goddard et al. (14) propose that orientation biases result from the surround facilitating or suppressing the central detector depending on their orientation difference. This model is based on known center-surround properties of neurons in early cortical areas (15, 16) and bypasses the need for higher-level involvement. By using a different approach, Schwartz et al. (17) propose that the sign and strength of interactions between the center and surround arise as a downstream consequence of their perceptual similarity. A measure of similarity could, in principle, be computed in early cortical areas (V1) where many neurons' tuning properties are subject to contextual effects (15, 18) and some can signal figure-ground segmentation (19), although it has been suggested that this information is conveyed to V1 neurons via feedback from higher areas (20).

We examined the need for separate mechanisms to account for the direct and indirect TI. We hypothesized that if observers were made blind to the surround (e.g., have no conscious access to its orientation) yet still displayed an indirect TI, this would indicate that the direct and indirect TI could be accounted for by a single mechanism acting before the conscious extraction of the surround orientation. We also sought to compare the time course of the TI and the indirect TI. Determining the occurrence and time scale of both biases under these conditions is a critical step in quantifying the interactions between orientation detectors, constraining models of early visual processing.

Results

Observers were instructed to report the orientation of a briefly presented (11.7 ms) vertical target that appeared every 2 s within a circular aperture at the center of an annular surround whose orientation was refreshed on every frame (Fig. 1*A*). On each appearance of the target, the observer reported whether it appeared to be tilted clockwise (CW) or counter-clockwise (CCW) of vertical. This allowed for an analysis of the surrounds' influence on the perception of the target as a function of surround orientation and timing: if the surround orientation had no effect on the observers' perception of the target, their responses should be random and the distribution of key presses will resemble a uniform distribution over the range of surround orientations. To determine if observers' responses were nonrandom over time, we examined 100 frames of the surround (corresponding to 1,170 ms) around the time of the target presentation.

Fig. 1 (*Lower*) shows the distribution of key presses as a function of surround orientation for naive observer RM over a 1,170-ms period, separated into CW responses (Fig. 1*B*) and CCW responses (Fig. 1*C*). In this format, the number of key presses are represented as a gray level in which white represents greater than average key presses and black represents fewer than average (scale bar gives index of absolute number of key presses). The data in Fig. 1*B* clearly illustrate the contextual

Author contributions: I.M. and C.W.G.C. designed research; I.M. performed research; I.M. and C.W.G.C. analyzed data; and I.M. and C.W.G.C. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

[1]To whom correspondence should be addressed. E-mail: isabelle.mareschal@sydney.edu.au.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1200952109/-/DCSupplemental.

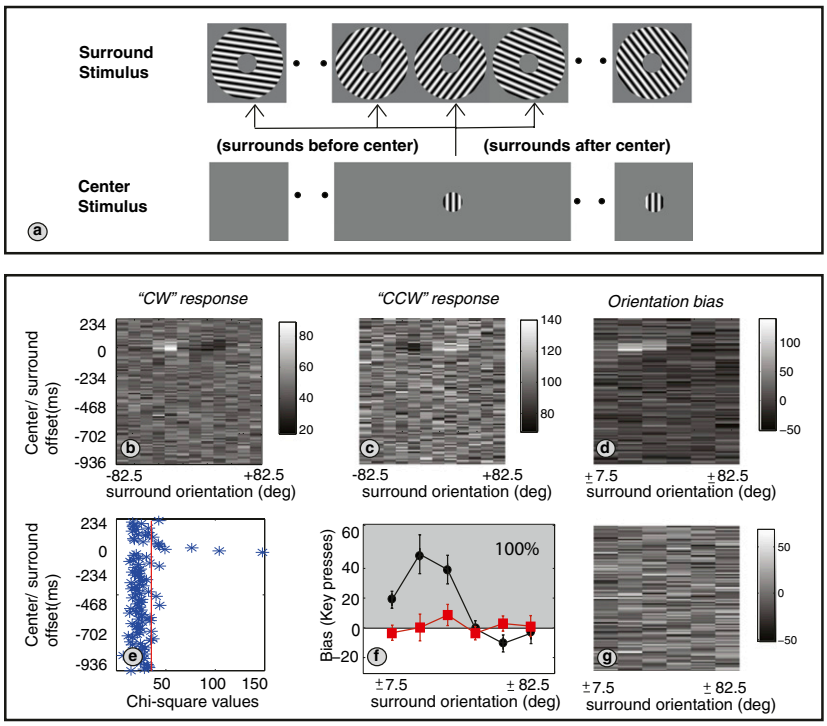

Fig. 1. Stimulus presentation and data analysis. (*A*) A rapid series of annular surround gratings was presented and, every 2 s, a target appeared in the center of the annulus. The observer reported the orientation of the target with a key press (*Lower*). (*B*) Observer R.M.'s CW response histogram, (*C*) CCW response histogram, and (*D*) orientation bias, as a function of the surround orientation over 1,170 ms. Time 0 represents the time of target onset. Surrounds before the target are assigned negative time values, and those after the target are assigned positive time values. Intensity is proportional to key presses. (*E*) An χ^2 analysis was performed on a frame-by-frame basis over the 1,170 ms. (*F*) Frames that reached significance from *E* were averaged to create an ROS and plotted as a function of the absolute orientation of the surround (black circles). Percentage is the number of targets the observer responded to over all trials. Error bars are ±1 SEM. Red squares and curve are the result of averaging over the same frames as the ROS in the control experiment and show that the observer could not access the orientation of the surround at any time within the experiment. (*G*) Orientation bias over the entire 1,170-ms sequence from the control experiment contains no discernable structure.

effects of the surround over a limited time range around the target. For example, the white patch near the center/surround offset of 0 ms demonstrates that the observer experienced a TI: negative (i.e., CCW) surrounds caused him to respond CW more than average. Orientation bias (Fig. 1*D*) combines the CW and CCW responses to account for any tendency an observer may have to prefer CW (or CCW) responses overall.

To estimate the effects of the surround as a function of orientation and timing, we performed a χ^2 analysis on a frame-by-frame basis over the 1,170-ms sequence (Eq. **1**). The frame with the largest χ^2 value is taken to represent the time when the influence of the surround on the target is maximal (Fig. 1*E*). We created a region of significance (ROS) by averaging frames around the peak χ^2 value. We used a liberal criterion for this ROS by including all frames within a three-frame "gliding window" (in both directions) that contained at least one frame significant at $P < 0.05$ (not corrected for multiple comparisons). The data from this ROS (Fig. 1*F*, black circles) show both the TI (i.e., positive bias key presses, data points in the gray region), which is largest with surround orientations at 22.5°, and the indirect TI (negative bias key presses, data points in the white region), which is greatest with surround orientations at 67.5°. This observer responded to all 1,740 targets (100%; Fig. 1*F*).

Fig. 1*G* plots the orientation bias for the control experiment conducted in the absence of a target in which the observer reported the orientation (i.e., CW or CCW) of the surround whenever cued at fixation. This bias plot lacks any discernable structure; over the 100-frame sequence, only three frames that were temporally separated from each other reached significance at $P < 0.05$ (uncorrected; occurring at times −94 ms, −667 ms, and −842 ms). To compare performance on the control task to performance on the orientation task, we averaged over the same frames as were used to create the ROS in the main experiment. The red squares in Fig. 1*F* plot data from the (nominal) ROS for the control experiment and shows that the observer's performance was at chance. The lack of a positive bias in this control experiment is consistent with subject reports that the surround orientations cannot be resolved at the 85-Hz refresh rate.

Fig. 2 plots the data in the same format for four additional observers who all show similar orientation biases (black curves), with three displaying an indirect TI [i.e., significant indirect TI at 67.5° across five observers; $t_{(4)} = 2.89$, $P < 0.05$]. For these observers, the significant frames lasted approximately 80 to 120 ms. The red curves confirm that none of the observers could identify the orientation of the surround in the control experiment.

We sought to determine how well the observers' performance could be accounted for by a linear transformation of the stimulus by their visual system. The response histograms in Fig. 1 *A* and *B* provide a linear estimate of the filter (Fig. 3*A*) applied to the input stream (stimulus; Fig. 3*B*) by the observer's visual system. Point-wise multiplying a 100-frame stimulus sequence that the observer saw by their filter produces a 100 × 12 matrix (Fig. 3*C*) that can be summed to obtain an index value whose magnitude correlates with how closely the stimulus matched the filter. Running this over all stimulus sequences produces a population of indices that can be separated into CW or CCW based on the observer's response. The decision criterion that maximizes the accuracy of predicting the observer's response from the value of the index can then be calculated (Eq. **3**). By using the CW response histogram over the data from which it was derived gives

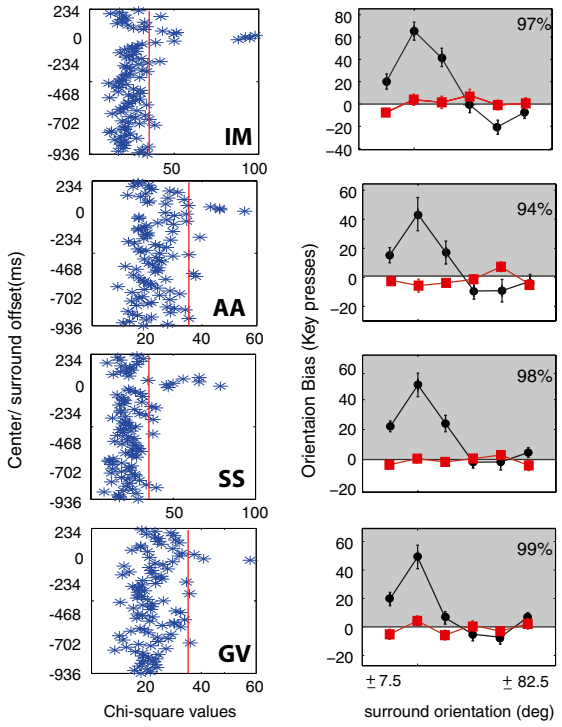

Fig. 2. Orientation bias from regions of significance for four additional observers. Only one observer (S.S.) does not get a clear indirect TI at a surround orientation of 67.5°. Error bars are ±1 SEM. Red squares are the control experiment.

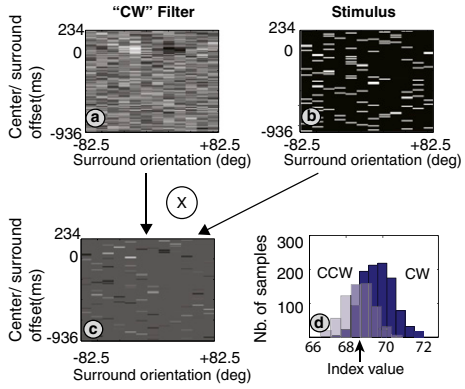

Fig. 3. Linear filter analysis. (*A*) The CW (or CCW) response histogram approximates the temporal linear filter. (*B*) Each 100-frame stimulus sequence can be represented as a 100 * 12 matrix in which the orientation present on a given frame is set to 1 and all other orientations are set to 0. (*C*) Point-wise multiplying the filter and the stimulus and summing across all 12 * 100 bins yields a single index denoting the time averaged response of the linear filter. (*D*) The value of the index is stored in the CW (dark blue) or CCW (pale blue) distribution according to the observer's response to the stimulus. The procedure is repeated for all 100-frame stimuli sequences, and a criterion (arrow) is calculated to optimize discrimination between the two distributions.

accuracies for the five observers between 71% and 82% for discriminating between CW and CCW responses when the decision criterion was optimized. This is not significantly different [$t_{(2)} = 1.4$, $P > 0.05$] than the maximum accuracy achievable given observer consistency as measured in a subsequent dual-pass experiment in which three observers were tested twice with the same sequence of surrounds [consistency, 75% (observer I.M.), 71% (observer S.S.), and 76% (observer R.M.)]. However, when the linear model (with decision criterion optimized on the original data) is applied to the new dual-pass data on which it was not trained, accuracy decreases to approximately 54% for I.M. and S.S. but is at 71% for R.M. Given that these new data were collected several weeks later, it is possible that the observers' decision criterion had shifted slightly. With the original linear filter but the decision criterion optimized for the new test data, accuracy reaches 71% for one observer (I.M.) and 74% for another (R.M.), but remains at 56% for the third (S.S.). This suggests that, for I.M. and R.M., their linear filter, if not their decision criterion, is stable over time.

Fig. 4 plots the time course of the TI and indirect TI. The upper panels (Fig. 4*A*) show biases averaged across the five observers and normalized by the maximum number of key presses over 164 ms of the stimulus. A video of the time course of these biases over the 40 frames around the target can be found in *SI Text* and Movie S1. Fig. 4*B* plots the key presses for the orientation that produced the greatest TI (22.5°) with the key presses for the orientation that produced the greatest indirect TI (67.5°) for data pooled across all observers and shows that the distributions of key presses deviate from their respective means at the same time. This can be seen more readily in Fig. 4*C*, which focuses on 40 frames around the target. Fitting Gaussian functions to four observers' time series data (observer S.S. had no indirect TI) reveals an insignificant difference of 2.3 ms between their peaks [average across observers of +16.4 ms for the TI and +18.7 ms for the indirect TI; $t_{(3)} = 0.15$, $P > 0.05$] as well as an insignificant difference of 20.5 ms in their bandwidths [average across observers of 85.8 ms for the TI and 62.0 ms for the indirect TI; $t_{(3)} = 2.14$, $P > 0.05$]. We also determined significance levels on the difference between the peaks for each observer by bootstrapping their individual data. We created response histograms by randomly selecting 60 (or 35) runs with replacement from the data set of 60 (or 35) runs for a given observer and used these "synthetic" data sets to build new CW, CCW, and bias-response histograms. Separate Gaussian functions were fit for the direct and indirect TIs, and the process was repeated 1,000 times. The bootstrapped differences between the peaks for all observers were nonsignificant (difference for I.M., +2.3 ms, $t = 0.23$, $P > 0.05$; difference for A.A., −6.3 ms, $t = -0.23$, $P > 0.05$; difference for R.M., −7.4 ms, $t = -0.16$, $P > 0.05$; difference for G.V., +4.2 ms, $t = 0.07$, $P > 0.05$).

In a final set of experiments, we assessed how the addition of a gap or a change in spatial frequency between the target and surround would affect the direct and indirect TI, as these manipulations have been reported to alter the magnitude of the direct TI (4, 5, 21) but not the indirect TI (8, 9). Fig. 5*A* plots data from the gap condition, averaged across four observers (R.M., I.M., A.A., and S.S.), and shows that the addition of a gap reduced the size of the direct TI and abolished the indirect TI. The peak response for the TI (taken as the average of the Gaussian fits to the four observers' peaks) occurred for surrounds at the time of target onset when a gap was introduced [peak, 0 ms (gap); peak, 15.2 ms (no gap)], although this temporal difference fell just short of significance [$t_{(3)} = 2.09$].

Fig. 5*B* plots the influence of the surround spatial frequency (blue circles, same spatial frequency as center; Fig. 5*B*). For the high spatial frequency surround condition (red squares, surround is two octaves higher in spatial frequency than the target), the viewing distance was increased to 114 cm, whereas for the low spatial frequency condition (green stars, surround is one octave lower; Fig. 5*B*), the viewing distance was kept at 57 cm. With the two different spatial frequency surround conditions, the size of the biases showed little variation [nonsignificant difference between the TI with low spatial frequency surround and the TI with the same spatial frequency surround over the seven frames of the ROS ($t_{(6)} = 0.92$, $P > 0.05$) and just significant difference with

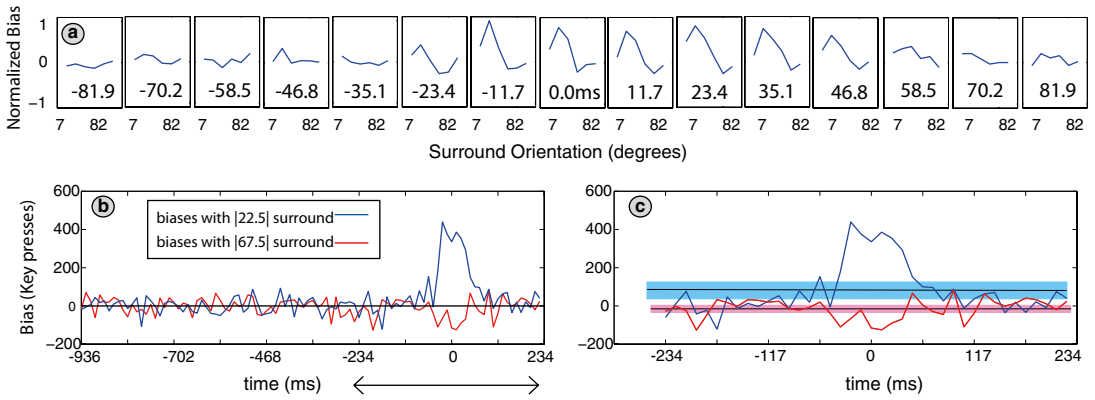

Fig. 4. Time course of orientation biases averaged across observers. (*A*) Distribution of key presses averaged across all observers and normalized by the maximum number of key presses for ±82 ms around the target. Note the clear emergence of the TI (positive peak) as well as the indirect TI (negative peak). *Lower*: Distribution of key presses for the TI and indirect TI summed across all observers plotted together over the 100 frames (*B*) and over 40 frames flanking the target (*C*). Shaded colored areas encompass the 95% confidence intervals centered on the means.

the TI obtained with a high spatial frequency surround ($t_{(6)}$ = 1.98), $P < 0.05$)]. It has been reported (4) that the size of the TI depends on the spatial frequency of the surround when targets are of a higher spatial frequency than 2.5 cycles/°. However, when lower spatial frequency targets were used, this dependency disappeared, a finding consistent with our results with a 2.0 cycles/° target.

Discussion

Reverse correlation techniques have been extremely useful in revealing the tuning properties of early visual processes (22–28). Here we develop a unique stimulus inspired by the work of Ringach (22) that increases the power of reverse correlation by providing measurements of bias in the absence of perceptual awareness of the context. This is achieved by the very rapid presentation rate of the surrounds that mask each other, rendering them inaccessible to conscious perception. We find a robust TI in all observers, consistent with an earlier finding (29) with the use of masked stimuli. Importantly, despite their inability to report the orientation of the surround at any given time, the majority of observers display a robust indirect TI. The introduction of a gap reduced the magnitude of the TI and abolished the indirect TI, but when the target and surround varied in spatial frequency, both biases persisted. These results are counter to earlier reports (8, 9) and may reflect the fact that our method eliminates any effect of observer response expectancy or bias. We also report that the TI and indirect TI peak in magnitude around the same time (when the surround is presented one or two frames after target onset). The duration of the direct TI is shorter than the reported temporal extent of effects measured previously (30) and is most likely caused by the fact that our surrounds mask each other. Taken together, these results could be interpreted in one of two ways: either that there is no need to invoke a separate, higher-level mechanism to account for the indirect TI; or, if higher levels are involved, they act at a stage before conscious extraction of the visual information. With functional MRI, it has been reported (31) that the earliest cortical area involved in visual processing (V1) encodes information that can be reliably used to predict the orientation of stimuli. Critically, this is the case even when subjects are made unaware of the stimulus orientation through masking, suggesting that V1 encodes unconscious stimuli. Taken in this context, the

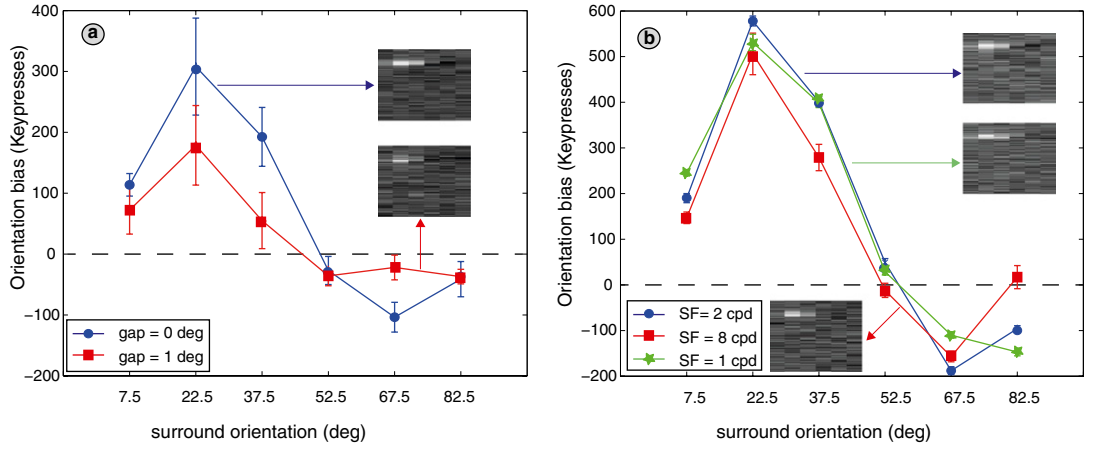

Fig. 5. (*A*) Orientation biases in the ROS measured in the presence (red squares) and absence (blue circles) of a 1° gap, averaged across four observers. Each ROS was symmetric and centered on the corresponding peak for the observer and condition. To combine data, the number of frames comprising the ROS was always the same (n = 5) and was determined by using the smallest ROS obtained by any observer in either condition. Error bars are ±1 SEM averaged within observers. Insets show orientation biases in the two conditions for the entire 100 frames. (*B*) Orientation biases with different surround spatial frequencies for observer I.M. The ROS for each condition contained seven frames (determined as described earlier), symmetrically centered on the peak for the given condition. Error bars are ±1 SEM.

most parsimonious explanation of our data would be one whereby both biases result from low-level inhibitory interactions between orientation-selective neurons. It has been shown that oriented surround stimuli can modify the response amplitude and bandwidth of V1 cells (18), changes that could underlie these perceptual biases. Regardless of the roles of the different cortical areas involved in visual processing, our results have clear implications for any model of orientation biases, namely that a single mechanism that does not rely on conscious access to the orientation of the surround can account for both types of bias.

Contextual effects abound in the perception of many fundamental visual attributes such as orientation (14, 21, 32), motion (33–35), texture (36–38), and contrast (39–41), although whether all of these effects arise preattentively or require awareness of the entire stimulus remains unclear. Unconscious processing of context has been used to examine orientation processing in the fovea (29), crowding in the periphery (42), and brightness illusions (43). We propose that our technique is a powerful tool that could be applied to contextual effects in other realms, such as color or depth. Distinguishing visual processes by their dependence on a complete, conscious representation of the stimulus is a critical step toward formulating a comprehensive model of cortical processing.

Methods

Observers. One of the authors and four naive observers served as subjects. All wore optical correction as necessary. A fifth naive observer failed to obtain a TI (no significant frames could be extracted from his orientation bias plot) and the data from this observer were discarded from the study.

Apparatus and Stimuli. A Dell Optiplex computer running MATLAB (MathWorks) was used for stimulus generation, experiment control, and recording of subjects' responses. The programs controlling the experiment incorporated elements of the PsychToolbox (44). Stimuli were displayed on a Diamond Digital monitor (resolution, 1,024 × 768 pixels, refresh rate, 85 Hz) driven by the computer's built-in Radeon graphics card. The display was calibrated by using a photometer and linearized by using look-up tables in the software. At the viewing distance of 57cm, one pixel subtended 2.1 arcminutes.

Stimulus and Procedure. Stimuli were viewed in the center of a black cardboard annulus that covered the edges of the display. The stimulus consisted of a concentric annular surround (outer diameter, 7.8°) containing a 2-cycles/° grating that could have one of 12 possible orientations (separated by 15°) abutting a circular patch (diameter, 1.8°), where a vertical 2-cycle/° grating was regularly presented. The surround was refreshed on each frame (11.7 ms), and each orientation had an equal probability of being selected. The center target was presented every 2 s for 11.7 ms and had one of four possible spatial phases separated by a quarter cycle. Both target and surround gratings were at 30% contrast. The observers' task was to report as quickly as possible after each target presentation whether it had appeared tilted CW or CCW of a subjective vertical using the response keys "L" for CW and "A" for CCW (Fig. 1*A*). A run lasted 1 min and observers completed a minimum of 60 runs (except for I.M., who completed 35 runs, each lasting 2 min). A movie of the stimulus can be found in *SI Text* and Movie S2. Please note that, because of limitations in the movie software, the stimulus is refreshed at 30 Hz, approximately three times slower than in the experiment.

Analysis. ***Linear filter estimation.*** At the end of the runs, the data were compiled and analyzed to determine how the sequence of surround orientations flanking the target presentation influenced observers' responses. CW and CCW response categories were calculated separately by determining observers' responses to targets, pooling across all spatial phases. Surrounds before the target presentation were assigned negative time values and those after the target were assigned positive values. In total, 100 frames around the target were analyzed: 20 frames (+234 ms) after the target presentation and 80 frames (−936 ms) before. Because the surround could have one of 12 possible orientations, we created 100 response histograms of the 12 orientations for the CW and CCW responses (see ref. 22). The observer's response to a given target determined which series of counter histograms was used. For each frame, the orientation of the surround at the given time was recorded and its counter incremented by 1.

If the surrounds had no influence on the perceived orientation of the target, every frame of the response histogram should resemble a uniform distribution. Determining when responses deviated from the uniform distribution will reveal the dynamics of lateral interactions underlying orientation biases. We performed a χ^2 analysis (Eq. **1**) on a frame-by-frame basis on the raw key press counts whereby the frequency of key presses at each orientation was compared with the mean number of key presses for the CW and CCW responses (e.g., ref. 28). This procedure deals with any small biases to report CW or CCW regardless of the surround orientation.

$$X^2 = \sum_{(i=1:24)} (\text{count}(i) - \text{meancount})^2 / \text{meancount} \quad [1]$$

where "meancount" was set to the mean count of CW when responses were CW and set to the mean count of CCW when responses were CCW.

Orientation bias histograms were created by adding corresponding frames from the flipped CCW response histogram to the CW response histogram. The response histogram was further reduced to six orientations by subtracting, frame-by-frame, the counts obtained with a positive (i.e., CW) surround orientation from counts obtained with the corresponding negative (i.e., CCW) surround orientation. On each frame, the following calculation was performed for the six orientations:

$$\text{Count}_\theta = [\text{CWcount}(-\theta) + \text{CCWcount}(\theta)] - [\text{CWcount}(\theta) + \text{CCWcount}(-\theta)] \quad [2]$$

In Eq. **2**, Count_θ represents the total key presses for the orientation θ, CWcount is the number of key presses in the CW response histogram for the orientations ($-\theta$) or (θ), and CCWcount is the number of key press in the CCW response histogram for the orientations ($-\theta$) or (θ).

By this procedure, positive key presses represent a TI and negative key presses represent an indirect TI. Note that, although orientation biases (i.e., the tendency to see true vertical as tilted in the absence of a surround) can occur, they result in a shift in the baselines (e.g., mean counts) of the CW and CCW histograms, and are discounted through the combinatorial procedure.

The very rapid presentation rate of the surround gratings ensured that observers were unable to identify the orientation on any given frame. As confirmation of this, the observers performed 60 trials of a control task in which the surround was refreshed on every frame but no targets were presented. Instead the observers fixated a central black square (2 pixels * 2 pixels) in the circular patch that changed to white every 2 s, cueing them to report whether the surround grating was CW or CCW of vertical. If the observers could correctly identify the surround orientations, the CW and CCW response histograms would resemble step functions in antiphase and combining responses following Eq. **2** would produce a positive bias in key presses across all orientations.

Dual-pass analysis. Each 100-frame response histogram can be likened to the linear filter underlying the observer's performance. To determine how well this characterization accounts for the data, we performed the following analysis. For every 100-frame sequence of surround orientations around the target presentation, we created a 100 * 12 surround matrix in which the orientation present on a given frame was assigned a value of 1, and all other surround orientations were set to 0. We then point-wise multiplied each surround matrix by the 100-frame CW linear filter and summed the output to obtain an index that was binned in the CW or CCW histogram depending on the observer's response. We determined the criterion producing the optimal accuracy (Eq. **3**) by summing the number of CCW indices that were less than the criterion (i.e., true negatives) with the number of CW indices greater than the criterion (i.e., true positives) and dividing by the total number of indices:

$$a = \left[\sum (\text{CCW} < \text{criterion}) + \sum (\text{CW} > \text{criterion})\right] \Big/ \left[\sum (\text{CCW}) + \sum (\text{CW})\right] \quad [3]$$

Three observers (I.M., R.M., and S.S.) completed an additional dual-pass experiment. In this procedure, the observer performs two runs of the experiment by using identical surround sequences such that any differences in responses reflects the observer's internal noise. In our dual-pass procedure each run consisted of 10 trials totaling 290 target presentations (note that the sequence of orientations in the 10 trials were different within the run, but identical across the two runs), allowing us to measure observer consistency and compare with the performance of the linear model.

ACKNOWLEDGMENTS. We thank Samuel Solomon for helpful discussions during these experiments. This research was funded by the Australian Centre of Excellence in Vision Science. C.W.G.C. was supported by a Future Fellowship from the Australian Research Council.

1. Gibson JJ (1937) Adaptation, after-effect, and contrast in the perception of tilted lines. *J Exp Psychol* 20:553–569.
2. Over R, Broerse J, Crassini B (1972) Orientation illusion and masking in central and peripheral vision. *J Exp Psychol* 96:25–31.
3. Carpenter RHS, Blakemore C (1973) Interactions between orientations in human vision. *Exp Brain Res* 18:287–303.
4. Georgeson MA (1973) Spatial frequency selectivity of a visual tilt illusion. *Nature* 245: 43–45.
5. Tolhurst DJ, Thompson PG (1975) Orientation illusions and after-effects: Inhibition between channels. *Vision Res* 15:967–972.
6. Wade NJ (1980) The influence of colour and contour rivalry on the magnitude of the tilt illusion. *Vision Res* 20:229–233.
7. Westheimer G (1990) Simultaneous orientation contrast for lines in the human fovea. *Vision Res* 30:1913–1921.
8. Wenderoth P, Johnstone S (1987) Possible neural substrates for orientation analysis and perception. *Perception* 16:693–709.
9. Wenderoth P, Johnstone S (1988) The different mechanisms of the direct and indirect tilt illusions. *Vision Res* 28:301–312.
10. Wenderoth P, van der Zwan R (1989) The effects of exposure duration and surrounding frames on direct and indirect tilt aftereffects and illusions. *Percept Psychophys* 46:338–344.
11. Blakemore C, Carpenter RH, Georgeson MA (1970) Lateral inhibition between orientation detectors in the human visual system. *Nature* 228:37–39.
12. Solomon JA, Felisberti FM, Morgan MJ (2004) Crowding and the tilt illusion: Toward a unified account. *J Vis* 4:500–508.
13. Clifford CWG, Wenderoth P, Spehar B (2000) A functional angle on some after-effects in cortical vision. *Proc Biol Sci* 267:1705–1710.
14. Goddard E, Clifford CWG, Solomon SG (2008) Centre-surround effects on perceived orientation in complex images. *Vision Res* 48:1374–1382.
15. Levitt JB, Lund JS (1997) Contrast dependence of contextual effects in primate visual cortex. *Nature* 387:73–76.
16. Cavanaugh JR, Bair W, Movshon JA (2002) Selectivity and spatial distribution of signals from the receptive field surround in macaque V1 neurons. *J Neurophysiol* 88: 2547–2556.
17. Schwartz O, Sejnowski TJ, Dayan P (2009) Perceptual organization in the tilt illusion. *J Vis* 9:19.1–19.20.
18. Gilbert CD, Wiesel TN (1990) The influence of contextual stimuli on the orientation selectivity of cells in primary visual cortex of the cat. *Vision Res* 30:1689–1701.
19. Zipser K, Lamme VAF, Schiller PH (1996) Contextual modulation in primary visual cortex. *J Neurosci* 16:7376–7389.
20. Angelucci A, Bullier J (2003) Reaching beyond the classical receptive field of V1 neurons: Horizontal or feedback axons? *J Physiol Paris* 97:141–154.
21. Kapadia MK, Westheimer G, Gilbert CD (2000) Spatial distribution of contextual interactions in primary visual cortex and in visual perception. *J Neurophysiol* 84: 2048–2062.
22. Ringach DL (1998) Tuning of orientation detectors in human vision. *Vision Res* 38: 963–972.
23. Neri P, Parker AJ, Blakemore C (1999) Probing the human stereoscopic system with reverse correlation. *Nature* 401:695–698.
24. Gold JM, Murray RF, Bennett PJ, Sekuler AB (2000) Deriving behavioural receptive fields for visually completed contours. *Curr Biol* 10:663–666.
25. Solomon JA (2002) Noise reveals visual mechanisms of detection and discrimination. *J Vis* 2:105–120.
26. Abbey CK, Eckstein MP (2002) Classification image analysis: Estimation and statistical inference for two-alternative forced-choice experiments. *J Vis* 2:66–78.
27. Mareschal I, Dakin SC, Bex PJ (2006) Dynamic properties of orientation discrimination assessed by using classification images. *Proc Natl Acad Sci USA* 103:5131–5136.
28. Wong EMY, Roeber U, Freeman AW (2010) Lengthy suppression from similar stimuli during rapid serial visual presentation. *J Vis* 10:14.1–14.12.
29. Clifford CWG, Harris JA (2005) Contextual modulation outside of awareness. *Curr Biol* 15:574–578.
30. Durant S, Clifford CWG (2006) Dynamics of the influence of segmentation cues on orientation perception. *Vision Res* 46:2934–2940.
31. Haynes J-D, Rees G (2005) Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nat Neurosci* 8:686–691.
32. Mareschal I, Sceniak MP, Shapley RM (2001) Contextual influences on orientation discrimination: Binding local and global cues. *Vision Res* 41:1915–1930.
33. Falkenberg HK, Bex PJ (2007) Contextual modulation of the motion aftereffect. *J Exp Psychol Hum Percept Perform* 33:257–270.
34. Spering M, Gegenfurtner KR (2008) Contextual effects on motion perception and smooth pursuit eye movements. *Brain Res* 1225:76–85.
35. Mareschal I, Clifford CWG (2012) Perceptual entrainment of individually unambiguous motions. *J Vis* 12:1–8.
36. Field DJ, Hayes A, Hess RF (1993) Contour integration by the human visual system: Evidence for a local "association field". *Vision Res* 33:173–193.
37. Caputo G (1996) The role of the background: Texture segregation and figure-ground segmentation. *Vision Res* 36:2815–2826.
38. Wolfson SS, Landy MS (1999) Long range interactions between oriented texture elements. *Vision Res* 39:933–945.
39. Cannon MW, Fullenkamp SC (1991) Spatial interactions in apparent contrast: Inhibitory effects among grating patterns of different spatial frequencies, spatial positions and orientations. *Vision Res* 31:1985–1998.
40. Snowden RJ, Hammett ST (1998) The effects of surround contrast on contrast thresholds, perceived contrast and contrast discrimination. *Vision Res* 38:1935–1945.
41. Solomon JA, Morgan MJ (2000) Facilitation from collinear flanks is cancelled by non-collinear flanks. *Vision Res* 40:279–286.
42. Wallis TS, Bex PJ (2011) Visual crowding is correlated with awareness. *Curr Biol* 21: 254–258.
43. Harris JJ, Schwarzkopf DS, Song C, Bahrami B, Rees G (2011) Contextual illusions reveal the limit of unconscious visual processing. *Psychol Sci* 22:399–405.
44. Brainard DH (1997) The psychophysics toolbox. *Spat Vis* 10:433–436.