How Perception Adheres Color to Objects and Surfaces:

Studies Using

Visual Illusions and Transcranial Magnetic Stimulation

Thesis by

Daw-An Wu

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

California Institute of Technology

Pasadena, California

2006

(Defended August 22, 2005)

© 2006

Daw-An Wu

All Rights Reserved

To my Parents

Abstract

The basic roadmap of the visual cortex starts with an early stage in which visual input is encoded primarily as edges. Information then flows through branching pathways into higher, more specialized regions of the brain. One such area is specialized in processing color, another area processes shape, etc. This encoding scheme poses some problems in explaining certain aspects of our daily experience of seeing. If the cortex processes the various visual features separately, how do we see all those elements integrated as unified objects? What mechanisms *bind* the features together? If the cortex encodes the visual scene in terms of its edges, then how do we see solid surfaces? What mechanisms *fill-in* the map of outlines?

This thesis investigates the problems of binding and filling-in using the techniques of visual illusion psychophysics and transcranial magnetic stimulation (TMS). TMS is known to cause the perception of a brief visual flash, or phosphene. We find that the appearance of the phosphene is found to depend on concurrent and previously viewed visual stimuli. In particular, TMS can cause an *instant replay* effect whereby recently presented visual stimuli are seen again. This serves as an effective probe for the visual information contained within hidden internal brain states.

TMS-induced instant replay is used to probe the mechanisms of Cai's asynchronous binding illusion, in which a color change in a moving object is perceived to occur later in the motion stream. After subjects view this illusion, TMS can sometimes cause them to see a replay of

iv

the color change, but in the correct position. This indicates that the color change was properly encoded in some portion of the visual system, but that this representation normally remains unconscious due to further processing. It appears that TMS can selectively reactivate this representation, revealing its visual content without the distortion caused by other processes.

Further investigations of binding demonstrate a number of cases in which visual features are decomposed and/or misbound. TMS-induced instant replay can cause the color of one object to be bound to the position and orientation of another. It can also separately replay the color and orientation of a grating. Finally, in a non-TMS experiment, we create a stimulus that induces a steady-state misbinding of color and motion. The illusion is a vivid, long-lasting misbinding effect ideal for neurophysiological investigation. These experiments confirm the separate encoding of visual features, and the existence of an active binding mechanism.

In an investigation of filling-in, a class of illusions derived from the artwork of Julian Stanczak is shown to defy the dominant model of color filling-in. Color is perceived to spread discretely among segregated patches of space. We present evidence that the extent of color filling is based on the high-level processing of globally defined perceptual surfaces rather than the low-level processing of locally defined retinotopic features. Based on a review of past electrophysiological studies and the phenomenology of the illusions presented here, it appears possible that the neural mechanisms of binding and filling-in might be intimately related, both of them highly integrated with the process of surface segregation.

V

Contents

Dedication		
Abstract		
Table of Contents		
List of Illustrations		
Chapter 1: Introduction		
Chapter 2: Development of TMS as a Probe—The "Instant Replay" Effect		
Background	8	
Experiment 1: The baseline phosphene	10	
Experiment 2: Interaction of the phosphene and visual stimuli	15	
Experiment 3: Interaction across time—instant replay	20	
Chapter 3: Object Color Perception—the binding problem		
Background	24	
Experiment 1: TMS probe of Cai's asynchronous binding illusion	27	
Experiment 2: TMS unbinds the visual features of a color grating	39	
Experiment 3: Inducing misbinding without TMS:		
Steady-state misbinding of color and motion	42	
Chapter 4: Surface Color Perception—Color Filling-In		
Background	48	
Experiment 1: Color jumps over gaps, filling into distant areas	53	

Experiment 2: Texture-defined edges block color-filling	58
Experiment 3: Overlaid color-filling processes are independent and selective	61
Experiment 4: Color-filling on perceptual surfaces in transparent motion	65
Discussion	69
Chapter 5: Conclusion	76
Appendix A: Detailed methods	80
References	89

List of illustrations

Figure 1.1	Classical visual illusions.	4
Figure 2.1	Sample phosphene percept at baseline	11
Figure 2.2	Phosphenes enhance overlapping visual stimuli	17
Figure 2.3	Instant replay	21
Figure 3.1.1	Cai's asynchronous binding illusion	28
Figure 3.1.2	"Clock" form of Cai's illusion, followed by TMS	32
Figure 3.1.3	Sample subject data	33
Figure 3.2	TMS-induced visual feature decomposition	40
Figure 3.3	Steady-state misbinding of color and motion	43
Figure 4.1.1	Basic filling-out stimuli and trial sequence	52
Figure 4.1.2	Basic filling-out results	55
Figure 4.2	Color-filling over texture-defined surfaces	59
Figure 4.3	Color-filling over two overlaid color gradients	63
Figure 4.4	Color-filling over perceptual surfaces defined by transparent motion	66

Chapter 1: Introduction

The modern roadmap for visual processing describes a branching network of brain areas, each performing their own particular transformations on the visual input. In the earliest cortical stages, cells respond robustly to spatial and temporal change, distilling our visual input into a map of edges (Hubel & Wiesel 1968). As neural activity travels to later stages, the paths branch and brain areas become specialized. For example, one area contains cells primarily sensitive to different directions of motion; another area, color (Livingstone & Hubel 1988; DeYoe, et al. 1994). As much as this roadmap lays the foundation for our understanding of visual computation, it also makes it difficult to explain certain aspects of our daily experience of seeing.

For one thing, we see colored shapes as solid and filled-in, even though the very first stage of visual cortex basically contains a visual map consisting of outlines. When a shape of solid color is viewed, most of the resulting neural activity corresponds to the edge of the shape. Almost no activity is found corresponding to the interior surface. Given this edge-based encoding of visual input, it seems there must be further neural mechanisms to decode the map, *filling-in* the surfaces.

Also, we see the color, shape and motion of objects as unified percepts, even though the processing of those individual features occurs in different regions of the brain. When we see a red ball rolling across the floor, its redness, its roundness, and its motion seem inseparably

bound together as features of an object. If all of these features are processed in different areas of the brain, then what mechanisms "bind" all the features together? What keeps objects unified and prevents features of one object from being associated onto another?

Many possible computational and neurobiological solutions to the problems of "binding" and "filling-in" have been proposed. These solutions fall into three basic categories: the explicit decoding solutions, the connection-based solutions, and the symbolic solutions.

The explicit solutions tend to result in anatomical activity patterns that "look like" the desired visual percept. For example, the filling-in of surfaces might be accomplished at some level if edge-based activity diffused laterally into neurons representing neighboring space (Cohen & Grossberg 1984). The binding of visual features might occur if the scattered processing streams for the features of an object are somehow routed into an even higher order set of cells which code for complete objects (Olshausen, et al. 1993; Olshausen, et al. 1995).

Meanwhile, the connection-based solutions propose that information might be encoded in the establishment of linked networks, the entirety of which would represent the percept. For example, the different features of a single object might be encoded by cells in disparate regions of the brain, but these scattered cells would maintain connections to each other, perhaps by synchronizing their activity patterns (Singer & Gray 1995).

Finally, the symbolic solutions argue that there really is not a problem to be solved. There is no reason why the anatomical activity patterns have to "look like" or be organized the way we experience them qualitatively (Garson 2001; von der Heydt, et al. 2003). Neurons that code for colors are, of course, not themselves physically colored, and the upside-down image that the retina receives does not mean that we see the world upside down. If we could scramble up all of the cells of the visual cortex, as long as they were to somehow maintain the same types of connections to the retina and to each other, we could imagine that the system would still work the same way. Similarly, particular patterns of activity can simply "stand for" a visual percept even if their anatomical organization does not look quite the same.

In this thesis, I address the aforementioned problems of binding and filling-in using the convergence and integration of two experimental techniques: The classical psychophysics of visual illusions, and transcranial magnetic stimulation.

Visual illusions have long been used to reveal the computational and neurobiological mechanisms of the brain. The systematic deviations in perception that are induced by certain stimuli reveal the basic visual cues that the visual system relies on to extract higher-order information. For example, manipulations of the Titchner and Mueller-Lyer illusions (figure 1.1a,b) demonstrate the contextual cues that the visual system uses in judging size and length. Other illusions can be related directly to the basic neurobiology of the brain. For



Figure 1.1: Classical visual illusions

(a) Titschner illusion. The central circles are the same size, but left one looks smaller

(b) Mueller-Lyer illusion. The vertical line segments are the same length, but the left looks longer.

(c) Hermann grid illusion. Gray spots are seen at the intersections of the gaps, a manifestation of lateral inhibition

example, the Hermann grid illusion (figure 1.1c) illustrates the mechanism of lateral inhibition between neighboring neurons.

In comparison to visual illusions, the use of transcranial magnetic stimulation (TMS) in neuroscience is fairly new. TMS is a noninvasive and safe method of altering brain activity in human subjects. Transient electrical field gradients are magnetically induced inside the brain by an insulated coil of wire, which is held above the subject's head. Because no current actually flows between the equipment and the subject, TMS avoids the pain associated with techniques such as electrical stimulation using scalp electrodes.

This thesis combines the manipulation of known visual illusions, the discovery and investigation of new visual illusions, and the use of TMS to explore the computational and neurobiological mechanisms underlying the binding problem and filling-in.

Chapter 2 of this thesis describes the discovery, development and exploration of a TMS paradigm that can be used as a probe into internal brain states. This probe extracts the visual information encoded in those brain states and reveals them in the visual percepts that the subject experiences. In an *instant replay* effect, the visual information that is revealed reflects not only the processing of concurrent visual inputs, but visual stimuli from the recent past—the subject re-perceives earlier visual stimuli.

Chapter 3 describes an investigation of the binding problem, using TMS and visual illusions, in combination and separately. First, instant replay is applied to Cai's *asynchronous binding* illusion. The replay of a situation where the perceptual experience diverges from the physical stimulus proves informative in deducing the mechanism underlying the illusion. In certain subjects, TMS caused the sudden appearance of a flash at a location corresponding to the real stimulus, rather than the previously seen illusory percept. This indicates that the stimulus was correctly encoded in some part of the visual system, casting doubt on the asynchronous binding account of the illusion. The result also shows that TMS does not simply induce a wholesale recapitulation of recent visual cortex activity, but that it preferentially activates certain components of the system, some of which may have been previously suppressed or altered by other components. This allows us to look more directly at certain component visual representations, without the effects of masking or distortion by other factors.

Chapter 3 also describes several instances of misbinding, providing evidence that there is indeed an active binding mechanism, arguing against a symbolic solution for the binding problem. TMS-induced instant replay is found, on occasion, to reflect decomposition and/or misbinding of features from previous stimuli. Furthermore, in the absence of TMS, an unchanging, steady visual stimulus is found to induce a long-lasting illusory misconjunction of color and motion. These properties of the illusion make it uniquely suitable for neurophysiological investigation, opening up the topic of visual feature binding to a search for neural correlates and mechanisms.

Chapter 4 describes an investigation of filling-in, the process whereby smooth surfaces of color are defined from early visual cortex's edge-based map of activity. Drawing from an effect seen in art by Julian Stanczak, we demonstrate a class of *discrete color-filling* illusions that defy the classical model of color-filling. Color is found to fill from one area into remote areas, jumping over intervening gaps. In manipulating the discrete color-filling effect, we show that the processing of globally-defined surfaces is involved in determining the local spread of color.

Although binding and filling-in are generally thought of as separate processes, our findings point to the possibility that the two are intimately related. The phenomena may be different faces of several interwoven mechanisms. In particular, the process of perceptual surface segregation seems to be central to both phenomena. The steady-state misbinding illusion relies heavily on how perceptual surfaces are interpreted from stimulus features. The color-filling illusion also operates on the basis of the perceptual surfaces that are parsed from the stimuli. In reviewing the neurophysiological investigations of surface segregation and classical color-filling, and then considering the computational requirements imposed by discrete color-filling, we propose the possibility that the mechanisms of color-filling resemble the binding of high-level surface color information to low-level edge information. Neurophysiological investigation of discrete color-filling should indicate whether this feedback model is valid, or if modifications of the classical local network model are sufficient.

Chapter 2: Development of TMS as a Probe—the Instant Replay Effect.

Background

Transcranial magnetic stimulation (TMS) is a safe, non-invasive method of stimulating human cortical tissue. The device relies on magnetic induction to create electrical field gradients around coils of wire. The procedure basically involves running a pulse of current through loops of wire held above the subject's head. Although no current runs between the device and the subject, the brain is subject to the effects of magnetic induction, whereby a transient electrical field gradient is created around the stimulus coil.

This method was first used to stimulate motor cortex, resulting in the activation of various muscle groups. Motor stimulation studies remain the bulk of basic TMS research, as the measurements of muscle potentials provides an objective and quantitative source of data in studying low level phenomena such as the interactions between paired TMS stimuli (Valls Sole, et al. 1992).

When applied over visual cortex, TMS has two types of perceptual effects, elicited under different stimulus conditions. A single pulse tends to elicit scotoma, a transient "blind spot" which is noticeable when TMS is administered with flashed visual stimuli (Amassian, et al. 1994; Kamitani & Shimojo 1999). A pair of pulses administered ~50 ms apart tends to elicit a phosphene, a brief flash of illusory light (Meyer, et al. 1991; Marg & Rudiak 1994).

Studies involving TMS-induced phosphenes generally take place in total darkness, in darkadapted subjects. The minimal stimulus strength needed to elicit the phosphene percept gives a measure of gross cortical excitability. These measure have, for example, been used to show a correlation between chronic migraines and elevated visual cortex excitability (Aurora, et al. 2003).

In the following studies, we induce phosphenes while the subject views visual stimuli. The perceptual interactions between the two stimuli are used to probe both the properties of phosphene induction as well as the mechanisms of visual processing.

Experiment 1: The baseline phosphene

We first assessed the phenomenology of the phosphene under baseline conditions – either in total darkness, or in a dimly lit room. Following established protocols, we delivered two pulses 50 ms apart through a figure-8 coil (See appendix A for detailed methods). When TMS is applied to optimally stimulate left visual cortex, the subject typically sees a brief flash of brightness (the phosphene) somewhere in the lower right visual quadrant. The phosphene is generally colorless or very pale in color. The texture of the phosphene is not smooth, but tends to have a finely speckled or pixilated quality. The spatial distribution of the phosphene varies widely between subjects. Borders are usually not clearly defined, except if the phosphene abuts the horizontal or vertical midlines, in which case there are clear, sharp and straight edges to the phosphene at these borders (see figure 2.1 for example).

The mapping of coil position to phosphene position is very poor. In general, the movement of the coil does not result in a systematic shift in the location of the phosphene. Rather, the movement of the coil away from an optimal position results in the weakening and shrinkage of the phosphene. Control of phosphene location is essentially limited to the stimulation of phosphenes in either the left or right hemifield, which is controlled by the polarity of the current rather than the position of the coil.

The anatomical locus of phosphene stimulation is highly debated. Most researchers have suggested V1 as a candidate, others have suggested extrastriate cortex or the optic radiation



Figure 2.1: Sample phosphene percept at baseline

The subject sits in a dimly lit room in front of blank computer screen. A Magstim figure-8 coil is placed over occipital cortex so that the current polarity points rightward, optimally stimulating left visual cortex. When two TMS pulses are delivered with 50 ms separation, the subject perceives a momentary brightness in the lower right quadrant of the visual field.

The phosphene's brightness is not uniform, and can have a noisy, pixelated quality to it. The edges of the phosphene are poorly defined except where the phosphene meets the horizontal and vertical midlines. The overall distribution of the phosphene varies from subject to subject, the above is a crude rendering of one instance.

(reviewed in Kammer, et al. 2005). Although this study is not aimed at resolving this issue, the reports by many subjects of sharp, straight borders at the horizontal midline may provide a clue as to cortical regions participating in the generation of the phosphene.

Such sharp borders are also found in quadrantic visual field defects, in which patients experience flashes of light or lack of visual awareness restricted to one of the visual quadrants. In one study, pre- and post-operative magnetic resonance images were taken of patients experiencing such quadrantic field defects. The location of cortical abnormalities indicated that the symptoms arose from defects or lesions in extrastriate areas V2 and/or V3. The researchers pointed out that a lesion that includes the anatomical border between V2 and V3 would be the most likely source of visual phenomena in which the percept stops at the horizontal meridian in a strict manner. This is because V2 and V3 do not contain continuous maps of the visual hemifield. Instead, each is divided into two anatomically separate portions. The portions representing the upper visual quadrant lie deep within the hemispheric fissure, while the portions representing the lower visual quadrant lie on the surface of the brain. Thus, a lesion could easily affect one quadrant without crossing over into the other. Furthermore, the retinotopic maps of V2 and V3 mirror each other, such that their representations of the horizontal meridian both abut the V2/V3 border. Thus, an irregularly shaped lesion that includes the V2/V3 border would result in a visual defect which is sharply limited at the horizontal meridian (Horton & Hoyt 1991).

Applying this reasoning to phosphene generation, the subject reports indicating that the phosphene was sharply limited to a lower quadrant of the visual field indicate the involvement of extrastriate cortex in phosphene stimulation. The limitation of the phosphene to the lower quadrant is consistent with the fact that the lower quadrant's V2 and V3 representations lie closest to the scalp. The phosphenes sharp border at the horizontal meridian would suggest that TMS is activating an area that includes both V2 and V3. This is consistent with a much more complete physiological explanation proposed by Kammer et al. (Kammer, et al. 2005).

When considering TMS as a method, it is tempting to view it as similar to techniques such as electrode-based microstimulation. One key difference, however, lies in the level of control one has in stimulating particular regions of the brain. It is not simply that the focus of stimulation is broader, but that the electromagnetic focus of stimulation does not seem to be the critical component in determining the spatial distribution of the phosphene. Our inability to systematically control the phosphene's position by moving the TMS coil is consistent with past reports (Kammer, et al. 2005; Marg & Rudiak 1994). It appears that the distribution of neural activity is more dependent on set internal factors such as neuroanatomy and local excitability.

This is consistent with theoretical modeling and experimental findings. Basic physics and single cell modeling of TMS both indicate that the length, shape, and orientation of neural processes is critical to the amount of activity induced in the cell by electric field gradients

(Kamitani, et al. 2001). A study using optical imaging and TMS of cat visual cortex also shows that the anatomical spread cortical activity contains components which are independent of TMS coil position (Valero Cabre, et al. 2004).

Overall, these lines of evidence indicate that the pattern of activity elicited by TMS is heavily dependent on factors internal to the brain, which determine any given cell or network's excitability with respect to TMS.

Experiment 2: Interaction of the phosphene and visual stimuli

Having established the phenomenology of the phosphene when TMS is administered to visual cortex in a baseline state, I investigated how the phosphene would interact with visual stimuli that overlapped in visual space. Perceptually speaking, would the whiteness of the phosphene area of overlap appear as if the phosphene "covered" the visual stimulus, or would there be a transparent overlay? Alternatively, might there be some kind of interaction between the two elements' brightness and/or color? Neuro-computationally speaking, would TMS drive the network into a set state in the affected area, overriding the input from the visual system, or would TMS linearly add its induced activity to that induced by the visual stimulus? Alternatively, could there be a non-linear interaction between the two inputs?

In this experiment, a colored disk was flashed on the screen and timed with TMS so as to evoke the simultaneous perception of an overlapping disk and phosphene (see Appendix A for detailed methods).

The overlapping region was perceived to be filled with a brighter and more saturated version of the disk's color. This enhanced region appeared as a predominantly solid surface of color, rather than a mosaic of varying color intensities. Where inhomogeneities were present, they manifested as overlaid bits of noise rather than fluctuations in the color of the disk's surface.

The degree of enhancement was assessed quantitatively by having the observers adjust an onscreen display after each trial to match their percept. The degree of enhancement varied widely across subjects and conditions, between 17% and 200%. However, responses were statistically significant within each condition for each subject (N = 3, 2 conditions each, all 6 t-tests p < 0.05).

We thus find that there is extensive interaction between the inputs from visual stimuli and TMS. The TMS-induced brightness percept, rather than simply overriding or adding to the visual stimulus, in some sense *multiplied* it. Reds became more red; greens more green; etc. To frame this in terms of a simplistic and reductionist neuro-computational model, it seems that a given neuron's excitability with respect to TMS is activity-dependent. When a neuron, such as a "red" neuron, is already being stimulated by a visual display, TMS causes a much larger increase in activity than if the cell were at baseline. This might result if the field induced across the cellular membrane is small relative to spike threshold. TMS would then be unlikely to push a resting neuron past threshold, but could induce additional spiking in a neuron already possessing other depolarizing inputs.

Taking a slightly wider view, visual perception is probably not based on the independent activity of single cells, but on the establishment of propagating network activity patterns characteristic of conscious perception. Given a random and sparse activation of neurons by TMS, this disorganized activity is likely to give rise to weak and fragmented visual percepts. However, the presentation of the colored disk would induce organized activity in visual



Figure 2.2: Phosphenes enhance overlapping visual stimuli

When a flashed visual disk is timed with TMS so that the phosphene and disk are perceived at the same time, the two percepts interact. The area of overlap between the disk and phosphene is seen as a more intense version of the disk's color. In contrast to the uneven quality of the baseline phosphene, the enhancement of this area is much smoother.

cortex, providing a patterned substrate for TMS to build on. This would explain why the enhancement interaction smoothly brightened the disk's color, rather than adding scattered bits of color in accordance with the uneven spatial patterning of the baseline phosphene.

The previous section established that the pattern of activity induced by TMS is highly dependent on brain-internal factors. This experiment shows that these factors include not only fixed elements such as neuroanatomy, but also dynamic factors such as activity due to concurrent inputs.

Similar findings exist in the studies of TMS to motor cortex. When a subject actively tenses a muscle which has been targeted for activation by TMS, the magnitude of the muscle contractions induced by TMS increases. Also, if a subject practices a particular movement using the targeted limb or digit, this biases the direction and magnitude of movement evoked by subsequently applied TMS (Abbruzzese, et al. 1996). These effects might be partially attributed to interactions at the level of peripheral nerve or muscle, but nevertheless provides a clear physiological analogy.

Furthermore, other motor studies find a similar phenomenon with a clearly cortical basis. It has been found that if a subject concentrates on an internal *intent* to move in a particular direction, then the movement evoked by TMS is also biased (Fadiga, et al. 1999; Rossi, et al. 1998).

Overall, these lines of evidence indicate that the phenomenal effects of TMS are dependent on dynamic internal states of the subject's brain. This holds out the promise that TMS could be used as a probe that prospectively reveals internal brain states. Just as TMS of motor cortex reveals local brain states in terms of the movements that those brain states encode, the perceptual "read-out" from visual cortex could reveal brain states in terms of their encoded visual images.

Experiment 3: Interaction across time—instant replay

During the investigation of the interaction between the phosphene and visual stimuli, it became apparent that the two percepts did not have to be simultaneous to show the interaction. When TMS was timed so that the subject would see the colored disk and phosphene separately, one after the other, the percept induced by TMS was not an ordinary phosphene. Instead, in the area where the phosphene and disk would have overlapped, the subject saw a repetition of that portion of the colored disk. Subjectively, this region looked much like the region of enhancement seen in Experiment 2, just displaced in time.

Under optimal conditions, this "instant replay" had a certain photographic quality to it: colors were smooth and edges were sharply defined. The color of the replay was the same as the color of the disk, even in contrast-inverted cases where a black disk was flashed on a bright background. This was useful in differentiating the phenomena from afterimages, which were of the disk's complementary color.

When stimulation intensity was lowered, or the delay between the disk and TMS was lengthened, the replay became weaker. The color around the edge of the disk was the most resilient to weakening, and the edge itself remained sharp so long as replay was visible. This is consistent with the notion that smooth surfaces of color are encoded mainly by their edges. This edge-based encoding is found mainly in early visual cortex, suggesting an entry site of TMS-induced activity.



Figure 2.3: Instant replay

If TMS is delivered a second or two after a flashed disk has already disappeared, subjects re-perceive a portion of the disk. This re-perception is qualitatively similar to the enhancement interaction described in Experiment 2, except displaced in time. The edge of the replayed disk is the strongest. When lower intensities of TMS and/or longer inter-stimulus delays are used to diminish the replay effect, the sharpness of the edge is the most resilient feature. The amount of delay that can be allowed between the visual stimulus and TMS can be seen in an experiment described more fully in chapter 3, Experiment 2. Even with as much as two seconds separation, subjects saw a weak replay effect, and reported percepts containing both the color and form of the preceding display.

At the time of TMS, the subject has already seen the visual display come and go, and now sits in a dimly lit room facing a blank screen. As such the external conditions at the moment of TMS are the same as the conditions in our first "baseline" experiment. The replay effect seen in this experiment indicates that TMS can indeed prospectively reveal internal brain states that are otherwise not reflected in conscious visual experience. Though subjects were not experiencing a live percept of the disk, TMS clearly revealed that the visual cortex was not really in a baseline state. In neurons and/or networks previously activated by the disk, hyper-excitability lingers on well after the stimulus and conscious experience have passed.

The lingering brain states revealed by this experiment correspond easily with some known phenomena, but not others. The time scale is in line with psychophysically observed phenomena such as visual short-term memory and certain forms of visual priming. On the other hand, the long delays at which replay can still be induced does not seem to match the typical data from physiological recordings (such as VEP). It may be that the residual brain states in question are not fully reflected in active neural patterns, such as elevated spike rates. Instead, the residual brain states might lie in nascent neural patterns, such as channel or

network excitability. These kinds of states would not be reflected in passive neural recording techniques, but would require active stimulation to be revealed.

The previous sections established that the perceptual effects evoked by TMS are highly dependent on both fixed, anatomically-determined contributions to excitability, and dynamic, activity-dependent contributions. The instant replay effect indicates that the dynamic contributions include not only the induced activity resulting from concurrently presented visual stimuli, but also organized residual excitability states which linger on long after the presentation of visual stimuli.

Thus it seems that TMS can indeed be used to probe for the presence of neural activity patterns that are not normally available to visual consciousness. The visual information hidden within these brain states reflects not only the processing of currently presented stimuli, but also contains the residual traces of activity patterns from the recent past.

Chapter 3: Object Color Perception – The binding problem

Background

Neurophysiological evidence indicates that visual processing segregates input into diverging streams leading to specialized brain areas which code for different stimulus features. In area MT, cells respond selectively for different directions of motion. In V4, the same is true for color, in IT the same is true for shape, etc. (Livingstone & Hubel 1988). Unlike early visual cortex, which contains cells responsive to very small regions of the visual field, these higher level areas have only a coarse mapping to very broad regions of space.

This leads to a question of object correspondence. In the complex scenes we regularly see, there are often multiple objects right next to each other possessing different shapes, colors, and/or directions of motion. Since the higher-order cells encoding these features specify rather large regions of the visual field, it is unclear how proper correspondence is maintained between all the different features of each object (von der Malsburg 1981; von der Malsburg 1995).

This issue is referred to as the "binding problem." The title can be a bit misleading, as the situation does not really seem to be much of a problem for the working brain. It sorts through the information effortlessly, robustly, and accurately. Though many researchers investigate the binding mechanism by inducing disruptions in binding, it has proven quite

difficult to induce clearly perceived instances of "misbinding" or "illusory misconjunction" of basic visual features.

Past efforts have used transient or rapidly changing stimuli, and/or the diversion of attention with difficult secondary tasks (Wolfe & Cave 1999). Under these conditions, it is possible to infer misbinding events from error patterns in observer reports. In an early study, researchers flashed colored letters to observers and found that patterns in their responses reflected occasional swapping or migration of colors between the letters (Treisman & Schmidt 1982).

The difficulty with data from these types of experiments is that they are confounded with issues of performance noise, expectation, task strategy (Butler, et al. 1991), and memory (Tsal 1989). Subjects, when faced with an uncertain percept, may produce a pattern of guesses that mimic misbinding events when they're actually just making logical inferences from limited information (Donk 1999; Donk 2001). Even as a first-hand observer, it is difficult to verify for oneself whether the mistake was due to a true misbinding event or a vague, incorrect memory.

Given the difficulties in showing a clear case of misbinding, some have argued that the distributed nature of the neural activity does not really pose a problem for the brain to actively solve. However widely activity is distributed, there is no theoretical necessity for any additional mechanism to bind it all together (Garson 2001). Speaking purely symbolically, a pattern of activity across any number of elements could directly code for the

unified percept of the object, just as pieces of a computer file can be physically scattered across a hard drive. Furthermore, others have proposed that unified objects can be specified by an ordinary feed-forward system of intermediate detectors for simple sets of features (Fukushima 1980; Perrett & Oram 1993; Wallis & Rolls 1997). Under these schema, the solution to the binding problem is not some additional neural mechanism. The feed-forward activation of higher order neurons is sufficient to specify unified objects. This direct correspondence would explain why it has not been possible to break the binding mechanism – no such mechanism exists.

The experiments in this section, by inducing events of unbinding and misbinding, provide positive evidence for an active binding mechanism, thus arguing against a direct feedforward, symbolic solution to the binding problem. TMS induces the piecemeal activation of isolated stimulus features, as well as several misbinding phenomena. Misbinding can also occur in the absence of TMS, caused by a stable, unchanging visual stimulus which creates a sustained, verifiable misbinding phenomena. The properties of this illusion make it suitable for neurophysiological investigation, making it possible to search for neural correlates of feature binding.

Experiment 1: TMS probe of Cai's asynchronous binding illusion

Having demonstrated the properties of TMS as a probe into visual activity patterns in cortex, this tool is now applied to a visual illusion related to the binding problem. This experiment probes the mechanisms that lie behind the illusion, and also provides further insight into the processes that lie behind TMS-induced instant replay. Both the illusion mechanisms and the general properties of TMS-induced activity are found to have direct bearing on the binding problem.

Visual illusions have long been used to deduce the way perceptual computation is organized. Here we combine the viewing of an illusion with TMS-induced "instant replay" to gain further insight into the processing that occurs in visual cortex. Given that replay reflects recent neural history, it would be informative to see what is replayed when the preceding stimulus and percept do not match. What kind of visual information is contained in the stimulated networks?

In Cai's "asynchronous binding" illusion (Cai & Schlag 2001), a sudden color change appears spatially displaced and at a different size, as if bound onto the wrong object (figure 3.1.1). In this stimulus, a bar moves across the screen, steadily changing in length. The following description uses the example of a growing green bar moving from left to right. In one frame of the animation, the color of the bar is changed to red. Observers viewing this stimulus perceive the brief color change erroneously. The color change is seen farther to the



Figure 3.1.1: Cai's asynchronous binding illusion

In this illusion, the position and size of a color change is misperceived. The stimulus shown here is a green bar moving from left to right and also growing in length. In one frame of the animation, the color is changed to red. Observers do not perceive a color change at that location, instead they perceive the color change to happen in the bar after it has moved and grown further. The redness is thus seen too far to the right and in too large a size.
right, and in a correspondingly larger size. The color change thus appears to have happened a little later in the animation sequence than it actually did, after the bar had moved and grown a little further.

Cai proposed that this was a particular twist in the heavily studied class of "flash-lag" illusions in which sudden events (flashes) are seen to lag behind moving objects (Krekelberg & Lappe 2001; Nijhawan 2002). Many have argued that this effect exists because predictable stimuli, such as moving objects, are processed by the visual system more quickly than are sudden stimuli, such as flashes. Thus, though a flash and moving object may be aligned in space and time, the differential processing delays cause the flash to be seen later relative to the moving object – the flash "lags."

The everyday ecological relevance of this phenomenon lies in the fact that we have to interact with moving objects in the world effectively, despite inherent delays in visual processing. In order to minimize neural conduction delays whenever possible, our visual system may facilitate activity in neurons coding for an object's predicted path of motion. Since no such prediction is possible in the case of a sudden event, the processing of flashes is at a disadvantage.

The added twist in Cai's illusion is that the predictable and sudden elements, rather than being from separate objects, are properties embedded in a single object. This brings the proposed differential delay mechanism within the framework of visual feature binding. Cai proposes that differential processing delay between sudden feature changes and predictable feature changes might create an error in binding. Since the sudden color change of the bar is processed more slowly than the predictably moving position and changing size of the bar, their temporal alignment becomes shifted. When these features are reintegrated into a coherent percept, or bound together, the color change becomes bound with a position and bar size that corresponds to a later point in time.

This "asynchronous binding" explanation for this illusion is controversial (Eagleman 2003), as is the differential delay explanation for the broader class of flash-lag illusions. In Cai's particular configuration, it has been argued that the illusion could be explained by an effect termed "flash-drag." It has been demonstrated that motion fields have the effect of "dragging" the perceived position of flashed objects in their vicinity. For example, a spinning wheel will cause nearby flashes to be perceived in a location shifted further forward in the wheel's rotational direction (Whitney & Cavanagh 2000). In Cai's configuration, the motion field created by the moving and expanding bar could both drag and expand the color flash.

What contribution can the use of TMS-induced instant replay make in explaining this illusion? If TMS is delivered shortly after the color change event, what kind of neural history will the replay percept reflect? The display has resulted in a visual percept that is decidedly different from the physical stimulus, so it is unclear what kind of visual information might be encoded in the networks stimulated.

The stimulus configuration used with TMS was modified from Cai's version to be more compatible with the phosphene location, to optimize the robustness and effect size, and to be easier to report accurately. The stimulus was a rotating clock hand, which changed color for a single frame (figure 3.1.2a). Though a variety of colors, velocities and positions of the color flash were used, the descriptions in this section assume a green clock hand rotating clockwise, which briefly changes to red at the four o'clock position. Observers perceived the color change to occur at an erroneous position, further forward in the motion (say, five o'clock). The redness was not only seen at a shifted position, but at the appropriate rotation and as part of the motion of the clock hand. This configuration robustly recapitulated Cai's illusion (N = 5, mean effect size 13 degrees, t-test p < 0.05).

Subjects viewed this illusion and received TMS shortly after the color change, such that they would see the illusion normally, and then see the TMS-induced percept afterwards. They made two types of reports. The first was of the position at which they saw the clock hand change to red (the color change prior to TMS). The second type was of the position at which they saw a red bar evoked by TMS. Detailed methods can be found in Appendix A.

The results from the first report reflected the typical illusory percept, the position being shifted forward in the rotation direction (N = 4, p < 0.05. Sample subject data in figure 3.1.1a and b, blue dotted line). This confirmed that the illusion itself persisted under the TMS paradigm.



(b) TMS effect type 1

In some trials, TMS replay caused the subject to perceive a red flash the behind the spot in which they had previously seen the hand turn red. This new location turns out to correspond to the true location of the flash. (c) TMS effect type 2 In other trials, TMS replay caused the subject to perceive the clock hand (which has been continuing to move on the screen) turning red again.



Figure 3.1.3: Sample subject data

Subject reports for perceived position of color change in the visual display, and for perceived position of TMS-induced "replay". Plots are of error size vs. physical position. Horizontal black line denotes zero error. Dotted blue line plots the reports for the visual display. Solid red line plots the reports of TMS-induced replay.

(a) Data from subject who consistently saw a type-1 effect (Fig. 3.1.2b).
Subject shows significant effects of the illusion (dotted blue line, t-test, p<.05).
Position of TMS replay corresponds to the physical position (solid red line, p>>.05).

(b) Data from subject who consistently saw a type-2 effect (Fig. 3.1.2c). Subject shows significant effects of the illusion (dotted blue line, p<.05). Reports are farther forward in position (solid red line, two-sample t-test, p<.05). The data from the second report, of the TMS-induced percept, reflected strong inter-subject variability. Some subjects saw nothing at all. Others, whose phosphenes were located peripherally, did not feel confident in reporting the position of a retrieved flash. The following represents the results from four out of the eight subjects, who saw a retrieved flash in the majority of trials and were confident in its position.

The second report fell into two diverging classes. Some subjects saw one of these classes of percepts exclusively, others varied from trial to trial. In the first class of reports, subjects reported that TMS caused them to see an additional red clock hand (Figure 3.1.2b). Compared to the location where they perceived the original color change, the replayed flash was seen farther back. This reported position was found to correspond with the actual location of the stimulus color change (N = 3, p >> 0.05, sample subject data in figure 3.1.3a, solid red line). In the second class of trials, the subjects reported that TMS caused the clock hand to turn red again as it continued to rotate (figure 3.1.2c). Accordingly, these reports were of locations further forward than both the physical and perceived locations of the initial color change (N = 3, two-sample t-test p < 0.05, sample subject data in figure 3.1.3b, solid red line).

The first class of trials offers insight into the mechanisms behind Cai's illusion. The second class of trials offers insight into the topic of feature binding, and what kind of visual activity is evoked during instant replay.

In the first class of trials, the fact that TMS elicited a flash in the position, orientation and color corresponding to the actual color change indicates that the visual system does indeed register the correct information at some level. TMS is interacting with residual brain states left over from the processing of the preceding event, and those states apparently encode the correct position of the color change.

At first glance, this result seems more consistent with the "flash-drag" account of the illusion. At the time of the color change, there is a motion field formed by the sweep of the clock hand. However, at the time of TMS, the clock hand has already passed, and so the reactivation of activity associated with the color change is no longer distorted by motion signals.

Unlike the effects found in earlier sections, TMS did not simply cause a recapitulation of visual cortex activity. Subjects saw a new event in response to TMS. Using the above explanation, it appears that TMS re-activated the color and raw position of the color change, but did not re-activate the associated motion signal. At the very least, TMS induced an activity pattern which was categorically different from the activity which underlay the initial illusory percept. It seems that TMS is able to selectively isolate and re-activate specific components of the visual processing system.

Keeping this in mind, we may not be able to entirely rule out the asynchronous binding account of the illusion. An alternative account would be that there are multiple representations of the color change's location, and that TMS interacts with only a subset of them. During the viewing of the stimulus, perception reflects the dominance of a representation of the illusory location, which is formed by a higher level binding processes. Afterwards, though, TMS might preferentially activate the representation of the actual location, perhaps because it has a slower decay time, or because its encoding scheme is somehow more suitable for activation by TMS.

In the second class of trials, subjects report that TMS induces an event at a position that corresponds neither with their prior percept nor with the physical location of the color change. Instead, TMS causes them to see the clock hand turn red again, as it continues to rotate. This is both quantitatively and qualitatively different from the first class of trials in which TMS induced an isolated flash *behind* the reported location of the original color change.

The induced percept contains redness in a location and orientation where it has no precedent, instead reflecting an amalgam of visual features from different events and times. It includes the color that was seen earlier, and it is combined with the new location and orientation of the rotating clock hand. This is an occurrence of misbinding.

One way to interpret this phenomenon is that TMS has successfully activated the color component of the earlier event, but not the other feature components. The color is then bound to a location and orientation which visual cortex is coding due to concurrent input from the ongoing visual display, namely, the position of the moving bar.

Viewing the two trial types separately thus far, two rather divergent interpretations have been given. For the first trial type, TMS is described as activating a compact representation of the actual stimulus, a representation containing a specific position, orientation and color. This is consistent with the organization of low-level, retinotopically-mapped visual cortex. For the second trial type, TMS is described as activating a partial representation containing only the stimulus color, whose position and orientation is not defined. This is consistent with a higher-level area such as V4, which processes color information, and has very coarsely defined retinotopy.

Viewing the two trial types in light of each other, a more unified explanation seems possible. Isolated activity in any single cortical area is unlikely to be sufficient to create a conscious visual percept. Thus, if TMS activates multiple scattered visual representations throughout visual cortex, the final percept probably depends on the propagation of that initial activity to form a global computational resolution of those components.

At the moment of TMS, there is a full representation of the still-present rotating bar, and a residual representation of the color change. TMS interacts with all of these, and from trial to

trial the balance of interaction may shift. In the first trial type, the orientation and location of the physical color change is sufficiently activated so as to be integrated into the evoked visual percept. In the second trial type, the orientation and location of the rotating clock hand dominates.

The use of TMS in conjunction with the asynchronous binding illusion has shed some light onto the mechanism behind the illusion, has caused some misbinding of its own, and revealed some of the complexity behind the instant replay effect. The results indicate that TMS does not simply recapitulate patterns of visual cortex activity, but that it interacts with some visual representations much more than others. This allows TMS to be used as a technique to break apart the components of visual processing, examining the information encoded in each piece more directly.

Experiment 2: TMS unbinds the visual features of a color grating

The combination of TMS and the asynchronous binding illusion produced effects where it seemed the different stimulus features of objects were activated independently and then reassigned. However, the complexity of the paradigm makes it difficult to completely isolate these events. Here, TMS is combined with a much simpler stimulus, a flashed color grating, in a search of simple and direct evidence for the independent activation of features and misbinding.

The following observations were made as a part of a larger study in which we assessed the effects of grating orientation and TMS delay on the brightness, saturation and spatial distribution of the TMS-induced replay. The subjective descriptions of the percepts, which are relavent to the binding problem, are described here. In this experiment, a large colored grating was flashed on the screen, followed by TMS at various timings. Subjects then drew regions in which replay was observed and adjusted the on-screen brightness and color saturation of a simulated percept. Although various stimulus parameters were used, the following description will assume a red diagonal grating. Detailed methods may be found in Appendix A.

Subjects saw replayed patches of red diagonal grating, solid red, and achromatic diagonal gratings. Often, they would see many of these elements arrayed in segregated patches of space (figure 3.2).



Figure 3.2: TMS-induced visual feature decomposition

(a) Stimulus sequence. A colored grating is flashed on the screen, followed by a dealy period of up to two seconds. TMS is then administered.

(b) Sample stimlus - A red grating with an orientation of 45 degrees

(c) Sample percept - from trial to trial, sujects saw regions of achromatic grating, red grating and/or solid red. The various percepts could be seen within a individual trial in different patches of visual space.

Thus, even given a simple flashed stimulus, the components of color and orientation could be broken down and independently activated. The percept reflects an uneven activation pattern across the visual field in which color can be strongly or weakly activated, as can be orientation or texture. The strength of activation for these two visual features were not tightly linked together. Even within a single trial, certain regions would reflect strong activation of color without orientation, and other regions would reflect strong activation of orientation without color.

This experiment clearly demonstrates that TMS can independently activate different visual features. This provides further confirmation of TMS' ability to isolate and activate components of visual processing. This also reflects an organization of the visual system in which color and orientation are separable elements. Thus, this provides further evidence that there is indeed an active binding process, which can be disrupted.

Experiment 3: Inducing misbinding without TMS: steady-state misbinding of color and motion

The previous experiments have shown several instances where TMS induces the independent activation of particular stimulus features, as well as the misbinding of those features. The following experiment demonstrates that misbinding phenomena can also occur in the absence of TMS. Here, a steadily presented, unchanging visual stimulus produces a long-lasting, verifiable instance of misbinding, suitable for use in neurophysiological experiments. [The following is an edited and expanded version of Wu et al. 2004.]

Here we report examples of a novel type of illusory feature misconjunction in which incorrectly bound perceptions persist during the continuous viewing of stable stimuli. Colormotion misconjunctions are induced by transparent motion stimuli (figure 3.3a) consisting of two sheets of random dots, where one sheet is moving up and one is moving down. The sheets contain dots of two colors such that the central and peripheral portions of the stimuli combine color and motion in opposite fashions (figure 3.3a). On the upward-moving sheet, dots in the center are red and dots in the periphery are green. On the downward-moving sheet, dots in the center are green and dots in the periphery are red.

Observers gazing at the center of the display perceive dots in the peripheral areas erroneously: they bind color and motion in the wrong combination. The entire display therefore appears to be covered by a sheet of red dots that are moving upwards and a sheet of



Figure 3.3: Steady-state misbinding of color and motion

(a) Sample transparent motion stimulus.

Two sets of random dots move rigidly in opposite directions. "Real" circles show details of the physical stimulus. In the upward-moving surface, dots near fixation are red, and dots beyond 60° of horizontal eccentricity are green; in the downward-moving surface, the reverse coloration applies. White bars mark the transition line. "Perceived" circles show details of the illusory percept. Observers incorrectly pair colour and motion in the periphery. Across the entire stimulus, upward-moving dots appear red and downward-moving dots appear green, forming two homogeneous surfaces of dots moving in opposite directions.

(b) Observer reports.

Proportion of trials in which observers (N=5, 3 naive) judged that most red dots beyond the white bars were moving upwards, plotted as a function of the actual percentage of red dots moving upwards; time to view stimuli was unlimited. In control trials (black line), the central region was blank, and responses followed the stimulus. When the central region contained upward-moving red dots and downward-moving green dots (red line), responses shifted predominantly to "up". When the central region contained the opposite motion (green line), responses shifted predominantly to "down" (MANOVA, F(10,16) = 35.2, p < 0.001). green dots that are moving downwards. The percept is dependent on the stimulus properties at the site of fixation. When observers shift their gaze to one of the peripheral regions, the color-motion relationship for the entire stimulus appears to reverse.

To quantify this effect, observers were asked to report whether the majority of peripheral red dots were moving up or down in a series of trials in which we varied the percentage of peripheral dots moving in either direction. In control trials, where the central portion of the stimulus was omitted, responses of observers (n = 5, 3 naive) roughly followed the physical stimulus (figure 3.3b, black line). But when the central portion was present, responses shifted markedly, and in the same direction as the central dots (figure 3.3b, red and green lines; MANOVA, F(10,16) = 35.2, p < 0.001).

To ensure that observers' responses reflected their perception and not a simple bias to respond to the central dots, we generated stimuli in which one peripheral region matched the center and the other contained the opposite motion. Observers (n = 4; 3 naive; all from the first experiment) perceived both sides to move with the center, performing poorly in reporting which side moved the other way (mean \pm s.e.m.: 47 \pm 8% correct; d' = -0.12 \pm 0.38). On the other hand, they performed well in reporting which side was which when the stimulus center was blank (86 \pm 2.5% correct; d' = 1.5 \pm 0.2, paired t-test p < 0.05)

We propose that the misbinding percept is driven by a foveal bias in the perceptual processing of broad surfaces. The control trials demonstrated that observers' peripheral

vision was sufficient to quite reliably judge the correct feature conjunctions in the periphery. However, because color signals and visual resolution are strongest in the fovea, it would normally be advantageous to use foveal information to resolve peripheral ambiguity if the physical source is likely to be a common, uniform surface. Uniform surfaces are strongly implied by the contiguous and precisely equal values of color and motion present across this stimulus. Therefore, conjunction information from the periphery is overridden, and peripheral features are bound in such a way as to accord with the usually more reliable center. This is consistent with the computational view of the brain as a Bayesian ambiguityresolving apparatus, which takes into account prior knowledge, such as statistical likelihood of inputs and the reliability of information channels, in order to arrive at a solution.

Previous experiments aimed at inducing the direct experience of illusory misconjunctions, or the erroneous binding of one object's features onto another object, disrupted visual attention by using stimuli with brief presentation times, rapidly changing features and extraneous distracting tasks. The misbinding phenomena produced were transient, ambiguous and sporadic. In those experiments, individual instances of binding failure were not separable from other sources of performance noise except in the aggregate statistics, and were subject to artifacts from memory, expectation, task strategy, and response strategy (Wolfe & Cave 1999).

The misbinding demonstration presented here avoids those confounds by allowing perceptual processing to reach a steady-state resolution during fully attentive, extended inspection.

Subjects were allowed to view the stimuli as long as they wished and made their judgments during live, active inspection. This is in contrast to previous experiments, in which subjects made judgments and responded based on their memory of a fleeting percept.

In combination with the TMS-induced occurrences of unbinding and induced misbinding, this experiment provides firm evidence that there exist active binding mechanisms beyond a direct symbolic correspondence. There is indeed a binding problem for the brain to solve.

Possible models for the binding process are constrained by the attentive conditions under which this illusion is induced. One theoretical framework used in the study of binding and visual search holds that it is the direction of visual attention that serves the central role of binding features together (Treisman & Gelade 1980). This was consistent with previous studies in misbinding, which prevented subjects from fully attending to the features of an object. However, the illusion presented here shows that feature conjunctions are fragile even under fully attentive states.

Because this illusion provides a directly identifiable misbinding event, it opens the binding problem and its mechanisms to neurophysiological investigation. Previous studies relied on statistical analysis to infer that a certain fraction of trials contained misbinding, and so could not definitively identify specific trials as either misbinding or generic mistakes. Here, the robust and verifiable nature of this illusion allows for occurrences of misbinding to be

individually identified and correlated to neural recordings. Furthermore, the extended duration of the effect brings it into a timescale suitable for fMRI studies.

Chapter 4: Surface Color Perception—Color filling-in

Feature binding leads to the association of a color to an object. This chapter deals with another problem in object color perception: how color becomes extended across the entire surface of the object. [The following is an edited version of a paper currently under review. The co-authors on this paper are Ryota Kanai, Frans Verstraten and Shinsuke Shimojo.]

Background

The representation of uniform color surfaces in the activity of early visual cortex is decidedly different from the intuitive assumption of isomorphic coding (such as what one might find in the CCD array of a digital camera). When a patch of uniform color is presented, little or no activity is evoked in neurons whose receptive fields lie directly on the patch's surface. Instead, vigorous activity is found in neurons whose receptive fields lie on the edges of the patch (Hubel & Wiesel 1968; Friedman, et al. 2003). In a corresponding set of results from psychophysical studies, the perceived color and/or brightness of a uniform surface has been found to depend on signals from the edges of the surface, rather than from its interior (Cornsweet 1970; Redies & Spillmann 1981; Montag 1997; Pinna, et al. 2001). How is this edge-based encoding scheme "decoded" into our perception of filled surfaces? The mechanisms whereby the edge signals inform the perception of the surface itself have been elucidated by the study of a broad class of visual phenomena, which can be described as "filling-in" effects.

Color filling-in, for example, is a visual phenomenon where the perceived color of a visual region is dictated, or filled-in, by the surrounding color (Pessoa, et al. 1998). This process is responsible for the invisibility of the blind spot (Ramachandran 1992), or other regions that lack bottom-up signals due to pathological scotomas (Bender & Teuber 1946; Sergent 1988; Zur & Ullman 2003). Filling-in is also observed in an extensively characterized process known as Troxler fading. After prolonged fixative viewing of a scene, small objects disappear from view, their locations "filled-in" by the surrounding color(Troxler 1804; Yarbus 1967; Gerrits, et al. 1966; Riggs, et al. 1953; Cornsweet 1970; Ramachandran & Gregory 1991; De Weerd, et al. 1998).

One characteristic found consistently throughout all known filling-in effects is that they are disrupted by luminance edges (Paradiso & Nakayama 1991; Redies, et al. 1984; Cole, et al. 1990; Eskew & Boynton 1987; Boynton, et al. 1977; Montag 1997; Cornsweet 1970; Todorovic 1987; Kingdom & Moulden 1988). The interposition of a luminance edge (e.g. a dark gap) between two colored regions blocks filling-in from one side to the other. Only after prolonged adaptation, whereby the luminance edge itself fades from awareness, can filling-in proceed.

This characteristic has been foundational to classical mathematical and neuromorphic models of visual processing (reviewed in Pessoa, et al. 1998; Cohen & Grossberg 1984). Under these interpretations, color and luminance information diffuses along the retinotopic map, starting and stopping at borders defined by local features such as high-contrast luminance transitions. Through this process, information encoded at surface edges comes to fill enclosed shapes. This decodes the edge-based scheme found in early visual cortex, resulting in the perception of uniform surfaces. This notion that sharp luminance transitions define surface borders has been used to explain the illusory spread of colors across areas lacking strong edge signals, why dark luminance gaps enhance the contrast of colors on either side, and why they block color-filling.

In this paper, we report a color-filling effect that defies the above conceptualization. While viewing the artwork of Julian Stanczak (Stanczak & McCleland 1998), we noticed strong phenomenological effects whereby color-filling jumped over clearly-defined luminance gaps. Unlike the classical demonstrations, which require that luminance gaps adapt and fade from conscious perception, this form of color-filling occurred without the adaptation of the gaps, leaving the luminance profile perceptually intact.

While Stanczak's art vividly shows counter-examples to the classical model, it is difficult to discern why a completely different type of color filling occurs in these situations. Most of the pieces contained a number of visual elements and effects, and used colors chosen for their tendency to dazzle or "vibrate" in the viewer's eye. We set out to verify that the effect could occur with the simple color gradients used in classical filling-in studies and a single set of strong, clearly presented luminance gaps, so that the effect could be parametrically characterized and investigated further.

Our simplified stimulus is depicted in figure 4.1.1a. A broad color gradient from red to green is displayed within an array of squares set against a black background. This stimulus successfully reproduced the discrete color-filling effect. When the center of this stimulus is gazed at steadily, the foveal color (red) appears to fill the peripheral squares, overwriting the original color (green). Notably, the black grid pattern remains perceptually prominent throughout the entire process. Cursory examination of a number of stimulus configurations revealed that this effect survives a wide range of gap and patch sizes, and does not require co-linearity of edges.

In a series of experiments, we show that filling-in occurs in a fashion highly selective to perceptual surfaces. Our study extends the classical view, and unites it with a wide literature showing the importance of surface parsing in other basic aspects of vision. We argue that the spreading and interaction of color information is governed not by the mere presence or absence of local, retinotopically defined features such as luminance contrast, but by their ability to affect the parsing of perceptual surfaces that are defined in a visually holistic context. These parsed surfaces likely have separate representations, because we find that multiple color-filling processes can proceed on retinotopically overlaid surfaces, in ways which indicate that each process is mediated independently of the other.



Figure 4.1.1: Basic filling-out stimuli and trial sequence

(a) Typical color gradient stimulus with a grid pattern

(b) The corresponding control stimulus, the same gradient without the grid. This configuration is known from classical studies to cause Troxler fading.

(c) Color profile of the stimuli. The percentage of a central color is defined by a cumulative Gaussian as a function of visual eccentricity. When the central color is red, 0% means 100% green.

(d) Trial sequence. A trial consisted of the observation of test stimulus until reponse or 30 second timeout. Particpants responded when the entire area of the stimulus became perceptually homogeneous. Inter-trial effects were minimized by presenting a color gradient of opposite polarity, and then a dynamic noise pattern, each for half the duration of the trial stimulus.

Experiment 1: Color jumps over gaps, filling into distant areas

In our first experiment, we created two sets of basic stimuli using elements described in the introduction. The first set of stimuli consisted of simple color gradients, taken from configurations known to induce classical Troxler fading (figure 4.1.1b). The color gradient was defined as a function of the visual eccentricity from a fixation marker and had a cumulative Gaussian profile as shown in figure 4.1.1c. These stimuli had differing values for the location parameter μ , which corresponds to the size of the central disk. The second set of stimuli was identical to the first, except for the superimposition of a dark grid pattern over each stimulus (figure 4.1.1a). In each trial, one stimulus was presented with a small fixation point at the center. When observers perceived the filling-in process to be complete, that is, when the stimulus area became perceptually homogeneous in color, they pressed one of three keys. This response indicated whether (1) the central color had filled into the peripheral squares, (2) the peripheral color had filled into the central squares, or (3) the direction of filling was ambiguous. After the response, observers viewed two types of counteradaptation stimuli to minimize cross-trial effects (figure 4.1.1d). Trials in which observers did not perceive a homogeneous color within 30 seconds were aborted.

Two measures were made in order to characterize the phenomenon. First, the time at which a keypress occurred (Time To Fill, or TTF) was used as a measure of how readily color-filling occurs in a given stimulus condition. Second, the direction of the filling-in was analyzed. These measures were used to characterize the effects of the size of the central color disk and the effects of superimposing the grid pattern.

The TTF results are shown in figure 4.1.2a & b. In the classical conditions without the grid pattern, the filling-in took 15 seconds or longer on average. Changing the size of the disk did not significantly change the TTF (figure 4.1.2a, dashed curve; F(2,18) = 0.701, P = 0.509). However, superimposing the grid significantly sped up TTFs overall (figure 4.1.2a solid curve vs. dashed curve; F(1,18) = 24.11, P < 0.001), by 3-7 seconds on average.

The direction of filling-in had an expected relationship to the size of the central disk. There was an increased tendency for the foveal color to fill outwards with larger disk sizes (figure 4.1.2b, F(2,18) = 21.94, P < 0.001). The presence of the grid overlay also drastically affected the direction of color filling (figure 4.1.2b, solid curve vs. dashed curve, F(1,18) = 17.65, P < 0.001). When the grid was present, the direction of color-filling was biased outwards. This was particularly clear in the conditions at the smallest radius of the central color, where the presence of grid caused a complete reversal of filling direction.

The finding that color filled outwards was initially surprising, as classical studies reported only the filling-in of objects by the background color. This experiment, however, shows that even classical configurations cause an object's color to fill into the background, when a color spot of sufficient size is centrally fixated.

The tendency for the grid texture to change the direction of color-filling might be explained by the receptive field sizes present in the fovea and in the periphery. Whereas the spatial



Figure 4.1.2: Basic filling-out results

(a) The effects of luminance gaps on the direction of color-filling. Mean results of the Directionality Index is plotted as a function of central disk radius (μ). A positive Index indicates a tendency of the color to fill outwards; a negative Index, inwards. An index of one would indicate all observed color-filling events were in the outward direction. Open circles with a dotted line indicate the conditions with gaps, and solid circles with a solid line the conditions without gaps. All error bars represent standard error of the mean (N=4, 40 trials per condition). The presence of gaps markedly shifts the percept from filling inward to filling outwards (F(1,18) = 24.11, p < 0.001). Effect of central disk radius was not significant.

(b) The effect of luminance gaps on the speed of color-filling. Mean results of the time required for the completion of color-filling is plotted as a function of central disk radius (μ). The presence of luminance gaps markedly reduces the time to filling (F(1,18) = 17.65, P < 0.001).

resolution of the fovea is sufficient to separately encode the color and gap regions, this is not true for the periphery. There, receptive fields are large, and individual elements tend to sum across larger regions of space. The grid texture, then, dilutes the strength of the color signal much more in the periphery than in the fovea. This may increase the tendency for color of the central disk to override the color of the periphery.

We leave this topic of filling direction for future study. Rather than investigating the factors that determine which color wins the competition for perceptual supremacy, we focus on finding the factors that define the regions within which color competition is allowed to occur.

For this, the TTF results provide the central finding. It is a surprise that color-filling should occur at all in the stimuli with the grid texture, let alone in an expedited fashion. Why would luminance edges promote color-filling in these stimuli, when they blocked color-filling so effectively in classical studies? Our hypothesis is that the spatial limits of color-filling are not determined by luminance edges per se, but by the surface segregation cues that they often give rise to. In accordance with the process of surface parsing, color-filling proceeds across all regions which are interpreted as arising from the same visual surface.

In the classical color-filling experiments, isolated luminance edges were presented in alignment with the transitions between color regions (e.g. a single dark ring around the color disk of experiment 1). This was originally interpreted to mean that luminance edges themselves block color-filling. We propose that the blocking is not a result of the edges per se, but the fact that they provide a strong cue that the regions on either side arose from separate surfaces. Both interpretations can explain the classical findings, because the contribution of locally defined features such as luminance edges is experimentally confounded with the contribution of globally defined surfaces.

The stimuli presented here eliminate that confound, and indicate that global surface segregation is the key step in defining the limits of color-filling. Speaking in terms of local features such as luminance edges, the grid texture creates more boundaries, breaking up the visual field. But speaking in terms of perceptual surfaces, the grid is clearly seen as a superimposed surface, and thus it *unifies* the visual field. The grid provides strong cues that all of the colored elements arise from a common surface. Our surface-based hypothesis would explain why the color-filling occurs with the grid, and why it occurs so much more rapidly.

Experiment 2: Texture-defined edges block color-filling

The first experiment showed that the color-filling process passes through luminance edges that do not provide segmentation cues. The next experiment provides positive support for the role of surface segmentation by showing the converse – that filling in is blocked by non-luminance-based surface segmentation cues. Here, we observed filling-in on texture stimuli consisting of oriented bar elements (Landy & Bergen 1991; Grossberg & Mingolla 1985). In texture stimuli, we can embed non-luminance-based segmentation cues, using differences in orientations to form second-order borders (figure 4.2a). While keeping the color gradient constant, we varied the strength of the texture segmentation cues. We compared a condition lacking texture edges (figure 4.2a, *left panel*) with conditions containing texture edges of three different levels of sharpness (figure 4.2a, *other panels*) by measuring the TTF for each condition. According to our hypothesis, this kind of segregation cue should block the filling-in phenomenon, especially when the texture border is sharply defined.

The TTF results are shown in figure 4.2b. In the condition without texture edges, the fillingin occurred across the entire stimulus almost instantly. In trials with texture edges, especially the sharpest condition, subjects observed that the color would fill-out relatively quickly to the textural borders, stop there, and then require much longer to cross over. These subjective reports were reflected in the quantitative behavioral data. A repeated measures of ANOVA on the factor of the conditions showed that the TTF significantly differs among conditions (F(3,9) = 6.079, p = 0.005). A post hoc (Tukey) test showed that the TTFs in the "no edge"



Figure 4.2: Color filling-out on texture defined by oriented bars

(a) The four types of stimuli: No edge condition in which the texture was homogeneous, Shallow edge condition ($\sigma = 5.0^{\circ}$), Medium edge condition ($\sigma = 3.3^{\circ}$), and Sharp edge condition ($\sigma = 0^{\circ}$).

(b) Effects of textural borders on color-filling. Mean results of time to filling for the four conditions. Error bars represent standard error of the mean (N=4, 20 trials each condition). The percentages of trials in which a participant did not respond before the 30 second time limit are shown in the brackets, and are excluded from the analysis. TTF differs significantly among conditions (repeated measures ANOVA, F(3,9) = 6.079, p = 0.005. See text for post hoc pairwise comparisons). The completion of the filling-out process was inhibited by the textural borders, most strongly by the sharpest borders. This indicates that the borders of the color-filling processes are determined by surface segmentation, not luminance edges.

condition were significantly smaller than in the sharp and medium edge conditions (p < 0.01, p < 0.05), but not the shallow condition (p > 0.05). As in experiment one, color-filling was not blocked by the many luminance gaps. Instead, it was blocked by textural edges, resulting in a longer duration of viewing required to complete the filling-in process.

This fully disassociates the roles of luminance edges and surface segregation in determining the limits of the color-filling process. Luminance edges are neither necessary nor sufficient in blocking color-filling. Instead, surface segregation is the critical component in defining the borders of the colored areas.

Experiment 3: Overlaid color-filling processes are independent and selective

The previous experiment explains how color information is restricted in its retinotopic extent, but does not explain the "discreteness" of the effect. That is, color-filling not only proceeds beyond the black boundaries, but seems to "jump" over them. If the transmission of color information is meant to cope with situations of multiple overlying surfaces, it could be that there exist multiple surface representations over which color information can flow independently. On the other hand, this can also be explained within the framework of a single retinotopic representation, with some kind of suppressive mechanism that prevents the color from being fully expressed in the dark gaps. Some support for this may be seen in the phenomenon of color induction, where black areas are sometimes seen with a "tint" of a neighboring color.

To examine these possibilities, we created stimuli in which two color gradients were overlaid in the stimulus. If the substrate for the color-filling process is a single retinotopic representation, then we would expect that multiple color-filling processes might not be able to cross each other in a region of visual space, or that they would otherwise mix or interfere with each other. On the contrary, we find in the following experiment that these stimuli give rise to two independent processes, indicating that the visual representations over which color information is flowing encode each surface separately.

We created a series of stimuli in which one color gradient was presented in discrete squares as in the first experiment, but now a second gradient was presented in the area previously occupied by the black grid (figure 4.3a). We varied the size of the central disk in the grid area, but held constant the gradient within the squares. In each trial, subjects made two keypresses. With the left hand, they pressed a key when the color gradient in the squares became homogeneous. With the right hand, they did the same task for the other gradient. Keypresses indicated whether the direction of the color-filling was foveal-outwards or peripheral-inwards. Data analysis was based on responses given within a 30 second time limit.

Both the qualitative and quantitative aspects of observers' responses indicated that two separate processes of color-filling occurred independently of each other across the overlaid gradients. Observers reported that color-filling occurred in both the array of squares and the grid, but they remained distinct and separate. For example, given the stimulus shown in figure 4.3a, an observer might see all the squares take on their foveal blue color and see the entire grid take on its peripheral red color (illustrated in figure 4.3b). The colors of the squares do not fill-in to areas occupied by the grid, or vice versa.

Quantitative analysis of the data also showed that two filling-in processes were acting separately and independently. First, the directions of filling for the two gradients were reported to be opposite in many trials (illustrated in Figs. 4.3a and 4.3b, data in figure 4.3c).



Figure 4.3: Color-filling over two overlaid color gradients

(a) Sample stimulus. Configuration is a modified version of stimuli in experiment 1, where the gaps between squares contain a second color gradient.

(b) Illustration of an eventual percept. Participants perceived color-filling to occur on both surfaces. Component colors remained restricted to their own surfaces. Here, the blue color of the central squares has filled out to the peripheral squares, and the red color of the peripheral grid has filled into the central grid.

(c) Participants often perceived color to spread in opposite directions in the grid and the squares. The percentage of trials in which filling proceeded in opposite directions is graphed as a function of the grid gradient's disk size.

(d) Color-filling occurred at independent rates on the two surfaces. The TTFs for the grid and squares are plotted against each other trial by trial. On average, responses for the grid and squares were separated by 5.7 seconds in either direction. Correlation between the two is poor: R² is 0.18, or 0.12 after correcting for inter-subject contributions. This indicates that color-filling processes are occurring separately on each perceptual surface.

The highest likelihood of opposite directions of filling (51%) occurred in the trials where the gradient in the grid pattern was set to the smallest disk radius.

Analyzing the TTF results trial by trial, we found that the timing of the two responses also reflected separate processes. Figure 4.3d is a trial-by-trial plot of the TTF for the squares vs. the TTF for the grid. Subjects usually reported completion of filling-in at different times for the grid and the squares, on average 5.7 seconds apart. Overall, the correlation between completion times for the two processes is extremely weak. In the raw data plot of figure 4.3d, $R^2 = 0.18$. Much of this is due to the contribution of systematic differences between subjects. When data is scaled to each subject's mean, $R^2 = 0.12$.

These spatial and temporal results indicate both qualitatively and quantitatively that visual processing segregates the input into perceptual surfaces and creates separate representations for each surface. It is within each of these representations that the color-filling process occurs. Thus, the flow of information is not determined by simple retinotopic adjacency or local borders, but by common surface origin.
Experiment 4: Color-filling on perceptual surfaces in transparent motion

Our final experiment tests the limits of this notion that color-filling is occurring across multiple perceptual surface representations. In the following stimuli, two surfaces are transparently overlaid, both occupying the entire display, so that every point in retinotopic space has an equal correspondence with either surface. We saw whether two color-filling processes could still remain separate, restricted to elements of their own surface.

We used a transparent motion stimulus (see figure 4.4a), in which multiple surfaces coexist at the same retinal position (Snowden & Verstraten 1999). Two surfaces consisting of random dots moved in opposite directions, each having a different color gradient from center to periphery. On one surface (say, moving upwards), the color of individual dots gradually shifted from red to green as the eccentricity increased, while on the other surface (moving downwards), the color of the dots gradually shifted from red to green.

We observed that after prolonged viewing, the color filling-in to peripheral dots occurred selectively for each motion-defined surface. Figure 4.4b illustrates the eventual percept for the stimulus diagrammed in figure 4.4a. After prolonged viewing, the color of the surface moving upwards appeared to become homogeneously red, whereas the color on the other surface appeared homogeneously green.



Figure 4.4: Selective color filling-out for transparently overlaid surfaces

(a) The transparent motion stimulus was created by superimposing two surfaces of random dots, each defined by common motion. On one surface, dot colors had a color gradient from red (center) to green (periphery), whereas on the other surface the opposing color gradient was used.

(b) After prolonged viewing of this stimulus, it was perceived as two homogeneously colored surfaces, the central color of each surface having filled-out to the peripheral elements sharing the same motion.

(c) Time required for completion of color-filling for single- and dual-surface stimuli of varying gradient steepness. (N=4, 16 trials per condition)

To parametrically characterize the effect, we created two sets of stimuli. The first set were stimuli constructed as above, with varying steepness of the color gradient (sigma). The second set consisted of single surfaces, in which only half the dots (say, the ones moving upwards) were displayed. In each trial observers made one or two keypresses to indicate when each of the moving surfaces appeared homogeneous in color.

As expected, response times increased as the steepness of the color gradient increased. Most importantly, no significant differences were found in the response times when comparing dual surface to single-surface stimuli of the same sigma settings (figure 4.4c).

The cases in which the color change was abrupt (sigma = 0), proved unique, and shed some light on the relationship of the filling-in effect with the steady-state misbinding found in chapter 3, experiment 3. Observer's responses to the single-surface stimuli of this type were as expected: response times were significantly longer than in the previous cases of smooth color change. Filling-in took as long as 2 minutes to initiate. However, when both surfaces were presented, the stimuli were immediately perceived as having homogeneous colors. This stimulus configuration is the same as the steady-state misbinding stimulus.

With shallow gradients, there arise yellow areas at intermediate eccentricity. In these regions, the color of green or red has to be created from the adjutant color signals through the fillingin process. On the other hand, with a steep color change, there is no such intermediate color area. In such a case, the color signals for both red and green are already available from the bottom-up input, and this can be exploited to create a percept of homogeneous surface color.

Thus, there seem to be two mechanisms to that lead to a homogenization of surface color, filling-in and misbinding. Our results indicate that they occur at different time scales. In the trials with smooth color gradient stimuli, the dissociation can be seen directly, as color misbinding can be seen in the extreme periphery before the color filling occurs into the yellow of the intermediate areas.

This provides further support for the idea that color-filling proceeds along multiple perceptual surface representations, rather than across a single retinotopic map. First, a colorfilling process occurring across a surface took the same amount of time whether or not a second surface was overlaid. Second, in these stimuli, any given point in retinotopic space would at times be occupied by elements of either or both surfaces. Yet the two color-filling processes remained separate and restricted to elements of their own perceptual surface.

Discussion

Taken as a whole, our experimental results show that the substrate for color-filling is not the classically conceived retinotopic map subdivided by locally-defined edges. Instead, surface segregation creates separate representations over which multiple color-filling processes can occur. The color interactions to which a particular region of visual space is subject, then, is determined by its membership in these surfaces. The processing of a region's color will be conducted in the context of the color information from other regions determined to arise from a common surface, even if they are disconnected. Conversely, color information from other surfaces, even though it may be in retinotopically intervening space, is essentially ignored.

Speaking in broadly theoretical and ecological terms, it would seem appropriate for color computations to be surface-specific. Our visual world is filled with occlusive relationships where one surface is partially occluded by another, as in an object behind a tree. It would be advantageous to allow visual processing to deal with the visible parts that object as if it were whole, with a minimal amount of interference from signals arising from the leaves and branches. But what kind of neural mechanisms could be involved in the transfer of color information across remote visual regions in such a specific manner?

The types of pattern, configuration, and common-motion based cues we used in the stimuli would point toward processing that occurs in higher visual areas. Higher visual areas, however, cannot offer a full account of the visual phenomena described here. While the spatial resolution of these areas is sufficient to encode changes in perceived color over the broad color gradients, it is not sufficient to explain the fine-grained surface specificity displayed in our experiments. Receptive field sizes are too large to encode different colors for all of the squares and the grid, let alone for the separate dots in the transparent motion stimulus.

It seems then, that a full account requires distributed, cooperative coding between high and low-level visual areas, one computing surface properties and relationships based on global patterns, and the other encoding fine-grained maps of surface correspondence. Results from a number of neurophysiological studies provide details for such a distributed computation. More generally, the results also lend support for the primary role of the global surface segregation process in determining surface color and visibility.

First we review single-unit neurophysiology studies of classical Troxler fading and its relationship to the processing of uniform color surfaces. Recordings in macaque V1 and V2 have been taken while the monkey viewed classical Troxler fading stimuli consisting of a solid disk of one color and a surrounding ring of a different color. It was found that neurons selective for the edge of the color disk declined in activity with a time-course that corresponded to the monkey's behavioral report of disk fading. Meanwhile, in neurons that were selective for the interior surface of the color spot, activity remained steady. In preliminary data, the same was found to be true in V4 (von der Heydt, et al. 2003).

This evidence indicates that early visual cortex does not "fill-in" perceptual color in the literal sense. Instead, von der Heydt et al. argue that perceptual color is directly encoded in low level "oriented edge signals" which are used in figure-ground segregation. In an earlier study from the same group, single-cell recordings from macaque V2 during stimulation with solid color patches indicated that edge-selective neurons encode not only the luminance polarity across an edge, but also border ownership – which side represented figure and which side ground (Zhou, et al. 2000). The time-course of activity suggested that feedback from higher visual areas supplied the global surface information, as border ownership signals appeared at a timing that strained the speed limits of lateral neural connectivity.

Under this kind of symbolic filling-in interpretation, these oriented edge signals would be the direct source of color information for higher order processing of form – there would be no "spread" of surface signals in the explicit, isomorphic sense. The inclusion of figure-ground components in the oriented edge signals could begin to explain some aspects of the phenomena described in our study. However, it does not account for how certain edges might be discounted in the context of global surface cues.

What is required in the illusions presented here is a neural signal present in early visual cortex that codes for a common surface across non-contiguous patches of color. This could provide the high-resolution map necessary to allow color information to interact across gaps. It would also provide the necessary information for higher-order color processing to

determine the set of color edge signals which are relevant to the surface, and which edge signals are to be ignored.

Such signals have been found in V1, in monkeys viewing texture-defined figures much like the ones we used in experiment 2. Neurons with receptive fields in the center of a texture surface experienced a late enhancement in their response if that surface was segregated in a figure-ground relationship (Zipser, et al. 1996). When feedback from higher visual areas is disrupted, either by anesthesia (Lamme, et al. 1998) or lesion (Lamme, et al. 1997), the figure-ground enhancement is eliminated, while other stimulus-response patterns remain the same.

These studies indicate that the edge signals in early visual cortex are multiplexed with a map of surface correspondences that is built based on feedback from higher visual areas. This parallels our psychophysical finding that edge signals contribute to the color of filled surfaces in a manner that is dependent on their role in globally segregated surfaces.

Other psychophysical studies in humans also indicate that surface segregation is a critical step for visual perception in general, and in the perception of surface color in particular. Studies in a number of different fields and techniques offer somewhat different interpretations as to how surface segregation enters the picture.

In the context of illuminated objects and surfaces, there has been a long tradition of debate as to whether there is explicit decomposition of stimulus luminance into surface reflectance and prevailing illumination, or if surface qualities are calculated based on local cue processing alone (Anderson 1997; Bergstrom 1977; Gilchrist, et al. 1999; Gilchrist 1979). Recently, Anderson and Winawer (Anderson & Winawer 2005) provided vivid support for the existence of explicit decomposition for luminance stimuli, showing that a configural presentation of transparently overlaid cloud-like formations greatly enhanced the magnitude of the typical simultaneous contrast illusion. Although color and luminance processing are known to diverge in many aspects, this shows that surface attribution can play a role in the appearance of basic surface properties.

Looking at color perception specifically, Moradi and Shimojo (Moradi & Shimojo 2004) conducted a study comparing the perception of colored random dot surfaces under conditions of rapid alternation or transparent motion. They found that surfaces in transparent motion were segregated pre-attentively. Color, on the other hand, required attention in order to successfully attribute color to each surface.

In a similar regard, it has been proposed that the color filling-in found in stabilized and/or equiluminant images is more consistent with a breakdown in surface segmentation and feature binding than with a decay in bottom-up visual input (Billock & Tsou 2004b). It has been found that optical stabilization of images on the retina causes color to be unbound from surfaces, inducing color migration or switching across borders, and eventually a complete

visual black-out (Billock, et al. 2001). Also, an extensive history of studies shows that phenomenal progression of image fading is based on the sequential disappearance of globally organized parts, rather than a progression along local features (reiewed in Billock & Tsou 2004a).

Neurophysiological investigation of activity during the viewing of stimuli of the kind presented here would be highly informative as to how colored visual surfaces are encoded in cortex. Up until now, evidence has been consistent with the notion that local edge signals correspond to perceived surface properties. In contrast, the stimuli presented here contain sets of local edges which are consistent with each other when analyzed locally, but are in competition when organized into globally segmented surfaces. Direct measurement of the neural signals corresponding to the color edges around local elements would reveal whether such competitive interactions are resolved at the level of early visual cortex.

If the local edge signals are found to somehow reverse in accordance with the perceived color, this would be consistent with a direct symbolic code such as the one proposed by von der Heydt, in which low-level oriented edge signals are the seat for a direct encoding of surface color used in higher-level form processing. If the local edge signals are found to persist in contradiction of the perceived color, this would be more consistent with a cortical binding model such as proposed by Billock and Tsou (Billock & Tsou 2004a; Billock & Tsou 2004b), where the appearance and visibility of surfaces is determined by how connections within and between various feature and surface maps are maintained.

An intriguing possibility is that one would find a selective effect on the figure-ground enhancement (feedback) component of activity, leaving the bottom-up feature selectivity component of activity intact. This would mirror the effect found in macaque V1 as a correlate for surface visibility (Super et al. 2002). Such a result would also be consistent with the participation of surface feature binding mechanisms in the process of color-filling, and would also suggest late feedback to V1 as the mechanism by which surface color is bound to low-level edge information.

Our findings argue that the processing of surface color is dependent on the formation of multiple perceptual surface representations, which are segmented based on global features and configurations. This processing seems to require both the wide receptive fields and global processing of higher visual areas and the high-resolution local feature map in lower visual areas. We propose that this process might involve surface processing in higher visual areas and a feedback loop, which causes fine-grained surface correspondences to be encoded in early visual cortex. There is neurophysiological evidence for such surface-encoding feedback in awake, behaving monkeys. This reverberant feedback between multiple higher level surface representations and lower level local features could signal the proper correspondences between local elements and their surfaces. Such a mechanism might explain how color processing can occur across distantly segregated elements of a common surface, independently of nearby elements of other surfaces.

Chapter 5: Conclusion

Chapter 2 described a method whereby TMS could be used as a probe to reveal dynamic internal brain states. The visual information contained within these brain states was revealed as a visual percept induced in the subject. The brain states were found not only to reflect the processing of concurrent visual inputs, but also residual activity from the processing of visual stimuli in the recent past. This resulted in the perception of an *instant replay* when TMS was delivered after a visual stimulus.

This technique complements existing neurophysiological techniques in terms of the type of information it provides. Whereas other techniques provide high levels of anatomical accuracy to map local units of computation, this technique induces system-wide activation to probe into global states. Whereas other techniques detect the physical phenomena occurring in the brain, this technique probes into the information that those phenomena encode.

The investigation of the binding problem in chapter 3 further revealed that TMS did not simply trigger a wholesale recapitulation of recent neural processing, but preferentially activated certain visual representations over others. This was seen in the first experiment, a combination of TMS-induced replay with Cai's asynchronous binding illusion. One of the effects induced by TMS was to induce percepts that were not simple replays of the preceding percept, but instead reflected the activation of an accurate representation of the stimulus display. This indicates that at least some components of the visual system had properly

processed the visual stimulus. TMS, by preferentially activating that component, was able to provide a more direct examination of that visual information by removing the distortion or masking that was occurring during the viewing of the illusion. The existence of this correct representation is inconsistent with the proposed asynchronous binding mechanism. It implies that the illusion has more to do with distortions caused by motion, and may not be related to binding.

That first experiment, as well as the following experiment with simple color gratings, demonstrated that TMS was able to trigger the unbinding and misbinding of visual features. TMS-induced replay could induce a percept that was the amalgam of the color of one object and the position and orientation of another. In grating stimuli, the replay effect produced patchy percepts reflecting the independent activation of the gratings' color and orientation. While in some patches, the percept reflected the activation and binding of both features, other patches reflected the activation of only one of the features. These effects provide evidence that the feature binding is indeed an active process of integrating a distributed encoding of visual features.

The final experiment in chapter three demonstrated a clear, long-lasting case of misbinding in the absence of TMS. The steady-state misbinding of motion and color seen in a transparent motion display represents the first individually verifiable, sustained misbinding illusion, which occurs under continuous, fully attentive viewing. These qualities of the display make

it uniquely suitable for neurophysiological study, opening up the binding problem to a search for neural correlates.

In chapter four, we investigated the mechanisms by which colored surfaces are perceived. Using an effect derived from artwork by Julian Stanczak, we demonstrated illusions that run counter to classical findings and models of the visual system. We found that color could easily jump over large gaps of space, overwriting the color in distant areas. Developing a more generalized model that accounts for both the old and new findings, we argued that the perceptual extent of surface color is governed, not by locally-defined features such as luminance edges, but instead by globally-defined features such as the edges defined by perceptual surface segregation. This would allow the visual system to compensate for situations such as partial occlusion, such as an object behind branches.

Furthermore, we show that color-filling does not operate on the basis of a retino-topographic map of the visual field. Multiple overlaid surfaces can each support their own independent color interactions. Even though the two surfaces may occupy the same retinotopic space, colors do not flow between the two surfaces as they spread. This is demonstrated both in a foreground-background configuration and in a transparent motion configuration.

In examining possible neurophysiological implementations of our model of filling-in, we find possible connections to feature binding mechanisms. Our experiments demonstrate that surface segmentation is critical to the process of color-filling. The locus of surface

segregation mechanisms has been alternately ascribed to oriented edge detectors in V2/V3 or to late activity in V1 resulting from feedback from higher cortical areas. Because the stimuli we present create unique conflicts between local edges and global surfaces, their use in neurophysiological investigations would be informative as to the mechanisms of color surface perception. Depending on the outcome of these experiments, there might be a unification of surface segregation, color-filling and feature binding mechanisms.

In sum, this thesis probes into the problems of feature binding and color surface perception using a convergence of classical visual illusion psychophysics and modern physiological techniques. We demonstrate the existence of active binding processes which can be disrupted, and provide a stimulus which is uniquely suited for future neurophysiological study of the binding problem. We demonstrate a link between surface segregation mechanisms and color-filling, and provide stimuli which are uniquely suited for the investigation of color surface perception. Finally, these investigations also develop a TMS methodology, which serves as a complement to existing neurophysiological techniques in extracting the information content encoded in system-wide brain states.

Appendix A Detailed Methods

Chapter 2

Experiment 1

Subjects sat 57 cm from a Sony Multiscan 20se CRT computer monitor in a dark or dimly lit room. All visual stimuli were generated using Matlab Psychophysics toolbox running under Mac OS 8.6. The stimulus program also controlled the TMS equipment via a National Instruments I/O board. Timing was synched to the 75 Hz refresh of the monitor. Paired pulses of TMS were delivered by a pair of Magstim 200 single-pulse units, their pulses sent to a figure-8 coil by a Magstim Bistim unit. The Bistim unit also controlled the pulse separation, set at 50 ms. The coil was positioned over occipital cortex with the handle oriented horizontally, resulting in either a leftward or rightward current direction at the focus. Rightward current preferentially excites the right hemisphere, leftward current, left hemisphere. For any given current direction, the coil position was chosen based to optimize phosphene perception.

Experiment 2

Methods same as experiment 1, with the addition of the following. A red disk with a diameter of 15 degrees of visual angle is flashed for 100ms. 70ms after the offset of the flash, TMS administered to stimulate left visual cortex. Following each trial, a reproduction of the disk is shown steadily on the screen, and the subject adjusts the intensity of the bottom right quadrant of the disk to match the earlier percept.

Experiment 3

Methods same as experiment 2 except that delay between flash and TMS was varied, up to 2 seconds in length.

Chapter 3

Experiment 1

Physical setup was the same as in the previous chapter. Visual stimulus is as follows. The clock hand is a one degree wedge extending seven degrees out from fixation. The hand rotates at 0.5 Hz. For a single frame of the animation, the hand is displayed as a different color. Color combinations varied, but were typically of a green or blue clock hand sweep with a brief change to red. TMS was delivered 300 ms after the display of the color flash. After each trial subjects reported the position of their percepts by adjusting a radial line whose angle was controlled by the mouse. Two types of reports were given, one of the perceived location of the color change, and another of the perceived location of the percept induced by TMS.

Experiment 2

The visual display consisted of a full screen 53 ms flash of an oriented red grating. Orientation of the grating varied between vertical, horizontal, and \pm 45 degrees. Spatial frequency was 3 cycles per degree. After a 0.5 to 2.0 second delay, TMS was administered. Subjects gave verbal reports, and used the mouse to draw outlines of the region in which they saw a replay of a colored grating. Then an oriented grating was displayed in this region, and subjects adjusted the brightness and saturation of the display. Horizontal movement of the mouse increased the amount of red, vertical movement increased the amount of gray. This thesis includes a qualitative account of the data only.

Experiment 3

Environment: Chinrest is 57 cm from screen, eye level with center of stimulus. Room lights are off, though there is some indirect light from an LCD monitor behind and to the side of the subject.

Hardware/software: Mac G4 1.25 GHz, LaCie electron 21 monitor, semi-flat (has horizontal curvature). Video at 75Hz, 1600x1200, projects 36x27.5cm image, (all cm measurements are flat, lined up with view from chinrest). MacOS 9.2.2, Matlab 5.2.1, Psychophysics Toolbox 7.1.4.

Stimlus - spatial: dot field is 880x440 pixels(20.5x10.5cm). Total dot density is 0.007, each dot is one pixel. Color-motion transition line is at 263 pixels from the center (6 cm). Fixation point is 8 pixel diameter. Transition marker lines are 8x50 pixels, vertically offset from the stimulus by 20 pixels.

Stimulus - motion: The movie is a loop of 40 animation frames. Animation is presented at one frame every 2 screen cycles, wherein each dot moves by one pixel distance. Each dot has a lifetime of 20 frames. Number of dots moving up and number of dots moving down are always balanced.

Stimulus - color: Based on the typical isolumanance reports (from full-screen heterochromatic flicker) for the setup, we set red dots at [255 0 0], Green dots at [0 150 0]. When a red and a green dot overlay each other, it is presented as [255 150 0]. Fixation point and transition markers are [250 250 250]. These settings were sent raw to the video board, not gamma-corrected. Number of red dots and number of green dots are always balanced.

Procedure: Subects were screened to ensure that they could discriminate motion direction in the peripheral regions under the empty-center control condition. Two subjects were rejected in this manner. In experiment 1, all combinations of peripheral percentage coherence and central red up/red down/control conditions (total 15 conditions, 15 trials each) were randomly interleaved and performed in three blocks. To maintain color/motion balance, if red dots were, for example, 75% up, 25% down, then green dots would be 25% up, 75% down. In experiment 2, peripheral red up/red down and center present/absent conditions (20 trials each, total 80 trials) were randomly interleaved and performed in three blocks.

Chapter 4

Participants & Apparatus Experiments 1, 2 and 4 each involved four participants, including two of the authors (RK and DW). Experiment three involved six participants, including three of the authors (RK, DW and SS). All had normal or corrected-to-normal visual acuity and normal color perception.

Stimuli were presented on a 22-inch monitor, a LaCie Electron21 with a screen image projecting 37.5 cm by 28 cm. Stimuli were generated in Matlab using Psychophysics Toolbox (Brainard 1997), running under Macintosh OS 9. The participants sat in front of the computer screen at a viewing distance of 57 cm, their heads partially immobilized in a chinrest. Chromaticity of colors was measured at viewing distance with a Minolta Chroma Meter CS-100A, which outputs CIE 1931 x, y units.

Stimuli & Procedure Prior to each experiment, the green value perceptually equiluminant to the maximum intensity of the monitor's red gun was obtained for each individual participant using heterochromatic flicker.

Experiment 1: We used radial color gradients where color gradually changes between red (CIE x, y = 0.629, 0.342) and green (0.287, 0.610), displayed at the equiluminant level for each subject (see figure 4.1). Intermediate colors in the gradient were created via linear interpolation after gamma correction. The gradient was defined as a function of the visual eccentricity from a fixation marker and had a cumulative Gaussian profile (figure 4.1c),

$$f(x) = \frac{1}{\sqrt{2\pi\sigma}} \int_{-\infty}^{x} \exp(\frac{-(t-\mu)^2}{2\sigma^2}) dt$$

where *x* is the eccentricity, and σ the scale parameter that determines the steepness of the gradient, and μ is the location parameter for the eccentricity at which the intermediate color (yellow) occurred.

In the first experiment, we used three radii for the central color disk ($\mu = 2.8^{\circ}$, 5.6° and 8.3°) with $\sigma = \mu/2$. On half the trials, we used red-center stimuli, and on the other half, greencenter stimuli. The stimuli were displayed within a circular aperture with a radius of 14.3°. On half of the trials, luminance (black) gaps (6.7 arcmin in width) were overlaid on the color gradient. Thus the color gradient was displayed only on regularly arranged squares (26.7 arcmin by 26.7 arcmin; see figure 1a). In the control condition, the same stimuli were presented without the gaps (figure 1b).

Participants viewed each stimulus until they judged the entire area to be homogeneous in color. At that point they would hit one of three keys, indicating whether the central color had filled into the periphery, vice-versa, or if the direction was too ambiguous to judge. Each trial was subject to a 30 second time limit.

Following each stimulus presentation, inter-trial effects were minimized by presenting the opposite color gradient, followed by a dynamic pattern mask consisting of all of the colors present in the stimulus. Both were presented for half the duration of the original stimulus. 20 trials were performed on each condition. There were thus a total of [3 (radii) x 2 (color

$$f(x) = \frac{1}{\sqrt{2\pi\sigma}} \int_{-\infty}^{x} \exp(\frac{-(t-\mu)^2}{2\sigma^2}) dt$$

where *x* is the eccentricity, and σ the scale parameter that determines the steepness of the gradient, and μ is the location parameter for the eccentricity at which the intermediate color (yellow) occurred.

In the first experiment, we used three radii for the central color disk ($\mu = 2.8^{\circ}$, 5.6° and 8.3°) with $\sigma = \mu/2$. On half the trials, we used red-center stimuli, and on the other half, greencenter stimuli. The stimuli were displayed within a circular aperture with a radius of 14.3°. On half of the trials, luminance (black) gaps (6.7 arcmin in width) were overlaid on the color gradient. Thus the color gradient was displayed only on regularly arranged squares (26.7 arcmin by 26.7 arcmin; see figure 4.1a). In the control condition, the same stimuli were presented without the gaps (figure 4.1b).

Participants viewed each stimulus until they judged the entire area to be homogeneous in color. At that point they would hit one of three keys, indicating whether the central color had filled into the periphery, vice-versa, or if the direction was too ambiguous to judge. Each trial was subject to a 30 second time limit.

Following each stimulus presentation, inter-trial effects were minimized by presenting the opposite color gradient, followed by a dynamic pattern mask consisting of all of the colors present in the stimulus. Both were presented for half the duration of the original stimulus. 20 trials were performed on each condition. There were thus a total of [3 (radii) x 2 (color

polarity; red-center vs. green-center) x 2 (with and without the gaps) x 20 trials] = 240 trials. The experiment was divided into 4 blocks and each took about 30-40 minutes.

Analysis: The directionality index, *I*, was calculated as I = [(number of filling-out trials) - (number of filling-in trials)] / [(number of filling-out trials) + (number of filling-in trials)].The value of*I*ranges between -1 and 1. A positive value of*I*indicates a tendency for filling-out to occur, and a negative value indicates a tendency toward filling-in.

Experiment 2 (texture defined by bar orientation): Stimuli subtended the entire display. The color gradient was defined in the same manner as in the first experiment except that it was created based on horizontal, instead of radial, eccentricity. μ and σ were fixed at 10° and 3.3°, respectively. Bar size was 3.3° by 0.56°. Orientation of the bar at a given horizontal eccentricity was determined in a way similar to the color gradient as follows;

$$\phi(x) = \frac{\pi}{4} \pm \frac{\pi}{2} \cdot f(x),$$

where $\phi(x)$ is the orientation of a bar at an eccentricity *x* and *f*(*x*) is the cumulative Gaussian function. The first term is the orientation of a bar displayed at the center. This orientation gradually changed to the orthogonal orientation through $\phi = 0$ (i.e. horizontal bars). The position of a bar, *x*, was based on the central point of the bar. The density of the bars was 2.25 bars/deg². For the gradient of the texture, we used a fixed μ of 3.3° and varied σ between 0°, 3.3°, and 5.0°.

As in experiment 1, trial types were balanced in terms of color polarity, and each trial was followed by a full-field opposing gradient and pattern mask. Each participant responded to 20 trials for each condition.

Experiment 3 (static dual-gradient stimuli): Stimuli subtended the entire display. Squares subtended 0.37°, and were spaced by 0.55°. The color gradient within the squares went from light-blue (CIE x, y = 0.168, 0.166) to green (0.274, 0.563). The color gradient within the grid went from red (0.642, 0.346) to yellow (0.544, 0.418). Spatial parameters of the gradient were fixed for the squares, with μ = 7.32 and σ = 3.66. Spatial parameters of the grid were varied across trials, with μ = 5.49, 7.32 or 10.98 and σ = $\mu/10$.

Trial sequence remained the same. Participants pressed keys when they perceived the grid and/or the squares to become homogeneous in color. Keypresses by the left hand indicated the time and direction of color-filling in the grid, keypresses by the right hand indicated the time and direction of color-filling in the squares. Each participant performed 10 trials per condition.

Experiment 4 (transparent motion): The transparent motion stimulus was created by superimposing two surfaces (upwards and downwards), each consisting of 600 random dots. Each dot was 0.178° in diameter and had a lifetime of 2 seconds. The velocity was 2.84°/s. The stimulus was displayed in an area subtending 14.2° vertically and 17.8° horizontally. The

color gradient for each surface was created in the same manner as for the shape-based stimulus. Here, μ was fixed at 3.75° and sigma varied between 0°, 2° or 3°.

The participants viewed the stimuli for unlimited duration of time and responded when they perceived the entire stimulus as homogeneous. Each participant performed 16 trials per condition.

REFERENCES

- Abbruzzese, G., C. Trompetto and M. Schieppati (1996). The excitability of the human motor cortex increases during execution and mental imagination of sequential but not repetitive finger movements. *Experimental Brain Research 111*, 465-472.
- Amassian, V. E., P. J. Maccabee, R. Q. Cracco, J. B. Cracco, M. Somasundaram, J. C. Rothwell, L. Eberle, K. Henry and A. P. Rudell (1994). The polarity of the induced electric-field influences magnetic coil inhibition of human visual-cortex implications for the site of excitation. *Electroencephalography And Clinical Neurophysiology 93*, 21-26.
- Anderson, B. L. (1997). A theory of illusory lightness and transparency in monocular and binocular images: the role of contour junctions. *Perception 26*, 419-453.
- Anderson, B. L. and J. Winawer (2005). Image segmentation and lightness perception. *Nature 434*, 79-83.
- Aurora, S. K., K. M. A. Welch and F. Al-Sayed (2003). The threshold for phosphenes is lower in migraine. *Cephalalgia* 23, 258-263.
- Bender, M. B. and H. L. Teuber (1946). Ring scotoma and tubular fields their significance in cases of head injury. *Archives Of Neurology And Psychiatry 56*, 300-326.
- Bergstrom, S. S. (1977). Common and relative components of reflected light as information about illumination, color, and 3-dimensional form of objects. *Scandinavian Journal Of Psychology 18*, 180-186.
- Billock, V. A., G. A. Gleason and B. H. Tsou (2001). Perception of forbidden colors in retinally stabilized equiluminant images: an indication of softwired cortical color opponency? *Journal Of The Optical Society Of America A-Optics Image Science And Vision 18*, 2398-2403.

- Billock, V. A. and B. H. Tsou (2004a). A role for cortical crosstalk in the binding problem: Stimulus-driven correlations that link color, form, and motion. *Journal Of Cognitive Neuroscience 16*, 1036-1048.
- Billock, V. A. and B. H. Tsou (2004b). What do catastrophic visual binding failures look like? *Trends In Neurosciences* 27, 84-89.
- Boynton, R. M., M. M. Hayhoe and D. I. A. Macleod (1977). Gap Effect—chromatic and achromatic visual-discrimination as affected by field separation. *Optica Acta 24*, 159-177.
- Brainard, D. H. (1997). The Psychophysics Toolbox. Spatial Vision 10, 433-436.
- Butler, B. E., D. J. K. Mewhort and R. A. Browse (1991). When do letter features migrate—a boundary-condition for feature-integration theory. *Perception & Psychophysics 49*, 91-99.
- Cai, R. and J. Schlag (2001). A new form of illusory conjunction between color and shape. Presented at: Vision Sciences Society Annual Meeting (Sarasota, FL, USA.).
- Cohen, M. A. and S. Grossberg (1984). Neural dynamics of brightness perception features, boundaries, diffusion, and resonance. *Perception & Psychophysics 36*, 428-456.
- Cole, G. R., C. F. Stromeyer and R. E. Kronauer (1990). Visual interactions with luminance and chromatic stimuli. *Journal Of The Optical Society Of America A-Optics Image Science And Vision* 7, 128-140.
- Cornsweet, T. N. (1970). Visual Perception (New York, Academic Press).
- De Weerd, P., R. Desimone and L. G. Ungerleider (1998). Perceptual filling-in: a parametric study. *Vision Research* 38, 2721-2734.
- DeYoe, E. A., D. J. Felleman, D. C. Van Essen and E. McClendon (1994). Multiple processing streams in occipitotemporal visual-cortex. *Nature* 371, 151-154.

- Donk, M. (1999). Illusory conjunctions are an illusion: The effects of target-nontarget similarity on conjunction and feature errors. *Journal Of Experimental Psychology-Human Perception And Performance 25*, 1207-1233.
- Donk, M. (2001). Illusory conjunctions die hard: A reply to Prinzmetal, Diedrichsen, and Ivry (2001). Journal Of Experimental Psychology-Human Perception And Performance 27, 542-546.
- Eagleman, D. M. (2003). Motion-biasing, not asynchronous feature binding, explains the feature-flash drag effect. Presented at: Vision Sciences Society Annual Meeting (Sarasota, FL, USA).
- Eskew, R. T. and R. M. Boynton (1987). Effects of field area and configuration on chromatic and border discriminations. *Vision Research* 27, 1835-1844.
- Fadiga, L., G. Buccino, L. Craighero, L. Fogassi, V. Gallese and G. Pavesi (1999).
 Corticospinal excitability is specifically modulated by motor imagery: a magnetic stimulation study. *Neuropsychologia 37*, 147-158.
- Friedman, H. S., H. Zhou and R. von der Heydt (2003). The coding of uniform colour figures in monkey visual cortex. *Journal Of Physiology-London* 548, 593-613.
- Fukushima, K. (1980). Neocognitron a self-organizing neural network model for a mechanism of pattern-recognition unaffected by shift in position. *Biological Cybernetics* 36, 193-202.
- Garson, J. W. (2001). (Dis)solving the binding problem. *Philosophical Psychology 14*, 381-392.
- Gerrits, H. J. M., B. Dehaan and A. J. H. Vendrik (1966). Experiments with retinal stabilized images. Relations between observations and neural data. *Vision Research* 6, 427-&.

- Gilchrist, A., C. Kossyfidis, F. Bonato, T. Agostini, J. Cataliotti, X. J. Li, B. Spehar, V. Annan and E. Economou (1999). An anchoring theory of lightness perception. *Psychological Review 106*, 795-834.
- Gilchrist, A. L. (1979). Perceptiono of surface blacks and whites. *Scientific American 240*, 112-&.
- Grossberg, S. and E. Mingolla (1985). Neural dynamics of perceptual grouping textures, boundaries, and emergent segmentations. *Perception & Psychophysics 38*, 141-171.
- Horton, J. C. and W. F. Hoyt (1991). Quadrantic visual-field defects a hallmark of lesions in extrastriate (v2/v3) cortex. *Brain 114*, 1703-1718.
- Hubel, D. H. and T. N. Wiesel (1968). Receptive Fields And Functional Architecture Of Monkey Striate Cortex. *Journal Of Physiology-London 195*, 215-&.
- Kamitani, Y., V. M. Bhalodia, Y. Kubota and S. Shimojo (2001). A model of magnetic stimulation of neocortical neurons. *Neurocomputing* 38, 697-703.
- Kamitani, Y. and S. Shimojo (1999). Manifestation of scotomas created by transcranial magnetic stimulation of human visual cortex. *Nature Neuroscience 2*, 767-771.
- Kammer, T., K. Puls, M. Erb and W. Grodd (2005). Transcranial magnetic stimulation in the visual system. II. Characterization of induced phosphenes and scotomas. *Experimental Brain Research 160*, 129-140.
- Kingdom, F. and B. Moulden (1988). Border effects on brightness: a review of findings, models and issues. *Spatial Vision 3*, 225-262.
- Krekelberg, B. and M. Lappe (2001). Neuronal latencies and the position of moving objects. *Trends In Neurosciences 24*, 335-339.

- Lamme, V. A. F., K. Zipser and H. Spekreijse (1997). Figure-ground signals in V1 depend on extrastriate feedback. *Investigative Ophthalmology & Visual Science* 38, 4490-4490.
- Lamme, V. A. F., K. Zipser and H. Spekreijse (1998). Figure-ground activity in primary visual cortex is suppressed by anesthesia. *Proceedings Of The National Academy Of Sciences Of The United States Of America* 95, 3263-3268.
- Landy, M. S. and J. R. Bergen (1991). Texture segregation and orientation gradient. *Vision Research 31*, 679-691.
- Livingstone, M. and D. Hubel (1988). Segregation of form, color, movement, and depth anatomy, physiology, and perception. *Science 240*, 740-749.
- Marg, E. and D. Rudiak (1994). Phosphenes induced by magnetic stimulation over the occipital brain - description and probable site of stimulation. *Optometry And Vision Science* 71, 301-311.
- Meyer, B. U., R. Diehl, H. Steinmetz, T. C. Britton and R. Benecke (1991). Magnetic stimuli applied over motor and visual-cortex - influence of coil position and field polarity on motor-responses, phosphenes, and eye-movements. *Electroencephalography And Clinical Neurophysiology*, 121-134.
- Montag, E. D. (1997). Influence of boundary information on the perception of color. *Journal* Of The Optical Society Of America A-Optics Image Science And Vision 14, 997-1006.
- Moradi, F. and S. Shimojo (2004). Perceptual-binding and persistent surface segregation. *Vision Research 44*, 2885-2899.
- Nijhawan, R. (2002). Neural delays, visual motion and the flash-lag effect. *Trends In Cognitive Sciences* 6, 387-393.

- Olshausen, B. A., C. H. Anderson and D. C. Van Essen (1993). A neurobiological model of visual-attention and invariant pattern-recognition based on dynamic routing of information. *Journal Of Neuroscience 13*, 4700-4719.
- Olshausen, B. A., C. H. Anderson and D. C. Van Essen (1995). A multiscale dynamic routing circuit for forming size-invariant and position-invariant object representations. *Journal Of Computational Neuroscience* 2, 45-62.
- Paradiso, M. A. and K. Nakayama (1991). Brightness perception and filling-in. *Vision Research 31*, 1221-1236.
- Perrett, D. I. and M. W. Oram (1993). Neurophysiology of shape processing. *Image And Vision Computing 11*, 317-333.
- Pessoa, L., E. Thompson and A. Noe (1998). Finding out about filling-in: A guide to perceptual completion for visual science and the philosophy of perception. *Behavioral and Brain Sciences 21*, 723-748.
- Pinna, B., G. Brelstaff and L. Spillmann (2001). Surface color from boundaries: a new 'watercolor' illusion. *Vision Research* 41, 2669-2676.
- Ramachandran, V. S. (1992). Blind spots. Scientific American 266, 86-91.
- Ramachandran, V. S. and R. L. Gregory (1991). Perceptual filling in of artificially induced scotomas in human vision. *Nature 350*, 699-702.
- Redies, C. and L. Spillmann (1981). The neon color effect in the Ehrenstein Illusion. *Perception 10*, 667-681.
- Redies, C., L. Spillmann and K. Kunz (1984). Colored neon flanks and line gap enhancement. *Vision Research 24*, 1301-&.

- Riggs, L. A., F. Ratliff, J. C. Cornsweet and T. N. Cornsweet (1953). The disappearance of steadily fixated visual test objects. *Journal Of The Optical Society Of America 43*, 495-501.
- Rossi, S., P. Pasqualetti, F. Tecchio, F. Pauri and P. M. Rossini (1998). Corticospinal excitability modulation during mental simulation of wrist movements in human subjects. *Neuroscience Letters 243*, 147-151.
- Sergent, J. (1988). An investigation into perceptual completion in blind areas of the visual-field. *Brain 111*, 347-373.
- Singer, W. and C. M. Gray (1995). Visual feature integration and the temporal correlation hypothesis. *Annual Review Of Neuroscience 18*, 555-586.
- Snowden, R. J. and F. A. J. Verstraten (1999). Motion transparency: making models of motion perception transparent. *Trends In Cognitive Sciences* 3, 369-377.
- Stanczak, J. and E. McCleland (1998). Julian Stanczak "A Retrospective: 1948-1998" (Youngstown, Butler Institute of American Art).
- Todorovic, D. (1987). The Craik-O'brien-Cornsweet effect new varieties and their theoretical implications. *Perception & Psychophysics 42*, 545-560.
- Treisman, A. and G. Gelade (1980). A feature integration theory of attention. *Cognitive Psychology* 12, 97-136.
- Treisman, A. and H. Schmidt (1982). Illusory conjunctions in the perception of objects. *Cognitive Psychology 14*, 107-141.
- Troxler, I. P. V. D. (1804). Über das verschwinden gegebener gegenstände innerhalb unseres gesichtskreises. In Ophthalmolgische Bibliothek, K. Himly and J. A. Schmidt, eds. (Jena, Friedrich Fromann), pp. 1-53.

- Tsal, Y. (1989). Do illusory conjunctions support the feature integration-theory a criticalreview of theory and findings. *Journal Of Experimental Psychology-Human Perception And Performance 15*, 394-400.
- Tsotsos, J. K. (1995). Toward a computational model of visual attention. In Early Vision and Beyond, T. V. Papathomas, ed. (Cambridge, MA, MIT Press).
- Valero Cabre, A., A. Pascual-Leone and L. J. Toth (2004). Patterns of TMS induced activity visualized with intrinsic signal imaging. Presented at: Society For Neuroscience Annual Meeting (San Diego, CA).
- Valls Sole, J., A. Pascual-Leone, E. M. Wassermann and M. Hallett (1992). Human motor evoked-responses to paired transcranial magnetic stimuli. *Electroencephalography And Clinical Neurophysiology* 85, 355-364.
- von der Heydt, R., H. S. Friedman and H. Zhou (2003). Searching for the neural mechanism of color filling-in. In Filling-In—From Perceptual Completion to Cortical Reorganization, L. Pessoa, de Weerd., P, ed. (New York, Oxford University Press), pp. 106-127.
- von der Malsburg, C. (1981). The correlation theory of brain function. In Abteilung für Neurobiologie, MPI für Biophysikalische Chemie (Göttingen).
- von der Malsburg, C. (1995). Binding in models of perception and brain-function. *Current Opinion In Neurobiology* 5, 520-526.
- Wallis, G. and E. T. Rolls (1997). Invariant face and object recognition in the visual system. *Progress In Neurobiology* 51, 167-194.
- Whitney, D. and P. Cavanagh (2000). Motion distorts visual space: shifting the perceived position of remote stationary objects. *Nature Neuroscience 3*, 954-959.

- Wolfe, J. M. and K. R. Cave (1999). The psychophysical evidence for a binding problem in human vision. *Neuron* 24, 11-17.
- Yarbus, A. L. (1967). Eye Movements and Vision (New York, Plenum Press).
- Zhou, H., H. S. Friedman and R. von der Heydt (2000). Coding of border ownership in monkey visual cortex. *Journal Of Neuroscience 20*, 6594-6611.
- Zipser, K., V. A. F. Lamme and P. H. Schiller (1996). Contextual modulation in primary visual cortex. *Journal Of Neuroscience 16*, 7376-7389.
- Zur, D. and S. Ullman (2003). Filling-in of retinal scotomas. Vision Research 43, 971-982.