The History of Neuroscience in Autobiography
Volume 13

Edited by Thomas D. Albright and Larry R. Squire
Published by Society for Neuroscience
ISBN: 978-0-916110-12-3

Charles Douglas Gilbert

pp. 268–293

https://doi.org/10.1523/hon.013007

Charles Gilbert

Charles Douglas Gilbert

Born:

New York, New York
January 15, 1949

Education:

Amherst College, Amherst, MA, BA, Summa cum laude, Biophysics (1971)
Harvard Medical School, Boston, MA, Neurobiology, MD, PhD (1977)

Appointments:

Teaching Fellow in Neurobiology, Harvard Medical School (1977–1979)
Principal Research Associate in Neurobiology, Harvard Medical School (1979–1981)
Assistant Professor, Department of Neurobiology, Harvard Medical School (1981–1983)
Assistant Professor, The Rockefeller University (1983–1985)
Associate Professor, The Rockefeller University (1985–1991)
Professor, The Rockefeller University (1981–present)
Arthur and Janet Ross Professor, The Rockefeller University (2004–present)

Honors and Awards:

Fellow, The Danforth Foundation (1971–1975)
Fellow, The Medical Foundation (1978–1980)
Weill-Caulier Award (1984)
Presidential Young Investigator Award (1984)
Rita Allen Foundation Award (1986)
Elected member, International Neuropsychological Symposium (1992)
Krieg Cortical Kudos, Cortical Discoverer Prize, Cajal Club (1993)
Fellow, American Association for the Advancement of Science (2000)
Elected, American Academy of Arts and Sciences (2001)
W. Alden Spencer Award (Columbia University College of Physicians & Surgeons) (2002)
Elected, National Academy of Sciences (2006)
Inaugural Fellow, American Physiological Society (2015)
Edward M. Scolnick Prize in Neuroscience (McGovern Institute, MIT) (2015)
Honorary PhD Degree, Amherst College (2022)

The objective of Charles Gilbert and his colleagues has been to understand the mechanisms of brain function at the circuit level. Gilbert has identified a component of the cortical circuit consisting of long-range lateral connections and has shown how connections between neurons operate dynamically. He has identified long-term changes in cortical connections that mediate perceptual learning and functional recovery after brain lesions, and short-term changes that enable neurons to change their function according to task demands. He has found that the functions of brain regions are under top-down influences of attention, expectation, and perceptual task. In effect, neurons are adaptive processors, capable of selecting subsets of inputs depending on the behavioral context. He has proposed a model in which this input selection, and the associated cortical dynamics, are achieved by an interaction between feedback to cortical areas and intrinsic connections within these areas. He is currently exploring how this mode of circuit interaction may account for perceptual dysfunction in neurological and behavioral disorders.

Charles Douglas Gilbert

In this memoir I hope not only to include an account of my past accomplishments but also to present a vision for a changing view of the nature of information processing in the brain. Although all of my work has focused on the modality of vision, the principles of circuit interactions we have developed in studies of visual cortical function may be more broadly applicable to all systems in the brain.

Family Life

I was born in New York City, although my parents were living in Princeton, New Jersey, at the time, where my father was a professor at Princeton University. My parents were both academics. My father was a clinical psychologist and professor of psychology and my mother was a schoolteacher. My father was fluent in German, so together with his training in psychology, he was able to apply his expertise as the psychologist during the Nuremberg war tribunals, where he interviewed all of the most senior Nazis who were being tried for war crimes. From this, he developed an interest in the psychology of dictatorship. When he sent his book *Nuremberg Diary* to a Princeton colleague, Albert Einstein, Einstein thanked him for the book and said, "It is truly an agonizing thought that people of such little sense of responsibility could rise to power with the consent of the people of a large and highly developed country. The greatest of political problems has made no real progress since the days of Aristotle." Later, my father identified the same tendencies at home when he publicly called on Joseph McCarthy to submit to a psychological examination, saying "we should like to know more about the psychodynamics of the personality that delights in deliberate deception, in sensational and reckless attacks." These thoughts remain eerily relevant. But that typified the liberal environment in which I grew up.

Our family moved several times, from Princeton to Michigan, where my father was a professor at Michigan State University, to Mexico, where he worked with psychologist Erich Fromm, to New York, where he returned to academia. I went to school on Long Island, and then went to Amherst College. At Amherst, I majored in biophysics, but perhaps following my family tradition I also had a parallel involvement in liberal politics. It was the Vietnam War era, and I chartered a fleet of buses to take me and my fellow students to a march on Washington to protest the war. Before the date of the protest we learned that the Federal Bureau of Investigation tried to induce the bus company, Peter Pan buses, to renege on our contract. We recruited the assistance of the American Civil Liberties Union, which

threatened to sue on our behalf, and the bus company backed down. But I always chuckle at the idea of suing Peter Pan.

My own interest in biomedical science started at an early age. At the age of four, I declared a wish to be a medical scientist. I was the stereotypical science nerd who had many scientific adventures and misadventures in grade school and high school. At one point, convinced that the continents must have at one time fit together, I got permission from my sixth-grade teacher to take apart a globe and reassemble it into a single continent. For a physics club project, I built a particle accelerator, complete with an oil vacuum pump and a Van der Graaf generator. When I started it up, it emitted so much radiation that my science teacher shut the whole thing down. In my teen years, I was an avid reader of both fiction and nonfiction, including a number of science-related works. One book that fed my interest in biomedical research was *Arrowsmith* by Sinclair Lewis. It portrays an idealistic medical researcher who comes to work at a research institute in New York City. Only after I came to The Rockefeller University did I learn that Arrowsmith's institute was modeled after the Rockefeller Institute for Medical Research, the forerunner of The Rockefeller University, where I work now. Some considerable irony that I ended up at the fictional institute that stoked my interest in medical science. I met my wife, Jennifer Jones, at Rockefeller, where she is a writer and fundraiser. We have two sons, Zachary and Jonah, both now adults and embarked on their own careers.

My family all share my passions for cooking and travel, probably because I exposed them to these from an early age. When Jennifer and I brought Zachary home from the hospital at two days old, I put him in his chair on the kitchen counter and proceeded to give him a lesson in the proper way to slice fennel. Now in a turnabout-is-fair-play, he peppers me with requests to purchase expensive cooking equipment, like a restaurant-grade vacuum sealer and a laboratory-grade water bath in order to properly sous vide a steak. And his younger brother (our son Jonah) instructs me on the proper method to slice ginger for his stir fry.

Great dining and the communal act of cooking together is an important part of our family life and life in the laboratory. Back when laboratories could have kitchens, complete with stovetops and ovens, all day (and night) experiments would be paused for a communal meal. The graduate students or postdocs would shop for ingredients (usually fresh fish), and I would prepare the meal. Jennifer and Zachary would join us, then head home while we resumed the experiment.

When Jonah came along we swapped our one-bedroom Rockefeller apartment for an old house in Brooklyn, a short commute to Rockefeller (although on some days in New York City traffic, that's arguable). The house is a large federal-style structure, on the federal and city landmarks list, sitting on about three acres of land. It is a rather unusual place to be found in the middle of downtown Brooklyn. The only reason we could afford

it was that the neighborhood at the time was in a rundown industrial area, surrounded by abandoned warehouses. Now, this part of Brooklyn has seen massive resurrection and is lively with people and even shops. But we spent many years in a largely deserted part of New York, and a house of this vintage has required a lot of attention. We are constantly designing and renovating parts of the house, if one wants to call it an extracurricular activity.

College and Medical School

While I was at Amherst College, I worked summers, and during the year, I worked at the Cold Spring Harbor Laboratory in the lab of Ray Gesteland and Robert Crouch. My project centered around the bacteriophage lambda. It was a great way to get my feet wet in experimental science, although my interests were gravitating toward studying more complex systems, and I developed an interest in studying the brain. On a fortunate day while at Amherst, I met with a visiting speaker, David Hubel, and thought his approach to studying the brain would be a good model for me to follow. I applied to, and was accepted to, the department of Neurobiology at Harvard Medical School. I wanted to pursue a combined medical and doctoral degree. Although there was no such combined formal program at Harvard at the time, I was able to put together my own program, taking the courses for the medical and doctoral programs simultaneously.

I worked in the lab of David Hubel (see volume 1) and Torsten Wiesel. Unlike much of the way labs work these days, they worked together on their own projects and let their students and postdocs (few in number at any one time) do their own thing, and they did not put their names on the papers that came from their student's efforts. Although I was new to the field, they provided me with an empty room and a catalog and told me to think of a project.

Early Work: The Anatomy of Cortical Circuits

After some initial poking around in not so interesting projects, I ended up doing a combined anatomical and physiological project to explore the roles of the different layers of the primary visual cortex (V1). I took advantage of the then-new technique of retrograde tracing of connections with horseradish peroxidase (HRP). In papers published in 1975 with Jim Kelly, a postdoc in David and Torsten's lab, I identified the projection targets of the different cortical layers. We showed that feedforward cortico-cortical connections originate from the superficial cortical layers, projections from layer five target the superior colliculus, and feedback cortico-cortical connections and cortico-thalamic projections originate from layer six (Gilbert and Kelly, 1975). Our findings establishing the principle that deep layers provide feedback connections and that superficial layers provide feedforward connections were subsequently used by others to map out a hierarchy of cortical areas subserving vision.

The fact that orders of magnitude more neurons project to the lateral geniculate nucleus (LGN) of the thalamus from the cortex than those originating from the retina posed the intriguing question of the role of recurrent inputs to the LGN and of the role of the LGN. Although the LGN had been characterized as a "relay" between retina and cortex, clearly much more was happening in the LGN to transform the information coming from the retina. I then did single-unit recordings from the different cortical layers, showing that the superficial layer neurons had small receptive fields, appropriate for the high-resolution function of the cortical pathway. Layer five neurons had large receptive fields selective for direction of movement, appropriate for the eye-tracking function of the superior colliculus (Gilbert, 1977). Although layer six had properties unique to that layer, it was not clear what those properties contributed to the function of the LGN. But in any event, this work led to my continued interest in mapping the circuitry of the cortex as an entry point to understanding cortical function.

At the end of my combined medical and doctoral training at Harvard, I was confronted with the decision of whether to continue with my medical training, going onto internship and residency and so on. I very much enjoyed working in the clinic, but I felt I could not do justice to either clinical work or research if I tried to do both, so I decided to continue exclusively on the research track. Still, I feel that my experience in the medical part of my training provided a valuable background for the translational possibilities of my research. It is also worth pointing out how much basic research needs to be done before its translational potential can be realized. Fortunately, we have had a glimpse of that potential, as outlined at the end of this memoir.

After completing my doctoral work on the function and efferent connections of the layers of visual cortex, I stayed on at Harvard for a period to do postdoctoral work. At that point, Torsten Wiesel and David Hubel developed research programs independent from each other, and Torsten joined my interest in doing a structure-function analysis of intrinsic cortical circuits. These experiments, which began in the mid-1970s, related the interlaminar and horizontal cortical connections to the functional properties of neurons in V1. This involved doing intracellular recording with micropipettes filled with a solution of horseradish peroxidase (HRP). After characterizing the receptive field properties of the recorded neurons, we injected the neurons with HRP, which then filled the entire soma, dendritic, and axonal arbors. This provided more complete labeling than the anatomical techniques that were available at the time—notably various silver impregnation techniques (the Golgi or Nauta stains), which provided only partial labeling, due at least in part to the lack of labeling myelinated axons. These were difficult and challenging experiments, often lasting late into the night. On occasion, Torsten would lapse into Swedish in the wee hours. From these experiments, we were able to outline a basic circuit of interlaminar connections, and further pursued the functional role of these connections by inactivating

specific cortical layers, such as layer 6 (Gilbert and Wiesel, 1979; Bolz and Gilbert, 1986). We described a systematic series of connections, starting from the LGN input to layer 4 and then sequentially to the superficial layers, then layer 5, then layer 6, and finally back to layer 4. This circuit was subsequently dubbed as the "canonical circuit," although the similarities and differences in the circuit between V1, where it has been most fully characterized, and other cortical areas remain to be further explored.

Long-Range Horizontal Connections and Contextual Influences

Our experiments (Torsten and mine) produced another result that seemed to run counter to how scientists viewed cortical circuits and receptive fields in V1. At the time, it was thought that cortical connections ran principally in the vertical direction, across the cortical layers, with little lateral extent parallel to the cortical surface, which was consistent with the columnar cortical functional architecture. Our experiments did reveal a systematic pattern of connections between the cortical layers, as outlined previously, but in addition to the interlaminar connections, we found, initially in the cat and then in the monkey, that cortical pyramidal cells formed long-range horizontal connections, extending up to about 8 millimeters parallel to the cortical surface, crossing many cycles of orientation change (Gilbert and Wiesel, 1979, 1981, 1983). The implication of this finding, if one considers the receptive field sizes and retinotopic organization of V1, is that neurons integrate input over a large area of cortex representing an area of the visual field that is much larger than their receptive fields. The resolution to this seeming contradiction is that the definition of the receptive field, and its extent, is stimulus dependent, and there are substantial nonlinearities in how neurons respond to complex stimuli. For example, although a neuron will respond to a line segment presented in the receptive field center and not at all when the segment is placed in an adjoining visual field position, the neuron can respond severalfold stronger when two line segments are presented in a collinear arrangement, with one "inside" and the other "outside" the receptive field. These interactions are sometimes referred to as "contextual influences" (see below). Of course, based on this and related observations, what is inside or outside loses its meaning. At the time the popular distinction was between the "classical" and "nonclassical" receptive field, although I found the terms unsatisfying. After all, even in his classical study of retinal ganglion cell receptive fields in the 1950s Stephen Kuffler defined the receptive field as "not only the areas from which responses can actually be set up by retinal illumination may be included in a definition of the receptive field but also all areas which show a functional connection, by an inhibitory or excitatory effect on a ganglion cell. This may well involve areas which are somewhat remote from a ganglion cell and by themselves do not set up discharges" (Kufffler, 1953, p. 45). The critical difference in

our study of receptive fields in the visual cortex is in the distinction between the selectivity of neurons to a simple stimulus, such as a single oriented line segment vs. its selectivity for complex stimulus configurations, which involve many components that interact nonlinearly. The lateral interactions and contextual influences operating in V1 also can account for perceptual interactions, such as the tilt illusion. The tilt illusion involves a change in the perceived orientation of a line when it is surrounded by other lines of a slightly different orientation. Surrounding a neuron's receptive field with a set of lines of a given orientation induces a change in its tuning to the orientation of a line placed in the receptive field center, potentially providing a cellular basis for the tilt illusion (Gilbert and Wiesel, 1990).

These contextual influences, seen in V1, are likely related to the Gestalt rules of perceptual grouping. Even in V1, neurons integrate information over large parts of the visual field, following the laws of perceptual grouping. One property of perceptual grouping, "good continuation" as described by the Gestalt psychologist Max Wertheimer, is consonant with the properties of the long-range horizontal connections, which can account for how object contours are perceptually assembled (Ts'o et al., 1986; Kapadia et al., 1995; Li and Gilbert, 2002; Stettler et al., 2002; Li et al., 2006). The responses of V1 neurons to global contours also reflect the statistical regularities of contours in natural scenes, as analyzed by then-graduate student Mariano Sigman (Sigman et al., 2001). As such, the cortex is tuned to pick out the more probable correlations in natural scenes, rather than the exceptions. As I will describe, the long-term change associated with perceptual learning is another example of how the visual cortex picks up new statistical regularities by repeated exposure to the same stimulus.

In studying the synaptic physiology of the horizontal connections, Judith Hirsch found a strong voltage dependence of inputs coming through long-range horizontal connections, potentially accounting for the nonlinearities in the responses to complex stimuli (Hirsch and Gilbert, 1991). The sort of interactions afforded by the horizontal connections have properties that underly an "association field" within V1 that mediates the Gestalt principle of good continuation. In particular, over the years with Torsten and my students and postdocs, we showed that the long-range horizontal connections connect columns of similar orientation preference. We demonstrated this in several ways as the technology for labeling the horizontal connections and visualizing the cortical functional architecture became more sophisticated. First, with Torsten, we combined labeling neurons forming horizontal connections with retrograde tracers and the orientation columns with 2-deoxyglucose autoradiography. Active neurons take up more radiolabeled deoxyglucose, but they can't further metabolize it, so it remains stably in the neurons that have been activated during its infusion, such as with a visual stimulus consisting of a grating of a particular orientation. The active cortical regions can then be visualized with autoradiography.

When the retrograde tracer was placed in a column of a given orientation, it was transported to adjoining columns preferring the same orientation. Then, with Daniel Ts'o, we measured the orientation dependency of lateral interactions by cross-correlation analysis between pairs of recorded neurons (Ts'o et al., 1986; Ts'o and Gilbert, 1988). Here again we found connectivity between similarly oriented neurons with widely separated receptive fields. With Dan Stettler and Aniruddha Das, we labeled the horizontal connections with virally mediated transfer of the GFP gene (using an adenovirus, in collaboration with Jean Bennett) and imaged the cortical architecture with intrinsic signal optical imaging (Stettler et al., 2002). This imaging technique was developed in Torsten's lab by Amiram Grinvald (Grinvald et al., 1986), where he found that active areas of cortex have lower reflectance than inactive areas, because of changes in local blood flow and oxidative changes in local chromophores.

The continuing emergence of new experimental approaches has greatly expanded the reach of what we were able to do. The linking of cortical columns of similar orientation preference along with their lateral extent across the cortical map matched psychophysical measures of perceptual interactions mediating contour saliency, as found by Wu Li (Li and Gilbert, 2002). Wu has been a longstanding collaborator, and he skillfully helped us transition from studying neuronal function in anesthetized animals to doing experiments with awake, behaving animals. This transition was centrally important in developing our theory of the relationship between cortical circuits and function in a behavioral context. The properties of the long-range horizontal connections are reflected in the contour dependent neural responses, which are facilitated by the presence of contours that extend over large parts of the visual field (Li et al., 2004, 2006). Yet, on top of this, as I'll discuss, these interactions are not constitutively active but rather are additionally dependent on top-down influences, with the contour-related facilitation being strongest when animals perform a contour detection task (Crist et al., 1997; Ito et al., 1998; Ito and Gilbert, 1999; Kapadia et al., 1999; Ito et al., 2000; Westheimer et al., 2001; Crist et al., 2001).

By working with awake, behaving monkeys trained to do visual discriminations tasks, we were able to link receptive field properties to perception and study the higher-order functions of early cortical areas and, critically, the role of top-down influences in mediating these properties. In this regard, I greatly benefited from my long-standing friendship and collaboration with Gerald Westheimer (see volume 5), a noted figure in the field of perceptual psychology. With his valuable input, we incorporated the tools of visual psychophysics in linking the receptive field properties of cortical neurons to visual perception. With Gerald and students and postdocs Mitesh Kapadia, Roy Crist, and Minami Ito, we made our initial findings on the interaction between contextual influences, perceptual learning, and top-down influences in early visual processing (Kapadia et al., 1994, 1995, 1999, 2000;

Gilbert et al., 1997; Crist et al., 1997; Ito et al., 1998; Westheimer et al., 2001). Another fruitful collaboration was with the highly imaginative perceptual psychologist Ken Nakayama, where, with a postdoc Jonathan Bakin, we showed how neurons in areas V1 and V2 conveyed information about three-dimensional surface configurations (Bakin et al., 2000).

Adult Cortical Plasticity

Another theme of my work that emerged from our investigation of the function of the long-range horizontal connections in the cortical circuit is that of adult cortical plasticity. From the early work of Torsten Wiesel and David Hubel, it was established that experience-dependent changes in cortical circuits were limited to an early period in postnatal life, lasting a few months, known as the critical period. Their observations were based on the property of ocular dominance, how the inputs from the two eyes are brought together in V1. Monocular deprivation leads to a shrinkage of the input from deprived eye afferents and expansion of the nondeprived eye afferents. This change could be induced or reversed only within the critical period, after which the thalamocortical projections became fixed.

Although there had to be a measure of cortical plasticity throughout life to encode new memories, or perceptual learning, it was generally believed that this took place in higher-level areas of the visual cortical hierarchy, far from V1. But the work on ocular dominance plasticity during the critical period was primarily a matter of changes in thalamocortical connections, the projection from the LGN to area V1. I felt that a different component of the cortical circuit, the long-range horizontal connections, was a promising candidate for adult plasticity, even in early visual cortical areas. To test this idea, we came up with an experimental model to induce experience-dependent changes in V1—making focal binocular lesions (also referred to as scotomas) with an infrared laser at corresponding locations in the two retinas. This removed functioning visual input from a restricted area of V1, which we refer to as the lesion projection zone (LPZ). Although initially this region was silenced after binocular lesions, over a subsequent period of weeks and months, the LPZ regained functioning visual input. In effect, the cortex was remapped, with a shrinkage of the representation of the retinal scotoma and an expansion in the representation of the normal, surrounding retina (Gilbert et al., 1990; Gilbert and Wiesel, 1992). Our finding of shifting receptive fields in the LPZ was confirmed by the application of new technologies for the study of neuronal populations, cortical functional architecture, and circuitry. These included optical imaging, recording with chronically implanted electrode arrays, labeling neurons with virally mediated gene transfer, and 2-photon imaging. The receptive fields of the neurons in the LPZ had shifted from their prelesion retinal locations to sites outside the retinal scotoma, having received input from cortical neurons outside

the LPZ through the long-range horizontal connections (Gilbert et al., 1990; Darian-Smith and Gilbert, 1994, 1995; Das and Gilbert, 1995; Yamahachi et al., 2009; Marik et al., 2014; Abe et al., 2015).

Although a small measure of remapping takes place in the LGN, this was insufficient to account for the more substantial reorganization occurring in the visual cortex, as I found with Corinna Darian-Smith (Darian-Smith and Gilbert, 1995). Cortical plasticity following retinal lesions was most compellingly demonstrated through the use of chronically implanted electrode arrays, which was accomplished in the lab by Hiroshi Abe (Abe et al., 2015). At each recording site, we could document the receptive field positions of single neurons before and after making the lesions. When the array was implanted in the future LPZ, we could track the changes occurring after cortical sites were initially silenced by the retinal lesion, and then we gradually recovered visually driven responses. After recovery, the receptive field positions had made large shifts to positions in intact parts of the retina, outside the retinal scotomas. This line of evidence showed that the remapping was not merely a matter of selection of neurons that were unaffected by the lesion but rather was mediated by neurons that were initially silenced and that subsequently developed new receptive fields that were substantially shifted relative to their original locations.

Another approach I applied in studying the functional recovery after retinal lesions, along with Aniruddha Das, involved the technique of intrinsic signal optical imaging, comparing the orientation maps before making the lesion and after recovery. There was a close match, supporting the idea that, because the horizontal connections linked orientation columns of similar orientation preference, they were the likely source of the reorganization (Das and Gilbert, 1995). A model, developed by a student Justin McManus in collaboration with the computational neuroscientist Shimon Ullman (see volume 13), showed how this form of reorganization could account for perceptual fill-in through the lesioned parts of the visual field and mediate a measure of recovery of function after CNS lesions (McManus et al., 2008). The model provided a mechanism for the perceptual fill-in, and perhaps more generally for how the brain can adapt to neurodegenerative disorders and stroke.

The cortical circuit responsible for this reorganization was the plexus of long-range horizontal connections. The experiments that led to this discovery benefited from my collaboration with Winfried Denk, who pioneered in developing the technique of 2-photon microscopy for deep tissue imaging in live animals. With Denk, we engineered the first 2-photon microscope for imaging in nonhuman primates. When combined with virally mediated expression of the green fluorescent protein (GFP), we were able to image, at high resolution, the somata, dendrites, and axons of neurons in V1 of Macaque monkeys. Our initial surprising finding was that, under normal circumstances, there is a continuing turnover of axonal boutons, the presynaptic sites of cortical circuits. With a student, Dan Stettler, along

with Aniruddha Das and Denk, we found a rate of turnover of 7 percent per week (Stettler et al., 2006). I'll have more to say about that later. We then examined what happened to the horizontal connections after making retinal lesions. Using the technique of longitudinal 2-photon imaging, we were able to track the changes occurring among identified axons over weeks. We found that the lesioning was followed by a period of exuberant outgrowth and pruning of the horizontal connections projecting from outside the LPZ to inside the LPZ, with the balance shifted toward a net increase (Yamahachi et al., 2009). This confirmed a result we initially obtained from postmortem studies (Darian-Smith and Gilbert, 1994) and accounted for how the LPZ neurons recovered visually driven input, shifting their receptive fields to positions overlapping with those of the neurons outside the LPZ from which they received input. Notably, the plasticity we observed occurred in adult animals, well after the end of the critical period for thalamocortical connections. Consequently, one can surmise that the limitations on cortical plasticity associated with the critical period are specific to certain cortical connections, most particularly thalamocortical connections, but other circuits, even those in V1, remain plastic throughout life.

With Sally Marik, I extended the players involved in adult plasticity from excitatory to inhibitory neurons, from the visual system to the somatosensory system, and from nonhuman primates to mice (Marik et al, 2010; Marik et al., 2014). Along with our study of the circuits engaged in cortical plasticity, we documented the expression of neurotrophins in the area undergoing topographic reorganization (Obata et al., 1999). With Marik, we revealed some of the other molecular pathways involved, notably the caspase pathway, which, rather than leading to apoptosis as it does in early development, is involved in axonal pruning in the adult (Marik et al., 2013, 2016).

Cortical Plasticity, Perceptual Learning, and Top-Down Influences

Our observations on cortical plasticity following retinal lesions may be relevant for the recovery of function after brain lesions or degenerative disease, but it seems unlikely that it evolved solely for that purpose. Instead, it was likely to be associated with experience-dependent changes in cortical function that occur throughout life under normal circumstances. Specifically, a form of implicit memory known as perceptual learning, which is broadly associated with the neocortex, seemed to us a likely candidate for neuronal plasticity in the visual cortex. Although it is established that different forms of memory are supported by different parts of the brain, with declarative memory in the medial temporal lobe and some forms of implicit memory elsewhere in the brain (including habit memory in basal ganglia and classical conditioning in the cerebellum; see Squire, volume 10), we have developed a theory of the circuit mechanisms of learning, and of systems consolidation of memory, that may be common to all forms of memory.

Clearly, it is necessary to have ongoing neuronal plasticity throughout adulthood as we continually undergo perceptual learning. One might have presumed that such plasticity would be the province of cortical areas high in the visual cortical hierarchy, with V1 being free of experience-dependent changes. In experiments with my colleague Wu Li, we found quite to the contrary that even V1 was capable of undergoing experience-dependent changes with learning, similar to what we observed following retinal lesions. Importantly, we found the neural mechanism of perceptual learning to be related to the role of top-down influences in neuronal function, which I will detail below in "The Neuron as an Adaptive Processor." A useful tool in measuring learning-dependent changes in neuronal function is that of mutual information—that is, how much the tuning of a neuron to a stimulus attribute is dependent on the value of that attribute. We found that as animals improve on a visual discrimination task, neurons carry increasing amounts of information about task- relevant stimulus characteristics and carry less information about task-irrelevant stimulus features. The design of the experiments that show this are as follows: while recording from hundreds of neurons in multiple cortical areas, we train animals (nonhuman primates) to shift between perceptual tasks with the identical visual stimulus. Some components of the stimulus are relevant to the task, and some irrelevant. Neurons carry information about, or are tuned to the properties of, the task-relevant stimulus components and their responses are not influenced by task irrelevant components. Changing tasks switches the relationship between what is relevant and what is irrelevant, and the information conveyed by each neuron's response changes accordingly. An example of a pair of tasks we use involves a set of five lines—a central line flanked by two parallel lines side-by-side with the central line and two colinear lines, either above or below the central line. On a trial-by-trial basis, the flanking lines are placed in different positions relative to the central line, so we can measure how each neuron's response changes as we alter the positions of the flanking lines. We can then construct a tuning curve that represents how each neuron encodes the positional information. Animals either perform a bisection task, reporting the proximity of the central line to either of the parallel flankers, or a Vernier task, reporting the position of the central line relative to the colinear lines. We find that when animals perform the bisection task neurons become sensitive to the position of the parallel lines, which are task relevant and suppress the influence of the colinear lines, which are task irrelevant. Conversely, when the animal performs the Vernier task, the same neurons become insensitive to the position of the task irrelevant parallel lines and are highly tuned to the position of the task relevant colinear lines. Thus, as I will elaborate, neurons are adaptive processors, in which a given level of firing conveys different information under different task conditions. We have documented these changes in various visual discrimination tasks and in different visual cortical areas, including V1 and V4

(Gilbert, 1994; Sigman and Gilbert, 2000; Crist et al., 2001; Li et al., 2004, 2008, Gilbert et al., 2009; Astorga et al., 2022).

The tuning of neurons to stimulus attributes changes during the course of perceptual learning. Along with the shift in the psychometric functions, reflecting improved performance on a task, there is a shift in V1 neurometric functions, reflecting an increased sensitivity of neurons to the property involved in the task. They become more sharply tuned to the stimulus attribute that shows perceptual improvement, or put another way, neurons increase the mutual information they carry about that attribute. For example, with the bisection task, the perceptual threshold in discriminating the position of the central line relative to the parallel line improves after hundreds of trials, and the neurons in different visual cortical areas increase the information they carry about the position of these lines and further suppress the information about the task-irrelevant lines. As with the changes in cortical circuitry that we saw following retinal lesions, during periods of perceptual learning and functional changes in neurons, the horizontal connections undergo parallel processes of sprouting and pruning (van Kerkoerle et al., 2018). As we have seen with the recovery following retinal lesions, one has to make the distinction between connections that are subject to a critical period and those that undergo change throughout life. Although it may have been presumed that the functional changes associated with perceptual learning would be limited to higher-order areas in the visual cortical hierarchy, it is particularly intriguing that even the earliest stage of visual cortical processing, area V1, shows such changes. This begs the question of how perceptual learning on one task doesn't interfere with performance on other tasks. The resolution to this problem, as I will elaborate, is that the functional properties of neurons, or their task-related tuning, are under top-down control. The underlying functional properties at the neuronal level are present only when the task is being performed, and these properties, for the same neurons, switch from task to task.

The dynamic changes in cortical circuits require a mechanism that can retain previously acquired information while encoding new information. We have seen an emerging theme: In different cortical areas, animal models and cell types, there is a high rate of synaptic turnover in the disappearance and appearance of axonal boutons. This has been seen by others in the form of dendritic spine turnover. We have documented this by longitudinal in vivo imaging of axonal boutons, the presynaptic sites. As described previously, even in the absence of experiential manipulation, such as when an animal is undergoing perceptual learning, boutons turn over at a rate of 7 percent per week. Although we were able to image for only a few weeks, there was no evidence for a privileged stable set of boutons. Even so, this amount of turnover raises the question of how information can be retained in the face of the circuit dynamics, because those circuits are the substrate of encoded information. The dynamics in connectivity may be related to the phenomenon

of representational drift, whereby neurons change their tuning or optimal stimulus on a day-to-day basis. (I will discuss this further in the section on dynamic connectivity.) In addition, the changes in connectivity are associated with the constitutive sharpening in neuronal tuning that occurs during the course of perceptual learning.

The Neuron as an Adaptive Processor

Working with awake, behaving animals opened up a world of new findings and new insights into the dynamic nature of cortical function and connectivity. We further benefited from the implementation of new technologies to record from large populations of neurons while animals perform complex visual discrimination and object recognition tasks. Over a series of experiments, we studied how neuronal responses were highly dependent on contextual influences, and these influences were in turn gated by top-down influences of perceptual task and expectation. By contextual influences, I mean that our perception, as well as neurons in V1, do not simply carry information about a single, local stimulus component but rather are highly influenced by the global characteristics of scenes in which that component is embedded. And these influences of stimulus context are in turn gated by top-down influences of perceptual task and expectation, as described earlier. We showed that a given level of response by a neuron does not signal the same information under all circumstances, but it can mean entirely different things when animals switch from task to task. This shows again that a neuron is an adaptive processor, carrying different information under differing behavioral contexts. This phenomenon also requires revisiting the idea of the receptive field. Neurons do not simply reflect the presence of a given sensory stimulus in a restricted part of the visual field. Rather, they carry information in the service of a particular perceptual or behavioral goal. Even V1 shows selectivity for complex stimulus configurations, and this selectivity shifts under different perceptual tasks. This contradicts the traditional notion that V1 is a mere preprocessor, carrying only simple stimulus characteristics, such as orientation or direction of local movement. Although there may be a measure of increasing complexity as one moves anterior in the ventral visual pathway, clearly V1 represents more complex stimulus configurations than single oriented line segments. And while it had been supposed that attentional influences increase along this pathway, we have seen top-down influences operating as strongly in V1 as in higher visual areas. We have observed the effects of top-down influences involving different perceptual tasks, including brightness discrimination, contour detection, shape detection, and object recognition (Ito et al., 1998; Ito and Gilbert, 1999; Gilbert et al., 2000, Li et al., 2004, 2006; McManus et al., 2008; Ramalingam et al., 2013; Chen et al., 2014; Liang et al., 2017; Astorga et al., 2022; Altavini et al., submitted).

It might be argued that we are simply talking about attentional influences. The use of the term "attention" has been applied in various ways, most commonly spatial attention, but also object oriented and feature based attention. Although one does not wish to get embroiled in a semantic argument, we prefer to use the term "top-down influences" because the information conveyed top-down is much richer than the mere locus of attention but instead includes the discrimination task being executed at the attended location. Top-down influences also play a role in setting the stimulus selectivity of neurons by object expectation and during perceptual learning. Importantly, we found that profound top-down influences operate not just in cortical areas at the top of the visual hierarchy but as far back as V1. The effects of the top-down influences are changes in the kind of information conveyed by neurons in the areas targeted by this feedback. A given level of firing may reflect very different stimulus characteristics when subjects perform different perceptual tasks or are operating under different expectations. Effectively, this means that a neuron does not represent a fixed "labeled line," but it represents different information on a moment-to-moment basis. Rather than compartmentalizing information for doing different tasks by different subsets of neurons, a large fraction of the entire neuronal population assumes tuning appropriate for every task. This raises the question of "read-out"—that is, how the response of a neuron is "interpreted" by the brain. The answer, presumably, is that the top-down instruction is for the target neuron to perform a particular calculation, and the area receiving the return feedforward input uses this response as the result of that instruction.

Dynamic Effective Connectivity—A Pointer Theory of Feedback

Underlying this task switching, we found that the task-dependent changes in cortical function are associated with changes in effective connectivity within the cortical network. With students Nirmala Ramalingam and Justin McManus, and my collaborator Wu Li, we studied the circuit mechanisms of top-down control using the chronically implanted electrode arrays referred to previously. With each electrode inserted into cortical sites representing different components of a complex stimulus, we were able to measure the effective connectivity between sites that were relevant to the task being performed vs. those that were task irrelevant. Effective connectivity was measured with either cross-correlation analysis of spiking activity or coherence between local field potentials (LFPs). The effective connectivity between sites switched on and off depending on the task, even though under all task conditions, all the recorded neurons were activated by the stimulus (Ramalingam et al., 2013). Although a neuron receives 10,000 inputs from other neurons, it is affected by only a small subset of these inputs at any given time. The cellular mechanism underlying this gating of specific

inputs remains a mystery, although one can speculate about possibilities. Whatever the mechanism, it is likely to involve an interaction between feedback connections, which carry information about the nature of the task, and local circuits, which carry information about the stimulus context. We have proposed that in effect the feedback serves as a pointer, dynamically linking the bits of information represented in the target area. Even under a fixed anatomical framework of connections, the matrix of effective connectivity within the network is constantly changing. Our theory is further supported by simultaneous recordings in areas V4 and V1, in which we showed that the lateral interactions in V1 are dependent on V4 input (Liang et al., 2017). One can imagine that the properties of any neuron result from nested dependencies of interactions, extending widely across the brain. The pointer idea also relates to our observations about functional changes occurring over extended periods of time during perceptual learning. What is required is a heterosynaptic interaction, which involves establishing the appropriate "addressing" of the relevant intrinsic connections by the feedback inputs to an area, allowing these inputs to be gated to conduct task-relevant information. This heterosynaptic interaction is distinct from classical models of cortical plasticity that involve Hebbian changes in synaptic weights that are permanently expressed. Instead, different influences of stimulus context are not constitutively expressed but instead change from moment to moment along with a changing behavioral context.

It is worth pointing out that this view puts additional constraints on what is required of connectomics. As challenging as it has been to reconstruct all the neurons and local connections in a cortical column, a large fraction of the inputs to that column come from distant sources through long-range horizontal connections within the area, including feedforward and feedback connections from multiple other areas and from the thalamus. In a sense, each neuron, taken together with all of its inputs, is a microcosm of the brain as a whole. A full description of the anatomical circuit affecting a neuron would require a connectivity map encompassing the entire brain. Perhaps that is why most of the progress in this area comes from smaller systems like *Drosophila* or *Caenorhabditis elegans*. But the greater challenge, as I have described, is that the connectivity matrix is dynamic, with changing effective connections within a fixed anatomical framework over short timescales. And with the synaptic turnover, even the anatomical framework is undergoing constant change.

When we did combined recordings in areas V1 and V4, research associate Guadalupe Astorga and I found an intriguing sequence in the functional changes occurring during perceptual learning. At early stages of learning, the information about the attributes involved in the visual task are better represented in V4 than in V1. Another important property seen in neuronal function is the suppression of the influence of stimulus attributes that are irrelevant to the task, again better reflected initially in the characteristics

of V4 neurons than V1 neurons. But over the training period, and in the following weeks, the increase in task-relevant information and suppression of task irrelevant information is seen in V1 (Astorga et al., 2022). This finding is reminiscent of an earlier study done with Mariano Sigman involving a visual search task with fMRI in humans, where a shift in cortical areas is activated by targets as they become more familiar from lateral occipital cortex to V1/V2 (Sigman et al., 2005). Continuing with the pointer idea, this reflects an engagement of pointers to V1, in addition to those previously established in V4. We have seen the phenomenon of task switching in multiple visual areas extending from V1 forward, and with different discrimination and recognition tasks. Importantly, learning in these paradigms involves not just increasing the level of mutual information for task-relevant attributes but increased suppression of task-irrelevant attributes, further supporting the idea of input gating.

With Astorga, we have made the intriguing observation that a small proportion of neurons change their task-related tuning over consecutive days, even when the animal is doing a given task. This phenomenon, known as "representational drift," has been seen elsewhere in the brain, most notably in the hippocampus. It is tempting to relate this phenomenon to our observation of the persistent turnover of synapses that we observed with longitudinal 2-photon imaging, where at the levels of connectivity and function, the brain is in a constant state of flux. Perhaps this is the mechanism underlying experience-dependent changes, such as those we have seen during perceptual learning. By having this underlying dynamic, learning can be accomplished by a process of selection of those neurons whose tuning carries more task-related information and less task-irrelevant information. As a result, the population as a whole becomes more adaptive for the trained behavior. Of course, this idea, and the drift itself, again raises the challenging issue of readout—how a particular neuron is linked to a particular percept.

Although it has long been the goal of systems neuroscientists to find and apply the optimal stimulus set for activating neurons, some looking for a "basis set," a fixed set of visual primitives that could be combined in various ways to represent full objects, our finding that neurons operate as adaptive processors posits a different framework. In effect, neurons take on properties that are appropriate for the current perceptual task, with their stimulus selectivity determined by top-down influences of expectation and perceptual task. Even their receptive field position shifts in accordance with that of the attended stimulus. So rather than thinking of a neuron having a fixed receptive field, it is more appropriate to think of it having a "perceptive field," adapted to the immediate perceptual demands, than reflecting only the bottom-up inputs arising from the retina. In current work with Tiago Siebert-Altavini, we have extended the functional switching to the process of object recognition, based on natural stimuli. We have seen that when animals are cued to expect different objects, neurons switch their stimulus

preferences to different objects or object components. We have seen these dynamics operating in several areas along the ventral visual pathway, including V1, which, as we have emphasized, has a more complex role in taking on task or expectation-dependent properties, not just passively encoding local stimulus features (Altavini et al., submitted). We suggest that, as we experience different environments, we are constantly operating under a set of expectations for the identity of objects in those environments, and these top-down influences interact with the incoming sensory information, leading to our ability to interpret the visual scenes we encounter.

As mentioned earlier, the pointer idea, and our findings of changes in neuronal tuning to task-relevant and task-irrelevant stimulus attributes, suggests a mechanism of synaptic plasticity different from the conventional idea of a Hebbian mechanism. Hebbian plasticity involves permanent changes in synaptic weights after concurrent activation of an input with a postsynaptic cell. Rather, what is required is a heterosynaptic mechanism, in which feedback connections develop the appropriate mapping, or "addressing" of local inputs that need to be gated during the execution of a task. Because different effective connectivities have to be set for different perceptual tasks, changing on a moment-to-moment basis, this doesn't involve any connection strength to be established on a permanent basis. Long-term changes may also occur, as long as the strengthened or weakened inputs can be dynamically gated. By this way of thinking, memories are not encoded as a change in weight of a small number of synapses but rather as a set of potential interactions distributed across a broad cortical network.

Although I now engage in a bit of speculation, I would like to propose that the pointer idea is applicable throughout the brain. For example, the hippocampus can be considered to be a pointer that links components of memories represented in the entorhinal cortex. As observed by Brenda Milner with patient H. M., when the hippocampus is removed, older memories are retained, but only those that were acquired several years before surgery. This process, requiring an extended period for memories to become independent of hippocampal input, is known as systems consolidation. We suggest that this consolidation involves a switch in pointers originating from the hippocampus to other areas, such as the perirhinal, parahippocampal, and prefrontal cortices as well as the thalamus. We have proposed that our observation of shifting pointers from those targeting V4 to those targeting V1 during perceptual learning reflects a similar process of systems consolidation for implicit memories. We further speculate that a similar role is played by the striatum, which can activate motor primitives by targeting premotor cortex via thalamic nucleus VL. This raises a potentially important function for the thalamus as a mediator of pointers. An additional feature of the striatum pointer idea is that it may contain the dimension of time, activating the motor primitives in a particular sequence in the execution of a complex movement. Here, too, systems consolidation may play a

role, with highly learned movements becoming independent of the striatum, depending on prefrontal inputs instead.

It is therefore worth considering whether this general theory of circuit interactions, involving pointers that gate neuronal inputs, expressing task-relevant inputs and suppressing task-irrelevant inputs, may be relevant to behavioral disorders that are not purely visual in character. This idea has driven my current research focus on autism. I've suggested that autism may be, at least in part, a perceptual disorder, in which there is a reduced ability to suppress task-irrelevant inputs. The result is a cacophony of stimulus inputs that would cause one to avoid interacting with the environment, which may be interpreted as a social deficit but is at root a perceptual one. We are currently testing this idea in animal models of autism. It is worth pointing out that other brain regions that are the source of pointers, such as the hippocampus and striatum, have also been implicated in autism, which therefore may be thought of as a disruption of a general mechanism of brain interactions that are broadly distributed, rather than a disfunction of any single brain region. In this regard, it is worth noting that hundreds of genes have been linked to autism and that these genes are ubiquitously expressed in the brain. How this pattern can account for the specific behavioral manifestations of autism is a puzzle, but perhaps the diverse set of signaling pathways linked to autism may have a common link at the level of circuits, such as those describe previously. Although we have studied some of the molecular mechanisms of experience-dependent axonal plasticity, the great challenge going forward is to understand how the diverse set of genes linked to autism and other behavioral disorders can produce the specific phenotypes of these disorders.

An emerging picture in how function is represented by the brain is one more of distributed interactive network, rather than the product of a feedforward hierarchy in which perception is represented at the top of the hierarchy after an amount of preprocessing from the lower levels. Experimentally, this requires simultaneous recording or imaging over many cortical and subcortical areas. One should consider that the function of feedback is to carry much more information than merely the locus of attention, but instead a variety of influences, including perceptual task, expectation, efference copy, and perceptual learning. Rather than thinking of perception and consciousness as being the products of the apex of a hierarchy of feedforward processing between cortical areas, they may lie in the mode of countercurrent interactions across the entire cortical network.

Cortical Dynamics

As we have pursued the themes of adaptive processing and perceptual learning, it is clear that the cortex is dynamic on a range of timescales. The information conveyed by a neuron switches from moment to moment

under different behavioral contexts and changes over weeks for a given task during perceptual learning. The seemingly random changes in tuning (representational drift) and synaptic turnover may be a useful substrate for enabling the experience-dependent changes associated with perceptual learning. One may consider the changes that are adaptive for perceptual learning result from a process of selection, keeping the neurons with tunings that fit the requirements of a task. Still, one must consider how the long-term changes associated with improvement in performance on a task can keep previously learned information intact. Moreover, here again we are challenged to understand how the system can "read-out" the meaning of a given neuron's response in the face of this constant flux in the information conveyed by each neuron under the same task conditions. This may be achieved by having cortical areas at higher levels of a hierarchy adapting to the changes at earlier levels and making the appropriate associations.

From our developing understanding of cortical circuit interactions, it is worth considering how they compare with neural network models that lack some of the key features we have unearthed. Most convolutional neural networks have a feedforward structure, lacking the lateral and feedback interactions we see in the brain. One may ask whether a neural network with enough layers can replicate cortical function without the dynamics seen experimentally, or whether certain functions are made possible by recurrent processing that are not achieved with neural network models. In other words, are current neural network models not in fact representative of real neural networks? It will be interesting to see what capabilities can be conferred to the models once true feedback and lateral interactions are incorporated.

With all this, the questions going forward proliferate. What is the mechanism, at the cellular or circuit level, of the input switching? How do feedback connections interact with intrinsic cortical circuits to enable the dynamic switching between tasks? What is the nature of the signal conveyed by feedback? What are the relative roles of different sources of feedback, for example, direct cortico-cortical connections vs. cortico-thalamo-cortical connections, or feedback to a given target cortical area from different cortical areas? In the face of synaptic turnover and shifting tuning, how is the previously learned information kept intact? Do these functional and anatomical dynamics help promote learning-dependent changes? How do circuit changes associated with learning prevent catastrophic loss of previously learned information? Even if the information is maintained, when individual neurons change their "line label," what is the mechanism of readout to stably guide behavior? What is the value of feedback/top-down interactions that operate in the brain relative to convolutional neural network models that operate by feedforward connections alone? How can one incorporate these interactions into neural network models, and what additional capabilities would this confer?

Although I began my career as an excursion into the nature of cortical circuits and the functional roles of specific cortical connections, our findings have led me over the years to a larger exploration of the more complex properties of visual cortical neurons, the short-term dynamics and long-term experience-dependent changes in adult cortical function and connections, and their potential relevance to behavioral disorders. At present, we are left with a profusion of questions still to be resolved, ones that will lead us to an even greater understanding of the basis of human cognition.

Chronological References

Kuffler, SW (1953) Discharge patterns and functional organization of mammalian retina. *J. Neurophysiol*. 16: 45.

Gilbert CD, and Kelly JP. (1975) The projection of cells in different layers of the cat's visual cortex. *J. Comp. Neurol.*163: 81–106.

Gilbert CD. (1977) Laminar differences in receptive field properties of cells in cat primary visual cortex. *J. Physiol.* 268: 381–421.

Gilbert CD, and Wiesel TN. (1979) Morphology and intracortical projections of functionally characterized neurons in the cat visual cortex. *Nature* 280: 120–125.

Gilbert CD, and Wiesel TN. (1981) Laminar specialization and intracortical connections in cat primary visual cortex. In *The Organization of the Cerebral Cortex*, ed. Schmitt FO, Worden FG, Adelman G, and Dennis SG, 163–191. Cambridge, MA: MIT Press.

Gilbert CD, and Wiesel TN. (1983) Clustered intrinsic connections in cat visual cortex. *J. Neurosci.* 3: 1116–1133.

Bolz J, and Gilbert CD. (1986) Generation of end-inhibition in the visual cortex via interlaminar connections. *Nature* 320: 362–364.

Grinvald A, Lieke E, Frostig RD, Gilbert CD, and Wiesel TN. (1986) Functional architecture of cortex revealed by optical imaging of intrinsic signals. *Nature* 324: 361–364.

Ts'o DY, Gilbert CD, and Wiesel TN. (1986) Relationships between horizontal interactions and functional architecture in cat striate cortex as revealed by cross-correlation analysis. *J. Neurosci.* 6: 1160–1170.

Ts'o DY, and Gilbert CD. (1988) The organization of chromatic and spatial interactions in the primate striate cortex. *J. Neurosci.* 8: 1712–1727.

Gilbert CD, Hirsch JA, and Wiesel TN. (1990) Lateral interactions in visual cortex. *Cold Spring Harbor Symposia on Quantitative Biology* 55: 663–677.

Gilbert CD, and Wiesel TN. (1990) The influence of contextual stimuli on the orientation selectivity of cells in primary visual cortex of the cat. *Vision Res.* 30: 1689–1701.

Hirsch JA, and Gilbert CD. (1991) Synaptic physiology of horizontal connections in the cat's visual cortex. *J. Neurosci.* 11: 1800–1809.

Gilbert CD. (1992) Horizontal integration and cortical dynamics. *Neuron* 9: 1–20.

Darian-Smith C, and Gilbert CD. (1994) Axonal sprouting accompanies functional reorganization in adult cat striate cortex. *Nature* 368: 737–740.

Gilbert CD. (1994) Early perceptual learning. *Proc. Natl. Acad. Sci.* 91: 1195–1197.
Kapadia MK, Gilbert CD, and Westheimer G. (1994) A quantitative measure for short-term cortical plasticity in human vision. *J. Neurosci.* 14: 451–457.
Darian-Smith C, and Gilbert CD. (1995) Topographic reorganization in the striate cortex of the adult cat and monkey is cortically mediated. *J. Neurosci.* 15: 1631–1647.
Das A, and Gilbert CD. (1995) Long-range horizontal connections and their role in cortical reorganization revealed by optical recording of cat primary visual cortex. *Nature* 375: 780–784.
Kapadia MK, Ito M, Gilbert CD, and Westheimer G. (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.
Crist RE, Kapadia MK, Westheimer G, and Gilbert CD. (1997) Perceptual learning of spatial localization: specificity for orientation, position and context. *J. Neurophysiol.* 78: 2889–2894.
Gilbert CD, A Das, M Ito, MK Kapadia, and G Westheimer. (1997) Cortical dynamics and visual perception. *Cold Spring Harbor Symposia on Quantitative Biology* 61: 105–113.
Ito M, Westheimer G, and Gilbert CD. (1998) Attention and perceptual learning modulate contextual influences on visual perception. *Neuron* 20: 1191–1197.
Ito M, Gilbert CD. (1999) Attention modulates contextual influences in the primary visual cortex of alert monkeys. *Neuron* 22: 593–604.
Kapadia MK, Westheimer G, and Gilbert CD. (1999) Dynamics of spatial summation in primary visual cortex of alert monkeys. *Proc. Natl. Acad. Sci.* 96: 12073–12078.
Obata S, Obata J, Das A, and Gilbert CD. (1999) Molecular correlates of topographic reorganization in primary visual cortex following retinal lesions. *Cerebral Cortex* 9: 238–248.
Bakin JS, Nakayama K, and Gilbert CD. (2000) Visual responses in monkey areas V1 and V2 to 3-D surface configurations. *J. Neurosci.* 20: 8188–8198.
Gilbert CD, Ito M, Kapadia MK, and Westheimer G. (2000) Interactions between attention, context and learning in primary visual cortex. *Vision Res.* 40: 1217–1226.
Ito M, Westheimer G, and Gilbert CD. (2000) Attentional modulation of contextual influences. In *Visual Attention and Cortical Circuits*, ed. J Braun and C Koch. Cambridge, MA: MIT Press, 89–102.
Kapadia MK, Westheimer G, and Gilbert CD. (2000) Spatial distribution of contextual interactions in visual perception and in primary visual cortex. *J. Neurophysiol.* 84: 2048–2062.
Sigman M, and Gilbert CD. (2000) Learning to find a shape. *Nature Neurosci.* 3: 264–269.
Crist RE, Li W, Gilbert CD. (2001) Learning to see: experience and attention in primary visual cortex. *Nature Neuroscience*, 4: 519–525.
Sigman M, Cecchi GA, Gilbert CD, and Magnasco MO. (2001) On a common circle: Natural scenes and gestalt rules. *Proc. Natl. Acad. Sci. USA* 98: 1935–1940.
Westheimer G, Crist RE, Gorski L, and Gilbert CD. (2001) Configuration specificity in bisection acuity. *Vision Res.* 41: 1133–1138.

Li W, and Gilbert CD. (2002) Global contour saliency and local colinear interactions. *J. Neurophysiol.* 88: 2846–2856.

Stettler DD, Das A, Bennett J, and Gilbert CD. (2002) Lateral connectivity and contextual interactions in macaque primary visual cortex. *Neuron* 36: 739–750.

Li W, Piech V, and Gilbert CD. (2004) Perceptual learning and top down influences in primary visual cortex. *Nature Neurosci.* 7: 651–657.

Sigman M, Pan H, Yang Y, Stern E, Silbersweig D, and Gilbert CD. (2005) Top-down reorganization of activity in the visual pathway after learning a shape identification task. *Neuron* 46: 823–835.

Li W, Piech V, and Gilbert CD. (2006) Contour saliency in primary visual cortex. *Neuron* 50: 951–962.

Stettler DS, Yamahachi H, Li W, Denk W, and Gilbert CD. (2006) Axons and synaptic boutons are highly dynamic in adult visual cortex. *Neuron* 49, 877–887.

Li W, Piech V, and Gilbert CD. (2008) Learning to link visual contours. *Neuron* 57: 442–451. [PMC2409109]

McManus J, Ullman S, and Gilbert CD. (2008) A computational model of perceptual fill-in following retinal degeneration. *J. Neurophysiol.* 99: 2086–2100. [PMID: 18199820]

Gilbert CD, Li W, and Piech V. (2009) Perceptual learning and adult cortical plasticity. *J. Physiol.* 587: 2743–2751. [PMC2718234]

Yamahachi H, Marik SA, McManus JN, Denk W, and Gilbert CD. (2009) Rapid axonal sprouting and pruning accompany functional reorganization in primary visual cortex. *Neuron* 64: 719–729. [PMC2818836]

Marik SA, Yamahachi H, McManus JNJ, Szabo G, and Gilbert CD. (2010) Axonal dynamics of excitatory and inhibitory neurons in somatosensory cortex. *PLoS Biology* 8(6): e1000395. [PMC2885981]

Marik SA, Olsen O, Tessier-Lavigne M, and Gilbert CD. (2013) Death Receptor 6 regulates adult experience-dependent cortical plasticity. *J. Neurosci.* 33: 14998–15003. [PMC3776054]

Ramalingam N, McManus JNJ, Li W, and Gilbert CD. (2013) Top-down modulation of lateral interactions in visual cortex. *J. Neurosci.* 33: 1773–1789. [PMC3711382]

Chen M, Yan Y, Gong X, Gilbert CD, Liang H, and Li W. (2014) Incremental integration of global contours through interplay between visual cortical areas. *Neuron* 82: 682–694.

Marik SA, Yamahachi H, Meyer zum Alten Borgloh S, and Gilbert CD. (2014) Large scale reorganization of inhibitory neurons following retinal lesions. *J. Neurosci.* 34: 1625–1632. [PMC3905138]

Abe H, McManus JN, Ramalingam N, Marik SA, Meyer zum Alten Borgloh S, and Gilbert CD. (2015) Adult cortical plasticity studied with chronically implanted electrode arrays. *J. Neurosci.* 35: 2778–2790. [PMC432540]

Marik SA, Olsen O, Tessier-Lavigne M, and Gilbert CD. (2016) Physiological role for amyloid precursor protein in adult experience-dependent plasticity. *Proc. Natl. Acad. Sci.* 113: 7912–7917. [PMC4948331]

Liang H, Gong X, Chen M, Yan Y, Li W, and Gilbert CD. (2017) Interactions between feedback and lateral connections in the primary visual cortex. *Proc. Natl. Acad. Sci.* 114: 8637–8642. [PMC5559040]

van Kerkoerle T, Marik SA, Meyer Zum Alten Borgloh S, and Gilbert CD. (2018) Axonal plasticity associated with perceptual learning in adult macaque primary visual cortex. *Proc. Natl. Acad. Sci.* 115: 10464–10469.

Astorga GL, Chen M, Yan Y, Altavini TS, Jiang CS, Li W, and Gilbert CD. (2022) Adaptive processing and perceptual learning in visual areas V1 and V4. *Proc. Natl. Acad. Sci.* 119: 36223395.

Altavini, T.S., Chen, M., Astorga, G., Yan, Y., Li, W., Freiwald, W. and Gilbert, C.D. (2024) Expectation dependent stimulus selectivity in the ventral visual cortical pathway. Submitted.