Visual Experience and Immediate Memory

J T Enns, University of British Columbia, Vancouver, Canada

To appear in: Banks, W. (Editor-in-Chief) (2008). Encyclopedia of Consciousness. Elsevier Inc.

Author Contact Information:

J. T. Enns
Department of Psychology
University of British Columbia
2136 West Mall
Vancouver, B.C., Canada V6T 1Z4

Email: jenns@psych.ubc.ca
Phone: 604-822-6634

Visual Experience and Immediate Memory

Word count = 5,938

Table of Contents

I. Introduction
II. Paradox One: Rich Experience from Impoverished Input
III. Paradox Two: Rich Experience but Limited Report
IV. Paradox Three: Experiences Evolve over Time from General to Specific
V. Conclusion

Key Words

Attention, Change Blindness, Inattentional Blindness, Visual Cognition, Visual Short-term Memory,

Glossary

Attention – any influence causing some information to be selected for more detailed processing at the expense of other information.

Cartoon advantage – improved recognition for pictures exaggerating distinctive features of a familiar object

Change blindness – being insensitive to a difference between two images

Comparison failure – failure to actively compare one region in a scene with its counterpart in another scene

Constraint – a physical regularity of the world in which our visual systems have evolved, e.g., gravity

Convexity bias – a tendency to assume, in the absence of contradictory information, that shapes with bulges are nearer to you than shapes with indentations

Edge alignment – a tendency to assume, in the absence of contradictory information, that aligned edges belong to the same surface

Ganglion cells – neurons in the eye that are organized in their circuitry to enhance edge detection

Heuristic – a rule of thumb that relates image interpretation to an expected constraint of the visual world

Lateral inhibition – circuitry found in ganglion cells that enhances edge detection for one ganglion cell by inhibiting neighboring ganglion cells as a direct function of its degree of activation

Own-race advantage (equivalently, the cross-race disadvantage) – improved ability to learn the faces of previously unfamiliar individuals from one's own race

Receptive field – a description of the behavior of a sensory neuron in terms of a filter or template

Subjective edge (equivalently, virtual edge) – the appearance of a faint edge at a surface boundary in the absence of any physical evidence for such an edge

Article Synopsis

Visual experience is a construction by the neurons in our brain. To understand how this is accomplished, scientists in a number of related disciplines study three paradoxes: (a) that our phenomenally-rich visual experience seems to derive from relatively impoverished visual input, (b) that although our experience may be rich in color, shape and texture, we are severely limited in our ability to make an immediate report on it to others, and (c) that our visual experiences seem to evolve over time in an order that is reversed with respect to the anatomical order of processing.

Introduction

Healthy people with fully functioning visual systems take many things for granted. But what they may take most for granted, even more than the air they breathe, is that the world they experience through vision is vivid in color, rich in texture and detailed in shape. There is no longer any doubt that this experience is created by the neurons in our brains. However, how that experience is created by our neurons is still largely a mystery. This article will review what scientists from a number of related academic disciplines currently understand about how visual experience is constructed by our mind when it confronts the pattern of light that reach our eyes. To help the reader understand some of the questions that are still unanswered in this process, the review will be organized around three primary puzzles that everyone faces when they try to understand how vision works. Each puzzle is in fact a **paradox,** because each one involves strong scientific evidence about our visual experience, on the one hand, that flies in the face of our common sense beliefs, on the other hand. These common sense beliefs are themselves presumably based on our own experiences, which only deepens the mystery.

Paradox One: Rich Experience from Impoverished Input

When we see a cat through a picket fence we have no doubt that we are experiencing the cat as an entire animate object, even though the light reflecting from the cat is only visible to us through the openings between the pickets. Moreover, we experience the cat as a solid object with volume, even though the individual light patterns between the pickets is in reality only two-dimensional at any point in time. And over time, as we move past the picket fence, or the cat moves behind it, we experience a unique cat of a given size, with a specific color pattern, and with a unique posture, even though the information from between the pickets is constantly changing.

As this example shows, our visual brain is able to take fragmentary, ever-changing patterns of two-dimensional light, and transform them into an experience of individual objects interacting with one another in a stable world. At no point in this process do we gain insight from our own experiences of the seen world about *how* our brain accomplishes this transformation. We simply have the experience of seeing the world, out there, and for the most part, as it really is. Our common sense understanding of vision, based on these experiences, thus vastly underestimates the real difficulty of the problem of converting fragmentary two-dimensional patterns of light into holistic three-dimensional visual experiences. An understanding of that only comes from treating other humans (and often even ourselves) as temporary *objects* of scientific interest. When we adopt this stance of the *experimenter*, we are able to conduct experiments on the transformation of light to experience by observing how human awareness and behavior (considered as the *output* of the scientific object) are related to the physical patterns of light (considered as the *input* to the object) that give rise to these experiences in us.

Experiments studying the relationship between the physical conditions of vision and the associated experiences in humans are traditionally referred as **psychophysics**. This branch of science, also known as experimental psychology, had its beginning over 100 years ago in the laboratories of Hermann Helmholtz, Wilhelm Wundt, and Gustav Fechner. In the twenty-first century, researchers who conduct psychophysical experiments often refer to themselves as **cognitive neuroscientists,** if they have a special interest in the brain mechanisms associated with the experiences they are studying, **cognitive scientists,** if they are interested in generalizing what they know about the problem of extracting information from light to non-biological systems, or **vision scientists**, if they are interested in comparing how different visual systems, both biological and non-biological, solve the problem of creating experience from light).

Edge enhancement. Psychophysical experiments over the past century have revealed that the tendency for the visual brain to *go beyond the information given*

applies equally well to its very earliest and simplest processes as it does to the highest levels of visual cognition and imagination. Consider one of the lowest level functions performed by the eye, that of registering the existence of an **edge** (the difference between light and dark) at a particular location in an image. The daylight neural receptors in the eye, known as **cones**, do not merely record edges as faithfully as they can. Rather, they are organized to deliberately exaggerate the contrast of edges they encounter, through a system of circuitry known as **lateral inhibition**. The net effect of this circuitry is to highlight the existence of edges that are detected, so as to help bridge and join edges in regions of the image in which the lighting conditions make their detection difficult.

The **Hermann grid** shown in Figure 1a provides first-hand experience of the effects of the lateral inhibitory circuitry in the eye. Notice that when you first glance at the grid, you see gray smudges at most of the intersections. However, when you fix your gaze directly on one of the intersections, the gray smudge you previously saw out of the corner of your eye disappears. The gray smudges seem to be present only at the intersections that are not currently at the center of your gaze.

Figure 1b illustrates schematically how the circuitry of the eye causes these smudges to appear in our experience. Many **cones** feed their information into a smaller number of **ganglion cells**, and the ganglion cells are organized to inhibit their neighbors as a direct function of how active each one is. The plus signs in Figure 1b represent activation by light; the minus signs inhibition. So the stronger the input to a given ganglion cell, the more that cell will reduce the activity in neighboring cells. The net effect is that patterns of light are registered by ganglion cells that are organized as small filters, optimally sensitive to a discrete spot of light that is surrounded by region of reduced light. This is illustrated in Figure 1b by the circular regions of plus signs, surrounded by an annulus of minus signs. When one of these filters is centered on an intersection in the grid, its net activity (the sum of plusses and minuses exposed to the light) will be *less* than when it is centered on a white

space between the squares. The result is that we experience a reduced sensation of light at the intersections.

But why do the gray smudges disappear every time we fix our gaze directly on them? This relates to the way in which visual acuity diminishes as move away from the center of the eye. Because ganglion cells have very small **receptive fields** (circular-surround filters) at the center of the eye, they are smaller than the size of the intersections and so illusory gray smudges are not seen. As we move away from the center, the receptive fields grow increasingly larger, ensuring that there will be a match between intersection size and the optimal size of receptive field at some distance from the center of gaze.

Edge interactions. Once the pattern of signals from the eyes has reached the brain, similar processes in the circuitry there operate to group edges together that belong to the same object and to segment edges away from each other that belong to different objects. Psychophysical experiments designed to study this process have used displays such as those shown in Figure 2. Both of these patterns contain the same number of short line segments arranged to form an S shape. Yet it is easy to see the S in the pattern in Figure 2a and almost impossible to see it in Figure 2b. The only difference between the two patterns is that the short segments that make up the S in Figure 2a are aligned with one another along their endpoints, whereas the same segments forming the S in Figure 2b are aligned along their midpoints, placing each one of them at right angles to the imaginary curve forming the S. You will be able to confirm their existence by looking at each line segment individually and comparing it to its neighbor, even if you are unable to see the large S-shape as a whole.

The visibility of the S-shape in Figure 2a points to the existence of cooperative mechanisms that group neighboring edges of similar orientation and alignment. The invisibility of the S-shape in Figure 2b points to the existence of competitive mechanisms that keep neighboring edges of different orientation distinct from one

another. Both cooperative and competitive mechanisms of this kind have been confirmed in electrophysiological studies of animals and in brain imaging studies in humans.

Filling in edges and colors. Everyone who looks at the shapes in Figure 3a says they see a white square shape in addition to some black round and oval shapes. Moreover, the square shape seems to lying on top (or in front) of the oval shapes. When it comes to the color of the white square, it seems to be even brighter than the whiteness of the surrounding paper, causing faint edges to appear where the square lies in front of the surrounding background.

The faint edges of the square are known as **subjective (or virtual) edges** because they exist only in your brain. They do not exist on the page, as you can prove to yourself by blocking your view of one of the ovals with another piece of paper. These edges illustrate the way the visual system is constantly trying to find a reasonable interpretation of an ambiguous pattern of input. To do this the visual system uses heuristics (rules of thumb that work most of the time) to assign plausible interpretations to what it sees. In this case, it seems to be using a **convexity bias** (in the absence of contradictory information, assume shapes with bulges are nearer to you than shapes with indentations) and **edge alignment** (in the absence of contradictory information, assume that edges that are aligned belong to the same surface).

Everyone who looks at Figure 3b says they see regions of orange separated by regions of white. Yet, when you look closely you can see that the pattern really only consists of two thin lines of equal thickness. One line is purple and much darker than the white of the background page; the other is orange and therefore more similar in its brightness to the background white. The illusion created by these two thin lines is that the color of the lighter line (the one more similar in brightness to the background) fills-in its region with a faint tinge of color whereas the darker of the two

lines creates a sharp edge, giving the appearance its regions may even be more intensely white than the background page.

This **watercolor illusion** occurs because the visual system follows several other heuristics when it comes to assigning color values to surfaces and to lighting conditions. In brief, sharp or abrupt changes in color or brightness are assumed to signal the edges of surfaces. More gradual changes are assumed to arise from surface pigmentation and lighting conditions, which includes coloration, shading, shadows, and other factors that contribute to the uneven distribution of light. The application of these assumptions then lead to the perception in Figure 3b of three surfaces with an orange-ish color, separated by two background regions of white separating these surfaces.

Using heuristics and constraints. There are many other heuristics used by the visual system in addition to those used in interpreting the boundaries and colors of surfaces. Figure 4a gives an example of a consequences of critical heuristic used to interpret whether a shaded region of an image is bulging out toward the viewer or is dented inward, away from the viewer. This is the heuristic that *light generally shines from overhead*. With this assumption, a disc that is lightly shaded on top, and gradually becoming darker toward the bottom, will be seen as a bump on a surface, whereas a disc with the reverse pattern of shading will be seen as a divot. The heuristic is effective, meaning it usually brings us to the correct interpretation, because we live in a world in which our most important source of light, the sun, shines from above. Thus, a heuristic can be thought of as our visual systems' taking into account of a **constraint** (i.e., a physical regularity) of the world in which our visual systems have evolved and in which we have spent a lifetime of experience.

Gravity is a constraint that is also used in a number of heuristics employed by the visual system to interpret the pattern of light given to the eye. Because of gravity, objects are rarely able to hover in mid air, unless they are lighter than air or have a source of their own energy that permits them to fly. A very useful heuristic then, is

that *objects are usually resting on or attached to other surfaces*. This heuristic prompts us at first glance to interpret Figure 4b as a small block resting on a larger block. Closer inspection will reveal that there is another perfectly reasonable interpretation — it could be a small block hovering over the front right hand corner of the larger block — but this interpretation is more rare simply because our visual system seems to take the constraint of gravity into account.

Feature exaggeration. Sketch artists and portrait painters have known for a long time that in order to communicate quickly and effectively with the brain of the viewer, it is best to make a drawing of a familiar face that is not quite faithful to its actual appearance. For best results, one should exaggerate those features of a face that distinguish it from the average or prototypical face. This means, for example, over-emphasizing the chin and beard of Abraham Lincoln and the protruding ears of Prince Charles. Pictures drawn with these distinctive exaggerations are actually identified more quickly and accurately by viewers than more realistic photographs of the same people, an effect referred to as the **cartoon advantage** in object recognition.

The cartoon advantage illustrates that the brain has some very efficient ways of storing and retrieving information in memory. Rather than keeping copies of all the image and scenes it encounters in their original form, it seems to store images in terms of their differences from some prototype or average for the category as a whole. One advantage of this type of coding is that it then benefits the task of recognizing a familiar object when it is seen from a novel vantage point or under unusual viewing conditions. This is because differences between any specific object and the average or prototype object will usually be preserved in the face of conditions of lighting or of viewpoint that affect the specific object and the average object in the same way.

Enhancing vision with knowledge. Constraints that arise from the physical conditions under which vision occurs are not the only factors that determine what we

see. Knowledge that we have acquired about the peculiar characteristics of objects and the contexts in which they are most likely to occur also play a large role in our immediate visual experience.

Experts see things at a glance that novices miss entirely or are only able to see with external guidance or more time. For example, experienced bird watchers and specialized dog breeders directly experience the specific species in a glance (in the case of bird watchers) and specific breeds in a glance (in the case of dog breeders) while novices directly experience only the higher-order category of *bird* or *dog* in a single glance. The same seems to be true for those trained in medical diagnosis, automotive repair, and rock climbing.

One domain in which most humans have expertise is that of face recognition. This means they tend to see individual identities and specific emotional expressions in a single glance rather than simply experiencing *a face*. But because humans also tend to live in groups of similar looking individuals, they also tend to have expertise that is uneven when it comes to the identification of individuals of races that differ from their own. This means that when individuals from one racial group are given the task of learning to recognize individuals they have never seen before from either their own race, or from another race with which they have less experience, they tend to learn new faces from their own race more quickly. This is known as the **own-race advantage** (or equivalently, the **cross-race disadvantage**) in face recognition. It is as though they experience faces of their own race at the level of individual identity whereas they experience faces of another race at the level of group identity. Training and experience with faces of another race can overcome the cross-race disadvantage, just as training and experience with different bird species can turn the novice into an expert bird watcher.

Paradox Two: Rich Experience but Limited Report

The previous section discussed how the visual system *goes far beyond the information given* in creating a rich visual experience from visual input that is, by comparison, much more impoverished and ambiguous than our corresponding experience of it. The present section will contrast this rich sensory experience with the surprisingly limited nature of what we are able to report about our experience at any moment in time.

The severely restricted nature of our ability to report what we see was first studied systematically by George Sperling. His experiments involved flashing an array of English letters on a screen (several rows and columns) and asking participants to report as many letters as they could. The main result was that their report was limited to only three or four letters, even though it felt to the participants that they could see all of them quite clearly. This aspect of participants' experience was confirmed in another study where they were asked to report only one row of the letters, depending on an auditory tone (high, medium, or low in pitch, corresponding to the top, middle or bottom row) that occurred immediately after the letter array had been erased from the screen. Under these conditions, participants could still report three or four letters from any given row, indicating that immediately following the flash of the letters they had conscious access to all of them. The problem with trying to report all of the letters was that in order to do so, some of the letters had to be committed to short-term memory while other letters were being reported. Short-term visual memory therefore placed restrictions on what participants could report about their immediate visual experiences.

Since the original experiments of Sperling, the nature and limits of short term visual memory have been studied in a large number of ways. Studies focusing on the question of its **capacity** (the size of the memory store) indicate that it is large (virtually limitless) for a period of about 200 milliseconds but then, following that brief period, it is limited to only three or four individual objects or events. Studies focusing

on the **content** (the nature of the representations in memory) suggest that each of the three or four items being held in short term memory are coded quite richly, meaning that several of their features, including size, color, and shape, can be accessed for report from any of these objects if required. However, detailed comparisons of how much can be reported when the system is stressed to its maximum reveal that there are still deficits in reporting multiple features from two objects when compared to reporting the same number of features from a only single object. Thus, our ability to talk about our immediate visual experiences is greatly limited by the requirement to hold this information in short term memory while we are generating utterances about them that unfold over time.

A popular laboratory task to study immediate visual experience that places only very minimal demands on the participant's ability to report their experience is the **change detection task**. In typical studies of change detection, two images are presented to the participant, either in temporal alternation, or beside each other in space, and the participant is asked to detect, locate, or identify any difference they can find in the two images. If the difference in the two images is detected immediately, then the experimenter assumes that the region in which the difference occurred in each image was represented in the visual experience of the participant. On the other hand, if the difference is only detected after some time, then the experimenter assumes that the region of the difference was not immediately accessible in the experience of the participant. Furthermore, the time required to detect the change can be used to index how long it takes before the changed region of the image has entered the consciousness of the participant.

Many studies of change detection have shown that participants are surprisingly insensitive to large physical differences in two images (e.g., a large passenger jet that is missing an engine at the center of one of the two images), being able to detect the difference only after dozens of image alternations, which can take upward of ten or twenty seconds. This profound insensitivity to change is therefore often

referred to as **change blindness** to help emphasize how long it takes for some aspects of the images to enter the consciousness of the participant.

The primary factor that predicts success in a change detection task is whether the region of change is at the focus of the participant's visual **attention**, where this term refers to all those influences causing some information to be selected for more detailed processing at the expense of other information. One obvious influence on attention is the content of the participant's short-term memory. If the participant has selected three or four specific objects in a scene to be the focus of their immediate experience, then a change to one of those objects will be detected much more quickly than a change to an object that is not currently being held in short-term memory. What is selected as the focus of attention can itself be influenced by a wide range of factors, including the capture of attention by an abrupt physical event in a specific region of the image, the interest of the participant in the content of the images, the relevance of the information in the images to primary task of the participant, and the expertise of the participant for the scenes that are depicted.

Many researchers believe that change blindness occurs because participants do not have the conscious access to the information in the scenes that they assume they have. In other words, change blindness represents a gap between our experience of a rich visual world and the reality of much more impoverished representations. Some have even referred to this gap as the **grand illusion of complete perception.** However, others caution that successful change detection requires both a representation of the scene before the change and a comparison of that representation with the scene after the change. This means that change blindness could occur, either because of a failure to represent one of the scenes (as has often been assumed) or because of a failure on the part of the viewer to actively compare a region in one scene with its counterpart in the other scene. As such, **comparison failures** may have been overlooked in previous theoretical accounts of change blindness. In support of this possibility, it has been reported that viewers can recognize a previously attended object on a memory test even when they have failed

to detect a change to the same object and that viewers can recognize both the unchanged and changed object at above-chance levels when they have failed to detect any change. If this interpretation is correct, change blindness may not imply a failure to form initial scene representations as much it points to a failure to make and test predictions between scenes. But so far, this is still only a conjecture that awaits systematic study.

Paradox Three: Experiences Evolve over Time from General to Specific

Vision researchers have been puzzled for many years by the type of information that participants can report when given only a very brief glimpse of a scene. Their puzzlement comes from knowing that the visual information processed by the brain proceeds in a certain order that seems at odds with the order in which participants have conscious access to the information. In brief, the brain proceeds in its analysis of a visual image by first processing very specific details at each location in the image, before coalescing and abstracting this information in order to form higher-order concepts. This means that first there are neurons in our eye and then other neurons in the early visual centers of our brain that respond simultaneously to the minute details of what we are seeing. Only higher up in the anatomical hierarchy of processing do the neurons respond to familiar objects and entire scenes. Yet, when it comes to our experience of an image, things seem to proceed in exactly the opposite direction. We seem to be aware of the general content of a scene and its overall meaning to us well before we are aware of any of the specific details that underlie the formation of this general concept. Some vision researchers have referred to this as our ability to “see the forest before the trees” Others discuss this puzzle in terms of the “reverse hierarchy” of visual experience.

Some of the first studies of this puzzle showed participants brief glimpses of naturalistic images and asked them a variety of questions about the scenes’ content.

Although questions about any of the specific details, such as the particular shapes or colors depicted could only be answered for the small region of the scene that the participant happened to focus on in any glimpse, questions about the **gist** (meaning) of the entire scene could be answered quite accurately. For example, participants were easily able to say whether the scene depicted a seashore, a jungle scene, or the interior of an office. Participants also had a very good idea of the general **layout** (spatial arrangement) of the large objects in the scene. They even had very accurate impressions of the emotional content that was depicted, being able to say whether the scene was generally an angry or a happy one. Only with additional time and sometimes only with hard won expertise were participants able to experience the specific details in a glance.

Another way this issue has been studied involves creating images that depict different content at the level of detail and at the level of the overall configuration. Some famous art involving this principle is shown in Figure 5a, while some displays from laboratory experiments are shown in Figure 5b. In both cases, it is easier to see the larger configurations (e.g., the face in Figure 5a and the letter S in Figure 5b) than it is to see the components of the configurations (e.g., the vegetables in Figure 5a and the many letter Es in Figure 5b). Experiencing these details takes additional time. This suggests that conscious awareness of what we are seeing begins high in the anatomical hierarchy of brain processes and is only able to access the component details that are represented lower in the hierarchy with additional time and mental effort.

Conclusion

Visual experience is a construction by the neurons in our brain. To understand how this is accomplished, scientists in a number of related disciplines study three paradoxes. These include (a) that our phenomenally-rich visual experience seems to derive from relatively impoverished visual input, (b) that although our experience

may be rich in color, shape and texture, we are severely limited in our ability to make an immediate report on it to others, and (c) that our visual experiences seem to evolve over time in an order that is reversed with respect to the anatomical order of processing, allowing us relatively quick access to the more abstract information conveyed in an image, but only slower access to specific details on which these abstractions are based.

See also

Attention, Change Blindness, Inattentional Blindness, Visual Cognition, Visual Short-term Memory,

Further Reading

Biederman, I. (1981). On the semantics of a glance at a scene. In M. Kubovy & J. R. Pomerantz (Eds.), Perceptual organization (pp. 213-253). Hillsdale, NJ: Erlbaum.

Bruner, J. (1973). Going beyond the information given. New York: Norton.

Crick, F. (1994). The astonishing hypothesis. New York: Charles Schribner's Sons.

Duncan, J. (1984). Selective attention and the organization of visual information. Journal of Experimental Psychology: General, 113, 501-517.

Enns, J. T. (2004). The thinking eye, the seeing brain. NY: WW Norton.

Enns, J. T., & Austen, E. L. (2007). Mental schemata and the limits of perception. In M. A. Peterson, B. Gillam, & H. A. Sedgwick (Eds.), In the Mind's Eye: Julian Hochberg on the Perception of Pictures, Film, and the World (pp. 439-447). NY: Oxford University Press.

Hochstein, S., & Ahissar, M. (2002). View from the top: Hierarchies and reverse hierarchies in the visual system. Neuron, 36, 791-804.

Luck, S. J. & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. Nature, 390, 279-281.

Navon, D. (1977). Forest before trees: The precedence of global features in visual perception. Cognitive Psychology, 9, 353-383.

Noë, A., & O'Regan, J. K. (2000). Perception, attention and the grand Illusion. Psyche, 6(15), < http://psyche.cs.monash.edu.au/v6/psyche-6-15-noe.html >

Pinna, B., Brelstaff, G., & Spillmann, L. (2001). Surface color from boundaries: a new 'watercolor' illusion. Vision Research, 41(20), 2669-2676.

Simons, D.J. & Ambinder, M.S. (2005) Change blindness: Theory and consequences. Current Directions in Psychological Science, 14(1), 44-48.

Sperling, G. (1960). The information available in brief visual presentations. Psychological Monographs, 74 (11, Whole No. 498).

Figure Captions

Figure 1. Hermann grid (a) and the differential effect of intersections versus spaces between columns on the receptive fields of ganglion cells in the eye (b). Illusory smudges are seen at the intersections of the grid because receptive fields centered there receive more light falling onto their inhibitory surrounds.

Figure 2. Identical S-shape configurations, consisting of short line segments aligned at their endpoints (a) or short line segments at right angles to these and aligned at their midpoints (b). The fact that the S-shape is more easily visible in (a) illustrates the cooperative effects between aligned edges and the inhibitory effects of neighboring edges of different orientation.

Figure 3. An illusory square shape formed by subjective contours (a) and three illusory orange-tinged surfaces formed using the watercolor illusion. The visual system uses heuristics to interpret ambiguous images.

Figure 4. Two shapes defined by shading to illustrate perceptual heuristics based on the constraints that light generally shines from overhead (a) and that objects are usually resting on or attached to other surfaces (b).

Figure 5. The portrait of a face composed of vegetables (a) and the letter S composed of numerous copies of the smaller letter E (b). Visual experience usually begins with the larger configuration and only has access to the component details with additional time and mental effort.

Permissions

All figures were original drawings made by the author for this publication, with the exception of Figure 5A, which is a copy of Giuseppe Arcimboldo's (1527-1593), *Vertumnus.* 1590-1591. (Oil on wood). To my best understanding, because of its age, this image is in the public domain.

Figure 1

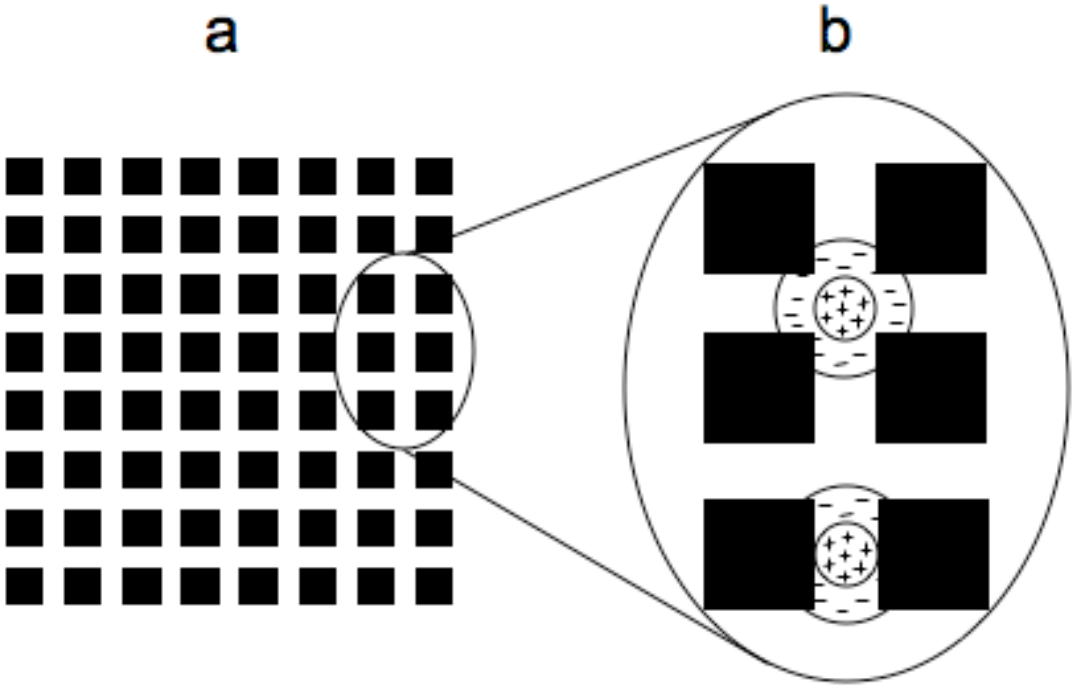

Figure 2

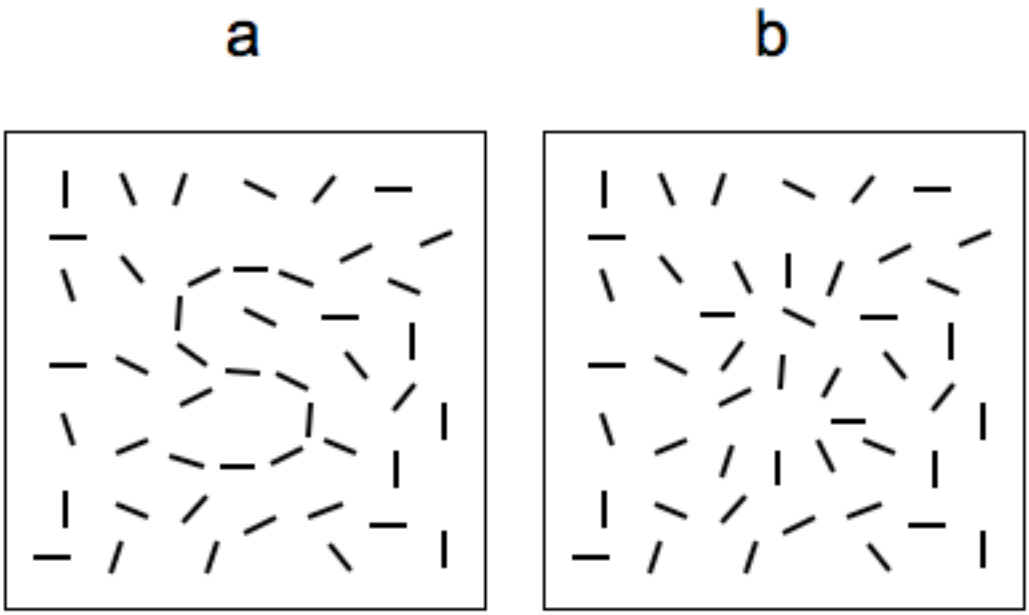

Figure 3

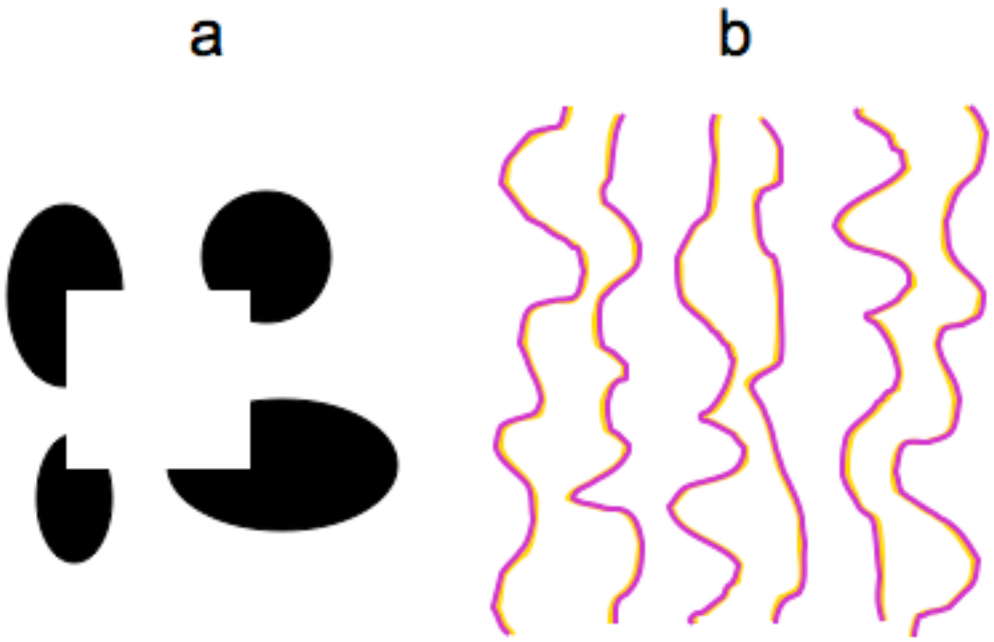

Figure 4

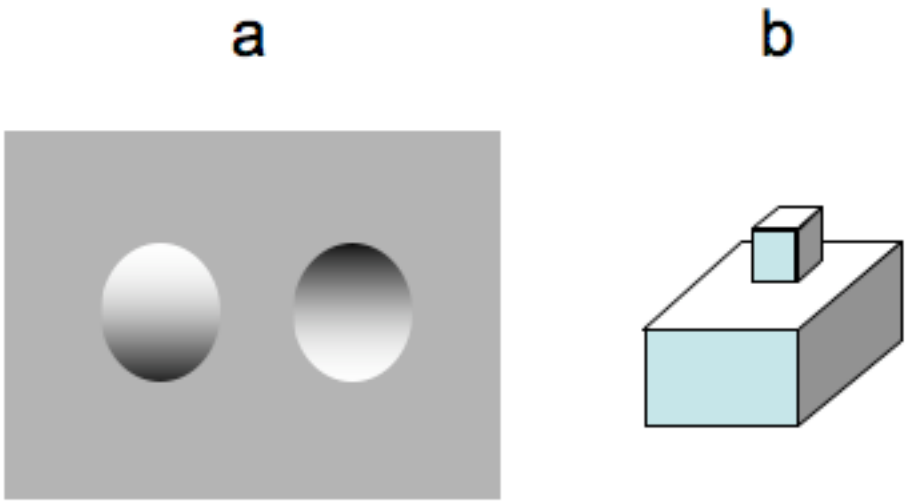

Figure 5

a

b

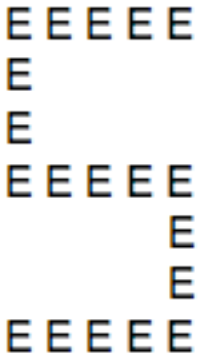

Biography

James T. Enns is a Professor at the University of British Columbia in the Department of Psychology and the Graduate Program in Neuroscience and the director of the UBC Vision Lab < http://www.psych.ubc.ca/~ennslab/ > A central theme in his research is the role of attention in human vision. This includes studies of how the visual world is represented inside and outside of focused attention, how attention changes the contents of consciousness, how perception changes with development, and how to design visual displays for optimal human performance. He has served as Associate Editor for the journals Psychological Science, Perception & Psychophysics, and Consciousness & Cognition. His research has been supported by grants from NSERC, BC Health & NATO. He has authored a book on visual cognition, The Thinking Eye, The Seeing Brain (W.W. Norton 2004), edited two volumes on the Development of Attention, coauthored a textbook, Sensation and Perception (Wiley 2004), and published numerous scientific articles. He was named a Distinguished University Scholar at UBC in 2003 and was the recipient of the Robert Knox Master Teaching Award in 2004. His PhD is from Princeton University and he is a Fellow of the Royal Society of Canada.

Color Photo of Author