

How we see art and how artists make it

by

Stephen Grossberg
Center for Adaptive Systems
Graduate Program in Cognitive and Neural Systems
Departments of Mathematics & Statistics, Psychological & Brain Sciences,
and Biomedical Engineering
Boston University
Boston, MA 02215

To appear in the multimedia book
Mobile Brain-Body Imaging and the Neuroscience of Art, Innovation, and Creativity
J. L. Contreras-Vidal, D. Robleto, J. G. Cruz-Garza, J. M. Azorin, and C. S. Nam (Eds.)
Cham, Switzerland: Springer Nature Switzerland AG

Corresponding author:
Prof. Stephen Grossberg
Center for Adaptive Systems, Room 213
Boston University
Boston, MA 02215
Email: steve@bu.edu
Web page: sites.bu.edu/steveg

Submitted: November 8, 2018

Abstract

This article provides an overview of an article entitled *How humans consciously see paintings and paintings illuminate how humans see* that I published in 2017 in the journal *ART & Perception* with my graduate student Lauren Zajac and then gave a lecture about at the 2017 conference about *Your Brain on Art* in Valencia, Spain. This article and lecture illustrate how the paintings of visual artists activate multiple brain processes that contribute to the conscious perception of their paintings. Different artists and artistic movements may focus on different combinations of brain processes to achieve their aesthetic goals. These distinct emphases allow viewers to identify the artist or artistic movement with considerable accuracy. The article and lecture summarized how neural models of how advanced brains consciously see have characterized these processes, and used the models to explain why paintings of Jo Baer, Ross Bleckner, Gene Davis, Charles Hawthorne, Henry Hensche, Henri Matisse, Claude Monet, Jules Olitski, and Frank Stella look the way that they do. Paintings by Rembrandt, Graham Rust, and Sean Williams were also mentioned along the way. These paintings illustrate brain processes that range from discounting the illuminant and lightness anchoring, to boundary and texture grouping and classification, through filling-in of surface brightness and color, to spatial attention, conscious seeing, and eye movement control. The current article provides a running commentary for a refinement of my previous lecture, with the simpler title *How we see art and how artists make it*, as a way to vividly summarize highlights of my archival article with Lauren.

Keywords: painting; neural networks; Baer; Bleckner; Davis; Hawthorne; Hensche; Matisse; Monet; Olitski; Rembrandt; Stella; complementary computing; perceptual grouping; boundary completion; surface filling-in; figure-ground perception; self-luminosity; lightness anchoring; neon color spreading; watercolor illusion; shading; texture; impressionism; Fauve; color field painting; chiaroscuro; trompe l'oeil

Introduction: From strokes to conscious percepts and back

Whenever an artist manipulates a canvas, say by applying a dab of color to a canvas, he or she immediately experiences a conscious percept of the result. This percept emerges from all the brain machinery whereby we consciously see and know about our visual world. Artists typically have no explicit knowledge about the brain processes that mediate between painterly manipulations and the resulting conscious percepts. Yet despite this intellectual chasm between action and percept, the particular interests and aesthetic sensibilities of different artists have led each of them to emphasize different combinations of these brain processes, and to thereby create their own artistic style. In the hands of a master, the results can be both astonishing and transcendently beautiful.

The corpus of works of art on two-dimensional surfaces, across time and culture, provide an incredible richness of issues that paintings elicit, both scientific and aesthetic. This article reviews several of these issues through a discussion of specific paintings by well-known artists that have been chosen to illustrate how different combinations of brain processes were used to achieve their aesthetic goals. Illustrative paintings or painterly theories by nine artists were given a unified analysis in Grossberg and Zajac (2017) using neural design principles and mechanisms that have been articulated and computationally characterized by the most advanced neural models of how advanced brains consciously see. This article also summarized, where possible, descriptions of an artist's stated goals, or summarized reviews of the artist's work written by art historians, curators, or critics.

The current article does not attempt to scientifically explain why a painting looks beautiful, or how it may arouse strong emotions. Such an analysis would require the study of how perceptual, cognitive, and *emotional* processes interact. Some promising approaches have been described whereby to understand aesthetic emotions by using mathematical models of the mind (e.g., Perlovsky, 2010). The current goal is to first try to better understand the brain mechanisms of perception and cognition whereby humans consciously see paintings, and whereby painters have achieved their aesthetic goals. Further studies of beauty and of aesthetic emotions may benefit for the considerable neural modeling literature about the brain processes that create coordinated conscious experiences of seeing, knowing, and *feeling* (e.g., Grossberg,

2013, 2017b). These more comprehensive theoretical insights would, in any case, need to build upon insights such as those described herein.

In addition, Grossberg (2017b) summarizes some (but not all!) of the basic brain processes that are needed to understand how we perceive and recognize music.

The current summary will provide comments about the numbered powerpoint slides in the lecture with the same title as the current article that can be found on this web page, along with a pdf of the lecture that I gave in Valencia in 2017.

Comments about numbered powerpoint slides

Slides 1-3. These slides introduce the lecture and raise the basic question of how various painters struggled to intuitively understand how they see in order to generate desired aesthetic effects in their paintings.

Slide 4. This slide notes that the lecture is possible due to neural modeling work that clarifies what goes on in each brain as it consciously sees, hears, feels, or knows something. The article Grossberg (2017b) provides a self-contained, non-technical summary of current modeling knowledge about how this happens. The current lecture focuses only on one aspect of how we consciously see. It also summarizes a claim concerning *why* evolution was driven to discover conscious states in the first place. This analysis begins with Slide 136. It proposes how conscious perception is used to close the loop between perception and action, in this case between manipulating a painting, seeing it, and then manipulating it again.

In brief, the lecture will explain how multiple processing stages overcome the incompleteness and ambiguities of the raw sensory data that reaches our brains. These sensory data are hopelessly inadequate for triggering effective actions that can enable us to survive in a changing world that is filled with potentially life-threatening challenges. After these processing stages do their work, the result is sufficiently complete, context-sensitive, and stable perceptual representations upon which to base effective actions. In civilized societies, these actions include the strokes that create a painting. The article hereby proposes that evolution discovered conscious states in order to mark, or “light up”, the sufficiently complete, context-sensitive, and stable perceptual representations that can support effective actions, notably feature-category resonances for consciously knowing about objects, and surface-shroud resonances for

consciously seeing them and triggering actions based upon them. These resonances will be defined and discussed below.

Slide 5. This slide summarizes some of the painters whose work will be discussed. The painters were chosen to demonstrate how different painterly effects emphasized different combinations of brain processes.

Slides 6-9. These slides emphasize that, since “*brain* evolution needs to achieve *behavioral* success,” neural models that hope to link brain to mind need to discover and model the level of brain processing that governs behavioral success. A half century of modeling has consistently shown that these are network and system levels, which is why we study neural networks.

In order to complete such a model, individual neurons must be designed and connected in networks whose emergent, or interactive, properties give rise to successful behaviors. Keeping all these levels in mind at once—behavior, network, neuron—requires an appropriate modeling language whereby to link them. Such a mathematical model makes it much simpler to understand how brains give rise to minds, not only by articulating appropriate brain design principles and mechanisms, but also by explaining the emergent properties that they generate when they interact together in response to a rapidly changing world. Unaided intuition cannot, by itself, understand these emergent properties.

Although rigorous mathematical modeling and computational analyses are needed to understand how brains give rise to minds in a way that feels inevitable, it is nonetheless possible to explain the ideas upon which these models are based using simple, self-contained, and intuitively understandable stories. That is what this lecture tries to achieve. In so doing, it clarifies that perhaps the hardest obstacle to understanding mind and brain is to know how to think about each problem. Once one is on the right path, the technical details can often be readily worked out. Finding such paths requires guidance from lots of data.

Slides 10 and 11. This perspective argues that, to deeply understand how brains work, you need to understand how evolution selects brain designs based on their behavioral success. That is why the modeling method and cycle that I have developed with many colleagues over the past 50 years always starts with behavioral data, often scores or even hundreds of experiments in a given

area of psychology. Having lots of data to guide ones thinking helps to rule out incorrect, but initially appealing, ideas.

The Art of Modeling consists in large part of figuring out how to understand these behavioral data, which one receives as static curves that plot one variable (e.g., number of correct responses at each position in a list after a fixed number of learning trials) against another (e.g., list position), as emergent properties of individual behaviors as they adapt autonomously in real time to a changing world. The results of such top-down analyses from behavioral data have always been the discovery of brain design principles that are translated into the simplest possible mathematical models. Then mathematical and computational analyses of these models are used to generate emergent behavioral properties that explain much more behavioral data than went into the hypotheses from which the model was derived. In this way, the modeling loop between behavior-to-design-to-model-to-behavior is closed.

In addition, and of critical importance, is the fact that the mathematical models always look like part of a brain. As a result, they explain a body of known brain data, as well as predict as yet unreported new brain data. Because this derivation proceeds from behavior-to-design-to-model-to-brain, it often proposes novel functional explanations of both known and unknown brain data.

Once the connection is made between behavior and brain, one can explain and predict lots of behavioral and brain data using the currently derived model. After the explanatory and predictive range of the model in its current form is understood, one can press both top-down from behavioral data, and bottom-up from brain data, to identify an additional design principle that the model does not currently embody. Then this new design principle is consistently added, “embedded”, of “unlumped” into an expanded model, and the cycle begins again, leading to a broader range of interdisciplinary data that can be explained and predicted.

This cycle has been repeated many times during the past 50 years. As a result, we now have models that can *individually* explain and predict psychological, neuroanatomical, neurophysiological, biophysical, and even biochemical data. In this sense, the classical mind/body problem is incrementally being solved.

Slides 12 and 13. After going through this modeling cycle, what is the result? Is the brain just a “bag of tricks” as even famous neuroscientists such as V. S. Ramachandran have claimed in the past? If that were the case, true theories would be impossible.

Instead, a small number of fundamental equations have sufficed to explain thousands of interdisciplinary experiments, just as in physics. A somewhat larger number of modules, or microcircuits, that are defined using these fundamental equations, are used in specialized forms to compute useful, but not universal, combinations of properties. These modules, in turn, are assembled into modal architectures for carrying out different kinds of biological intelligence. The word “modal” stands for different modalities of intelligence, such as vision, audition, cognition, emotion, and action. None of them computes all possible computable functions in the manner of a modern von Neumann computer. However, each of them is general-purpose within its own modality of intelligence, can respond adaptively to wide range of environmental challenges, and can seamlessly interact with other modal architectures to generate autonomous adaptive intelligence as we know it.

Slides 14 and 15. What principles determine how modal architectures are designed? It is here that the novel computational paradigms, and corresponding design principles, that underlie brain computing play a critical role in ensuring that we can autonomously adapt to rapidly changing environments that are filled with unexpected events. Two of these paradigms are called Complementary Computing and Laminar Computing. Together they also imply a third fundamental brain design that I call the Hierarchical Resolution of Uncertainty. It is this latter design that requires multiple processing stages before our brains can compute perceptual representations that are adequate for generating effective actions. It is because such representations need to be selectively used to generate action that conscious states “light them up” for this purpose. These are the processing stages that enable a painter to apply paint to a canvas and consciously see and appreciate his or her handiwork.

Slides 16-20. Complementary Computing asks what is the nature of brain specialization. It provides an alternative to the earlier idea that brains compute using independent modules. There are lots of specialized brain regions in the visual cortex, and at least three parallel cortical processing streams with which to activate them. However, independent modules should compute

each property—such as luminance, motion, binocular disparity, color, and texture— independently of the others. In reality, huge perceptual and psychophysical databases show that there are strong interactions between these various perceptual qualities.

Complementary Computing explains how such specialization coexists with, and indeed requires, these interactions by providing a very different answer to the question: What is the nature of brain specialization? Complementary Computing identifies new principles of uncertainty and complementarity that clarify why multiple parallel processing streams exist in the brain, each with multiple processing stages to realize a hierarchical resolution of uncertainty.

There are analogies to computationally complementary properties, such as a key fitting into a lock, and puzzle pieces fitting together, but these analogies do not explain the dynamism that is required to carry out Complementary Computing. In particular, computing one set of properties at a processing stage prevents that stage from computing a complementary set of properties. These complementary parallel processing streams are balanced against one another. This kind of balance is reminiscent of classical ideas about Yin and Yang, but again not explained by them. Instead, prescribed interactions between these streams, at multiple processing levels, overcome their complementary weaknesses and support intelligent and creative behaviors. They do so, in particular, by creating conscious visual states that can be used to guide looking and reaching behaviors, including those used to create and see paintings.

Slides 21-23. Each row in Slide 21 summarizes a pair of computationally complementary processes and the cortical streams in which they are proposed to occur. This list is not, however, exhaustive of all the complementary processes in our brains.

When one puts together the first four of them, one is led to an emerging unified theory of visual intelligence, starting at our photosensitive retinas and ending at the prefrontal cortex, or PFC. Each box in the slide functionally describes a basic process that occurs in the corresponding part of the brain, and both the What and Where cortical streams are included. The modeling work that I and my colleagues have carried out over the years to explain hundreds of interdisciplinary experiments support my hypothesis that the bottom-up, horizontal, and top-down interactions between these various processes help to overcome complementary processing deficiencies that each process would experience if it had to act alone.

Slides 24-26. These slides begin to show what it means for visual boundaries and surfaces to be complementary. Much psychophysical evidence has supported my prediction that 3D boundaries and surfaces are the basic functional units in natural vision. This prediction was first made in Grossberg (1984) and was supported by computer simulations of perceptual and psychophysical data in Grossberg and Mingolla (1985a, 1985b) and Grossberg and Todorovic (1988). I began to extend it in Grossberg (1987, 1994) to explanations and simulations of data about 3D vision and figure-ground perception using the Form-And-Color-And-DEpth (FACADE) model, and its 3D LAMINART model extension to the laminar circuits of visual cortex, with multiple PhD students and postdocs, including Rushi Bhatt, Yongqiang Cao, Nicolas Foley, Gregory Francis, Alan Gove, Simon Hong, Piers Howe, Seungwoo Hwang, Frank Kelly, Levin Kuhlmann, Jasmin Leveille, John Marshall, Niall McLoughlin, Steven Olson, Luiz Pessoa, Rajeev Raizada, William Ross, Aaron Seitz, David Somers, Karthik Srinivasan, Guru Swaminathan, Massimiliano Versace, James Williamson, Lonce Wyse, and Arash Yazdanbakhsh. The vision models were complemented by the SACCART, SAC-SPEM, TELOS, and lisTELOS models of the saccadic and smooth pursuit eye movements that occur during visual perception and planning, and invariant object category learning. A parallel but distinct line of work also developed the 3D FORMOTION model of visual motion perception, with its extensions to visually-based navigation and target tracking. See my web page sites.bu.edu/steveg for many such archival articles and https://en.wikipedia.org/wiki/Stephen_Grossberg for a list of the names of available models and the areas to which they contribute.

Visual boundaries are emphatically not just edge detectors. Rather, boundaries can form during texture pop-out, 3D shape from texture, figure-ground separation, and visual illusions, among other contexts. This versatility spares our brains from having to use specialized detectors for each of these types of stimuli, only to have to figure out at a later processing stage how to put all the information together. Such specialization cannot, in any case, work in response to natural scenes if only because edges, shading, texture, and figure-ground properties are often overlaid at the same perceptual positions in a scene.

Neon color spreading is one of the visual illusions that provides lots of useful information about the complementary properties of visual boundaries and surfaces. A typical neon-inducing image is constructed of black and blue arcs, where the blue contrast relative to its white background is smaller than that of the black contrast. When these arcs are properly arranged,

both boundary completion and surface filling-in of a neon color spreading illusion are caused. The boundary completion generates the illusory square that passed through the positions where the blue and black arcs touch. The surface filling-in causes the square to be filled with a bluish hue.

Note that three properties of boundary completion and surface filling-in are shown. The first two boundary properties are that boundaries are completed between pairs of inducers in an *oriented* and *inward* fashion. If outward completion were possible, then a single dot in an image could cause a radial proliferation of boundaries which could seriously obstruct vision. By comparison, the spread of the blue color through the square is generated by small breaks in the blue boundaries where they touch the more contrastive black boundaries. The blue color can then spread in an *unoriented* manner *outward* in all directions until it hits the square illusory boundaries. These boundary and surface properties are manifestly complementary: oriented vs. unoriented; inward vs. outward.

Slide 27. This slide shows that boundaries are completed within the interblob cortical stream from the lateral geniculate nucleus, or LGN, through V1 interblobs, V2 interstripes, and V4. The surfaces are completed in the parallel blob cortical stream through V1 blobs, V2 thin stripes, and V4. These are two of the brain's computationally complementary processing streams.

Slides 28 and 29. What does the third boundary completion property of “insensitive to direction-of-contrast” mean? This has to do with the classical distinction between seeing vs. knowing, or seeing vs. recognition. For example, in Slide 29, the lower left image shows an Ehrenstein Figure that is generated by blue lines pointing towards the center of an imagined disk. One can both see and recognize this disk because its interior is brighter than its background. This brightness difference is due to filling in of “brightness buttons” that are generated just beyond the line ends, whence this brightness spreads within the illusory circle that is also generated through the line ends.

In contrast, in response to the Offset Grating to the right of the Ehrenstein Figure, a vertical boundary is generated that passes through the line ends of the horizontal blue lines. We can *recognize* this vertical boundary, but we cannot *see* it: It is not brighter or darker, or nearer or

further, from the rest of the background. This percept shows that one can consciously recognize objects that one cannot see. There are hundreds of such amodal percepts.

One plausible answer to the question “Why do we see?” is that “We see things to recognize them”. However, we can recognize the vertical boundary that is generated by the Offset Grating without seeing it. This is thus a counterexample to the hypothesis that we see things in order to recognize them, because we can recognize this vertical boundary without seeing it. This conclusion does not deny that seeing objects does often help to recognize them, but it shows that there must be a different answer to the question “Why do we see?”

I earlier noted that, due to hierarchical resolution of uncertainty, our brains seem to have created conscious states of seeing so that we can selectively use those perceptual representations upon which to base actions like looking and reaching.

Slides 30-32. Slide 29 shows that some boundaries are invisible. Slide 30 provides one of several reasons why *all* boundaries are invisible, at least within the interblob cortical stream that generates boundaries. In particular, let us move along the circumference of the gray disk in the right figure of this slide. One passes from gray-to-white, then gray-to-black, then grey-to-white, etc. contrasts all along the circumference. These reversals of relative contrast are often found when an object is seen in front of a textured background.

If our brains only had separate boundaries that compute dark-to-light contrasts (e.g., gray-to-white) or light-to-dark contrasts (e.g., gray-to-black), then each type of boundary would have big holes in it. Brightness and color could spread through these holes during the filling-in process and thereby seriously degrade vision.

Slide 31 shows that, indeed, boundary computation begins with oriented local contrast detectors, called simple cells, that individually can respond to either a dark-to-light oriented contrast, or a light-to-dark oriented contrast, but not to both. If boundary processing ended here, there would, indeed, be big holes in the resulting boundaries.

Instead, at each position, pairs of like-oriented simple cells that are sensitive to opposite contrast polarities input to cells at the next processing stage that are called complex cells. Each complex cell can respond to both dark-to-light and light-to-dark contrasts at, and close to, its preferred position and orientation. Thus, by the time complex cells respond to the gray disk image in Slide 30, they would build a boundary at every position around its circumference.

However, because they pool signals from both polarities—that is, are insensitive to direction-of-contrast—complex cells cannot represent visual qualia like differences in relative luminance or color. Said in another way: *All boundaries are invisible!* We can experience how salient boundaries may be, but strong boundary salience does not imply a visible difference of qualia.

Slides 33-35. Despite being invisible, boundaries are extremely useful in helping us to recognize objects, especially objects that are partially occluded in a three-dimensional scene. The dashed red lines in Slide 34 illustrate where amodal boundaries of partially occluded objects may be created in order to help to recognize these objects. The abutting three rectangles in the right image of Slide 35 gives rise to a compelling 3D percept of a vertical rectangle that is partially occluding, and in front of, a horizontal rectangle. Even though we “know” that the horizontal rectangle is “behind” the vertical rectangle, we do not see it.

This property of figure-ground separation is exploited in all pictorial art, movies, and TV that use a 2D image to generate representations of 3D objects. For example, in the famous Mona Lisa of Leonardo da Vinci, her face partially occludes the background of the scene, which can nonetheless be amodally completed behind her, at least in the upper part of the painting.

Slides 36 and 37. Slide 36 shows a side view of the interior of an eye. After light passes through the lens of the eye and the retinal fluid that helps to maintain the eye’s shape, it needs to go past the nourishing retinal veins and all the other cell layers in the retina before it hits the photoreceptors. The photoreceptors that are activated by the light send signals along axons via the optic nerve to the brain.

Slide 37 shows a top-down view of the retina. It includes the fovea, which is the part of the retina that is capable of high acuity vision. Eye movements focus the retina upon objects of interest. There is also a blind spot that is as big as the fovea. Here is where the axons from the photoreceptors are bundled together to form the optic nerve. No light is registered on the blind spot.

Slides 38-40. Slide 38 illustrates how a blue line may be occluded by a retinal vein and the blind spot in multiple positions. This picture shows that even a simple image like a line may not be

completely registered on our retinas. This state of affairs raises several questions. For one, why do we not see retinal veins and the blind spot? This is true because our eyes rapidly jiggle in their orbits, even when we think that they are not moving. This jiggle generates transient visual signals from objects in the world. These transients refresh the neural responses to these objects. The veins and blind spot do not, however, generate such transients because they move with the eye. They are thus stabilized images. Hence, they fade. You may have noticed in an ophthalmologist's or optometrist's office your own retinal veins or blind spot when he or she moves a small light alongside your eye in order to examine it. That motion can create transients with respect to the borders of the veins and blind spot and makes them momentarily visible.

Another important question is this: How do we see even images like a line if they can be occluded in multiple positions? Figure 39 shows that boundary completion completes boundaries within occluded regions and surface filling-in spreads colors and brightnesses from surrounding regions to complete the surface percepts of the occluded regions within these boundaries.

The percepts that are generated across the occluded regions are constructed at higher brain regions. Because they are not provided directly by visual inputs to the retinas, they are, mechanistically speaking, visual illusions. On the other hand, we often cannot tell the difference between the regions on the line that receive their signals directly from the retina, and those that have completed boundaries and filled in colors and brightnesses. Both kinds of regions look equally "real". This raises the question in Slide 40: What do we call a visual illusion? I believe that we tend to call illusions those combinations of boundary and surface properties that look unfamiliar or unexpected, as in the case of the invisible vertical boundary that is generated by the Offset Grating in Slide 29.

Slides 41 and 42. If boundaries are invisible, then how do we consciously see? Slide 41 suggests that we see the results of surface filling-in after boundaries define the compartments within which lightness and color spread. The pdf of Slide 21 summarizes the fact that the stimulus that generates the percept called the Craik-O'Brien-Cornsweet Effect has the same background luminance, but a less luminous cusp abutting a more luminous cusp in the middle of the image (see the red line labeled stimulus). These two regions are surrounded by a rectangular black frame. The percept is, however, one of two uniform gray regions (see the blue line labeled percept). This percept may be explained by the fact that the boundaries which surround the gray

regions restrict filling-in to each of them. Then filling-in of the less luminous cusp in the left region leads to the percept of a uniformly darker gray region than does the filling-in of the more luminous cusp in the right region. A more complete explanation, and simulations, of this percept is given in Grossberg and Todorovic (1988), as well as of the very different percept that is seen when the black region is replaced by a gray region that matches the gray of the stimulus background. Many other brightness percepts are also explained and simulated within that article.

Slides 43 and 44. We can now understand the last computationally complementary property of boundary completion and surface filling-in that is shown at the bottom of Slide 43. As I earlier noted, “insensitive to direct-of-contrast” can also be summarized by the statement that “all boundaries are invisible”. “Sensitive to direction-of-contrast” can be recast as “filling-in of visible color and lightness” since filled-in surfaces are what we can consciously see. Slide 44 can now summarize my prediction from 1984 that all boundaries are invisible in the interblob cortical stream, whereas all visible qualia are surface percepts in the blob cortical stream.

Slides 45-49. We can now begin to apply these ideas to provide a better mechanistic understanding of the aesthetic struggles of various painters. Let us start with Henri Matisse. Slide 45 raises the provocative question: Did artists like Matisse know that all boundaries are invisible? Consider his painting, *The Roofs of Collioure*, from 1905 to understand a sense in which the answer to this question is Yes. Note that Matisse constructed much of this painting using patches of color to suggest surfaces. As noted in Slide 46, Matisse went through a life-long struggle to understand “the eternal conflict between drawing and color”. He wrote that “Instead of drawing an outline and filling in the color...I am drawing directly in color”. The bottom image in this figure illustrates what this means. Using color patches on canvas, Matisse created amodal boundary webs that organized these color patches in surfaces. By not “drawing an outline” to define these surfaces, Matisse ensured that he did not darken these colors. Generating vivid colors in their paintings was one of the goals of the Fauve artistic movement to which some of Matisse’s paintings contributed.

Thus, as Slide 48 notes, when discussing *The Roofs of Collioure* with your friends, you can tell them that it illustrates Complementary Computing in art because it generates so many invisible boundary representations to define its colorful surfaces.

Another Matisse painting from 1905, the Open Window, Colloure, is illustrated in Slide 49. This painting brilliantly combines surfaces that are created with sparse surface color patches, as well as surfaces that are rendered with continuously applied paint. Both types of surfaces blend together into a single harmonious scene.

Slides 50-52. Many artists have experienced Matisse's struggle to be "drawing directly in color". They include plein air painters such as the founder of the Cape Cod school of art, Charles Hawthorne, and his most famous student, Henry Hensche. These slides summarize their approaches to painting in their own words.

Slides 53-56. The next slides illustrate the same theme with a quote from the famous French Impressionist painter, Claude Monet, and the well known painting of Georges Seurat, Femmes au bord de l'eau. This latter painting illustrates that there are both large-scale boundaries that group regions of this painted scene, and small-scale boundaries that surround the individual color patches with which the painting was created. We can see both scales as our attention focuses upon different aspects of the painting.

Slides 57-59. The same brain processes of boundary completion and surface filling-in that enable use to appreciate Impressionist paintings also enable use to process natural images and images that are derived from artificial sensors. Slides 57 and 58 illustrates this by showing how these processes can transform a Synthetic Aperture Radar, or SAR, image into one that can be easily interpreted by human observers. SAR is the kind of radar that can see through the weather, and is thus very useful in remote sensing and international treaty verification applications where SAR images are taken from satellites and other airborne observers. The Input image in the upper left corner contains five orders of magnitude in the radar return. This huge dynamical range is hard to represent on a powerpoint slide, and much of the image is darkened relative to its maximal intensity due to the existence of sparse, but very high intensity, pixels in it. The Feature image results from a process of "discounting the illuminant", or compensating for variable intensities or gradients of illumination that could otherwise prevent the extraction of information about object form. This process normalizes the Input image without distorting its relative intensities. Despite

this normalization process, the resulting images still exhibit its individual pixels, just as in the painting by Seurat.

The Boundary image in the lower left corner detects and completes boundaries around and between sets of pixels with similar contrasts. Then the Feature image fills-in within the Boundary image. This result is the Surface Filling-In image in the lower right corner. One can here see a road that runs diagonally downward from the middle of the top of the image towards its lower right. One can also see individual posts along this road, the highway that runs beneath it, and the trees and shadows that surround the roads. The pixels have been largely replaced by shaded object forms.

Slide 59 shows that the filled-in surface representation in Slide 58 is the result of processing the Input image using three different spatial scales: small, medium, and large. The small boundary scale detects local image contrasts best, such as the individual posts on the road. The large boundary scale detects more global features, such as the collinear structure of the road. Surface filling-in occurs within the completed boundaries at each of these three scales. The final image is a weighted sum of the three Surface Filling-In images in the bottom row.

Slides 60-62. Slide 60 asks *how* do our brains compute boundaries inwardly and in an oriented fashion between pairs or greater numbers of approximately collinear inducers with similar orientations?

Slide 61 proposes that the cortical cells which complete boundaries obey a property that I have called the *bipole* property. This name describes the fact that these cells receive signals from nearby cells via receptive fields that have two branches, or poles, on either side of the cell body. Suppose, for example, that a horizontal edge, as in one of the pac men of a Kanizsa square stimulus, activates such a cortical cell (shown in green). It then sends excitatory signals via long-range horizontal connections (in green) to neighboring cells. These signals do not, however, activate these neighboring cells because inhibitory cells (in red) are also activated by the excitatory signals. These inhibitory cells inhibit the cells that the excitatory cells are trying to excite. The excitatory and inhibitory signals are approximately the same size, so the target cell cannot get activated. It is a case of “one-against-one”.

Slide 62 shows the case in which an entire Kanizsa square is the stimulus. Now there are two pac men that are like-oriented and collinear on each side of the stimulus. Consider the pac

men at the top of the figure. Each of them can activate a cell whose long-range excitatory connections try to activate intervening cells. As before, they also activate inhibitory interneurons that try to inhibit these target cells. Why, then, does not the total inhibition cancel the total excitation, as before?

This does not happen because the inhibitory interneurons also inhibit each other (see red connections). This recurrent inhibition converts the network of inhibitory interneurons into a recurrent, or feedback, competitive network. I proved in Grossberg (1973) that such a network tends to normalize its *total* activity. Thus, no matter how many inhibitory interneurons get activated, their total output remains approximately the same. The total inhibition to the target bipole cell thus does not summate like the excitatory signals do as more inhibitory cells are activated. This is thus a case of “two-against-one” so that the bipole cell can get activated if two or more approximately like-oriented and collinear neighboring cells send signals to it. This explains why boundary completion occurs inwardly and in an oriented manner from two or more neighboring cells, as noted in Slide 29.

Slide 62 also includes, at its upper right corner, a schematic way to represent the longer-range excitatory (in green) and shorter-range inhibitory (in red) effects on a bipole cell’s firing.

Slides 63-66. I predicted the bipole property in 1984. That same year, a famous article was published in *Science* by von der Heydt, Peterhans, and Baumgartner (1984) that provided experimental support for the prediction in cortical area V2; see Slide 27. Slide 63 shows that, either direct excitatory inputs to a bipole cell body, or similarly oriented excitatory inputs to both “poles,” or receptive fields, of a bipole cell, are needed to activate it. Moreover, an input to a receptive field is still effective in activating the cell if it is moved around within it.

Slide 64 shows that additional evidence for this kind of horizontal activation of cells in cortical area V1, which is the cortical area that feeds into V2, and which itself receives inputs from the Lateral Geniculate Nucleus, or LGN; see Slide 27. Both the longer-range excitatory influence (in blue) and the shorter-range inhibitory influence (in red) were found both in psychophysical and neurophysiological experiments by Kapadia, Ito, Gilbert, and Westheimer (1995). These excitatory effects are, however, of shorter range than they are in V2, and typically modulate, or sensitize, V1 cells, rather than fire them.

Slide 65 shows some of the anatomical evidence for cells with long-range oriented horizontal connections.

The top left image in Slide 66 shows the oriented bipole cell receptive field that Ennio Mingolla and I used to simulate boundary grouping and completion properties in 1985. The dot at the center represents a bipole cell body and the short lines to either side represent the relative strengths of inputs to the cell's two receptive fields. The orientation of a line in one of these receptive fields corresponds to the relative strength of the connection with that orientation and that position to the bipole cell. Note that inputs can be received from both collinear and nearly collinear positions and orientations, with the most collinear positions and orientations receiving the largest inputs. The upper right image represents psychophysical data of Field, Hayes, and Hess (1993) that support bipole cell properties. The two images in the bottom row represent the bipole receptive fields that were used in modeling studies by two sets of other authors.

Slides 68-70. This part of the talk introduces some of the main concepts and mechanisms of Laminar Computing which, as Slide 68 notes, tries to clarify why are neocortical circuits are organized into layers of cells, often six characteristic layers in perceptual and cognitive cortices. Slide 69 depicts a simplified diagram of the circuits in cortical layer 2/3 that carry out perceptual grouping using long-range, oriented, horizontal excitatory connections, supplemented by short-range disynaptic inhibitory interneurons, as summarized in Slides 61-66. Slide 70 asks what happens before layer 2/3. In particular, how do inputs reach layer 2/3?

Slides 71-73. Slide 71 notes that oriented local contrast detectors, called simple cells, are the first cortical stage at which cells fire in response to preferred orientations at their preferred positions and spatial scales. In addition, each simple cell can respond to either an oriented dark-to-light contrast or an oriented light-to-dark contrast, but not both. Slide 72 notes that simple cells are not sufficient, as I already noted when discussing Slide 30. As already noted in Slide 31, Slide 73 reminds us that simple cells of like orientation and position, but opposite contrast polarities, add their output signals at complex cells.

Slides 74-80. Slide 74 notes that complex cells are also not sufficient because they do not respond adequately at line ends or corners. Indeed, as Slide 75 remarks, multiple processing stages are needed to accomplish this hierarchical resolution of uncertainty.

Slide 76 illustrates what goes wrong if only simple and complex cells process line ends. At a bar end, these oriented cells can respond at each position, as illustrated by the red lines in the left image. However, they cannot respond at a line end, as illustrated by the gap in the red boundary there. This problem occurs for every choice of simple cell scale. One just needs to choose the width of the line accordingly. Slide 77 asks: Who Cares? Why is this a problem in the first place?

Slide 78 shows that it is, in fact, a very serious problem because color could flow out of every line end during the process of surface filling-in, thereby leaving the scenic representation awash in spurious color.

Slide 79 summarizes the problem that needs to be solved: Somehow the brain needs to create a line end, called an *end cut*, after the stage where complex cells act. After the end cut forms, color will be contained within the line end. Slide 80 emphasizes that the process which creates end cuts carries out a context-sensitive pattern-to-pattern map, not a pixel-to-pixel map, since it would be impossible, looking just at a pixel with no boundary, to decide if it needs to be part of an end cut, or just left alone because nothing is happening in the scene at that pixel.

Slide 81. Slide 81 depicts a circuit of simple, complex, and hypercomplex (or endstopped complex) cells that are capable of generating end cuts. The hypercomplex cells respond in two stages. The first competitive stage is defined by an on-center off-surround, or spatial competition, network. Using this network, each complex cell excites like-oriented hypercomplex cells at its position while inhibiting like-oriented hypercomplex cells at nearby positions. In addition to receiving these excitatory and inhibitory inputs, these hypercomplex cells are also tonically active; that is, they are activated even in the absence of external inputs, due to an internal source of activation.

In the absence of inputs from the first competitive stage, firing of the hypercomplex cells due to their tonic activation is inhibited by the second competitive stage, which is realized by a competition between hypercomplex cells at the same position that are tuned to different orientations. Maximal inhibition is delivered between hypercomplex cells that are preferentially

tuned to perpendicular orientations. When all the hypercomplex cells receive only tonic activation, they can inhibit each other equally using this orientational competition.

Slide 82. Slide 82 explains how end cuts are created at the end of a vertical black line on a white background. Near the end of the vertical line, its vertical edges can activate vertical complex cells which, in turn, can activate vertical hypercomplex cells at its position, and inhibit vertical hypercomplex cells at nearby positions, including positions beyond the end of the line. The inhibited vertically oriented hypercomplex cells remove inhibition from other oriented hypercomplex cells at the same positions, and do so maximally at perpendicular orientations. The activities of these other cells are thereby disinhibited. That is, their tonic activation can now drive them to fire. An end cut can hereby form.

Slide 83. Slide 83 shows a simulation of how complex and hypercomplex cells respond to a line end. The line end is shown in gray in both images. The lengths of the oriented lines are proportional to the responses of the corresponding cells with those positions and orientational preferences. Note in the complex cell responses in the left image that, despite strong vertically, and near vertically, oriented responses along the vertical sides of the line, there are no responses at the bottom of the line. This is due to the elongated shape of oriented simple and complex cells. In the current simulation, the receptive field size is shown by the dark dashed lines.

The hypercomplex cell responses in the right image of Slide 83 show a strong end cut that is perfectly aligned with the bottom of the line end (hyperacuity!) but also generates responses at multiple nearly horizontal orientations (fuzzy orientations). These near-horizontal hypercomplex cell responses result from the near-vertical complex cell responses.

Slides 84-86. These slides all illustrate some of the consequences of the need to create end cuts, and of their resulting properties.

Slide 84 notes that some kinds of fonts build in their own end cuts, in the form of serifs.

Slide 85 notes that the fuzzy orientations that occur as part of end cuts allow lines that are not perfectly parallel to nonetheless generate emergent boundaries by cooperation among their end cuts.

Slide 86 notes that the global grouping that forms through line ends may, or may not, go through their preferred perpendicular orientations. The boundary that ultimately forms has the most support from all the inducers with which it can group.

Slides 87 and 88. Slide 87 reminds us that all of these possibilities are due to the fuzzy receptive fields of individual bipole cells. This state of affairs raises the question: Why are not all the groupings that form using fuzzy bipole cells themselves fuzzy, which would cause a significant loss of acuity if it were true?

Slide 88 suggests that a fuzzy band of possible groupings often does form initially, and that this is a good property: If bipole cell receptive fields were too sharply defined, then there would be a close-to-zero probability that a grouping could ever get started. Having gotten started, then the challenge is to choose the grouping with the most evidence, while suppressing weaker groupings. This is done using another hierarchical resolution of uncertainty.

Slides 89-92. Slide 89 notes that sharp boundaries emerge from fuzzy bipole cells due to interactions within the larger network of which bipole cells form a part.

Slide 90 illustrates this property. Images (a), (c), (e), and (g) represent the inputs to such a network. Each line in these images is proportional to the size of the input to a cell centered at the middle of the line and with the orientational preference of the line. Thus, every input is composed of a “bar” of vertical features. The inputs differ only in whether or not the bars are aligned in rows, columns, or both. In (a), only the columns are aligned. In (c), both columns and rows are aligned. In (e), only the rows are aligned. And in (g), the rows are aligned and closer together.

Images (b), (d), (f), and (h) depict the steady-state responses of the bipole cells in this network. In (b), vertical boundaries are created between the bars. In (d), vertical and horizontal boundaries are created. In (f), horizontal boundaries are created. And in (h), both horizontal and diagonal boundaries are created, even though there are no diagonal orientations in the inputs. These simulations illustrate that the network is sensitive to the colinearity and orientations of input inducers, and that sharp boundaries can be completed using fuzzy bipole cell receptive fields. The simulation in (h) also shows how emergent diagonals can be created if there is enough evidence for them in the input inducers, just as they are in response to the bottom display

in Slide 86. The rows needed to be brought closer together for this to happen so that they fell within the span of the diagonally oriented bipole cell receptive fields.

Slide 91 includes images that induce percepts which illustrate the properties of the simulations in Slide 90. In response to the upper left image of an E that is composed of smaller A's, the top horizontal boundary of the E groups diagonal orientations of the A boundaries. The top horizontal boundary of the S emerges from the perpendicular line ends of the H's, whereas the right vertical boundary of the S emerges from collinear grouping of the right sides of the H's.

These properties have inspired works of art. Slide 92 shows a typography portrait of Sean Williams in which all the facial features and the hair exploit these properties of boundary completion.

Slides 93-98. The next few slides show how the processes that have already been reviewed can explain the percept of neon color spreading. Slide 94 depicts a neon color spreading image that is composed of black crosses abutting red crosses. In this image, the contrast of the red crosses with respect to the white background is smaller than the contrast of the black crosses with respect to the white background. In response to this image, one of several percepts can be perceived. One can either perceive red neon color filling local shapes around the individual red crosses, such as diamonds or circles, or one can perceive diagonal streaks of color passing through a collinear array of red crosses.

Slide 95 depicts how neon color can appear to spread beyond a red cross and be contained by the illusory circle that is induced where the black and red regions touch. Let us now see how the first steps in generating a neon percept are caused in the simple-complex-hypercomplex network of Slide 96.

Slide 97 considers what happens where a pair of collinear black and red line ends touch. Vertically oriented complex cells respond along their vertical boundaries. Because the black-to-white contrast is larger than the red-to-white contrast, the complex cells that are along the black line end become more active than those along the red line end. The black vertical complex cells inhibit red vertical hypercomplex cells more than conversely near where the two line ends touch, via the first competitive stage. As a result, these red boundaries are inhibited, or at least significantly weakened, thereby causing a hole, or weakening, in them that is called an *end gap*. Red color can spread outside the red crosses through these end gaps during surface filling-in.

In addition to causing end gaps, the weakening of the red vertical hypercomplex cell activities disinhibits other oriented hypercomplex cells at those positions, especially horizontal hypercomplex cells, thereby creating an end cut, just as in the case of the line end in Slides 82 and 83.

After these end cuts form, they can use the bipole cells that they activate to create an emergent boundary that best interpolates the end cuts, as illustrated by Slide 98. The red color that spreads outside the red crosses is blocked from spreading beyond this circular illusory boundary.

Slides 99-101. We can now apply these insights to better understand how various paintings look, starting with the paintings of Jo Baer. Slide 100 shows a group of three of Jo Baer's paintings side-by-side. All of them have a black border. Within this border is a less contrastive border with a specific color: red, green, or blue, from left to right. The percepts show redish, greenish, and bluish hues spread throughout the intervening canvas. How does this percept happen?

The main effect can be explained by the spatial competition of the first competitive stage, followed by surface filling-in. The black-to-white and black-to-red contrasts are larger than the red-to-white contrasts in the leftmost image. As a result, the red-to-white boundary is weakened, so red color can spread through the interior of the canvas. The same holds true for the green and blue contrasts.

A more vivid version of this effect was developed by Baingio Pinna, who calls it the watercolor illusion. In the image in Slide 101, there are four closed regions in which a dark blue line abuts a light blue line, which encloses a white interior region. The percepts within these regions is one of light blue color filling their interiors. This happens for the same reason that the Joe Baer effects do, because the dark blue contrast with respect to both the white background and the light blue contrast, is larger than the light blue contrast with respect to the white background. The effect is made stronger by using corrugated, or curvy lines, whose surface area relative to the surrounded white interiors is much larger than straight lines would allow, thereby creating many more positions at which light blue color can flow within the weakened boundaries to fill the white interiors.

Slides 102-104. Slide 102 calls attention to the fact that the bluish regions also seem to bulge slightly in front of the white backgrounds that surround them. This may be explained as a special case of how cells with multiple receptive field sizes, or spatial scales, influence how we see objects in depth. Slide 103 shows more examples of this using shaded images that create compelling percepts of objects in depth. These techniques are called *chiaroscuro* and *trompe l'oeil*. Slide 104 notes that similar effects make many shaded and textured objects in 2D pictures appear to have a 3D shape, and that responses of receptive fields with multiple sizes can create form-sensitive webs of boundaries that control filling-in of surfaces at multiple depths.

Slides 105 and 106. Slide 105 describes one factor that helps to explain how this happens. As an object approaches an observer, it gets bigger on the retina. As a result, other things being equal, a larger retinal image is closer. Slide 106 notes that smaller scales can respond better to small scales, whereas larger scales can respond better to larger scales so that, other things being equal, bigger scales can be associated with nearer depths during perception-action cycles. A big image on the retina is not, however, always due to a nearer object because a very larger object far away, and a smaller object nearby, can both generate retinal images of the same size. How this confound is resolved is explained in Grossberg (1994, 1997).

Slides 107-113. This series of slides describes some of the processes that enable an object like a shaded ellipse in a 2D picture to generate a compelling percept of a 3D ellipsoid. Slide 107 notes that, if boundaries were just edge detectors, there would be just a bounding edge of the ellipse. Slide 108 shows how the ellipse would then look after filling-in occurs. It would have a uniform gray lightness and would look flat. We know, however, from Slide 71 that simple cells are oriented local contrast detectors, not just edge detectors.

Because of this, as Slide 109 notes, different size detectors generate dense form-sensitive webs of boundaries, that I have called “boundary webs” for short, at different positions and depths along the shading gradient. Slides 110-112 show that increasingly large receptive fields are sensitive to broader bands of shading, starting from the bounding edge and working towards the ellipse interior. Other things being equal, the small scales signal “far”, larger scales signal “nearer”, and the biggest scales signal “near”, other things being equal.

Slide 113 notes that the boundary web corresponding to each scale captures the gray shading in the small form-sensitive compartments of that web. We see this pattern of shading as it is distributed across all the depths, thereby generating a shaded percept in depth.

Slide 114. This slide notes that this kind of process has succeeded in quantitatively simulating psychophysical data about human judgments of depthfulness using percepts of shape-from-texture as illustrative examples. Although all of the five disks in this slide are composed of spatially discrete black shapes on a white disk, the ones to the left appear to have a rounded shape in depth, whereas those to the right appear to be increasingly flat. These percepts can be quantitatively simulated using multiple-scale boundary webs and the surfaces whose filling-in they control.

Slides 115 and 116. Coming back in Slide 115 to the watercolor illusion, we can now explain its bulge in depth as a consequence of a multiple-scale boundary web, albeit one that is generated by just a few abutting corrugated lines of decreasing contrast. The images in Slide 116 also generate multiple-scale boundary webs but use gradual changes in contrast to induce them, so that more scales can be involved, leading to more gradual perceived changes in depth.

Slides 117-120. This next group of slides proposes why paintings by Monet of the Rouen cathedral at different times of day lead to different conscious percepts. In Slide 118, the cathedral was painted at sunset when lighting was almost equiluminant across most of the pointing. As a result, color, rather than luminance, differences defined most of the boundaries, which were correspondingly weakened. Fine architectural details were not represented, so that coarser and spatially more uniform boundary webs were created, thereby leading to less perceived depth in the painting.

Slide 118, in contrast, shows the cathedral in full sunlight that is very non-uniform across the painting, thereby creating strong boundaries due to both luminance and color differences. Due to the increased amount of detail, the boundary webs that form are finer and more non-uniform, leading to a more depthful percept.

Slide 119 emphasizes another consequence of full sunlight by marking some of the T-junctions that are now clearly visible in the painting, leading to additional cues to perceiving

relative depth, as in the percept of a partially occluded rectangle shown in red in this slide, and further discussed in Slides 34 and 35.

Slides 121-123. This group of slides considers quite a different combination of painterly properties, as illustrated by the color field paintings of Jules Olitski. Slide 122 summarizes four of these “spray” paintings, so called because of the method that was used to create them. Slide 123 contrasts the percepts created by these spray paintings with those of Monet and other Impressionists. In the spray paintings, there are no discrete colored units (or at least very few), and no structured color or luminance gradients. Instead, diffuse boundary webs are spread over the entire surface. When they fill in, the resulting percepts are of a space filled with a colored fog and a sense of ambiguous depth. The quote of Olitski summarizes his intention to create this kind of effect.

Slides 124-135. These slides explain how Ross Bleckner managed to create self-luminous effects in his paintings. Slide 126 claims that at least two different processes can create these effects: Boundary web gradients and lightness anchoring.

Slide 127 presents some examples of how a picture can seem to glow if boundary web gradients exist; that is, if the shading that creates boundary webs varies systematically across space in a way that enables brightness to spread across them into regions where it can be trapped. The four images in the upper left corner illustrate this. These four images, working from left to right in the top row, and then from left to right in the bottom row, have increasingly steep boundary web gradients. The steepest gradients inhibit stronger boundaries to more completely inhibit the weaker boundaries near to them, allowing more brightness to flow beyond them. This brightness summates in the interior square, thereby creating an increasing bright result that, in the final square, appears self-luminous.

The right column of Slide 127 shows a similar effect in its top row with the example of the double brilliant illusion. The rows beneath that summarize computer simulations using the Anchored Filling-In Lightness Model (aFILM) that I developed with my PhD student, Simon Hong. More will be said about aFILM in the next few slides, since it can explain the brightening effects due to boundary web gradients, as well as those due to lightness anchoring.

In the left pair of images in the bottom row, two vases are shown side by side. The rightmost vase looked matte, or dull. A highlight was manually attached to this dull vase to create the vase in the left image. Now the entire vase looks glossy. This can be explained by the fact that the highlight includes luminance gradients that match the shape of the surrounding vase. The resulting boundary web is thus assimilated into these vase boundaries, thereby allowing brightness to spread across the vase. Beck and Prazdny (1987), who reported this percept, also rotated the highlight and removed its luminance gradients. Both effects prevented the rest of the vase from looking glossy, as would be expected because the brightness could not flow into other shape-sensitive boundary webs of the vase.

Slide 128 asks what is lightness anchoring, while Slide 129 reminds us that we have only considered so far how discounting the illuminant preserves the *relative* activities of luminance values, without saturating, as they are converted into perceived brightnesses. Lightness anchoring demands more.

As Slide 129 notes, it asks how the *full dynamic range* of a cell is used, not just how the cell can compute relative activity values. Another way of saying this is to ask: How do our brains compute what is perceived to be white in a scene?

One hypothesis about how white is perceived, noted in Slide 130, was proposed by the great American psychologist, Hans Wallach, who suggested that the highest luminance in a scene is perceived as white, the so-called HLAW rule. Slide 131 shows that this rule sometimes works, as in the top row of images. However, the bottom row of images shows that, if there is a very intense light source in a scene, renormalizing it to make the light source white can drive the rest of the scene into darkness.

Simon Hong and I realized that if one, instead, computes the *blurred* highest luminance as white (BHLAW), then that problem can be avoided, as shown by the simulations in Slide 132.

Slides 133 and 134 illustrate how the BHLAW rule works. Slide 133 shows a cross-section of a luminance profile in green, and the spatial kernel that defines the BHLAW rule in red. In this situation, the width of the luminance step is considerably narrower than that of the blurring kernel. As a result, when this scene is anchored to make the blurred highest luminance white, the maximal brightness of the step is more intense than white. It therefore appears to be self-luminant.

In contrast, if as shown in Slide 134, if the luminance step in a scene is at least as wide as the blurring kernel, then when the scene is anchored to make the blurred highest luminance white, the entire luminance of the step is seen as white.

Returning now to look at different examples of Bleckner's paintings, we can see that the small bright regions look self-luminous because of lightness anchoring, whereas larger spatial luminance gradients look self-luminous due to escape of brightness, as in the images of Slide 127.

Slides 136-138. None of the above results would be as compelling as they are if we had no idea what goes on in each of our brains when we consciously see something, whether it is another person or a painting. Fortunately, there has been considerable progress during the past 40 years to incrementally understand how and why we become conscious. These slides summarize a definition of the Hard Problem of Consciousness that expresses these issues. My Grossberg (2017b) article about the Hard Problem was published Open Access and is also available on my web page sites.bu.edu/steveg.

Slides 139-144. These slides summarize the hypothesis that our brains "light up" to embody a conscious state when they go into a resonant state. Slide 142 proposes that, more generally, "all conscious states are resonant states". As Slide 143 notes, not all brain dynamics are resonant, so consciousness is not just a "whirl of information processing. Slide 144 provides a non-technical definition of what a resonant state is.

Slide 145. This slide notes that conscious states are part of adaptively behavioral capabilities that help us to adapt to a changing world. Conscious seeing, hearing, and reaching help to ensure effective actions of one kind or another. This lecture does not explain all the brain machinery that explains *why* evolution may have been driven to discover conscious states. Grossberg (2017b) does attempt to do this.

In brief, that article argues that evolution was driven to discover conscious states in order to use them to mark perceptual and cognitive representations that are complete, context-sensitive, and stable enough to control effective actions. This link between seeing, knowing, consciousness, and action arises from the fact that our brains use design principles such as

complementary computing, hierarchical resolution of uncertainty, and adaptive resonance. In particular, hierarchical resolution of uncertainty shows that multiple processing stages are needed to generate a sufficiently complete, context-sensitive, and stable representation upon which to base a successful action. Using earlier stages of processing could trigger actions that lead to disastrous consequences.

Slides 37-39 already illustrated this problem in the case of visual perception. How, for example, could you look at a part of a scene that was occluded by the blind spot? As summarized in Slide 39, processes like boundary completion and surface filling-in at higher processing stages are needed to overcome these occlusions. This is one example of hierarchical resolution of uncertainty. After a sufficiently complete surface representation is generated, a resonance develops that marks this representation as an adequate one upon which to base looking and reaching.

Slides 146 and 147. Slide 146 focuses on this question for the case of seeing and reaching. Slide 147 asks: What is this resonance? It proposes that a surface-shroud resonance “lights up” surface representations that are proposed to occur in prestriate visual cortical area V4. Surface-shroud resonances are predicted to occur between V4 and the posterior parietal cortex, or PPC, where a form-fitting distribution of spatial attention occurs in response to an active surface representation, and begins to resonate with it in the manner that I will explain in Slides 154-157.

Slide 148. This slide proposes that, just as a surface-shroud resonance supports conscious seeing of visual qualia, a feature-category resonance supports conscious recognition of, or knowing about, visual objects and scenes.

Slides 149-153. These slides briefly describe how feature-category resonances are proposed to be generated using mechanisms and circuits of Adaptive Resonance Theory, or ART, which models how we can learn to attend, recognize, and predict objects and events in a changing world without being forced to forget things that we already know just as quickly. In other words, ART proposes a detailed mechanistic solution of the brain processes whereby our brains solve the *stability-plasticity dilemma* that is summarized in Slide 149; namely, how can we learn quickly without being forced to forget just as quickly? Because ART works, it has been used in

many large-scale applications to engineering and technology that need these properties. Some of these applications are listed in Slide 151, including the use of ART by the Boeing company in a parts design retrieval system that was used to design the Boeing 777.

Slide 152 notes that top-down attentive feedback encodes learned expectations that dynamically stabilize learning and memory. ART models the neural networks that embody how top-down attention works, and how top-down expectations are learned.

Slide 153 summarizes how a feature-category resonance develops between an attended pattern of features, called a *critical feature pattern* (depicted in light green), and an active recognition category at the next processing stage. The reciprocal bottom-up and top-down excitatory signals synchronize, amplify, and prolong cell activations. During such a resonance, the adaptive weights, or LTM traces, in both the bottom-up adaptive filters and the top-down expectations can learn to selectively fire the active critical feature pattern and category when a similar input pattern is experienced in the future. It is because such a resonance triggers learning that I have called the theory *Adaptive Resonance Theory*.

Slides 154-158. These slides provide some basic information about surface-shroud resonances. In particular, Slide 155 notes that an attentional shroud is a surface-fitting distribution of spatial attention. I have predicted that such a shroud enables learning of view-invariant object categories by modulating category learning within the What cortical processing stream. How this is proposed to happen is explained Grossberg (2017b). Here I focus on related issues.

Slide 156 illustrates a one-dimensional cross-section of a simple scene in which two luminous bars occur, the left one a little more luminous than the right one. Both bars send topographic excitatory bottom-up signals to the spatial attention region, where they trigger a widespread spatial competition for attention.

In addition, as Slide 157 summarizes, the active spatial attention cells send topographic excitatory top-down signals back to the surfaces that activated them. The totality of these interactions defines a recurrent on-center off-surround network whose cells obey the membrane equations of neurophysiology, also called shunting interactions. I mathematically proved in Grossberg (1973)—see also the review in Grossberg (1980)—how such a network can contrast-enhance the attention focused upon the more luminous bar while also inhibiting the attention focused on the less luminous one. Indeed, such a network tends to normalize the total activity

across the network, so by increasing attention to one bar, attention to the other bar is automatically diminished.

The net effect of these recurrent interactions is a surface-shroud resonance. Due to the top-down excitatory signals, the attended surface appears to have greater contrast, a property that has been reported both psychophysically and neurophysiologically.

Slide 158 summarizes the claim that an active surface-shroud resonance means that sustained (via the resonance) spatial attention is focuses on the object surface.

Slide 159. I want this slide to stand alone because it summarizes the claim, which is critical to this section of the talk, that a surface-shroud resonance supports conscious seeing of an attended object, in particular a painting. The talk does not summarize the large amount of psychological and neurobiological data that are consistent with this claim, but Grossberg (2017b) does do this.

Slides 160-162. Slide 160 summarizes what kinds of resonances support knowing vs. seeing. A surface-shroud resonance, with the shroud in posterior parietal cortex (PPC), supports conscious seeing, whereas a feature-category resonance, with the category in inferotemporal cortex (IT), supports knowing. We can know about a familiar object when we see it because both resonances can synchronize their activities via shared circuits in prestriate visual cortical areas such as V2 and V4.

Slide 161 notes that, if the knowing resonance is damaged, then patients with visual agnosia can nonetheless accurately reach towards an object even if they cannot describe the orientation or other properties in space of the object that they are reaching.

Slide 162 emphasizes a dual use of PPC in doing this. First, there is the top-down attention from PPC to V4 that focuses spatial attention as part of a surface-shroud resonance. In addition, there is a bottom-up command from this attentive focus to motor control networks further downstream that carries out an intended movement to the attended object. Attention and intention are well known to both be parietal cortical functions. The theory clarifies why this so from the perspective of explaining how and why we become conscious of visual qualia.

Slide 163. This slide summarizes some of the brain designs that this lecture has used to explain properties of how we consciously see and know things, and how these processes help to guide

artists in making visual art. These designs clarify that our brains compute very differently than traditional computers, and the currently popular algorithm in machine learning and AI called Deep Learning. Adaptive Resonance Theory has also been used in machine learning and AI applications, as Slide 151 has illustrated. ART can thus shed light upon the artistic process as well as provide algorithms for large-scale applications in engineering and technology that require autonomous adaptive intelligence in response to rapidly changing environments that may be filled with unexpected events. ART has also been used to provide mechanistic neural explanations of mental disorders that afflict millions of individuals, such as Alzheimer's disease, autism, Fragile X syndrome, ADHD, schizophrenia, visual and auditory neglect, and medial temporal amnesia (Franklin and Grossberg, 2017; Grossberg, 2000, 2017a, 2017b; Grossberg and Kishan, 2018; Grossberg and Seidman, 2006). Deep Learning cannot do any of these things. I therefore welcome artists, as well as scientists and technologists, to further study ART and to help develop its ability to provide new insights and applications in all of these fields.

REFERENCES

- Franklin, D. J., and Grossberg, S. (2017). A neural model of normal and abnormal learning and memory consolidation: Adaptively timed conditioning, hippocampus, amnesia, neurotrophins, and consciousness. *Cognitive, Affective, and Behavioral Neuroscience*, 17, 24-76.
- Grossberg, S. (1973). Contour enhancement, short-term memory, and constancies in reverberating neural networks. *Studies in Applied Mathematics*, 52, 213-257.
- Grossberg, S. (1980). How does a brain build a cognitive code? *Psychological Review*, 87, 1-51.
- Grossberg, S. (1984). Outline of a theory of brightness, color, and form perception. In E. Degreef and J. van Buggenhaut (Eds.), *Trends in mathematical psychology*. Amsterdam: North-Holland, pp. 59-85.
- Grossberg, S. (1987). Cortical dynamics of three-dimensional form, color, and brightness perception, II: Binocular theory. *Perception and Psychophysics*, 41, 117-158.
- Grossberg, S. (1994). 3-D vision and figure-ground separation by visual cortex. *Perception and Psychophysics*, 1994, 55, 48-120.
- Grossberg, S. (1997). Cortical dynamics of three-dimensional figure-ground perception of two-dimensional figures. *Psychological Review*, 104, 618-658.
- Grossberg, S. (2000). The imbalanced Brain: From normal behavior to schizophrenia. *Biological Psychiatry*, 48, 81-98.
- Grossberg, S. (2013). Adaptive Resonance Theory: How a brain learns to consciously attend, learn, and recognize a changing world. *Neural Networks*, 37, 1-47.
- Grossberg, S. (2017a). Acetylcholine neuromodulation in normal and abnormal learning and memory: Vigilance control in waking, sleep, autism, amnesia, and Alzheimer's disease. *Frontiers in Neural Circuits*, November 2, 2017, <https://doi.org/10.3389/fncir.2017.00082>.
- Grossberg, S. (2017b). Towards solving the Hard Problem of Consciousness: The varieties of brain resonances and the conscious experiences that they support. *Neural Networks*, 87, 38-95.
- Grossberg, S., and Kishnan, D. (2018). Neural dynamics of autistic repetitive behaviors and Fragile X syndrome: Basal ganglia movement gating and mGluR-modulated adaptively timed learning. *Frontiers in Psychology, Psychopathology*. <https://doi.org/10.3389/fpsyg.2018.00269>.
- Grossberg, S. and Mingolla, E. (1985). Neural dynamics of perceptual grouping: Textures, boundaries, and emergent segmentations. *Perception and Psychophysics*, 38, 141-171.
- Grossberg, S. and Mingolla, E. (1985). Neural dynamics of form perception: Boundary completion, illusory figures, and neon color spreading. *Psychological Review*, 92, 173-211.

Grossberg, S. and Seidman, D. (2006). Neural dynamics of autistic behaviors: Cognitive, emotional, and timing substrates. *Psychological Review*, 113, 483-525.

Grossberg, S. and Todorovic, D. (1988). Neural dynamics of 1-D and 2-D brightness perception: A unified model of classical and recent phenomena. *Perception and Psychophysics*, 43, 241-277.

Grossberg, S., and Zajak, L. (2017). How humans consciously see paintings and paintings illuminate how humans see. *Art & Perception*, 5, 1-95.

Kapadia, M. K., Ito, M., and Gilbert, C. D. (1995). Improvement in visual sensitivity by changes in local context: Parallel studies in human observers and in V1 of alert monkeys. *Neuron*, 15, 843-856.

Perlovsky, L. (2010). Musical emotions: Functions, origins, evolution. *Physics of Life Reviews*, 7, 2-17.

Von der Heydt, R., Peterhans, E., and Baumgartner, G. (1984). Illusory contours and cortical neuron responses. *Science*, 224, 1260-1262.