



The History of Neuroscience in
Autobiography
Volume 12

Edited by Thomas D. Albright and Larry R. Squire
Published by Society for Neuroscience
ISBN: 978-0-916110-11-6

Richard A. Andersen
pp. 0–37

<https://www.doi.org/10.1523/hon.012001>

Richard Aron



Richard A. Andersen

BORN:

New Kensington, Pennsylvania
October 27, 1950

EDUCATION:

University of California, Davis, B.S. Biochemistry (1973)
University of California, San Francisco, PhD Physiology (1979)
Johns Hopkins Medical School, Postdoctoral Fellow (1981)

APPOINTMENTS:

Assistant Professor, Salk Institute (1981–1986)
Adjunct Assistant Professor, University of California, San Diego (1982–1987)
Associate Professor, Salk Institute (1986–1987)
Associate Professor, Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology (1987–1990)
Professor, Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology (1990–1993)
James G. Boswell Professor of Neuroscience, Division of Biology and Biological Engineering, California Institute of Technology (1993–present)
T & C Chen Brain-machine Interface Center Leadership Chair, California Institute of Technology (2017–present)

HONORS AND AWARDS (SELECTED):

Sloan Foundation Fellow (1982)
McKnight Foundation Scholars Award (1983)
Spencer Award, Columbia University (1994)
Fellow, the American Association for the Advancement of Science (2000)
Fellow, the American Academy of Arts and Sciences (2002)
Visiting Professor, College de France (2005)
Member, National Academy of Sciences (2005)
Member, National Academy of Medicine (2007)
A Recipient of UCSF's 150th Anniversary Alumni Award (2015)

Richard Andersen's research uses neurophysiological, psychophysical, neuroanatomical, and computational methods to understand the processing of sensory and motor signals in the cerebral cortex. These studies have included examining the senses of sight, hearing, balance and touch, the neural mechanisms of action, and the development of neural prosthetics. He discovered gain fields, the method the brain uses to transform signals between spatial representations. He also discovered neural signals of intention in the posterior parietal cortex, proving that they are not sensory in nature but rather reflect the planning of the actor. He applied this discovery of intention signals to advance research in brain-machine interfaces, showing that paralyzed patients' intentions can be decoded from brain activity to control assistive devices, such as robotic limbs and computers. He used intracortical electrical stimulation of somatosensory cortex to show that natural sensations of touch and body position can be restored to patients with cervical spinal cord lesions that have left them otherwise insensate to somatosensory stimulation below the lesion.

Richard A. Andersen

Family Background

I was born in New Kensington, Pennsylvania, in 1950. My mother Norma was born in Queens, New York, and my father John in Brooklyn, New York. The name Andersen comes from my grandfather on my father's side who was Danish. My grandmother on my father's side was from Germany. John was a first-generation American. My mother's parents were American and their parents came from Germany, making my mother a second-generation American. My brother Cliff was born in 1954 in Buffalo, New York.

My father joined the U.S. Army at the beginning of World War II and was in officer training, coincidentally here in Pasadena where I now work. He and several soldiers contracted polio from the swimming pool at Pasadena City College, literally across the street from Caltech. His case was severe and he spent many months in a hospital. Although he recovered, his lower legs were weakened, and later in life, he suffered from post-polio syndrome, which greatly hampered his mobility. Interestingly, I am now studying ways to assist people with paralysis. My mother was his primary caregiver, with me as backup, so I have a firsthand appreciation of the suffering from loss of mobility.

My parents met through family friends in New York City. My mother was a secretary in Manhattan and my father was an engineering student at Brooklyn Polytechnic (now the NYU Tandon School of Engineering).

My father designed and oversaw the building of chemical plants, mostly for Kaiser Chemical. As a result, our family moved quite a bit, including Buffalo, New York; Columbus, Ohio; Walnut Creek, California; Baton Rouge, Louisiana; and Moraga, California. He enjoyed talking about different chemical processes and engineering, and so I was getting exposed to the world of science at an early age. In high school, I enjoyed doing science projects, including building a fuel cell and studying pollutants in the San Francisco Bay. So it was natural that when I went to college at the University of California, Davis (UC Davis), that I was a chemical engineering major for the first two years, which was a good choice for getting a quantitative education. My brother Cliff took a different career route and became a musician, mostly playing jazz guitar.

In college, I was naturally interested, as many students are, in one's own psychology. At the time, psychology had the biggest classes and was a very popular major. But my leanings were more toward understanding the biological basis of brain function. At the time, there was not a major for neuroscience, so I pieced together a degree in biochemistry that would prepare me best for studying the brain. I was particularly lucky in my last

two years to be hired to do research in Robert Scobey's lab. Scobey studied the processing of visual information by retinal ganglion cells, and I participated in single neuron recording experiments in the retinas of cats, examining the response of these cells to small displacements of visual stimuli.

I also enjoyed nonacademic pursuits in college; maybe too much—but it was, after all, the late 1960s and early 1970s, which were socially dynamic times. I started a student newspaper and was a disc jockey at the college radio station, hosting a jazz and blues radio show. One of my friends went to work at a radio station in the Bay Area, and he hoped I would join him. DJ or scientist? I did go to San Francisco, but chose the latter career path.

Graduate School

After receiving my B.S. from UC Davis I entered the Physiology Graduate Program at UC San Francisco (UCSF). Again there was no formal neuroscience or neurobiology program, but UCSF had a number of young and now famous neuroscientists on the faculty at that time.

During my first year at UCSF, I worked in the lab of Ben Libet (see Volume 1). He was collaborating with a neurosurgeon on the timing of the awareness of sensation in patients. Later Libet did a groundbreaking experiment in which he found that the awareness of movement intentions came after the movement plans had already been formed. He interpreted the results as showing that the sense of free choice was not causal to action planning. He wanted me to work on his other main topic, which was neural transmission in the superior cervical ganglion by way of second messengers. Although these were interesting experiments, I decided to transfer to another lab, that of Michael Merzenich (see Volume 7).

Merzenich came from the Hopkins school of neurophysiologists, being a graduate student with Vernon Mountcastle (see Volume 6). At Hopkins, Merzenich studied the somatosensory system and, as postdoctoral fellow with Jerzy Rose from the University of Wisconsin, subsequently studied the auditory system. Merzenich was one of the primary developers of the cochlear prosthetic. Seeing the neural plasticity in humans who learned to interpret the rudimentary artificial signals from the prostheses, led him to become intensely interested in brain plasticity. He next performed famous experiments with Jon Kaas and colleagues on the plasticity of cortical somatosensory maps after deafferentation from peripheral nerve lesions. Merzenich was a tremendous and inspirational mentor who emphasized studying big ideas.

Thesis Research

In Merzenich's lab, I studied the anatomy of the auditory system in cat, detailing the topography of the cortical-cortical, cortico-thalamic, and

cortico-collicular connections. From these experiments, I found two parallel systems. Highly tonotopically organized cortical fields were connected to corresponding tonotopic regions of the thalamus, whereas another, nontotopically organized cortical area was connected to more complex and multimodal areas of the thalamus. These two pathways, which I called “tonotopic” and “diffuse,” suggest that the former system is very much like visual cortex with its strong topography, and the latter is highly integrative (Andersen, Knight, and Merzenich 1980).

Perhaps the most remarkable finding from my thesis research was that all cortico-cortical and thalamocortical pathways are highly reciprocal. We were using new anatomical tracers, one that was transported retrogradely and another anterogradely. By mixing the anterograde and retrograde tracers and making punctate cortical injections, a technique developed by one of my fellow graduate students Steve Colwell, the resulting labeling showed nearly identical patterns of projection for both tracers. This result indicated a precise reciprocity of connections between cortical areas and between the cortex and thalamus. Although other track tracing techniques had by that time shown that these connections were reciprocal, the newer tracers that we employed, horseradish peroxidase and tritiated amino acids, provided the first detailed demonstration of the spatial precision of this reciprocity. These results were similar to Colwell’s earlier findings in the rabbit visual cortico-thalamic system. They bear out the notion that all cortico-cortical and cortico-thalamic connections are precisely reciprocal, which to my knowledge is a rule to which subsequent research has not found an exception.

First Interest in Posterior Parietal Cortex

UCSF, which at the time was located only at the hospital on Parnassus Avenue, had a tremendous medical library. It was there that I would read books and papers by famous neurologists like Critchley, Luria, Geschwind, Holmes, and Balint. I found the deficits from posterior parietal lesions particularly fascinating. For instance, Balint’s syndrome, first described by the Hungarian neurologist Rezső Balint in 1909, results from bilateral lesions of the posterior parietal cortex (PPC). These subjects have difficulty making voluntary eye movements, misreach for objects in the peripheral visual fields, and cannot perceive multiple objects simultaneously. Another example deficit is neglect, which results from unilateral lesions, particularly of the right hemisphere of PPC. Patients with neglect are not blind but have an attentional deficit in which they are seemingly unaware of visual space contralateral to the lesion. I had the opportunity to go on rounds at a hospital while a graduate student and observed patients with PPC lesions. I was struck by how similar these patients’ deficits were to those described in the literature.

Postdoctoral Study

Around the time of my graduation, a few labs were using the new technique of single neuron recording in behaving non-human primates (NHPs). Luck would have it that one of the labs was Vernon Mountcastle's, and that Mountcastle was examining higher cortical functions in the PPC of NHPs. Following Merzenich's advice, I contacted Mountcastle and he accepted me for postdoctoral study. Mountcastle was a giant in the field, having discovered the cortical column, a shaft of neurons arranged vertically within cortex with similar response properties. These columns are the basic building blocks of much of cortex. Mountcastle was the world leader in quantifying the relation of neural activity in the somatosensory system to perception. I had heard stories that he was a tough taskmaster and, since I was a bit of a night owl, Merzenich thought the discipline of working in Mountcastle's lab would be good for me. However, on my arrival in Baltimore I found Mountcastle to be an amazingly devoted, inspiring, and exacting teacher. He had gathered at Hopkins top researchers in NHP research, including faculty Apostolos Georgopoulos, Mahlon DeLong, and Gian Poggio. Postdoctoral fellows beside myself in Mountcastle's lab were John Kalaska, Roberto Caminiti, and Brad Motter, who all went on to stellar careers in systems neuroscience. Around the time of the move I married Carol Ahern, an audiologist. We met at a seminar on the cochlear hair cells of the inner ear at UCSF. We moved together to Johns Hopkins in Baltimore, where she entered the master's program in audiology at Hopkins.

Mountcastle was interested in how perception leads to action along cortical pathways. This led him to the PPC because it is the major association cortical area bridging between somatosensory cortex and premotor cortex. His experiments on this sensory-to-motor transformation required recording from neurons while animals experienced perceptual goals and directed actions toward these goals. With the help of Ed Evarts from NIH, who was recording from the motor cortex of behaving NHPs, Mountcastle retooled his lab to record from PPC in behaving NHPs. What he found has provided the foundation for research in PPC to this day. In landmark papers in 1975 and 1977, Mountcastle and his colleagues described neural responses to specific actions and high-order sensory phenomena. Numerous cells were selective to the particular voluntary oculomotor behaviors of saccades, smooth pursuit, and fixation. Others responded to reach or hand manipulation of objects. High-level visual sensory properties included modulation by attention, and responses to global motion patterns that are indicative of flow fields during locomotion. Mountcastle proposed a command hypothesis in which PPC integrated sensory information and emitted command signals for action, which were elaborated by downfield motor areas. In 1978, Michael Goldberg (see Volume 10) and colleagues from NIH published a paper in NHPs in PPC. They replicated many of the findings of Mountcastle, but they

proposed a different interpretation in which the observed changes in PPC activity were better explained by visual attention. It is not too surprising, in retrospect, that an area involved in visuomotor transformations would have both high-level motor and sensory properties.

Gain Fields

My main research at Hopkins involved the discovery of “gain fields.” The PPC has long been implicated in spatial awareness, as evidenced by spatial-perceptual deficits in humans with lesion to this area. Among the spatial deficits, people fail to be aware of visual stimuli in the contralateral space, show spatial memory defects, are aware of only half an attended object, and have difficulty in navigation. From these observations, I reasoned that the PPC would have visual receptive fields anchored to the outside world. In healthy subjects, this hypothesized anchoring could account, for instance, for the world appearing to be stable in spite of the many eye movements we make, which constantly change the images on the retinas. Likewise, a coordinate transformation from retinotopic receptive fields in early visual areas is required for observers to make visually guided body movements. An example is that observers can reach to an object accurately, independent of gaze direction.

What Mountcastle and I found was that the visual receptive fields in PPC remained retinotopic—that is, anchored to the retina. But we also found that many of the visually responsive neurons are also modulated by the direction the eyes are looking in the orbits of the skull—that is, they convey eye position signals (Andersen and Mountcastle 1983). Moreover, while some neurons carried only eye position or retinal signals, it was most common that individual neurons carried both signals. This convergence of different eye position and retinal signals on individual neurons in PPC could conceivably represent locations with respect to at least the head, invariant of eye position. Mountcastle felt rather that the systematic modulation of the gain of the visual response by eye position had a role in attention, not coordinate transformations, and so my first paper on gain fields had essentially no discussion.

Faculty Years: Parallel Paths in Research

Over my past 40 years of research, and still going, I followed a few selective lines of research. I was a faculty member at three institutions, first the Salk Institute, then MIT, and then Caltech. I will follow the paths of research through these great institutions in chronological order, but my major research interests spanned these sites. From my postdoc days, I continued my interest in spatial representations and coordinate transformations. A second interest was, and is, the encoding of intent in the posterior

parietal cortex. And a third is studying structure from motion processing in the pathway leading from primary visual cortex into the PPC.

Along the way, I transitioned from primarily using the behaving NHP paradigm to performing clinical studies in humans (although we are still doing NHP studies). This path was made possible by applying what we learned about the PPC in NHPs to translate to neuroprosthetics for tetraplegic participants. We use brain-machine interfaces (BMIs) which record brain activity from electrodes implanted in the brain and translate the signals to control external devices such as robotic limbs or computers. BMIs typically have used motor cortex activity for control. We refer to our approach as cognitive neuroprosthetics because we implant the electrodes in PPC, a high-level/cognitive brain area, and we record the intent of the individual. As one might predict from association cortex, recording in PPC enables access to a plethora of different signals, from intended movements of all parts of the body to high-level cognitive signals, such as those related to semantics and memory decisions. This has been a thrilling ride—that is, to do basic research in NHPs and then to translate that research directly in clinical studies. Moreover, the work in humans has led to new scientific discoveries that contribute to new prosthetic applications.

Needless to say, BMIs allowed us to sail into rather uncharted territory, as we had never seen the human cortex functioning at such fine temporal and spatial scales. And, unlike animal studies in which inner mental processes are at best inferred, we can ask the participants how they did a task, what they felt with microstimulation, or what they imagined introspectively. This new clinical direction has left me at this point in my career excited every day about what new results we will find and how they will contribute to helping individuals with paralysis.

The Salk Institute

In some ways, my disagreement with Mountcastle about gain fields left the field of coordinate transformations open to me, and I continued studying this topic when I moved to my first faculty position as an assistant professor at the Salk Institute in 1981. Mountcastle felt it would be a good move for me to go to the Salk. I could concentrate on research without teaching or administrative duties, as the Salk was small and almost entirely research oriented. Mountcastle told me his first few years as a faculty member at Hopkins were similar in that he had few teaching and administrative responsibilities. Carol and I had our son Michael at the Johns Hopkins hospital and three months later we moved to San Diego.

At the Salk Institute, I was within the division directed by Max Cowan, a well-known neuroanatomist (see Volume 4). The process of setting up my behaving monkey lab and collecting the first neural data would take some time. So, to get an early start, I took advantage of working in Cowan's

state-of-the-art neuroanatomy lab to begin examining the neural connectivity of the PPC. In Mountcastle's papers, he had found a clumping of response properties for saccades, pursuit, and reaching. He speculated that a columnar organization in PPC was made up of mixed saccade, pursuit, and reach columns.

I reasoned that the best way to visualize the saccade columns in PPC was to make tracer injections into the frontal eye field (FEF). The FEF is a cortical region in the frontal lobe specialized in producing voluntary saccades, and a presumed target of saccade neurons in PPC. Instead of finding columns in PPC, the major reciprocal connections with FEF were mostly with a single cortical area within the lateral bank of the intraparietal sulcus, which we named the lateral intraparietal area (LIP) (Andersen, Asanuma, and Cowan 1985). LIP has become popular for the study of early eye movement processing, including movement planning, attention, and decision making. Other columns postulated by Mountcastle in PPC also turned out to be cortical areas, including the medial superior temporal area for pursuit eye movements and the parietal reach region for reaching.

With Jim Gnadt, a postdoctoral fellow at the time and now a program director at the National Institute of Neurological Disorders and Stroke (NINDS), we studied the saccade properties of LIP neurons. We found that if (1) a saccade target was briefly flashed in the visual response field of an LIP neuron, and (2) the monkey had to withhold a saccade to that target location for a delay period, the neuron remained active throughout the delay period. To show that this persistent activity was not a memory signal for the location of the visual target, we also used a task in which the flashed saccade target (the cue) was delivered outside the neuron's response field. The monkey was then instructed to make a saccade that brought the cued location into the response field. In this case, the delay period activity emerged when the previously cued saccade target location was now in the response field of the neuron. Jim and I reasoned that the activity represented the intended movement of the animal, and not the memory of the location of the target (Gnadt and Andersen 1988).

My early description of gain fields with Mountcastle was limited to only a few eye positions and retinal visual stimulus locations. With postdoctoral fellows Ralph Siegel and Greg Essick, we continued the examination of gain fields in PPC with much more thorough mappings of the visual receptive fields at different eye positions. These new findings reinforced our earlier interpretation that the response fields were strongly retinotopic, and the modulation by eye position could be best modeled as the product of an eye position gain factor on a retinotopic visual receptive field. Likewise, these new results strengthened the interpretation that the eye position-dependent modulation produced a tuning for locations in head-centered space (Andersen, Essick, and Siegel 1985).

As mentioned, a strategy I used to jump-start my lab at the Salk was to do anatomical experiments, which were sure to render quick results, while

building up a recording lab for behaving monkeys. Another strategy was to pick two major lines of research as a precaution in the event that one did not succeed. The first, outlined earlier, was to study the visual-motor properties of PPC neurons. The second was to study motion perception. The motion pathway had already been well established anatomically as proceeding from layer 4b neurons in the primary visual cortex (V1), to the middle temporal area (MT), to the medial superior temporal area (MST) in PPC. In the end, both lines of research were successful and certainly kept me busy.

The three-dimensional (3D) structure of objects can be perceived from two-dimensional (2D) motions of random moving dots that are arranged as if they were dots painted on a rotating object projected onto a 2D screen. A nice feature of this stimulus is that it is ambiguous, and sometimes the object is seen to rotate in one direction, but then it flips to appear to be rotating in the opposite direction. This stimulus can be made unambiguous by adding disparity to the dots. Thus, I started investigating structure-from-motion (SFM) along the motion pathway in NHPs. Beginning at the Salk with Ralph Siegel, we showed in psychophysical experiments that monkeys perceive SFM and that both monkeys and humans integrate spatial and temporal parameters of SFM similarly (Siegel and Andersen 1988).

A major influence on me during my time at the Salk Institute was Francis Crick. I had read James Watson's book, *The Double Helix*, which begins with the famous line "I have never seen Francis Crick in a modest mood." It was thus with some initial trepidation that soon after I arrived at the Salk, Francis walked into my office, extended his hand, and introduced himself. To my delight, he became a supportive and brilliant mentor.

Crick had taken an interest in consciousness as the other major question in biology from the one he had already tackled successfully: what is the genetic code? To this end, Crick; V. S. Ramachandran, a vision scientist from UC San Diego (UCSD); and Gordon Shaw, a physicist from UC Irvine (UCI), organized the Helmholtz Club. There were about 20 members from around southern California. Besides those of us from San Diego, regulars included Stan Schein, Joaquin Fuster (see Volume 7), John Schlag, and Madeleine Schlag-Rey from UC Los Angeles (UCLA) and John Allman and David Van Essen (see Volume 9) from Caltech. Every month or so, we would have a meeting at the UCI faculty club in which two renowned speakers from around the world gave extended talks with much interaction followed by dinner at a nearby restaurant. These were lively meetings. I recall that Richard Gregory, a famed British psychophysicist, barely got past his first slide before there was a vigorous, long debate with Francis. Of course, this was a bit of an exception from most presentations, but the long format provided a tremendous avenue for presenting research in depth with lively discussion.

One nascent field that was emerging from these meetings, from neural network researchers at UCSD, and from artificial intelligence (AI) researchers

at MIT, was computational neuroscience. This move toward more theoretical neuroscience helped me greatly to formulate my own research with more of a theoretical perspective.

I was particularly fortunate that right across the street from the Salk, at UCSD, there was the Parallel Distributed Processing group, a collection of mathematical psychologists, including David Rumelhart, Geoffrey Hinton, and James McClelland, who were advancing connectionist models of cognition. In 1986, this group developed backpropagation, a method for training multilayered neural networks by propagating error back through the networks. In 1988 David Zipser, a member of this UCSD group, and I published a paper in *Nature*, which was one of the first uses of neural networks to model a brain computation and to compare the network to neural data (Zipser and Andersen 1988). We trained a three-layer neural network that received retinotopic targets and eye position signals as inputs and, as outputs, produced receptive fields that coded the targets in head coordinates. After training the network with backpropagation, we found that the middle layer units developed gain fields, with retinotopic visual fields modulated by eye position very similar to the recorded neurons in PPC. Thus, the network model demonstrated that gain fields could be a method of computation for making coordinate transformations.

San Diego was a great environment for me and my family. Our daughter Kristen was born there. The Salk Institute was a great place to focus on basic research. We lived in Encinitas, a beautiful beach town just north of San Diego. I even learned to surf when I was there. In part through Crick and the Helmholtz Club, I also got to know several of the faculty in the Brain and Cognitive Sciences Department at MIT. MIT was the center for AI. At the time, AI referred to more algorithmic/symbolic theories of brain function. Many neuroscientists in my generation were deeply influenced by the posthumously published book *Vision* by David Marr, one of the luminaries in the MIT AI group. His book emphasized three levels for understanding the visual system: computation, algorithm, and hardware. The MIT faculty in Brain and Cognitive Sciences was an impressive group of neurophysiologists, cognitive scientists, and AI researchers. I was honored to receive a job offer from them, and it was an opportunity that was too great to pass up. So in 1987, we left the West Coast behind and moved to Boston. My family and I settled outside of Boston in Wellesley.

MIT

At the Salk, I had the only behaving monkey lab. At MIT, there were two giants of behaving monkey research. Emilio Bizzi was one of the founders of the field of motor control, studying how the motor system controls the trajectories of the limbs. He was also the chair of the department and was very supportive, patient, and a great leader of the faculty (see Volume

6). Peter Schiller was an eminent visual neurophysiologist who has made many seminal findings of the cortical processes of vision and attention (see Volume 7). Ann Graybiel was the world expert on the anatomy and functioning of the basal ganglia, and Sue Corkin was a renowned neuropsychologist. I mentioned the AI component of the department that included the giants in the field, among them Tomaso Poggio (see Volume 8) and Simon Ullman. It was a most fertile ground for the study of systems neuroscience.

Structure from Motion

MIT was a perfect place to study SFM. Faculty members Shimon Ullman and Ellen Hildreth were world experts on the topic. At the time, there was a running challenge in the literature to see what algorithms could determine whether a rotating object was rigid or not from viewing the fewest dots and the fewest frames.

We chose the topic of how the perception of surfaces is produced from SFM. When viewing a few dots on the surface of an otherwise invisible cylinder, one gets the clear impression that there is a surface between the dots. With postdoctoral fellow Masud Husain and graduate student Stefan Treue, we used random dot patterns that had limited lifetimes before they were repositioned to another location on the surface of the object defined by motion. We found that humans would integrate the information from such displays spatially and temporally across several point lifetimes, providing strong evidence that the brain constructs perceived surfaces from transient velocity information to produce more global surfaces. In fact, points could be left out of a quarter of an object's surface and the interpolation process would still fill in a perceived surface. Ellen Hildreth, collaborating with us, showed that a computational model of surface interpolation could account for all of the aspects of SFM perception that we had found.

The SFM displays we used had no disparity depth cues, being presented as moving dots on a flat screen. Yet, even viewed with just one eye, the motion fields produced vivid percepts of a 3D object. How is it possible for the nervous system to separate these directions into separate surfaces, since each patch of the display contains motions moving in different directions? Particularly problematic is the mechanism of motion perception itself, which is a result of inhibitory interactions between neurons. The question has even broader implications about how we can perceive transparency. For instance, we can make out both the reflection on a window and the scene behind it. We postulated that there would be two stages in the visual pathway. The early stage would measure local motions. The later stage would integrate the motion signals from the earlier stage more globally to construct the perception of surfaces.

Ning Qian, a postdoc in the lab, had the brilliant idea of presenting exactly balanced motion signals at every sample location on the surface.

This idea was motivated by the observation that counter-phase sine-wave gratings, which are mathematically equivalent to two sine wave gratings drifting in opposite directions, produce no perceived motion but rather static standing wave oscillations. We paired each moving dot with another dot moving in the opposite direction. This display destroys the perception of moving transparent surfaces, and rather appears like a field of flickering points.

Our recording experiments in areas V1 and MT of NHPs explained the phenomena of transparent surface perception. V1 neurons measured local motions with their small receptive fields, and did not reliably distinguish between the paired, perceptually nontransparent stimuli and the unpaired, perceptually transparent stimuli. In contrast, the directionally selective MT neurons, with their large receptive fields, responded better to the nonpaired stimuli than to the paired stimuli. Ted Adelson is a computational neuroscientist at MIT who introduced the concept of motion energy to explain neural interactions that produce direction selectivity. With Ted, we built a computational model and tested it with the same stimuli we used in the NHP experiments. The first stage models V1, which measures local motions and does not distinguish between transparent and nontransparent motions. The second stage models MT, in which motion energies of opposite direction but the same spatial frequency and disparity contents suppress one another. Like the MT neural responses, this second stage distinguishes between transparent and nontransparent motions. These results were published in three companion papers in the *Journal of Neuroscience*, which beautifully explained transparent motion perception based on psychophysical measurements and neural recording data, and conceptualized the process with a computational model (Qian, Andersen, and Adelson 1994a, 1994b; Qian and Andersen 1994).

Optic Flow

Another important function of motion is determining the direction of locomotion for navigation. The next stage in the V1-MT pathway is the dorsal division of the medial superior temporal area (MSTd). This area has large, bilateral receptive fields that are selective for patterns of motion that one sees when moving through the environment. Neurons were described that were sensitive to expansion/contraction (like a field of stars while moving through space in a science fiction movie), rotation, and translation. These different motion patterns could be decomposed by the nervous system into three channels (expansion/contraction, translation, and rotation) for analyzing the optic flow field. A possible advantage of such a decomposition is that translation motion due to eye movements could be subtracted from expansion components due to locomotion to determine the focus of expansion, which would correspond to the direction of heading.

Mike Graziano, Bob Snowden, and I tested this decomposition hypothesis by constructing a 2D spiral motion space, which had on the horizontal axis clockwise/counterclockwise rotation, and expansion/contraction on the vertical axis. A neuron tuned to 45 degrees in this space, for instance, would be selective to a clockwise outward spiral. We found many MSTd neurons that were tuned to spirals, indicating that a simple three channel decomposition is not correct, and rather there is a continuum of motion pattern selective neurons. We also found that the tuning was invariant for patterns over these large receptive fields. Thus, single neurons cannot determine accurately the direction of heading from expansion and a course coding at the level of populations of neurons would have to account for the precision of heading direction from optical flow patterns (Graziano, Andersen, and Snowden 1994). In a subsequent experiment with Bard Geesaman, we showed that the tuning of MSTd neurons to complex motions was invariant for the type of stimuli, whether they were moving dot patterns, solid squares, or a figure/ground display (Geesaman and Andersen 1996). Thus, MSTd neurons would appear to be involved in both the analysis of object motion as well as optic flow from observer translation.

Saccades and LIP

After having discovered LIP at the Salk Institute, two postdoctoral fellows in my lab at MIT, Shabtai Barash and Leonardo Fogassi, and graduate student Martyn Bracewell, began looking at LIP response properties in detail. LIP could be differentiated from adjoining area 7a of PPC by saccade timing. We found that the saccadic responses of most LIP neurons began just before saccades, whereas responses in 7a were mostly postsaccadic. The response fields in LIP were well defined, and their visual, delay, and saccade responses were spatially overlapping. Peter Their, who was doing a sabbatical in the lab, showed that microstimulation of LIP produced saccades, although requiring higher currents than in the frontal eye field in the frontal lobe. In experiments later carried out at Caltech, we further showed that pharmacological inactivation of LIP produced deficits in the accuracy and dynamics of saccades, particularly for saccades to remembered locations and primarily in the visual field contralateral to the inactivation. The combination of evidence from neurophysiology, neuroanatomical connectivity, microstimulation, and inactivation led us to conclude that LIP is a posterior eye field for saccades in parietal cortex, much like the frontal eye field is for saccades in the frontal lobe.

Gain Fields

While at MIT, we also extended research on gain fields. Previously, we had shown gain fields in area 7a of PPC, but in further study, we found them to

exist in LIP as well. Large shifts in gaze combine eye and head movements. Interestingly, we discovered that neurons in LIP had both eye and head gain fields and that these gain fields were aligned (Brotchie et al. 1995). Thus, the gain fields in LIP are, in fact, modulating visual responses by the gaze direction with respect to the body.

We also extended our theoretical studies of gain fields. We found that microstimulation of different locations in LIP produced saccades in different directions toward the contralateral field. However, the initial eye position fixation modulated the amplitude of the saccades, similar to simulations of microstimulation of middle-layer units in the neural network models of PPC. As a serendipitous control, we found that microstimulation of cortex in the intraparietal sulcus just ventral to LIP produced complicated and often head-centered goal-directed saccades when varying eye position. These stimulation results were consistent with retinotopic-eye position modulated neurons in LIP, which could converge onto the more ventral location in the intraparietal sulcus to code goals in head-centered coordinates.

A common criticism of neural network modeling of brain processes at the time was that backpropagation learning is not biologically plausible. With Pietro Mazzoni, a graduate student in the lab, and Michael Jordan, a faculty member in our department at MIT, we showed that a biologically plausible reinforcement model produced gain fields when training neural networks to convert stimulus locations from retinal to head centered, and with approximately the same number of training cycles as with backpropagation (Mazzoni, Andersen, and Jordan 1991). Thus, our networks were consistent with long-term/skill learning with biologically realistic constraints.

Caltech

When I was up for tenure at MIT, which I received, I was also offered positions at Rockefeller University and the California Institute of Technology. My wife Carol was a fourth-generation Californian on one side of her family and fifth on the other. MIT was an excellent academic environment for me, and New York City is in my roots, but for Carol, the desire to get back to California was very strong. Needless to say, Caltech and MIT have numerous similarities, particularly in my areas of interest of the visual system and computational neuroscience. Given these excellent opportunities it was not an easy decision, but family won out and we moved to California in 1993. My son was in sixth grade at the time and my daughter was in third. The transition for them was pretty smooth. My son and his wife are now graphic designers living in Brooklyn, back near where my parents grew up, and my daughter is a pediatrician in Redondo Beach, just south of Los Angeles.

An advantage of moving is that one can design a new lab fitted to exploring new research directions. In my case, I wanted to add vestibular and

auditory stimuli to the arsenal of sensations we could study. This opportunity was particularly important because the PPC is an association cortical area, and these senses, along with somatosensation and vision, all converge on this multimodal area (Andersen et al. 1997).

Caltech: Monkey Days

Intention

Intention and attention are intertwined. It was argued by some that the delay activity we saw in LIP was a correlate of attention. In fact, when one plans to look somewhere, visual attention is also directed to that location. But attention is a sensory process. If you plan to reach to a location, then attention is likewise directed to that location. But the plans, to saccade or reach, are very different since the action involves separate effectors. In experiments with Larry Snyder and Aaron Batista, we reasoned that LIP would be more active for planned eye movements than for planned reaches if it was coding intent, but would show similar activities if it was encoding spatial attention (Snyder, Batista, and Andersen 1997). In fact, we found the former. Serendipitously, we also found that regions within the intraparietal sulcus on the medial bank showed the opposite pattern of activity, being selective for reaches but not saccades. We termed this area the parietal reach region (PRR), which included the medial intraparietal area and area V6. Another area in the intraparietal sulcus that is anterior to LIP is the anterior intraparietal area (AIP). Hideo Sakata and colleagues found that single neurons in AIP coded entirely different grasps. These findings of eye movement, reach movement, and grasp movement fields in the intraparietal sulcus led Chris Buneo and me to propose that there is a map of intentions within PPC (Andersen and Buneo 2002).

In further experiments, we showed that PPC can plan simultaneously a sequence of reaches. Also, if NHPs do not know early in a trial whether a reach or a saccade will be required to obtain a target, they will plan movements to both; in other words, PPC can form intended movement plans that can be scrubbed later.

Gain Fields and Multimodal Integration

In the move to Caltech, we installed an anechoic chamber for auditory experiments, which has no reflections of sound off the walls of the chamber and is ideal for sound localization. Another addition was a 3D vestibular chair that could move the NHPs around all axes in three dimensions and record eye position at the same time. Included in our arsenal of new paradigms was training monkeys to rotate their heads on their bodies without changing gaze direction.

Not surprisingly, given that PPC is an association area that integrates different sensory modalities, we found that PPC neuronal activity was affected by vision, audition, and vestibular inputs, but in very specific ways. When passively stimulated with auditory and visual stimuli, in which the stimuli had no behavioral significance, LIP neurons respond only to visual stimuli. Thus, in terms of passive stimulation, LIP would be considered a visual and not an auditory sensory area. This preferred selectivity for visual stimulation is not surprising given its role in eye movements. However, graduate student Jennifer Linden and postdoctoral fellow Alexander Grunewald found that when saccading to sounds in space, LIP is also activated by the auditory stimulus. Auditory targets are initially encoded in head coordinates in early stages of the auditory system, given that locations of sounds in space are derived from intra-aural cues of phase and intensity at the two ears. However, the auditory stimuli in LIP were represented toward eye coordinates, although many were intermediate between eye and head coordinates. Thus, the auditory target representation is shifted toward a common coordinate frame with visual stimuli in LIP for making saccades (Cohen and Andersen 2002). Eye position modulated the auditory responses, suggesting that gain fields also account for the coordinate transformation of auditory space from head to eye coordinates.

Another way that gain fields, and presumably coordinate transformations, are handled in PPC came from studying how head position signals are represented in LIP and area 7a. By using various combinations of head position with respect to the body and world using a vestibular chair, we found that neck proprioceptive signals produced gain fields in LIP, whereas vestibular signals also accounted for gain fields in area 7a (Snyder et al. 1998). Thus, LIP, using the sense of the position of the head on the body, could code gaze location (eye plus head position) with respect to the body. Area 7a projects to the presubiculum of the hippocampal formation and, using the vestibular sense of the position in the world, could code visual stimuli in the world (so-called allocentric coordinates). This result is consistent with the finding of place and grid cells, which code in allocentric coordinates, in the hippocampal complex.

The Parietal Reach Region

The natural coordinate frame for coding a visually guided reach is for the visual target to be in the coordinates of the hand—that is, the movement vector from the hand to the target. But surprisingly, we found that PRR codes reaches, but in eye coordinates, not in hand coordinates (Batista et al. 1999)! Again gain fields were present in PRR, and they included modulations for gaze and limb position. Thus, the signals were all there for converting visual targets to hand coordinates further along in the sensorimotor pathway for reaching.

Gain fields for eye position have been measured by a number of groups since our initial findings in PPC and in a large number of cortical areas. These findings suggest that it is a general method for making coordinate transformations. However, to sample even two variables, for instance, eye and target position, in two dimensions (horizontal and vertical) in different combinations, requires a large number of trials. This experimental version of the “curse of dimensionality” is even greater if a third variable, hand position, is included. Chris Buneo came up with a novel idea, which was to use a matrix design for plotting pairs of the three variables (gaze, target, and hand position) in three matrices, and then to calculate the gradients of the matrices for the firing rate changes. This technique could distinguish gain fields from coordinate frames. For example, when plotting the firing rate of a neuron as a function of horizontal limb position and horizontal retinal stimulus location, the activity will vary only as a function of limb position for a gain field. For a visual stimulus in hand coordinates, the activity will vary for both hand position and retinal location to maintain selectivity for a location with respect to the hand. With this technique, we reaffirmed that PRR codes visual targets in retinal coordinates, but moving forward in the sensorimotor pathway to area 5d in PPC, anterior to PRR, we found a transition to representing the target in hand coordinates (Buneo et al. 2002). Interestingly, we also found that the coordinate representations in area 5d were dynamic. In the task, initially only cues of where to look and where to initially place the hand were present. In that case, the neurons coded only the hand position relative to the gaze direction (hand-eye coordinates). When the target for the reach appeared, the reference frame switched to coding the hand relative to the target (Bremner 2014). In our subsequent PPC recordings in humans, we have found similar coding of imagined hand movements to visual targets to be in hand coordinates, suggesting that the PPC recordings in humans are further downstream, like area 5 of the NHPs. We use imagined hand position because the subjects are tetraplegic and cannot actually move their limbs. We also found that the dynamics in the coordinate frames depend on the tasks. If the human subjects are making eye movements, the PPC neurons code the target in retinal coordinates, and if they are imagining hand movements, the same population of neurons code the target in (imagined) hand coordinates.

Perhaps the most direct evidence suggesting that the gain fields are important for coordinate transformations comes from the inactivation of PRR in NHPs. Although the representation of targets in PRR are in retinal coordinates, modulated by eye position, we found that inactivation of PRR produced one of the hallmarks of Balint’s syndrome, optic ataxia, which is the misreaching to visual targets. These reach inaccuracies were not due to a general sensory effect, since saccade deficits were not found. Thus, the inactivation of PRR interferes with the spatial transformation of the retinotopic locations of stimuli to hand coordinates for spatially accurate reaches.

In other words, PRR plays a critical role in coordinate transformations (Hwang et al. 2012).

It is remarkable that so much of the cortex is affected by eye position. What is the source of this signal? It could be derived from proprioceptors in the extraocular muscles. Alternatively, it could be derived from an efference copy of the command to move the eyes. One study suggested that the gain fields in LIP are too slow to follow the rapid number of eye movements we make, whereas another found that their dynamics in LIP match the psychophysics of perceived stability of the world across eye movements. Using statistical modeling, we found that there is even a prediction in the population response of LIP neurons of where the new position of the eyes will be prior to an eye movement (Graf and Andersen 2014). Because proprioceptive and efference copy signals are prominent in cortex, it is likely that both contribute to gain fields.

Internal Models

In studying the dynamics of reach movement-induced activity in PPC of NHPs, we found that the responses were too fast (occurring before reaches) to be sensory feedback, and too slow to be the command signal to move the limbs, since PPC activity began after M1 activity. This finding we interpreted as consistent with an internal model of the state of the limb (Mulliken, Musallam, and Andersen 2008). Internal models are popular in the motor control literature, as they can take out the lags introduced in the nervous system due to brain processing. These internal models are updated by sensory inputs as a consequence of movement as well as monitoring movements with internal copies of the commands (efference copies) sent out by motor cortex. But in PPC, they are also likely to be more general and provide a model of the position of the body in the world. This idea of PPC being a site for an internal model has been proposed by neuropsychologists studying deficits in the PPC, in which patients are slow or unable to correct movements. Interestingly, in our human experiments, we found that receptive fields for reach can be changed even when the tetraplegic subjects, who cannot move their limbs, imagine their hands at different locations in space. Presumably, imagination is able to access the internal models.

Free Choice

A large body of research focuses on decision making in NHPs. Generally, animals choose between alternatives that vary the amount, probability, or type of reward. In orbital frontal cortex, decisions are made in selecting between potential types of rewards. In PPC, decisions are made between actions. Alex Gail at Gottingen showed that potential action plans are

represented in PPC for two potential targets, followed by a decision favoring one plan. This result is consistent with what we found with decisions between effectors. In our study with He Cui, we found that a potential plan for reach is made in PRR and for a saccade in LIP, and when the animal is cued to make a free choice between saccading and reaching, one plan remains and the other disappears (Cui and Andersen 2007). This action choice mechanism, and preplanning, affords rapid decisions and actions over short timescales between competing alternatives.

We asked whether PPC is actually part of the circuit involved in decision making, or whether it just reflects decisions made elsewhere. The effector selectivity of PRR and LIP allowed us to address this question. NHPs were given the choice to move to two targets, one in each visual field, to obtain equal reward for either choice. Postdoctoral fellows Vasileios Christopoulos, James Bonaiuto, and Igor Kagan showed that PRR inactivation biased the animals to make more ipsilesional choices for reaches, but not saccades. The reverse was found for LIP, biasing saccades more than reaches. The dorsal pulvinar is a major thalamic input to PPC. Postdoctoral fellows Melanie Wilke and Igor Kagan found that inactivation of the pulvinar produced a bias for ipsilesional saccade choices when saccade targets were simultaneously presented in the two visual hemifields. If, however, rewards were increased for the contralesional saccades, this bias was largely erased. Thus, the animal could see the target in the contralateral visual field, but it became less appealing and thus was chosen less often. Increase in the salience of the contralesional target could not explain the magnitude of the reward effect, because increasing the contrast of the contralesional target produced much less modulation of the decision bias than increasing the reward.

A major deficit of PPC lesions, particularly to the right hemisphere, is neglect in which patients fail to be aware of stimuli located in the contralateral space. What has been thought previously to be a milder form of neglect is extinction. For extinction, subjects can see stimuli in the contralesional visual field. However, when a competing stimulus is simultaneously presented to the ipsilateral space, the patients report only the ipsilesional stimulus. Our experiments with changing reward between competing saccade targets after pulvinar inactivation suggest that an element of extinction may be a deficit in decision-making circuitry. These inactivation data explain why patients with PPC lesions tend to explore the contralesional space with eye movements less often.

If PPC is involved in action decisions, how do other areas in the frontal lobe coordinate the action plan with PPC? Using simultaneous recordings from PRR and dorsal premotor cortex (PMd), we found that correlations between these two areas are stronger during free choices than during instructed movements. This finding suggests that there is a subpopulation of neurons in the frontal and parietal cortex that orchestrates activity between these areas during free choice (Pesaran, Nelson, and Andersen 2008).

Visual Motion

At Caltech, we made great strides in understanding the neural mechanisms for SFM perception and optic flow navigation. These studies were made in behaving NHPs with single neuron recordings from MT and MSTd. Much of this research was headed by a brilliant postdoctoral fellow, David Bradley.

Structure-from-Motion

At MIT we had found that V1 neurons measured local motions for transparent two-surface random dot patterns and that the responses of MT neurons to the transparent stimuli were suppressed by the two surfaces moving simultaneously in the preferred and antipreferred directions, compared with neuronal responses to single surfaces moving in the preferred direction. Neurons in MT are also selective for binocular disparity. Natural transparent stimuli usually would contain other surface cues, such as depth from binocular disparity. When we presented the two surfaces at different depths using disparity cues, we found that the inhibition occurs mainly for transparent stimuli at the same depth, which is consistent with the idea that this inhibition exists to reduce spurious motion signals (noise). Thus, transparent surfaces are represented independently in MT when binocular disparity cues are present to disambiguate the depth ordering of surfaces (Bradley, Qian, and Andersen 1995).

In a follow-up study, we used the illusion of a rotating transparent cylinder composed of random dots. This display is perceptually bistable, much like a Necker cube and will appear to flip front and back surfaces, and thus, the direction of the perceived rotation of the cylinder. We trained NHPs to indicate the surface order they perceived, even though the visual stimulus was always the same. Many neurons in MT exhibited changes in activity that coincided with the reversals in perception, indicating that MT has a basic role in SFM perception (Bradley, Chang, and Andersen 1998). In a subsequent study, it was found that MT neurons were much more affected by the surface order of these perceived 3D stimuli than were V1 neurons. These results suggest that V1 is not directly involved in SFM perception but MT is.

Optic Flow and the Direction of Heading

We had studied the global motion fields in spiral space at MIT. The expanding and contracting stimuli have an important role in the perception of direction of locomotion. For instance, while walking or driving, the motion that is generated on the retina, for moving forward, is an expansion (contraction if you are moving backward), and the focus of expansion indicates the direction of heading when the eyes are still. However, when the eyes also

move, as happens when moving in one direction but fixating on an object to the side using smooth pursuit eye movements, this introduces an additional translation motion on the retinas, shifting the focus of expansion. Marty Banks and colleagues at UC Berkeley showed that humans are still able to perceive the correct direction of heading and this correct percept requires eye movement information. Marty and I were at a meeting together and discussed trying his paradigm while recording from area MSTd, which has expansion selective neurons.

We simulated locomotion with random dot fields. We measured the tuning of MSTd neurons for the location of the focus of expansion and found that this tuning shifted toward correctly registering the heading direction during pursuit eye movements. This tuning shift is possible only if MSTd neurons receive either visual or motor signals about pursuit. We eliminated the visual signal hypothesis by showing that the same neurons did not exhibit a compensatory tuning shift when the display simulated the retinal stimulus that occurs during pursuit, but with the eyes still (Bradley et al. 1996). Thus, MSTd neurons compensate for the motor signals that drive eye movements and register the true heading direction. Further experiments led by Krishna Shenoy, a postdoctoral fellow in the lab, found that this was the case when gaze was shifted by vestibular signals during whole body rotations. These results indicate that the compensation is a result of gaze rotation, whether it is the eyes moving in the head (pursuit) or the head rotating in the world (vestibular-ocular reflex). In another study he found that the amount of compensation was correlated with the speed of pursuit. Brian Lee, a graduate student, found there was also a retinal component of adjustment, which was apparent when the rate of visual expansion was changed, simulating different speeds of locomotion in the environment.

My studies on visual motion, spanning my faculty positions at the Salk, MIT, and Caltech, addressed fundamental questions about SFM perception and heading perception. Personally I felt I was ready to pivot to explore a new topic. A new direction was catching my interest: to explore brain-machine interfaces in NHPs and later in human clinical trials.

Brain-Machine Interfaces

As I was investigating the sensorimotor transformations in cortex, I began to think about how this knowledge might be used clinically. One straightforward application would be to design neural prosthetics that would help people suffering from paralysis. A colleague at Harvard, Marge Livingstone (see Volume 9), told me I was already talking about this when I was still at MIT, so at least 27 years ago.

The technology used for neural prosthetics is referred to as brain-machine interfaces (BMIs) or brain-computer interfaces (BCIs). A BMI is a device that interfaces a machine with the brain. It can write in or read out

information. Write-in BMIs generally use electrical stimulation to modulate brain activity. Examples are already available clinically, including cochlear implants for restoring hearing, deep brain stimulators for movement disorders, and cortical stimulators to disrupt epileptic seizures.

Read-out BMIs are still in development. The challenge of read-out BMIs is to record signals from the brain and decode them with mathematical algorithms to produce control signals for operating external devices, such as computers and robotics. One important use of such a read-out system is the development of neural prosthetics systems to assist patients with severe paralysis. Thus, a person who was paralyzed from neurological injury or disease affecting parts of the motor system could still generate signals from less affected parts of the brain to control external, assistive devices. Examples of patients that can benefit from a BMI are those suffering from spinal cord lesion, stroke, anterolateral sclerosis, multiple sclerosis, or Duchenne muscular dystrophy.

This field of read-out BMIs was established before neurophysiologists like myself, using single neuron recordings in NHPs, became involved. The signals for driving the BMIs were recorded primarily from electroencephalogram (EEG) recordings. Subjects could train their EEGs to control external devices like the movement of a cursor on a computer screen. An advantage of EEG is that it is noninvasive. However, a disadvantage is the spatial resolution is very low, averaging brain signals from centimeters of cortex. To overcome this spatial resolution limitation, ideally one would record the activity of single neurons. Importantly, not only would one record from single neurons, but also a number of single neurons simultaneously. Thus, populations of single neurons are recorded all at once. Arrays of microelectrodes have been developed over the past two decades that allow this simultaneous recording from populations of single neurons. The array that we use, the Utah array, was originally developed by Richard Normann at the University of Utah and refined by two companies, Cyberkinetics and later Blackrock Neurotech. It is a small four- by four-millimeter flat surface with 100 protruding microelectrodes that looks a bit like a bed of nails. Each of these arrays typically records from 100– to 200 neurons, generally more than one array is implanted in each subject, and the recordings are viable in humans for about five years.

Brain Control in Monkeys

We began using brain signals to control computers in earnest at Caltech in the late 1990s. This idea of single-neuron population BMIs had occurred independently at a few NHP neurophysiology labs, including Miguel Nicolelis at Duke, John Donoghue at Brown, and Andy Schwartz at Arizona State University and later the University of Pittsburgh. These labs concentrated on primary motor cortex with the exception of Nicolelis who studied

several cortical areas. I had been working on the high-level intent signal in PPC and so had concentrated on making a more cognitive neuroprosthetic—that is, one that specified the intended goal of a movement. These first experiments were in NHPs and were necessary to establish the success of brain control with single-neuron population recording in animal studies before human clinical studies. These animal studies were followed by studies in humans with severe paralysis by the Donoghue lab, the Schwartz lab, and our lab.

Our first paper on cognitive neural prosthetics with NHPs was published in *Science* in 2004 (Musallam et al. 2004). In this study, implants were made in PPC and the dorsal premotor cortex, another high-level pre-motor area. The study focused on decoding the intended goals of a movement instead of the trajectory to a goal, which was the approach of other NHP BMI studies. Decoding is a necessary step between the raw neural data and the control of assistive devices. Training data are collected at the beginning of an experiment that determines the correlation between patterns of neural activity and motor outcomes. The training data teach the decoder to subsequently predict the intended motor outcome from patterns of activity generated by single trials to drive the prosthetic device.

We found that determining goals could be accomplished very quickly, within 200 ms, and much quicker than the time that would be required for guiding the trajectories. Performance improved over a period of months suggesting that learning was at play. The desirability of the goals also could be decoded by indicating the value of each trial, which was cued at the target location. The value manipulations were amount, probability, and type (juice or water) of the reward. These cognitive variables reflecting the desirability of goals are factored into decision making. This finding of desirability indicated that very high-level signals related to the state of the NHP could be decoded. Conceivably these high-level signals can be used in clinical studies for assessing the motivation and preferences of the patient. They also suggest that implants, particularly those introduced in centers more involved with emotion and mood, could be used in clinical studies of neuropsychiatric disorders, such as depression, and could provide feedback to guide neuromodulation with, for instance, electrical stimulation.

Next, graduate students Markus Hauschild and Grant Mulliken showed that a trajectory signal also could be decoded from PPC. The trajectories could be decoded for targets in 2D and even in 3D, where NHPs moved a cursor to locations in different depths combined with horizontal and vertical movements. Graduate students Bijan Pesaran, John Pezaris, and Maneesh Sahani; postdoctoral fellows Hans Scherberger and Murray Jarvis; and colleague Partha Mitra showed that local field potentials (LFPs) are prominent and predictive of behavioral state in PPC (Pesaran et al. 2002; Scherberger, Jarvis, and Andersen 2005). Postdoctoral fellow EunJung Hwang showed that LFPs recorded from the microelectrodes also could be

used for brain control. Adding LFP information to the spiking information improved performance. LFPs average neural activity over larger volumes of cortex, resulting in more redundant information being recorded across electrodes. Spike recordings reflect activity in much more localized volumes of cortex, resulting in much more independent information recorded across electrodes and thus leading to better decoding performance. However, an important feature of the LFPs is that they can be recorded from most microelectrodes in an array, whereas spikes are more limited in the number of electrodes that are close enough to single neurons to record spikes. Also, over time, the spiking signals fade. LFP signals still remain and thus can be used to extend the time implants are viable.

While we were doing these NHP studies to confirm the feasibility of implanting PPC in humans, we were also in the process of preparing for the human studies. We built a multidisciplinary team of scientists and physicians from Caltech, the Keck Medical School of the University of Southern California, the Geffen Medical School at UCLA, the Rancho Los Amigos National Rehabilitation Center, and the Casa Colina Hospital and Centers for Health Care. We worked with regulatory bodies, including the U.S. Food and Drug Administration, and the institutional review boards at Caltech and at our collaborators' institutions to receive the go-ahead. It took about 15 years of NHP studies, building the infrastructure, and obtaining approvals before we began our first human clinical trials.

Caltech: Human Days

Neuroprosthetics Studies in Humans

Participants in neuroprosthetic studies are true heroes. At the time of my writing this autobiography, there have been only about 20 participants who have entered this type of clinical trial among the groups doing this research worldwide. The participants are akin to astronauts who have landed on the moon; an elite group with tremendous bravery and resolve to enter this new world of brain-controlled neuroprosthetics. They volunteer selflessly, as there is no direct medical benefit for them personally in enrolling in the study; only the knowledge that their groundbreaking work is contributing to medical advances that someday will benefit a future generation of people suffering from severe paralysis.

Our first participant was Erik Sorto. He was the victim of a gunshot that severed his spinal cord when he was 21 years old, rendering him tetraplegic. When Sorto entered the study he had already been paralyzed for 10 years. The implant surgery took place at the Keck Medical School at the University of Southern California in April 2013. Two microelectrode arrays were placed in the PPC. The surgeons were Charles Liu and Brian Lee, and the surgery went flawlessly. Going back to the planetary space analogy, my colleagues at

the Jet Propulsion Laboratory (JPL) in the Pasadena area refer to the seven minutes of terror when a Mars rover enters the atmosphere of Mars before it lands. For us, it was two weeks of terror after the implantation, during the healing period, before Sorto had his first session. We had every reason to believe that the implant would work based on our NHP studies. However, the monkeys were healthy and Sorto had been paralyzed for 10 years. Would the motor-related signals in PPC still be there after 10 years of disuse? No one had ever implanted human PPC with microelectrode arrays. We used functional magnetic resonance imaging preoperatively to help guide us to locations in PPC that were active when Sorto imagined a reach or a grasp. But this technique is indirect and monitors increases in blood flow when a brain area is active. Would we record single neurons in PPC of Sorto, and would he be able to modulate them with his thoughts when controlling a robotic limb or a computer cursor?

Our team was led by Spencer Kellis, a senior scientist at Caltech. During the first session, we isolated only a few neurons. This is not uncommon as yields of neurons become much larger over the first month. Also we had to adjust the unit isolation on the channels with the data collection software. In the next session, something amazing happened. As more neurons came into view, Sorto was asked to try to control their firing rates. Some of the neurons were activated when he imagined rotating his paralyzed wrist. We used the firing rates of these neurons to control the wrist of a robotic arm. He was able to use his thought to rotate the robot hand to gesture shaking the hand of a graduate student. This result was amazing. First, it showed that, even after 10 years of disuse, the intrinsic circuits for motor control were still intact. Second, it required no practice on Sorto's part; the brain control worked right out of the box. Third, Erik was ecstatic. This was the first time in 10 years that he had been able to move a limb and with a social gesture.

We ask participants at the beginning of the study what they hoped to achieve. In Erik's case, he wanted to be able to drink a beer on his own. A year into the study, he achieved this goal by controlling a robotic limb. The first time he accomplished this feat, there were shouts of joy from Erik, staff, and lab members present. Our second subject, Nancy Smith, was a high school teacher who had played the piano before a car accident six years earlier that left her tetraplegic. The implant surgery was performed by Nader Pouratian at UCLA, and the research team was led by Tyson Aflalo, a senior scientist at Caltech. Surprisingly, we found that she had a motor representation of all 10 individual fingers in PPC and was able to play a virtual keyboard piano with imagined finger strokes of one of her hands.

We have now had five participants over the past eight years, Erik Sorto, Nancy Smith, DL, FG, and JJ. In every case, it leaves me breathless to see them control their physical environment with just their thoughts. For me, it has been a high point in my career, being able to do both the basic

research of studying neuroscience, and then being able to apply what we learn directly to the clinic. These studies with humans reveal new scientific findings that are then used to further advance the clinical aspects of the study. The team is truly interdisciplinary. No single person can do this kind of project, and it requires physicians, scientists, engineers, and computational neuroscientists. The patients are truly members of the team. Unlike conventional clinical studies that involve large numbers of people for very short periods of time, these studies involve single individuals over extensive time periods. Each person does two to three sessions per week lasting about four hours per session. The total length of time a person is in a clinical trial is about five years.

As mentioned earlier, when we obtain consent from our participants, we indicate that they will not directly benefit from the study. However, I have found that this type of study is tremendously beneficial to them, and to us, in ways that became apparent only over the course of the studies. People with paralysis often do not leave their home environment and are depressed. The clinical studies enable them to have a job, get out of the house, meet interesting people, and have a purpose. Some come to engage more in outside interests, including giving lectures and being interviewed by media. From our team's side, we have direct contact with members of the population we want to help, being able to better understand their needs. We get to know their families and, before the Covid pandemic, we would have holiday and birthday parties with the participants, their family members and caregivers, and our team of scientists and physicians. It is a truly unique experience for all of us.

In 2015, we published our first human study with Sorto in *Science* (Aflalo et al. 2015). This first-ever use of PPC single-neuron population recordings for prosthetic control showed that indeed PPC provides high-level movement-related activity through motor imagery. Intended goals of movement could be decoded very rapidly, within 200 ms, similar to our findings in NHPs. Also the trajectories could be decoded. Moreover, goal signals were effector specific, and separate neurons encoded the contralateral limb or ipsilateral limb, referenced to the implanted hemisphere. Thus, this was the first demonstration of a cognitive neural prosthetic in humans.

Our *Science* paper received considerable press. A couple, Tianqiao Chen and Chrissy Luo, saw a news story on our findings on television in Singapore. They had founded a company, Shanda in 1999, which became the biggest online entertainment developer in China. They subsequently founded the Tianqiao and Chrissy Chen Institute for philanthropy targeting brain science. Tianqiao visited me at Caltech, and a month later, both he and Chrissy visited again. Their interests in our research and in Caltech led to their making a gift of \$115 million in 2016. This was the largest gift given to Caltech by a donor that was not an alumnus or previously connected or known to the Caltech administration. The gift provided \$50 million

toward the construction of the Chen Neuroscience Research Building and the remainder to support the Tianqiao and Chrissy Chen Institute for Neuroscience at Caltech. This new neuroscience institute included several endowed centers. One of these is the T&C Chen Brain-Machine Interface Center, for which I am currently the director.

“It Was the Best of Times, It Was the Worst of Times.”

This Charles Dickens quote (*A Tale of Two Cities*) captures the next few years of my life. Research was going fabulously, and building up the BMI Center was crucial in sustaining our BMI research with humans and to also building interest in BMI research by fellow Caltech faculty through seed grants from the BMI Center (Andersen 2019). But in early 2017, my wife Carol was diagnosed with an incurable disease. That year, I became her caregiver until her passing at home with our family in January of 2018. I lost my lifelong love and emotional foundation. Not long after her passing came the pandemic, still with us two years later as I write these words. My fantastic children, lab, and colleagues have helped me to weather these tragedies and also to keep my research moving forward.

Partially Mixed Selectivity

We were pleased that even after years of spinal cord injury the intent of the subjects could be recorded from PPC. As we explored the activity of the populations of recorded neurons, we found that a large number of other cognitive variables were present and decodable. These included mental strategy (imagined versus attempted movement), finger movements, decisions involving memory recall, hand shaping for grasp, observed actions of others, and action verbs, such as “grasp” or “drop.” From a few hundred neurons, we could decode much of what a person intends to do. How is this possible?

Recall that the gain fields we found in NHPs mixed gaze position-tuned activity with retinotopically tuned visual stimulus positions. In the Zipser-Andersen model, we showed that as few as nine neurons could represent all of visual space with respect to the head. Recent animal studies from other labs have shown a mixing of variables. For example, recordings from prefrontal cortical neurons in NHPs showed a mixing of two types of memory task and different visual objects. The key is that one neuron can respond to more than one variable and can be involved in more than one circuit. This type of coding was coined randomly mixed selectivity. However, in our studies, we found that, although there is mixing, there is also structure in the statistics of the representation. We referred to this as partially mixed selectivity.

We studied these statistical structures formally in two experiments. In the first, we examined the representation of task strategy (imagine/attempt), side of the body (left/right), and effector (hand/shoulder). Graduate student

Carey Zhang, along with Tyson Aflalo, found that whereas the effectors were randomly mixed on PPC neurons, the strategy and body side were correlated (Zhang et al. 2017). Thus, if a neuron responded to an imagined right hand movement, it was more likely to respond also to an imagined left hand movement. One could not predict, however, if that same neuron would respond to an imagined shoulder movement. This structure makes functional sense since, for instance, something learned with one hand is transferable to the other hand. However, the mechanics of the shoulder and hand are very different, and thus it is useful to have their activities more separated in the neural population.

A second experiment led by Tyson Aflalo and colleague Guy Orban examined selectivity for observed actions and action verbs. Participants viewed videos of five actions of another person manipulating objects. These included drag, drop, grasp, push, and rotate. Each action was shown from different views of the person performing the action. Besides the action videos, some trials showed only the text of the action verbs and the participants were asked to read the text silently. It was again found that all of these variables were mixed in the population, but there were correlations of the videos showing the same action from different points of view and even reading the verb for the same action (Aflalo et al. 2020).

Partially mixed selectivity may explain why so many roles have been attributed to NHP LIP. As mentioned, we had found the area and ascribed it to have a role in saccades. Subsequent studies emphasized its role in decision making, attention, and categorization. In NHP experiments, one selects an area of the cortex to study, usually based on preliminary data or prior literature suggesting its function. To train the NHPs to perform a behavioral task takes months, and collecting data is also very time consuming with the NHPs performing many thousands of trials. Not surprisingly, researchers find what they are looking for and then try to reconcile their results to other functions ascribed to the same brain area by different labs using different tasks. With humans, you can ask them to perform a task, so there is almost no training involved. Data collection is from array recordings so each daily yield is a large data set. With language, you can instruct many different types of tasks. With humans, you can also obtain their verbal reports of sensations for stimulation studies and mental strategies for behavioral tasks.

Although areas of the PPC show a great diversity of action variables recorded from single small patches, I still believe that there are core local functional differences within PPC. We have seen differences between human area 5, which appears to be reach selective, and AIP, which appears to be more involved in shaping the hand during grasp. The inactivation studies in NHPs produced a specific reach deficit for PRR, and a saccade deficit for LIP. I believe that what we record from an implant in PPC are the core local computation at the implant site as well as the information coming into the

area from the many other brain areas that project into the local region. This idea emphasizes that each node is a part of a network.

Stimulation

The people with tetraplegia from spinal cord injury not only are unable to move their limbs, but also are unable to feel them since the main touch and body position signals travel through the spinal cord to the brain. This deficit is problematic for two reasons. One is a sense of embodiment. When Nancy Smith imagined moving the avatar's fingers to play a virtual piano, she did not receive touch feedback when she struck the piano keys. Second, somatosensation is necessary for manual dexterity. Under visual feedback our participants are able to move a robot limb and grasp an object using visual feedback for guidance. However, once the robot hand has grasped an object, somatosensation is necessary for manually manipulating the object.

We and others have used one method to evoke somatosensation by delivering current through electrodes imbedded in a grid (referred to as electrocorticography or ECoG) that sits on the surface of the primary somatosensory cortex. The arrays are placed under the dura but do not penetrate the cortex, and the electrodes themselves can be small with diameters around two millimeters. Even so, currents to evoke sensation are high, in the order of milliamps, and the sensations are unnatural and have been reported as tingling or electric. On the positive side, the grids can be large and cover extensive areas of the brain, and the somatotopic map in S1 can be determined by stimulating different electrodes in the grid. Thus, this method can provide touch sensation at localized regions of a robotic hand and may improve dexterity.

Intracortical microstimulation can be used to stimulate sensation through the microelectrodes in the microelectrode arrays. Because the tips are in the cortex rather than on the surface, much smaller currents can be delivered to evoke a sensation. Rather than stimulating large numbers of neurons as in the case of ECoG, only a small number of neurons are activated around the tips of the electrodes. Using this technique, the laboratory of Rob Gaunt at the University of Pittsburgh implanted the hand region of S1, and the participant reported natural-like feelings of pressure. Our team, led by Michelle Armenta Salas, Luke Bashford, and Spencer Kellis, implanted the arm region of one of our participants and for the first time, he reported natural touch sensations, such as tap, squeeze, and vibration (Armenta Salas et al. 2018). Also for the first time, we also found proprioception, in which stimulation of some sites produced the sense of the arm moving. Interestingly, as we increased the current delivered, we could change the sensation elicited from a touch sensation to a proprioceptive sensation for many of the electrodes, providing preliminary evidence that the sensations can be manipulated by changing the properties of the electrical stimulus.

Spatial Representation

Using the relative position task and analysis we previously used in NHPs, we have been exploring the representation of space in humans. In our electrode placements, we see that the coordinate transformations are often largely complete, for instance, an imagined reach is coded in hand coordinates. One reason for this difference, I suspect, is that our human participant implants are recording from locations more anterior in PPC than in the NHPs. We began seeing hand coordinate frame neurons in area 5 of NHPs. Another interesting feature we saw in area 5 of NHPs, and now in PPC of humans, is a switching of neural representations. In area 5, when the initial hand and eye position were obtained in the task, the neurons coded the relative position of the eye and hand. When a reach cue was provided, the representation shifted to coding the location of the target relative to the hand. In humans, we have found a similar contextual shifting of coordinate frames. When the target is for a saccade, the coding is the gaze relative to the target (eye coordinates). When the target is for a reach, the coding in the population shifts coordinates to coding the target with respect to the hand (hand coordinates). We also are finding not only egocentric spatial representations but also allocentric representations. An egocentric representation is where an object is with respect to a part of the body. An allocentric representation is where objects in the environment are placed with respect to each other. Interestingly, both egocentric and allocentric deficits in spatial perception are seen with PPC injury in humans.

Although many of our spatial representation experiments are currently ongoing, graduate student Matiar Jafari and Tyson Aflalo led a study, published recently, in which we looked at reach activation and coordinate frames in the population of S1 neurons (Jafari et al. 2020). There are two remarkable findings from this study. The first is that even though the patient is paralyzed, we designed a task in which he imagined different initial hand positions as well as imagined reaching to the target. In this task, there is a visual cue phase, a delay phase when motor planning occurs, and then an action phase in which a reach is imagined. The first finding was that during the visual cue and motor planning phases, which produce robust activation in PPC neurons, there was no activity in S1. Thus, vision and planning activity are not variables that are represented in S1 but are present in a comparable task in PPC. Even if there is mixed selectivity in S1, it does not share the same variables as PPC. Second, during the imagined reach the activity was in hand coordinates, the appropriate coordinates one would expect for an actual reach.

Internal Models and Imagination

A third interesting feature of the S1 recording results was the activation of S1 for imagined reach. Reach activity has been recorded in S1 of NHPs

before an actual reach, although the coordinates were not determined. The fact that the reach activation in human S1 is in hand coordinates, and there was no activity in the cue or planning phase, suggests that the source of the reach activation is computed in another brain area and then sent to S1. This finding suggests that the activation is an efference copy and a component of an internal model.

As mentioned, the timing of reach responses in NHP PPC suggested that they were efference copies for updating an internal model in the sensorimotor system. With our human participants, we have been able to examine task strategies in which the patient imagines moving, attempts to move (for paralyzed parts of the body), or actually moves (for nonparalyzed body parts). It is rather amazing how much of the brain is activated by imagination. For example, S1 and PPC neurons can be activated for touch or imagined touch. This extends not only to the part of the body of the imagined touch, but also to the quality of the touch, such as a squeeze or a tap. The results from the coordinate frame studies suggest that PPC keeps a representation of the world and the relation of the body in the world. Injury to PPC further supports its role in this internal model framework, since these patients have great difficulty adjusting and correcting their movements to changes in the environment.

Awareness of Intent

Harkening back to the Libet experiment, we designed a study in which participants were free to choose when to move, or both what to move and when to move. They were asked to retrospectively report the time they felt the urge to move. Similar to Libet's findings, and subsequent findings by others, PPC activity rises several hundred milliseconds before subjects report the initial urge to move. Moreover, our BMIs can decode the subject's intended movement before they are aware of it, and they report that a brain-controlled action often occurred before they intended it. To provide better embodiment, we found that we could design decoding algorithms that predict when the subject would be aware of the intent in order to synchronize intention awareness with the movement output of the neuroprosthetic. In the bigger picture, these results indicate that a lot of movement computation is occurring implicitly or preconsciously in PPC.

Stability and Learning

There is an extensive literature documenting changes in cortical representations of the body in somatosensory and motor cortex; in fact, many of the seminal studies were made by my mentor Mike Merzenich. Perhaps less studied is the functional organization that remains intact after injury. The study of BMIs in tetraplegics has provided a unique opportunity to

determine remaining function. BMI studies from our group and others have found considerable preserved structure even years after the initial injury. The motor cortex is active and this activity can be immediately harnessed in brain-control tasks. Microstimulation of S1 produces the same somatotopy of sensation one would expect in a healthy individual, and the machinery for natural perception is still intact. Mixed selectivity, which is found in association cortex of laboratory animals, is also present in PPC of the human participants.

Learning, in which the control of the prosthetic improves over months, has been observed in most studies. Very fast learning, on the scale of one day, often involves cognitive strategies using existing intrinsic patterns of activity within an area. For instance, postdoctoral fellows EunJung Hwang and Sofia Sakellaridi have led studies that show a rotation of targets will lead to re-aiming, documented in NHPs and humans (Hwang et al. 2013; Sakellaridi et al. 2019). In the case of humans, they can tell you what cognitive strategy they are using. Studying long-term learning is currently more problematic because of the difficulty in recording from the same neurons over many days with the current recording technologies. The long-term learning that has been observed is in the domain of the cortical areas implanted—in other words, a sharpening of the skill of the area. Further work on plasticity with BMIs is needed. But at least in viewing what we know today, it is important to select areas of the brain that naturally function in the domain of the desired BMI output. For instance, one would likely be more successful to implant language cortex for decoding language than to implant visual cortex for language decoding.

Functional Ultrasound

No doubt there will be improvement in materials and electronics as BMIs move to the clinic. The gold standard for BMIs is to record from single-neuron populations. But to record the extracellular electric field of a neuron requires the microelectrode to be within a few tens of microns from the neuron. Thus, at least within the foreseeable future, this sort of spatial resolution will always require an invasive brain implant. As the technology becomes refined, this invasiveness may not be a major stumbling block as current neurotechnologies like cochlear implants and deep brain stimulators are invasive and yet readily available from the clinic.

Another avenue is to produce minimally invasive technologies. One such technology we have been working on is functional ultrasound. This technique detects changes in blood flow with the activation of brain areas. It has a very high spatial resolution, in the order of a 100 microns, but it is slow as it is limited temporally by the hemodynamic response. In a collaboration with Mikhail Shapiro's lab at Caltech and Mickael Tanter's lab in Paris,

postdoctoral fellow Sumner Norman and graduate student Whitney Griggs have shown that functional ultrasound can be used in NHPs for single-trial decoding of the direction of movement and by the effector used (Norman et al. 2021). This technique could be used as a minimally invasive BMI in humans by replacing a small bit of skull with an acoustically transparent window. It would not require surgery to the dura or penetration into brain tissue. In addition to BMI applications, this technology may be useful for decoding the state of the brain. It has a very wide field of view and could detect epileptic seizure locations as well as changes in more global brain activity that accompany neuropsychiatric diseases.

Final Thoughts

I have had extremely good fortune in my life and career. I have had the privilege to work with the brightest and best students and colleagues from around the world. I was fortunate to be at the beginning of what is now the huge and evolving field of neuroscience and to be trained by some of the giants in the field. I again find myself at the beginning stages of a new and expanding field: BMIs. I thank my family who have supported me all these years and for the great joy they have brought me. I thank the many brilliant students, postdoctoral fellows, and staff who studied with me. Many I have mentioned in the text and all have contributed tremendously to the research coming out of the lab. The following table lists our lab members and the decade in which they graduated or are currently in the lab for the decade beginning 2020. Finally, I am still doing research, which seems to me to be at a level comparable to early in my career, and it certainly bringing me the same exhilaration of adventure.

Lab Members

Postdoctoral Fellows	Graduate Students	Research Support
1980s	1980s	
Chiko Asanuma	Martyn Bracewell	
Greg Essick	Michael Craziano	
Ralph Siegel		
Peter Brotchie		
Gene Blatt		
Jim Gnadt		
Masud Husain		
Roger Erickson		
Leonardo Fogassi		
Peter Thier (faculty on sabbatical)		

Postdoctoral Fellows	Graduate Students	Research Support
1990s Robert Snowden Shabtai Barash Ning Qian Larry Snyder Sohaib Kureshi David Bradley Jim Crowell Jing Xing Philip Sabes	1990s Stefan Treue Bard Geesaman Pietro Mazzoni Brigette Stricanne Ray Li Maneesh Sahani Jennifer Linden	1990s Gail Robertson Catherine Cooper Betty Gillikin Grieve Viktor Shcherbatyuk (to present) Cierina Reyes Marks
2000s Jyl Boline Yale Cohen Kenneth Grieve Alex Gruenwald Krishna Shenoy Murray Jarvis David Dubowitz Brian Corneil Hans Scherberger Shau-Ming Wu Rodrigo Quian-Quiroga Jason Connolly Chris Buneo Bijan Pesaran Bradley Greger Alexander Gail Sam Musallam Boris Breznen Marina Brozovich Daniel Rizzuto Axel Lindner Elizabeth Torres He Cui	2000s John Pezaris Aaron Batista Kyle Bernheim Daniella Meeker Matthew Nelson Brian Lee Michael Campos Daniel Baldauf Asha Iyer Grant Mulliken Hillary Glidden Rajan Bhattacharyya	2000s Kelsie Pejsa (to present) Tessa Yao
2010s Melanie Wilke Claudia Wilimzig Igor Kagan Carlos Pedreira Xoana Troncoso EunJung Hwang James Bonaiuto Lindsay Bremner Markus Hauschild Arnulf Graf Christian Klaes	2010s Boris Revechkis Carey Zhang Matiar Jafari HyeongChan Jo	

Postdoctoral Fellows	Graduate Students	Research Support
Michelle Armenta Salas Vasileios Christopoulos Dan Kramer Spencer Kellis		
2020s (current lab members) Tyson Aflalo (member of professional staff) Luke Bashford David Bjanec Jorge Gamez Sumner Norman Sofia Sakellaridi Xinyun Zou Liang She	2020s (current lab members) Srinivas Chivukula Whitney Griggs Charles Guan Kelly Kadlec Isabelle Rosenthal Sarah Wandelt	

Selected Bibliography

- Aflalo, T., S. Kellis, C. Klaes, B. Lee, Y. Shi, K. Pejsa, K. Shanfield, S. Hayes-Jackson, M. Aisen, C. Heck, C. Liu, and R. A. Andersen. 2015. "Decoding motor imagery from the posterior parietal cortex of a tetraplegic human," *Science*, 348: 906–10.
- Aflalo, T., C. Y. Zhang, E.R. Rosario, N. Pouratian, G. A. Orban, and R. A. Andersen. 2020. "A shared neural substrate for action verbs and observed actions in human posterior parietal cortex," *Science Advances*, 6.
- Andersen, R. 2019. "The intention machine," *Scientific American*, 320: 25–29.
- Andersen, R. A., C. Asanuma, and W. M. Cowan. 1985. "Callosal and prefrontal associational projecting cell populations in area 7A of the macaque monkey: a study using retrogradely transported fluorescent dyes," *J Comp Neurol*, 232: 443–55.
- Andersen, R. A., and C. A. Buneo. 2002. "Intentional maps in posterior parietal cortex," *Annual Review of Neuroscience* 25: 189–220.
- Andersen, R. A., P. L. Knight, and M. M. Merzenich. 1980. "The thalamocortical and corticothalamic connections of AI, AII, and the anterior auditory field (AAF) in the cat: evidence for two largely segregated systems of connections," *J Comp Neurol*, 194: 663–701.
- Andersen, R. A., and V. B. Mountcastle. 1983. "The influence of the angle of gaze upon the excitability of the light-sensitive neurons of the posterior parietal cortex," *J Neurosci*, 3: 532–48.
- Andersen, R. A., L. H. Snyder, D. C. Bradley, and J. Xing. 1997. "Multimodal representation of space in the posterior parietal cortex and its use in planning movements," *Annu Rev Neurosci*, 20: 303–30.
- Andersen, R. A., G. K. Essick, and R. M. Siegel. 1985. "The encoding of spatial location by posterior parietal neurons," *Science*, 230: 456–58.

- Armenta Salas, M., L. Bashford, S. Kellis, M. Jafari, H. Jo, D. Kramer, K. Shanfield, K. Pejsa, B. Lee, C. Y. Liu, and R. A. Andersen. 2018. "Proprioceptive and cutaneous sensations in humans elicited by intracortical microstimulation," *eLife*.
- Batista, A. P., C. A. Buneo, L. H. Snyder, and R. A. Andersen. 1999. "Reach plans in eye-centered coordinates," *Science*, 285: 257–60.
- Bradley, D. C., G. C. Chang, and R. A. Andersen. 1998. "Encoding of three-dimensional structure-from-motion by primate area MT neurons," *Nature*, 392: 714–7.
- Bradley, D. C., M. Maxwell, R. A. Andersen, M. S. Banks, and K. V. Shenoy. 1996. "Mechanisms of heading perception in primate visual cortex," *Science*, 273: 1544–7.
- Bradley, D. C., N. Qian, and R. A. Andersen. 1995. "Integration of motion and stereopsis in middle temporal cortical area of macaques," *Nature*, 373: 609–11.
- Bremner, L. R., and R. A. Andersen. 2014. "Temporal analysis of reference frames in parietal cortex area 5d during reach planning," *Journal of Neuroscience*, 34: 5273–84.
- Brotchie, P. R., R. A. Andersen, L. H. Snyder, and S. J. Goodman. 1995. "Head position signals used by parietal neurons to encode locations of visual stimuli," *Nature*, 375: 232–5.
- Buneo, C. A., M. R. Jarvis, A. P. Batista, and R. A. Andersen. 2002. "Direct visuomotor transformations for reaching," *Nature*, 416: 632–636.
- Cohen, Y. E., and R. A. Andersen. 2002. "A common reference frame for movement plans in the posterior parietal cortex," *Nature Review Neuroscience*, 3: 553–62.
- Cui, H., and R. A. Andersen. 2007. "Posterior parietal cortex encodes autonomously selected motor plans," *Neuron*, 56: 552–9.
- Geesaman, B. J., and R. A. Andersen. 1996. "The analysis of complex motion patterns by form/cue invariant MSTd neurons," *J Neurosci*, 16: 4716–32.
- Gnadt, J. W., and R. A. Andersen. 1988. "Memory related motor planning activity in posterior parietal cortex of macaque," *Experimental Brain Research*, 70: 216–20.
- Graf, A. B. A., and R. A. Andersen. 2014. "Inferring eye position from populations of lateral intraparietal neurons," *eLife*, 3: e02813.
- Graziano, M. S., R. A. Andersen, and R. J. Snowden. 1994. "Tuning of MST neurons to spiral motions," *J Neurosci*, 14: 54–67.
- Hwang, E. J., P. M. Bailey, and R. A. Andersen. 2013. "Volitional control of neural activity relies on the natural motor repertoire," *Curr Biol*, 23: 353–61.
- Hwang, E. J., M. Hauschild, M. Wilke, and R. A. Andersen. 2012. "Inactivation of the parietal reach region causes optic ataxia, impairing reaches but not saccades," *Neuron*, 76: 1021–29.
- Jafari, M., T. N. Aflalo, S. Chivukula, S. S. Kellis, M. A. Salas, S. L. Norman, K. Pejsa, C. Y. Liu, and R. A. Andersen. 2020. "The human primary somatosensory cortex encodes imagined movement in the absence of sensory information," *Communications Biology*, 11.
- Mazzoni, P., R. A. Andersen, and M. I. Jordan. 1991. "A more biologically plausible learning rule for neural networks," *Proc Natl Acad Sci U S A*, 88: 4433–37.

- Mulliken, G. H., S. Musallam, and R. A. Andersen. 2008. "Forward estimation of movement state in posterior parietal cortex," *Proc Natl Acad Sci U S A*, 105: 8170–7.
- Musallam, S., B. D. Corneil, B. Greger, H. Scherberger, and R. A. Andersen. 2004. "Cognitive control signals for neural prosthetics," *Science*, 305: 258–62.
- Norman, S. L., D. Maresca, V. N. Christopoulos, W. S. Griggs, M. Demene, M. Tanter, M. G. Shapiro, and R. A. Andersen. 2021. "Single trial decoding of movement intentions using functional ultrasound neuroimaging," *Neuron*, 109: 1–13.
- Pesaran, B., M. J. Nelson, and R. A. Andersen. 2008. "Free choice activates a decision circuit between frontal and parietal cortex," *Nature*, 453: 406–09.
- Pesaran, B., J. Pezaris, M. Sahani, P. M. Mitra, and R. A. Andersen. 2002. "Temporal structure in neuronal activity during working memory in Macaque parietal cortex," *Nature Neuroscience*, 5: 805–11.
- Qian, N., R. A. Andersen, and E. H. Adelson. 1994a. "Transparent motion perception as detection of unbalanced motion signals. I. Psychophysics," *J Neurosci*, 14: 7357–66.
- Qian, N., and R. A. Andersen. 1994. "Transparent motion perception as detection of unbalanced motion signals. II. Physiology," *J Neurosci*, 14: 7367–80.
- Qian, N., R. A. Andersen, and E. H. Adelson. 1994b. "Transparent motion perception as detection of unbalanced motion signals. III. Modeling," *Journal of Neuroscience*, 14: 7381–92.
- Sakellaridi, S., V. N. Christopoulos, T. Aflalo, K. W. Pejsa, E. R. Rosario, D. Ouellette, N. Pouratian, and R. A. Andersen. 2019. "Intrinsic variable learning for brain-machine interface control by human anterior intraparietal cortex," *Neuron*, 102: 694–705.
- Scherberger, H., M. R. Jarvis, and R. A. Andersen. 2005. "Cortical local field potential encodes movement intentions in the posterior parietal cortex," *Neuron*, 46: 347–54.
- Siegel, R. M., and R. A. Andersen. 1988. "Perception of three-dimensional structure from motion in monkey and man," *Nature*, 331: 259–61.
- Snyder, L. H., A. P. Batista, and R. A. Andersen. 1997. "Coding of intention in the posterior parietal cortex," *Nature*, 386: 167–70.
- Snyder, L. H., K. L. Grieve, P. Brotchie, and R. A. Andersen. 1998. "Separate body- and world-referenced representations of visual space in parietal cortex," *Nature*, 394: 887–91.
- Zhang, C. Y., T. Aflalo, B. Reveckis, E. R. Rosario, D. Ouellette, N. Pouratian, and R. A. Andersen. 2017. "Partially mixed selectivity in human posterior parietal association cortex," *Neuron*, 95: 697–708.
- Zipser, D., and R. A. Andersen. 1988. "A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons," *Nature*, 331: 679–84.