Changing objects lead briefly flashed ones

Bhavin R. Sheth, Romi Nijhawan and Shinsuke Shimojo

California Institute of Technology 139-74, Pasadena, California 91125, USA

Correspondence should be addressed to B.R.S. (bhavin@caltech.edu)

Continuous, predictable events and spontaneous events may coincide in the visual environment. For a continuously moving object, the brain compensates for delays in transmission between a retinal event and neural responses in higher visual areas. Here we show that it similarly compensated for other smoothly changing features. A disk was flashed briefly during the presentation of another disk of continuously changing color, and observers compared the colors of the disks at the moment of flash. We also tested luminance, spatial frequency and pattern entropy; for all features, the continuously changing item led the flashed item in feature space. Thus the visual system's ability to compensate for delays in information about a continuously changing stimulus may extend to all features. We propose a model based on backward masking and priming to explain the phenomenon.

An object flashed at the instant another, continuously moving object arrives at the same location is perceived to spatially lag the moving object (flash-lag effect)[1,2]. In the case of motion, an object's visual location varies in time, but an object can change features other than location (luminance, color, shape and so forth) over time. An important issue then is whether the flash-lag effect is found exclusively in motion[3], or whether it is also observed in other feature domains[4] when the object remains motionless while another feature changes smoothly. Does an object increasing in intensity appear brighter than a flashed object of equal luminance? Does a grating becoming coarser over time appear more coarse than a flashed grating of equal spatial frequency? Answers to such questions are significant because a phenomenon that crosses feature boundaries would reflect a fundamental property of the brain. If the flash-lag effect were found for other visual features, and perhaps other sensory modalities, then mechanisms invoked to explain the phenomenon must be general across feature domains. Here we report that continuously changing objects appeared to lead flashed objects in feature space whether the changing feature was color, luminance, spatial frequency or pattern entropy.

Results

We first tested a stationary disk that began either as red or green and then continuously changed color (to become more green or red, respectively). Halfway through the sequence, we briefly flashed a second disk (**Fig. 1a**). In a two-alternative forced-choice task, the observer judged which of the two disks appeared greener. A typical trial (**Fig. 1b**, top) and the typical perceived relationship between the two (bottom) is shown. Thus, although the flashed disk was of the same color as the synchronously presented frame in the sequence of the changing color disk (synchronous frame), the continuously changing disk seemed to be of a subsequent color (**Fig. 1b**, bottom). Results for this color flash-lag effect are shown for six observers (**Fig. 1c** and **d**). Responses were combined over all observers (n = 6; 1 author and 5 naive observers) for both red→green and green→red transitions (**Fig. 1c**). The magnitude of the flash-lag effect is halfway between the two psychometric curves fit to the data. For all observers, the flashed disk lagged the continuously present, gradually changing color disk in color space (**Fig. 1d**). Converted into time units, the mean flash-lag magnitude in the color task was 394 ms (**Fig. 1d**).

Next, we tested other features, namely, luminance and spatial frequency. In the luminance task, a disk appeared on one side of the fixation point (FP) at the start of each trial and gradually became brighter or dimmer. On the opposite side of the FP, a second disk of random luminance was briefly presented for one frame. The observer judged which of the two disks appeared brighter (**Fig. 2a**). For the spatial frequency task, a patch of square-wave vertical grating was used. The spatial frequency of the grating changed incrementally while, midway through the sequence, a second vertical grating patch of randomly chosen spatial frequency was flashed. Again, the observer had to judge which of the two patches was higher in spatial frequency (**Fig. 2b**). For all observers (n = 3 for luminance, 2 naive observers; n = 4 for spatial frequency, all observers naive), there was a clear effect of the flashed object perceptually trailing the continuously changing one. The average time lag in the case of luminance (37 ms) was slightly less than half that in the case of spatial frequency (83 ms; **Fig. 2a** and **b**, gray bars on right).

Thus far, all the features tested have known neural correlates in visual cortex. For example, color is believed to be processed by chromatic wavelength-tuned neurons in V1 and V4 (ref. 5), and on- and off-center retinal ganglion cells[6,7], and center–surround cells in the LGN respond differentially to different light intensities[8]. Orientation-selective neurons biased to respond to a limited range of spatial frequencies are found in V1 (refs. 9, 10), and exquisitely direction-tuned cells are found all along the motion pathway, including areas V1 (ref. 11), MT and MST[12].

Demonstrating a flash-lag effect for an attribute with no known correlate would show that the effect is indeed due to properties that are widespread throughout the brain. Therefore we next tested observers on pattern entropy—a feature with no well documented neural correlate. The visual stimulus consisted of a square patch containing a fixed number of dots. The dots were initially arranged in a regular grid (**Fig. 2c**, left). The range of dot-position scatter, measured with respect to a given dot's grid position, was the same for all frames. However, the percentage of dots allowed to stray from their positions on the grid was gradually increased from 0% to 100%, causing the total pattern entropy of the dots to increase (**Fig. 2c**). The dots were enclosed

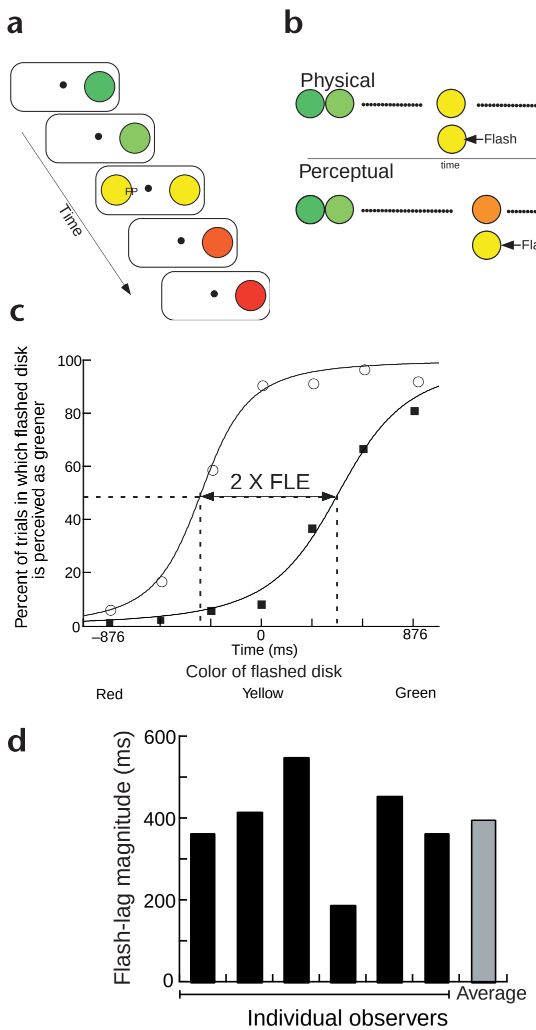

Fig. 1. The flashed object lags behind in color space. (**a**) A disk appearing either left or right of a central fixation point (FP) continuously changed color from green to red or from red to green while maintaining equiluminance[37] (see Methods). While it was in the middle of its smooth color transformation, a second disk of a randomly chosen (with equal probability from a set of seven predetermined red/green color combinations), fixed color briefly flashed on the opposite side of the FP. (**b**) Top, color sequence of the continuously changing disk across time (green→red) and the color of the flashed disk (yellow). Here the continuous and flashed disks were of the same color (yellow) when the flashed disk was on. Bottom, typical perception of this sequence. Because of neural transmission delays, the observer perceived the flashed disk some time after its image reached the retina. At that moment, the observer typically perceived the continuous disk as more red and less green than the flashed disk. (**c**) The grouped responses of all observers (n = 6) over each of the seven predetermined colors of the flashed disk for both red→green (filled squares) and green→red (open circles) color sequences. The abscissa gives the color of the flashed disk relative to that of the color sequence at the time of the flash, with the middle (fourth from left) representing the same color for both. The ordinate gives the percentage of trials the observer perceived the flashed disk as greener at the time of the flash. Each point is the mean of 120 trials. Intersections with the vertical dotted lines define the threshold estimates (T_{50}) of two psychometric curves fit to the data, and mark the points at which the flashed disk appeared more green than the continuously present disk as often as it appeared less green. (**d**) Magnitudes of the flash-lag effect for each of the 6 observers and the mean (394 ms; right, in gray). We found no major differences in effect magnitude between the first three observers, for which we applied gamma correction, and the last three, for which we used their respective perceptual equiluminance points obtained earlier[38].

by the invisible boundaries of the square area, which held the mean luminance of the patch constant. Starting with the inverse, a frame in which all the dots were free to occupy locations not on the grid, the percentage of such dots was gradually decreased to zero. Experiments using the increasing and decreasing pattern-entropy sequences (n = 3 observers; **Fig. 2c**) yielded results similar to other feature domains: namely, the flashed object trailed the continuously changing object (mean, 94 ms; **Fig. 2d**). In another experiment, the pattern entropy was increased (decreased) by gradually and uniformly increasing (decreasing) the range of scatter for all the dots. The mean flash-lag magnitude for three different observers using this set of stimuli was similar (95 ms; data not shown).

To our knowledge, cells with tuning curves for stimuli used here have not yet been found. Complex cells in V1 and V2 do respond to unoriented dot constellations, but it is not their optimal stimulus[13]. Moreover, other cortical and subcortical areas such as the LGN could also be stimulated by these patterns. Because the stimuli shown in Fig. 2c are unlikely to be processed by any one brain locus or set of neurons specialized for such stimuli, the combined results of all features described above suggest that the neural substrates of the flash-lag effect may be distributed over multiple cortical areas.

The generality of the flash-lag phenomenon clearly begs the question as to its underlying mechanism. One might argue that attention is an appropriate general-purpose, high-level process that could account for the effect. According to this account, the flash acts as a 'pull' cue to 'grab' the observer's attention[14] away from the continuously changing stimulus. Attention is then voluntarily switched some time later from the flashed stimulus back to the continuously changing one. During this delay, the changing stimulus has progressed beyond its value at the time of the flash, producing the flash-lag effect[15]. To test this account, the flashed and continuously changing disks were both turned on simultaneously at the beginning of each trial (**Fig. 3a**). The sequence of frames was identical to the sequence following the flash in the original color experiment, and so was the duration of each frame (see Methods). Because both stimuli were turned on simultaneously, the flashed stimulus was not alone in its abrupt onset, and so attention could not be exclusively pulled toward it. Consequently, one should not observe a flash-lag effect in this task assuming this attention hypothesis. Contrary to the attention hypothesis, this so-called 'flash-initiated cycle'[16] yielded the same effect as before: the continuous disk led the flashed disk in color space (**Fig. 3b**). The mean magnitude of the flash-lag effect across 3 observers was 334 ms. Next, we tested the attention account on a second task. In a temporal-order judgment task, the flashed and continuously changing stimuli appeared at different times relative to one another (**Fig. 3c**), and observers (n = 3) judged which of the two appeared first. If the time required for switching attention from the flashed to the gradually changing stimulus is on the order of several hundred milliseconds, then the flash would be perceived to have come on first, even when it did not. There was no significant shift of the curve (**Fig. 3d**). When the two came on concurrently (stimulus onset asynchrony, or SOA, was 0), observers were equally likely to judge either one as having appeared first, and for 75 ms SOAs (much

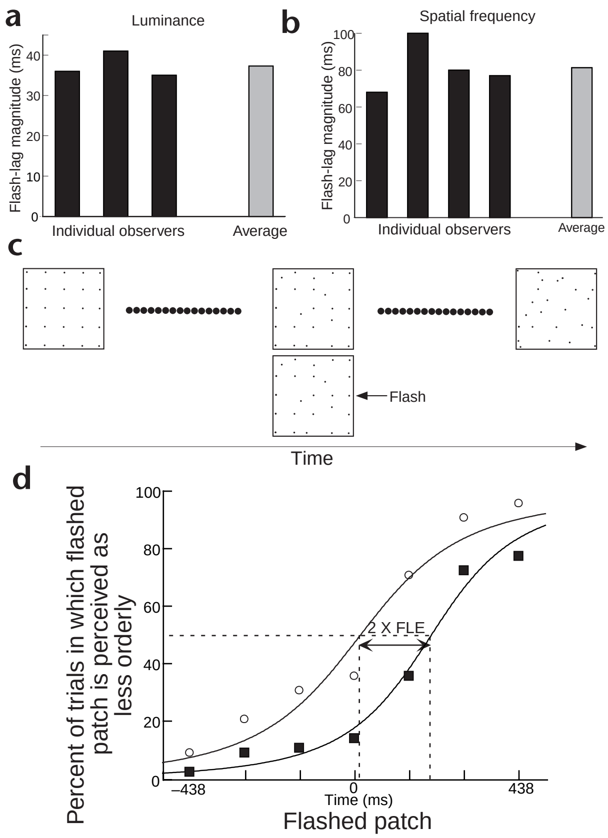

Fig. 2. A flashed object in luminance, spatial-frequency or pattern-entropy space perceptually trails a continuously present and smoothly changing one. (**a**) Luminance flash-lag effect magnitudes for three observers and the mean (gray, right). Mean luminance flash-lag effect, 36 ms. (**b**) Spatial-frequency flash-lag effect magnitudes for four observers and the mean (gray, right). Mean spatial-frequency flash-lag effect magnitude, 83 ms. (**c**) Stimulus used to study the flash-lag effect in pattern entropy, showing the sequence of the continuous square patch over time and the flashed patch at the instant it was presented. In this illustration, the continuous and flashed patches had identical dot locations at the moment the flashed square patch was turned on. We reversed the sequence of frames for the decreasing entropy sequence. (**d**) Psychometric curves (n = 3 observers) for both increasing (filled squares) and decreasing (open circles) entropy sequences. Each point is the mean of 60 trials. The difference in threshold estimates (T_{50}) between the two curves was highly significant ($p < 0.001$). The mean flash-lag effect magnitude was 94 ms.

less than the flash-lag effect with color), observers were very accurate—indicating no effect. Combining the results from both experiments, we conclude that attention cannot be the main process underlying the flash-lag effect, although it may have a small, yet critical role (see Discussion).

Adaptation (and the resulting aftereffect) is another mechanism proposed to explain the effect. An adapter of fixed feature value causes a subsequently presented test of high (low) feature value to appear even higher (lower) in feature value as compared with the adapter[17]. In the green→red color task, for example, a series of primarily green adapting frames presented early in a given trial may cause a subsequent yellow frame in the sequence presented simultaneously with a (yellow) flashed stimulus to appear redder: the continuously changing stimulus 'jumps ahead' of the flash in perceptual color space. Thus adaptation seems to account for the flash-lag phenomenon in the color, luminance and spatial-frequency domains, but the strength of adaptation required (hundreds of milliseconds in the color domain, for example) does render the adaptation explanation implausible. Also, adaptation would not account for the effect in the case of motion, as the aftereffect could occur either in a direction opposite to that of the adapting stimulus[18–20], yielding a 'flash-lead' effect, or in the same direction as the adapting stimulus[21], yielding a flash-lag effect. Nonetheless, we performed an experiment to further address the adaptation account in the color domain. The color of the continuously changing object was gradually changed from reddish brown (proportion of green to red slightly less than 1) to red (ratio of green:red, 0), or *vice versa* (**Fig. 4a**). Midway through the trial, a second object was flashed. If the continuously present stimulus (reddish brown→red) actually matched the color of the flashed stimulus at the time of the flash, the flash-lag effect should make it appear redder than the flash. However, because each frame in the sequence is reddish (green/red ratio<1), adaptation following prior viewing of the continuously changing stimulus leading up to the flash should make the red in the sequence appear desaturated. Hence, in the reddish-brown→red transition, adaptation should cause the continuously present stimulus to appear less red than it truly is, reducing the perceptual lag of the flash in color space. If adaptation were the major cause, the flash-lag effect should even be abolished. On the other hand, in the reverse red→reddish-brown condition (**Fig. 4b**), adaptation to red should enhance the effect. At the very least, if adaptation were to have any role at all, the net effect in the red→reddish-brown transition should be greater than in the reddish-brown→red transition. For data combined across three observers (1 author, 2 naive observers; **Fig. 4d**), magnitudes of the net effect in the red→reddish-brown (256 ms) and reddish-brown→red (317 ms) transitions show the opposite trend, contrary to the predictions of the adaptation explanation. In our spatial-frequency experiments, the sequence was presented for a total of 300–600 ms (see Methods)—at least a couple of orders of magnitude less than the time required to obtain a robust spatial-frequency aftereffect[22]. Thus adaptation cannot explain the flash-lag effect in spatial frequency, either. To summarize, adaptation, in the form of a color aftereffect[17], brightness aftereffect, spatial frequency aftereffect[22] or motion aftereffect[18], is not sufficient to explain the flash-lag effect.

An alternative general-purpose account is based on visual persistence. According to this account, the observer estimates the median color of the continuously changing object during the 100-ms interval that a flashed object would visibly persist, if left unmasked[23,24]. Because the median exceeds the flash in feature value, this would constitute the flash-lag effect. If the persistence account is correct in explaining the flash-lag phenomenon, then decreasing either the intensity or the duration of the flashed stimulus should enhance the effect[25]. First, we reduced the intensity of the flashed stimulus to one-eighth of the baseline (1.6 cd per m^2) while the luminance of the continuously changing stimulus was kept at baseline (12.7 cd per m^2). We compared the strength of the effect across three observers (one author, two naive observers); combined data and psychometric curves are shown for baseline conditions (**Fig. 5a**) and with reduced flash intensity (**Fig. 5c**). The mean flash-lag effect magnitude in the baseline condition and with reduced flash intensity was 466 ms and 406 ms, respectively (**Fig. 5b** and **d**; gray bars). If anything, the flash-lag effect decreased with reduced flash intensity, contradicting the persistence account. This result is also incompatible with the differential latency account, which claims that the flash-lag effect is due to a longer

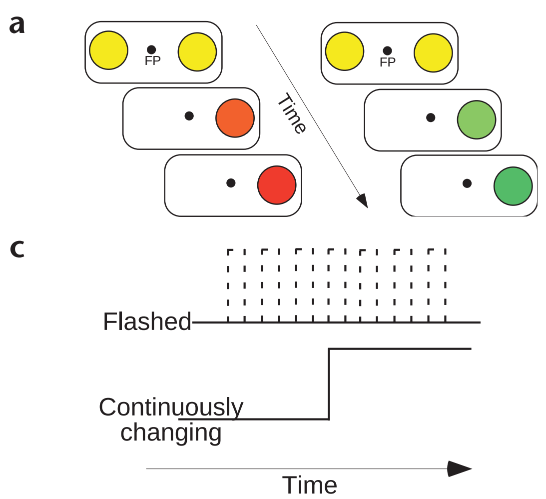

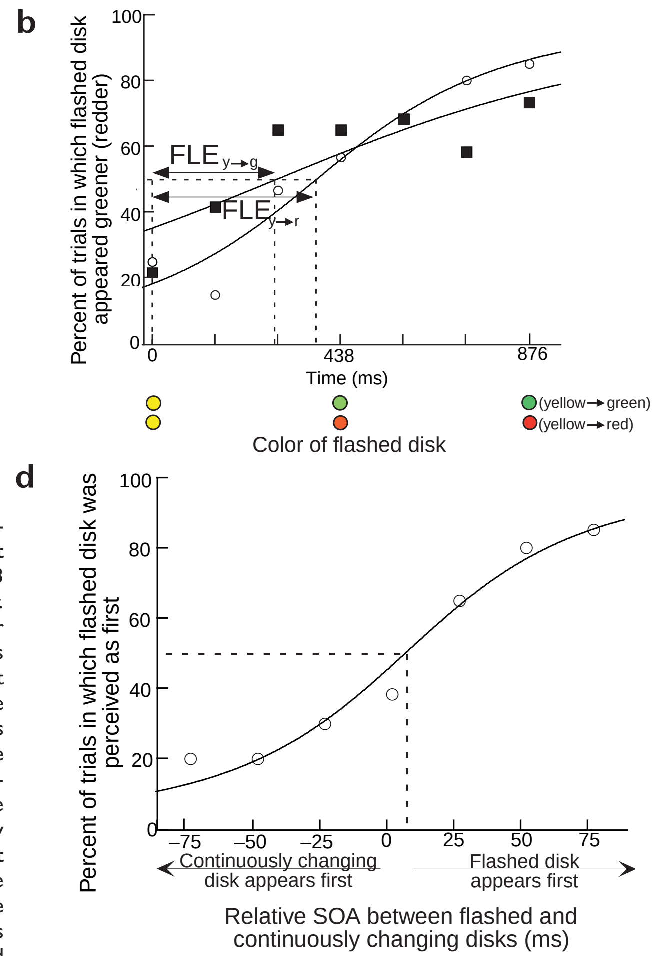

Fig. 3. Attentional switching effects cannot account for the flash-lag phenomenon. (**a**) Each trial began with two disks presented simultaneously at equal distance from the FP. The flashed disk disappeared after one frame (73 ms), whereas the second stayed on and gradually became redder or greener. (**b**) Grouped (*n* = 3 observers) psychometric curves for the two color sequences (yellow→green, filled squares; yellow→red, open circles) in this color flash-initiated task. Each point is the mean of 60 trials. Flash-lag effect was measured from the trial onset, as the flashed disk was presented at the start of each trial. Horizontal distance (converted into time units) from this leftmost point to the point along the psychometric curve at which the flashed disk appeared redder (greener) in 50% of the trials gives the magnitude of the effect; mean flash-lag effect magnitude was 334 ms. (**c**, **d**) The possible timing of the onsets of the two disks (flashed and continuously changing) in the two-alternative, forced-choice temporal-order judgment task. (**d**) Percentages of trials in which the observer (*n* = 3) perceived the flashed disk as having appeared first (ordinate) versus SOA (abscissa). The horizontal coordinate of the point at which the psychometric curve crosses the dotted line gives the SOA at which the observer perceived the flashed and continuously changing disks as having appeared first equally often (50%).

latency of the flashed stimulus[26,27]. According to this account, reducing flash intensity should increase its latency and thereby increase the magnitude of the flash-lag effect—in direct contradiction to our data. In a second experiment, we decreased the duration of the flash to a fourth of the 73 ms baseline duration, but the total time the continuously changing stimulus remained visible was the same as in the baseline condition (see Methods). The mean effect with the shorter, 18-ms flash (352 ms; **Fig. 5e** and **f**, gray bar), was not greater than the mean effect in the baseline condition (466 ms; **Fig. 5a** and **b**), again counter to the persistence account. Therefore, neither persistence nor differential latency provide a satisfactory explanation for the flash-lag effect.

Discussion

We found that the flash-lag effect described earlier for motion also exists for other features, such as color, luminance, spatial frequency and pattern entropy. The continuously changing object was perceived to be ahead of the flashed object in feature value, even though the feature values of both objects at the time of the flash were physically identical. Qualitatively speaking, the effect was the same across all features tested so far, but there was a wide range in the strength of the effect across attributes. These differences perhaps reflect differences in the way different attributes are processed in the brain.

These findings are somewhat paradoxical, as a linear-systems analysis should predict the opposite result. The flash should evoke a classic, transient impulse response, whereas the ramp of the gradual stimulus change should be low-pass filtered and attenuated, leading to a delayed signal. Thus, the flash should appear to lead the continuously changing object, contrary to our observations.

Several hypotheses are offered to explain these paradoxical results. One of these is attention, but our flash-initiated and temporal-order judgment findings (**Fig. 3**) showed that attention alone is not sufficient to explain the effect. Adaptation is also proposed as a mechanism for the features studied here, but our results disputed this hypothesis as well, because the magnitude of the effect was unaffected by adaptation (**Fig. 4**). Another account based on persistence of an unmasked stimulus was also proven to be inadequate: decreasing the intensity or duration of the flash did not enhance the effect, contrary to the persistence account (**Fig. 5**). To reiterate, across all features tested—color, spatial frequency, luminance and pattern entropy—the effect was qualitatively the same, and it must, at least in part, result from a compensation mechanism. We believe that any theory attempting to explain the flash-lag effect in motion should not ignore the remarkably similar findings in other feature domains. In this regard, modality-general mechanisms that are present across many or, perhaps, all object features and visual areas have an advantage.

One such model is based on two modality-general mechanisms—priming (B.R.S., S.S. & R.N. *Invest. Opthalmol. Vis. Sci.* **40**, 45, 1999) and backward masking(B.R.S. *et al.*; ref. 28,29). Both priming and masking of a target stimulus are enhanced

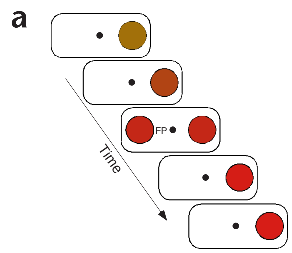

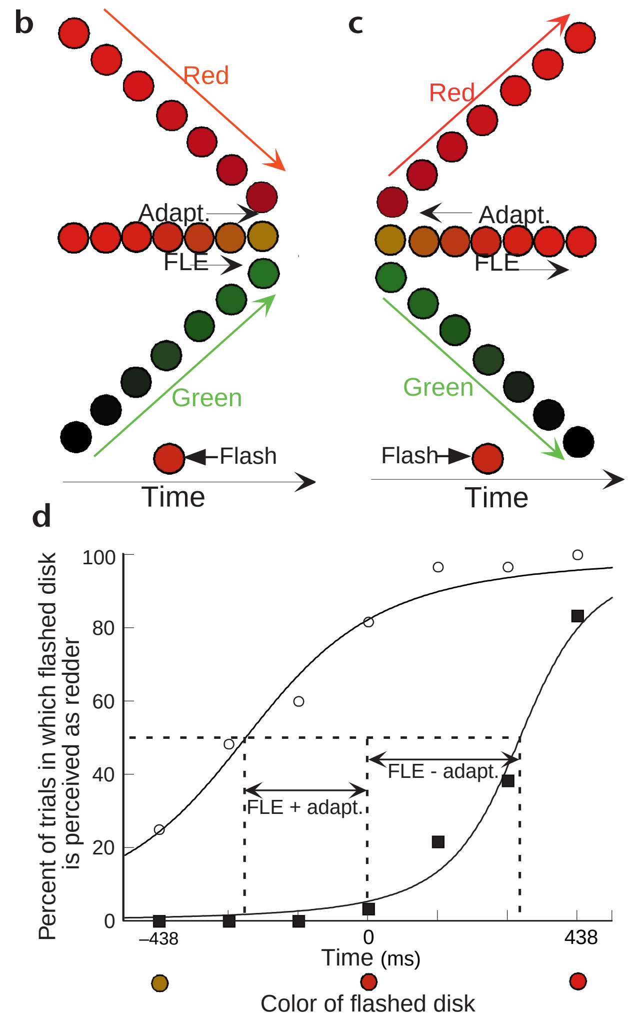

Fig. 4. Color aftereffect (adaptation) cannot explain the color flash-lag effect. (**a**) On alternate trials, a continuously present disk gradually changed color from red to reddish brown, or reddish brown to red. Halfway through, a second disk was flashed briefly on the opposite side of the FP. The color of the flashed disk was one of the colors of the continuously present disk at a particular time frame. (**b**, **c**) Time sequence of colors of the continuously present, changing disk and its component shades of red (top) and green (bottom). Red is more saturated than green throughout. (**b**) In the transition from red→reddish-brown, the direction of color adaptation is indicated (same direction as the flash-lag effect; FLE + adaptation). (**c**) In the reddish brown→red transition, the directions of color adaptation and flash-lag effect are opposite (FLE – adaptation). (**d**) Grouped responses and combined, fitted psychometric curves (n = 3) for both reddish brown→red (filled squares) and red→reddish brown (open circles) color sequences.

by other stimuli similar to the target in feature value and physical location. These conditions are also necessary to obtain the present flash-lag effect (B.R.S. *et al.*). We explain our results as follows: because of transmission delays[29], it took some time after the physical onset of the flashed stimulus before it was perceived. During this delay, the frame within the continuously changing sequence that occurred at the same time as the flash had already been succeeded by subsequent frames. These frames appeared soon after this synchronous frame (SOA in our experiments was 43–73 ms, well within the 100–150 ms time frame for masking), were very much like it in feature value, and were presented in the same spatial location. These conditions are optimal for backward masking[30–32]. Actually, for any given frame in the sequence, succeeding frames that appear within 100–150 ms are likely to mask it. In particular, the synchronous frame will be masked by later frames, preventing its signal from crossing perceptual threshold. On the other hand, stimuli not otherwise perceived can nonetheless improve detectability of similar stimuli presented thereafter[33,34]. This phenomenon of priming lasts longer than 1–2 seconds—the duration of the continuously changing sequence in our experiments. Therefore, we suggest that this subthreshold activity nonetheless facilitates activity of succeeding frames. In the case of the synchronous frame, attention is transiently diverted to the flash[14], which probably attenuates accumulated priming due to frames in the sequence before the synchronous frame. As a result, the signal corresponding to the synchronous frame is not able to overcome the masking from later frames in the sequence. Following the flash, however, the subthreshold activity gradually accumulates. (The priming corresponding to successive frames is successively greater.) When the subthreshold activity combines with direct input from one or more[35] frames of the sequence that physically occur after the flash, the activity crosses threshold. As a result, when the signal corresponding to the flash is perceived with it, a post-flash frame is simultaneously perceived. Because this post-flash frame in the sequence is ahead of the flash in feature space, one observes the flash-lag effect. Masking explains why the frame physically synchronous with the flash is not perceived, whereas priming causes subthreshold activity to accumulate, thereby allowing a subsequent frame to be perceived. Thus, in our view, both priming and masking contribute to the compensation of neural delays and to the flash-lag effect.

Methods

All stimuli were presented on a Sony Trinitron monitor (Tokyo, Japan; 75 Hz refresh) under control of a Mac PowerPC running Matlab (MathWorks, Natick, Massachusetts) and the Psychophysics Toolbox extensions[36,37]. From a pool of ten observers (one author, nine naive observers), a subset was chosen for each task. Each observer's head was partially immobilized by means of a chinrest (Handaya, Tokyo, Japan) placed 28.5 cm from the computer monitor (screen dimensions, 75° × 60°). Room lights were turned off, and the observer was asked to fixate a gray dot (0.2° × 0.2°) in the center of the screen throughout the experiment. Viewing was binocular. Specific verbal instructions were given before each experiment. Observers had to respond at the end of each trial by pressing one of two adjacent keys on a computer keyboard. No feedback was provided. A typical experiment lasted eight to ten min. In separate trials or in separate trial blocks, each stimulus sequence was run both forward and backward. This eliminated any selection bias the observer might have in the task.

Color task. For three observers, we measured the physical luminance of every frame using a photometer (Minolta LS-100, Minolta Corporation, Ramsey, New Jersey) and equalized the physical luminance of the continuously changing disk across all the frames in the sequence (12.7 cd per m^2). For the remaining three observers, their respective equiluminance points were obtained using the classical flicker-fusion

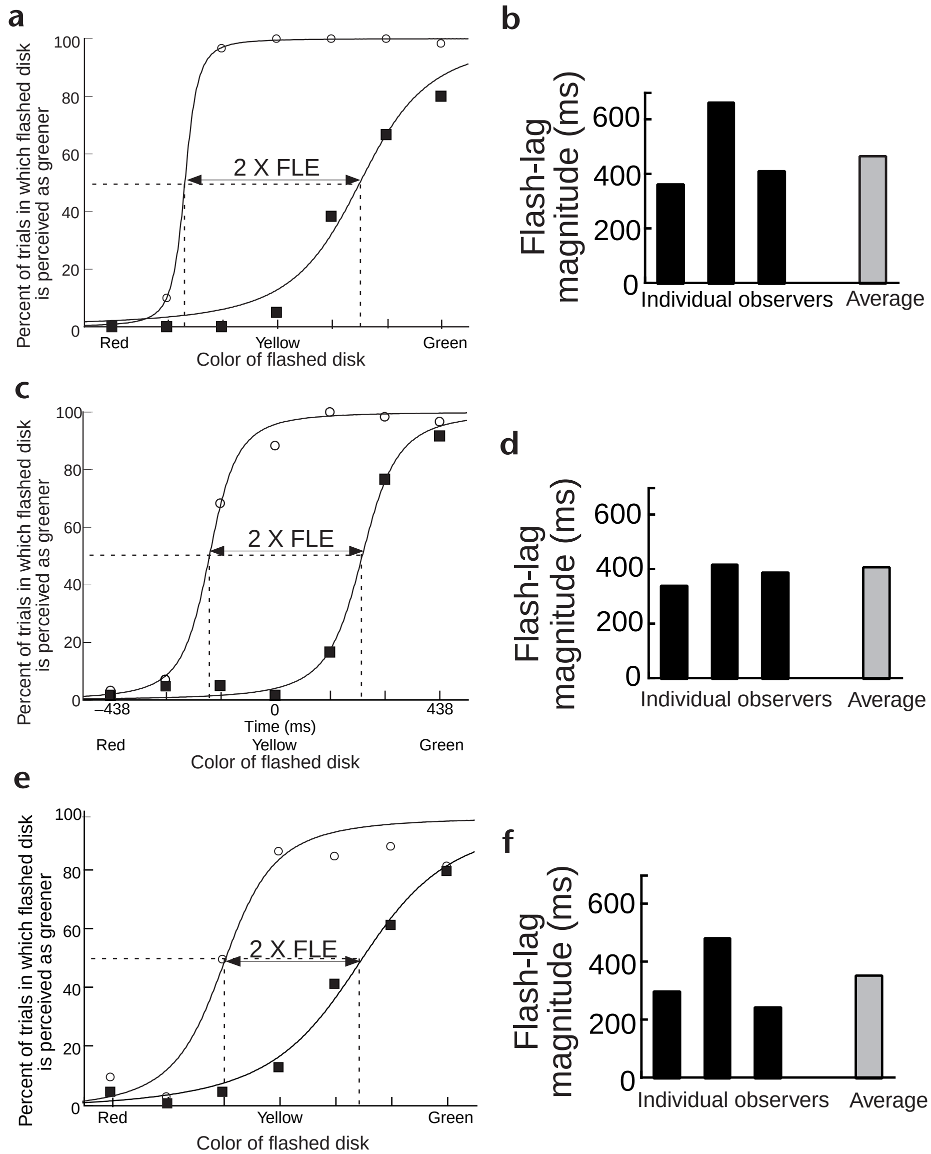

Fig. 5. Reducing the intensity of the flash or its duration did not enhance the flash-lag effect. (**a**) The combined data (n = 3) and fitted psychometric curves for red→green (filled squares) and green→red (open circles) sequences in the baseline condition. (**b**) Flash-lag effect magnitudes are shown. Mean magnitude of the color flash-lag effect, 466 ms. (**c**) Combined data and fitted psychometric curves for red→green (filled squares) and green→red (open circles) transformations in the reduced flash intensity condition. (**d**) Flash-lag effect magnitudes in the reduced flash intensity condition. The mean color flash-lag effect magnitude was 406 ms. (**e**) The pooled data for the same three observers and fitted psychometric curves for red→green (filled squares) and green→red (open circles) sequences with shorter flash duration (18.25 ms). The continuously changing disk remained on for the same amount of time as in the baseline (**a**; 2200 ms). (**f**) Flash-lag effect magnitudes with shorter flash duration. Mean flash-lag effect, 352 ms.

technique[38]. A circular disk (diameter, 2.9°), centered 3.5° directly above/below the FP, appeared at the beginning of each trial. Disk location (above or below the FP) was varied from trial to trial. The disk continuously changed color from green (CIE coordinates, x = 0.310, y = 0.572) to red (CIE coordinates, x = 0.602, y = 0.344) or from red to green. Each trial consisted of 32 or 64 discrete frames, and each frame was presented for ~73 ms on the monitor. Exactly halfway through the trial, a second disk appeared briefly (for ~73 ms) and then disappeared. The color of this flashed disk was randomly chosen from a set of seven predetermined, equally spaced, red/green color combinations and was one of the colors of the first, continuously present disk at some fixed time following trial onset. At the end of each trial, the observer had to respond by pressing the appropriate key on a keyboard, as to which of the two disks (continuously present or flashed) was greener at the instant the flashed disk was seen. Trial duration was approximately 2.2 or 4.4 s.

Luminance task. A circular disk (diameter, 2.9°), centered 3.5° directly above/below the FP, appeared at the beginning of each trial. The disk was bright or dim at the outset and gradually became dimmer or brighter, respectively. Each trial consisted of 64 discrete frames, and each frame was presented for 13 ms on the monitor. Halfway through the trial, a second disk of a randomly chosen fixed luminance appeared for one frame only (13 ms), and then disappeared. At the end of each trial, the observer had to respond as to which of the two disks (continuously present or flashed) was brighter at the time of the flash by pressing the appropriate key. Each trial took approximately 830 ms.

Spatial-frequency task. A vertical square-wave grating patch (11.7° × 0.6°) appeared at trial onset whose spatial frequency increased from a starting value of 0.7 cycles per degree to a final value of 6 cycles per degree, or *vice versa* (decreasing spatial frequency). The patch was centered 0.6° directly above/below the central FP and alternated location on every trial. Halfway through each trial, a second patch of equal area and distance from FP was briefly presented. For 2 observers, the flash duration was 43 ms, and the trial time was 300 ms; for the remaining 2 observers, the flash stayed on for 87 ms, and the trial lasted for 600 ms. The spatial frequency of the patch was randomly chosen from a set of seven values with equal probability. At the end of each trial, the observer responded as to which of the two rectangular patches (continuously present or flashed) was of a higher spatial frequency at the time of the flash.

Pattern-entropy task. At the beginning of each trial of the increasing entropy sequence, a grid of equally spaced dots (11 × 11 grid of dots, each of size 0.05° × 0.05° and same luminance) was inscribed inside a square area (6.4° × 6.4°) that provided an invisible boundary for all the dots. The sequence consisted of 13 frames, each with a progressively more irregular arrangement of dots. On each successive frame of the sequence, an increasing number of randomly chosen dots were allowed to deviate from their location on the grid. The range of positions each such chosen dot could occupy was constant throughout. On each frame, the fraction of dots selected to wander from their position on the grid were chosen randomly and independently of whether or not they were chosen in other frames. Each frame stayed on for 73 ms, and the trial lasted for approximately 1530 ms. Halfway through (~730 ms after trial onset), a second such configuration of dots (partially regular grid of dots) was briefly presented for 73 ms. The arrangement of dots in this flashed stimulus was randomly chosen to be exactly identical to one of the 13 discrete frames of the continuously present stimulus with equal probability. The observer had to choose the stimulus that appeared to have a more disorderly arrangement of dots at the time of the presentation of the flashed configuration. For the decreasing entropy sequence, the sequence of frames was reversed.

Color flash-initiated task. Two circular disks (diameter, 2.9° each), centered 3.5° above and below the FP, appeared at the beginning of each trial. One disk stayed on for 73 ms and was randomly chosen to be one of seven predetermined colors, and the second disk was initially yellow and gradually changed in color to green (red) over 1100 ms. Each frame of the continuously changing sequence lasted 73 ms, and all the frames were equiluminant. In the yellow→red (green) condition, the color of the flash was intermediate between the initial yellow and final red (green) values of the continuously changing disk. At the end of each trial, the observer had to choose the disk (continuously present or flashed) that appeared greener at the start of the trial.

Color adaptation task. On alternate trials, a continuously present disk gradually changed color from red (CIE coordinates, $x = 0.602$, $y = 0.344$) to reddish brown (CIE coordinates, $x = 0.491$, $y = 0.430$), or reddish brown to red. The intensities of red and green in each frame of the sequence were measured separately by means of a photometer, and the sum of the intensities was kept constant across frames (12.7 cd per m^2). The intensity of red was greater than that of green in all frames. The timings were identical to that in the color flash-lag experiment. Each trial consisted of 32 discrete frames, and each frame was shown for 73 ms. The stimulus sequence lasted for 2.2 s. Midway through, a second disk of a randomly chosen fixed color was flashed for one frame duration (73 ms). At the end of each trial, the observer had to choose the disk (continuously changing or flashed) that appeared redder at the instant the flashed disk was seen.

Statistical analysis and computer simulations. The points for each sequence (red→green, green→red) in Figs. 1c and 2d were fitted with psychometric curves using probit analysis[39,40]. Intersections with the vertical dotted lines define the threshold estimates (T_{50}) of the two curves, and mark the points at which the flashed disk appeared more green than the continuously present disk as often as it appeared less green. In the absence of a flash-lag effect, the curves should be identical and should project onto the same point on the abscissa [$T(\text{red} \rightarrow \text{green})_{50} - T(\text{green} \rightarrow \text{red})_{50} = 0$]. The magnitude of the flash-lag effect is given by half of the distance between their projections, or threshold estimates, converted into time units. Computer simulations using the bootstrap method[41] to test the significance of the difference in the two thresholds enumerated all possible pairs of psychometric functions from the pooled distribution (red → green and green → red averaged) and weighted them by their binomial probability. Each of the possible pairs of psychometric functions yielded a pair of threshold estimates. The difference in threshold between the two psychometric functions [$T(\text{red}\rightarrow\text{green})_{50} - T(\text{green}\rightarrow\text{red})_{50}$] was compared with the distribution of the threshold differences between the pairs of psychometric functions generated using the bootstrapping method. The upper and lower threshold-difference values of the distribution that would exclude the upper and lower 0.05% of the bootstrapped population, respectively, were taken as the 99.9% confidence limits. Using this technique, we found that the thresholds for both the curves fell outside of the confidence limits ($p < 0.001$ for both).

Acknowledgements

This research was supported by NSF grant SBR-9710116 and a Caltech Division of Biology fellowship awarded to B.R.S. We thank Johanna Weber for reviewing our manuscript, Christian Scheier for providing the statistics software, Gabriel Kreiman for help with printing and Chris Tyler for suggestions on improving the pattern-entropy experiment.

RECEIVED 17 SEPTEMBER 1999; ACCEPTED 20 MARCH 2000

1. Nijhawan, R. Motion extrapolation in catching. *Nature* **370**, 256–257 (1994).
2. Nijhawan, R. Visual decomposition of colour through motion extrapolation *Nature* **386**, 66–69 (1997).
3. Berry, M. J., Brivanlou, I. H., Jordan, T. A. & Meister, M. Anticipation of moving stimuli by the retina. *Nature* **398**, 334–338 (1999).
4. Cavanagh, P. Predicting the present. *Nature* **386**, 19–20 (1997).
5. Zeki, S. The representation of colour in the cerebral cortex. *Nature* **284**, 412–418 (1980).
6. Hartline, H. K. The response of single optic nerve fibers of the vertebrate eye to illumination of the retina. *Am. J. Psychol.* **57**, 243–249 (1938).
7. Kuffler, S. W. Discharge patterns and functional organization of mammalian retina *J. Neurophysiol.* **16**, 37–68 (1953).
8. Hubel, D. H. & Wiesel, T. N. Integrative action in the cat's lateral geniculate body. *J. Physiol. (Lond.)* **155**, 385–398 (1961).
9. Campbell, F. W., Cooper, G. F. & Enroth-Cugell, C. The spatial selectivity of the visual cells of the cat. *J. Physiol. (Lond.)* **203**, 223–235 (1969).
10. Schiller, P. H., Finlay, B. L. & Volman, S. F. Quantitative studies of single-cell properties in monkey striate cortex. III. Spatial frequency. *J. Neurophysiol.* **39**, 1334–1351 (1976).
11. Schiller, P. H., Finlay, B. L. & Volman, S. F. Quantitative studies of single-cell properties in monkey striate cortex. I. Spatiotemporal organization of receptive fields. *J. Neurophysiol.* **39**, 1288–1319 (1976).
12. Albright, T. Direction and orientation selectivity of neurons in visual area MT of the macaque. *J. Neurophysiol.* **52**, 1106–1130 (1984).
13. Hammond, P. & MacKay, D. M. Differential responses of cat visual cortical cells to textured stimuli. *Exp. Brain Res.* **22**, 427–430 (1975).
14. Jonides, J. & Yantis, S. Uniqueness of abrupt visual onset in capturing attention. *Percept. Psychophys.* **437**, 346–354 (1988).
15. Baldo, M. V. C. & Klein, S. A. Extrapolation or attention shift? *Nature* **378**, 565–566 (1995).
16. Khurana, B. & Nijhawan, R. Extrapolation or attention shift? *Nature* **378**, 565–566 (1995).
17. von Helmholtz, H. *Handbuch der Physiologischen Optik* **II** 240–245 (Verlag von Leopold Voss: Hamburg und Leipzig, 1909).
18. Wohlgemuth, A. On the aftereffect of seen movement. *Br. J. Psychol. Monogr. Suppl.* **1**, 1–117 (1911).
19. Snowden, R. J. Shifts in perceived position following adaptation to visual motion. *Curr. Biol.* **8**, 1343–1345 (1998).
20. Nishida, S. & Johnston, A. Influence of motion signals on the perceived position of spatial pattern. *Nature* **397**, 610–612 (1999).
21. Suzuki, S. & Cavanagh, P. Focused attention distorts visual space: An attentional repulsion effect. *J. Exp. Psychol. Hum. Percept. Perform.* **23**, 443–463 (1997).
22. Blakemore, C. & Sutton, P. Size adaptation: a new aftereffect. *Science* **166**, 245–247 (1969).
23. DiLollo, V. Temporal integration in visual memory. *J. Exp. Psychol.* **109**, 75–97 (1980).
24. Coltheart, M. Iconic memory and visible persistence. *Percept. Psychophys.* **27**, 183–228 (1980).
25. DiLollo, V. & Bischof, W. F. Inverse-intensity effect in duration of visible persistence. *Psychol. Bull.* **118**, 223–237 (1995).
26. Purushothaman, G., Patel, S. S., Bedell, H. E. & Ogmen, H. Moving ahead through differential visual latency. *Nature* **396**, 424–424 (1998).
27. Whitney, D. & Murakami, I. Latency difference, not spatial extrapolation. *Nat. Neurosci.* **1**, 656–657 (1998).
28. Eagleman, D. M. & Sejnowski, T. J. Motion integration and postdiction in visual awareness. *Science* **287**, 2036–2038 (2000).
29. Maunsell, J. H. & Gibson, J. R. Visual response latencies in striate cortex of the macaque monkey. *J. Neurophysiol.* **68**, 1332–1344 (1992).
30. Turvey, M. T. On peripheral and central processes in vision. *Psychol. Rev.* **80**, 1–52 (1973).
31. Breitmeyer, B. G. *Visual Masking: An Integrative Approach* (Oxford Univ. Press, New York, 1984).
32. Bachmann, T. *Psychophysiology of Visual Masking* (Nova, Commack, New York, 1994).
33. Bartram, D. J. The role of visual and semantic codes in object naming. *Cognit. Psychol.* **6**, 325–356 (1974).
34. Marcel, A. J. Conscious and unconscious perceptions: An approach to the relations between phenomenal experience and perceptual processes. *Cognit. Psychol.* **15**, 238–300 (1983).
35. Efron, R. Conservation of temporal information by perceptual systems. *Percept. Psychophys.* **14**, 518–530 (1973).
36. Brainard, D. H. The Psychophysics Toolbox. *Spat. Vis.* **10**, 443–446 (1997).
37. Pelli, D. G. The Video Toolbox software for visual psychophysics: tranforming numbers into movies. *Spat. Vis.* **10**, 437–442 (1997).
38. Boynton, R. M. & Kaiser, P. K. The additivity law made to work for heterochromatic photometry with bipartite fields. *Science* **161**, 366–368 (1968).
39. Finney, D. J. *Probit Analysis* (Cambridge Univ. Press, Cambridge, UK, 1971).
40. McKee, S. P., Klein, S. A. & Teller, D. Y. Statistical properties of forced-choice psychometric functions: Implications of probit analysis. *Percept. Psychophys.* **37**, 286–298 (1985).
41. Foster, D. H. & Bischof, W. F. Thresholds from psychometric functions: Superiority of bootstrap to incremental and probit variance estimators. *Psychol. Rev.* **109**, 152–159 (1991).