


Glucose-Powered Neuroelectronics

by

Benjamin Isaac Rapoport



ARCHIVES

Bachelor of Arts in Physics and Mathematics, Harvard University, 2003
Master of Arts in Physics, Harvard University, 2003
Master of Science in Mathematics, Oxford University, 2004
Master of Science in Physics, Massachusetts Institute of Technology, 2007

Submitted to the
Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2011

© Massachusetts Institute of Technology 2011. All rights reserved.

Author
Department of Electrical Engineering and Computer Science
May 9, 2011

Certified by ..
Rahul Sarpeshkar
Associate Professor
Thesis Supervisor

Accepted by
Professor Leslie A. Kolodziejski
Chairman, Department Committee on Graduate Students

Glucose-Powered Neuroelectronics

by

Benjamin Isaac Rapoport

Submitted to the
Department of Electrical Engineering and Computer Science
May 9, 2011
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy

Abstract

A holy grail of bioelectronics is to engineer biologically implantable systems that can be embedded without disturbing their local environments, while harvesting from their surroundings all of the power they require. As implantable electronic devices become increasingly prevalent in scientific research and in the diagnosis, management, and treatment of human disease, there is correspondingly increasing demand for devices with unlimited functional lifetimes that integrate seamlessly with their hosts in these two ways. This thesis presents significant progress toward establishing the feasibility of one such system: *A brain-machine interface powered by a bioimplantable fuel cell that harvests energy from extracellular glucose in the cerebrospinal fluid surrounding the brain.*

The first part of this thesis describes a set of biomimetic algorithms and low-power circuit architectures for decoding electrical signals from ensembles of neurons in the brain. The decoders are intended for use in the context of neural rehabilitation, to provide paralyzed or otherwise disabled patients with instantaneous, natural, thought-based control of robotic prosthetic limbs and other external devices. This thesis presents a detailed discussion of the decoding algorithms, descriptions of the low-power analog and digital circuit architectures used to implement the decoders, and results validating their performance when applied to decode real neural data. A major constraint on brain-implanted electronic devices is the requirement that they consume and dissipate very little power, so as not to damage surrounding brain tissue. The systems described here address that constraint, computing in the style of biological neural networks, and using arithmetic-free, purely logical primitives to establish universal computing architectures for neural decoding.

The second part of this thesis describes the development of an implantable fuel cell powered by extracellular glucose at concentrations such as those found in the cerebrospinal fluid surrounding the brain. The theoretical foundations, details of design and fabrication, mechanical and electrochemical characterization, as well as *in vitro* performance data for the fuel cell are presented.

Thesis Supervisor: Rahul Sarpeshkar
Title: Associate Professor

Acknowledgments

In writing these words of thanks at the end of my doctoral studies, I am conscious of what a great privilege it has been to spend nearly a quarter of my life in productive pursuit of a dream. That dream has lost none of its luster for me since I entered the M.D.–Ph.D. Program at Harvard and MIT in July 2004, to work on microelectronic interfaces with the human brain and nervous system. In the intervening years, many of the family, friends, colleagues, teachers, advisors, and others I wish to thank here have remarked that this field, my work within it, and the prospect of translating thought into action through brain-implanted electronics, all still seem to belong to the domain of science fiction. It is true. And so in completing this phase of my scientific work, I find myself thinking of the fantastic adventurers of science fiction I have grown up admiring who, realizing they have been dreaming, awaken to find themselves in yet another dream. The joy and adventure of scientific research and innovation continue, even as these words of appreciation mark a formal end. I wish a similar joy of pursuit to any present or future students who may happen to read the work I present here.

More than anyone else, my parents, Sandra and Samuel Rapoport, understand the dreams I have for myself and for the worlds of science and medicine. First and always my teachers, mentors, and trusted supporters, they have inspired and guided me, and made it possible for me to dream—and then to pursue and realize my dreams and some of theirs. With my deepest love, respect, admiration, and gratitude, the best of this work is dedicated to you.

I thank my thesis advisor, Rahul Sarpeshkar, for our exciting and productive collaborations, and for what he has taught me through example about rigorously pursuing and developing avant-garde science and technology. During those same years I have had the good fortune of working closely with especially talented members of our research group: Scott Arfin, Soumyajit Mandal, Lorenzo Turicchia, and Woradorn (Pok) Wattanapanitch. So often I have been humbled by my ignorance and saved by their generosity and wealth of knowledge. Our hours of debate and conversation,

days and nights of side-by-side collaboration, weeks of travel together, and months and years of jointly experienced frustrations and successes, have taught me and helped sustain me over the course my of graduate work. I thank you four, and hope that these past years have seeded the beginning of enduring friendships.

For what has been an enriching final year, my research has involved a collaboration with the Microfluidics Group in the Advanced Silicon Technology Group at Lincoln Laboratory, led by Jakub Kedzierski, who has been a knowledgeable and kind supervisor. Damiano (Desi) Biasella of the Microelectronics Laboratory at Lincoln Laboratory has been absolutely instrumental in the fabrication of the glucose fuel cell described in the second part of this thesis. I thank him for his consistent technical excellence, insight, astute observations, and good-natured forbearance as we modified and iterated some trial steps dozens of times before finalizing our fabrication process. Many others at Lincoln Laboratory treated me as a colleague and helped make my time there valuable, successful, and enjoyable: Shaun Berry, Vladimir Bolkhovsky, Sue Cann, Chenson Chen, Denise Holohan, Craig Keast, Jeff Knecht, Gayatri Perlin, Steven Rabe, Todd Thorsen, Thaddeus (T. J.) Wlodarczak, and Donna Yost.

A number of other collaborators have directly or indirectly made important contributions to my work over the past several years: Richard Andersen and Sam Musallam of the Division of Biology at the California Institute of Technology; Ed Boyden of the Media Laboratory and the Departments of Brain and Cognitive Sciences and Biological Engineering at MIT; Derek Lovley of the Department of Microbiology at the University of Massachusetts, Amherst; and Hector L. Penagos and Thomas J. Davidson, both formerly of the Wilson Laboratory in the Department of Brain and Cognitive Sciences at MIT, who provided electrophysiologic data I used in my work on neural decoding. With kindness and enthusiasm, Thomas Tachovsky of the Technology Licensing Office at MIT has guided me through the legal filings associated with intellectual property I and collaborators have developed during the past five years [75, 80, 79, 62].

For whatever inspired the truly wild gamble that led to my admission, I thank the faculties and administrations of the Harvard–MIT M.D.–Ph.D. Program and the

Harvard–MIT Division of Health Sciences and Technology (HST). Linda Burnley has been a source of support since the day I was admitted. She even conspired with me as early as April 2004, making it possible for me to run the Boston Marathon immediately before a set of interviews with prospective research supervisors, and to set a precedent that eventually led to my work on the physiology of endurance running. Rick Mitchell has been a gifted teacher and valuable advisor throughout my years with the HST Program. Martha Gray, during her tenure as Director of HST, offered me sound advice at a critical moment in my academic career. My thanks, also, to M.D.–Ph.D. Program Directors Chris Walsh and Steve Blacklow.

I doubt that my wonderful, first academic advisor at MIT, the late Professor Michael Feld, realized how important his guidance and support were to me. I sincerely hope that in completing the degree he helped me start, I honor his memory and his legacy as an advisor.

Separately from the work presented here, I have enjoyed a set of wonderful and unexpected research collaborations during my years of graduate training. William Kettyle inspired my work on the quantitative physiology of endurance running, a labor of love that has unexpectedly enjoyed some popular attention these past several months [60]. Joseph Madsen and Stan Anderson have been mentors to me in neurosurgery, and it has been my privilege to work with them in exploring a set of interesting clinical problems [22, 17].

After leaving my first home in New York City, I was able to make Lowell House my home at Harvard for almost a decade, first as an undergraduate, and then, for five years, as a resident tutor. House Masters Diana Eck and Dorothy Autin have become like family to me, and my life has been enriched immeasurably by friendships, conversations, responsibilities, and immersion in an endlessly fascinating community of gifted, worldly, creative scholars—not to mention the absolutely unique adventures I have had in my role as keeper of the Russian bells at Lowell House during a historic period of transition. I thank the students at Lowell House it has been my privilege to advise, knowing how much I have learned from working with people who share some of my passions, and who will eventually be my colleagues. I also thank Dr. Maurice

and Mrs. Kitty Pechet, true role models in so many ways, whose confidence in me since my undergraduate years has been validating, and who have graciously treated me as a member of their extended family. Livingston Taylor, my masterfully talented former neighbor at Lowell House, taught me everything I know about speaking for audiences.

It could be argued that I have enjoyed more than my fair share of adventures since beginning my graduate work. Some of my closest friends have shared those adventures with me. In particular, Craig Mermel has been my partner in the theory and practice of endurance running; and Adam Cohen and I, first childhood friends and now colleagues, have traveled and lived like brothers while working in Liberia.

Some very special individuals—friends, family, and colleagues—have helped me during my years of graduate work in ways that are difficult to explain concisely: Marlys Fassett, Sahar Nissim, Alex Papanastassiou, Sarah Rapoport, Lizzy Rossin, Sol Schulman, Lauren Zander. I thank you all.

I am grateful to the granting agencies that awarded the fellowships that have supported my work and afforded me an added level of intellectual freedom while pursuing my doctoral studies: the Graduate Research Fellowship Program at the National Science Foundation, the Medical Scientist Training Program at the National Institutes of Health, the Hugh Hampton Young Fellowship at MIT, the CIMIT–MIT Medical Engineering Fellowship from the Center for Integration of Medicine and Innovative Technology, and the Defense Advanced Research Projects Agency (DARPA) through a grant to Lincoln Laboratory.

Finally, I thank the members of my thesis committee: Dennis M. Freeman, Professor of Electrical Engineering, MIT Department of Electrical Engineering and Computer Science; Emad N. Eskandar, Associate Professor in Surgery (Neurosurgery), Harvard Medical School and Director of Stereotactic and Functional Neurosurgery, Massachusetts General Hospital; and Jongyoon Han, Associate Professor of Electrical Engineering, MIT Department of Electrical Engineering and Computer Science.

Contents

1	Introduction	13
1.1	Overview	13
1.2	Neural Decoding in General	15
1.3	Micropower Neural Decoders for Implantable Brain–Machine Interfaces	16
1.4	Energy Harvesting for Bioimplantable Electronics	20
1.5	A Glucose-Based Bioimplantable Fuel Cell	21
1.6	Literature Reviews	24
1.7	Structure of This Thesis	26
2	Biomimetic Algorithms and Micropower Circuit Architectures for Real-Time Neural Decoding	27
2.1	Overview	27
2.2	A Biomimetic Adaptive Algorithm and Low-Power Architecture for Implantable Neural Decoders	28
2.3	A Biomimetic Adaptive Algorithm and Micropower Circuit Architecture for Implantable Neural Decoders	33
3	Efficient Universal Computing Architectures for Decoding Neural Activity	75
3.1	Overview	75
3.2	Efficient Universal Computing Architectures for Decoding Neural Activity	76

4	Bioimplantable Glucose-Based Fuel Cells	101
4.1	Overview	101
4.2	An Implantable Glucose Fuel Cell for Brain–Machine Interfaces . . .	101
5	Conclusions	141
5.1	Motivations	141
5.2	Summary of Work	143
5.2.1	Biomimetic Algorithms and Circuit Architectures for Micropower Analog Neural Decoding	143
5.2.2	Universal Computing Architectures for Micropower Digital Neural Decoding	143
5.2.3	Biologically Implantable Glucose Fuel Cells for Brain-Powered Brain–Machine Interfaces	144
5.3	Future Outlook	145
A	Literature Reviews	147
A.1	Brain–Machine Interfaces and Neural Decoding	148
A.1.1	Brain–Machine Interfaces	148
A.1.2	Microelectronics for Brain–Machine Interfaces	148
A.1.3	Neural Decoding	149
A.2	Bioimplantable Device Engineering and Bioimplantable Power Sources	150
A.2.1	Bioimplantable Device Engineering	150
A.2.2	Bioimplantable Power Sources	152
A.3	Microbial Fuel Cells	153
A.4	Symbiosis	155
	References	157

List of Figures

1-1	Conceptual Design of a BrainMachine Interface System	14
1-2	Conceptual Design of a Neural Decoding System	17
1-3	Conceptual Design of an Implantable Glucose Fuel Cell	24

Chapter 1

Introduction

1.1 Overview

Imagine treating severe and currently irreparable brain damage by neurosurgically implanting synthetic neural circuitry—microelectronic devices that interface with neural tissue to restore and augment neurologic function. Imagine, too, being able to offer hope for recovery to victims of strokes, severe trauma, seizure disorders, dementia, and even psychiatric illness—through implantable devices that learn and compute like biological neurons, harvest power from their surrounding environments, and operate on the tiny scales of living cells.

The work presented in this thesis is part of an effort to develop a set of platform technologies for brain–machine interfaces, with a particular emphasis on *neuromotor prosthetics*: brain-controlled prosthetic limbs for paralyzed people, amputees, and others who have sustained damage to regions of the nervous system that control movement.

Neuromotor prosthetics and other brain–machine interface systems seek to restore and augment neural function by facilitating direct communication between the brain and external devices. These systems use brain-implanted electrode interfaces to extract neural signals associated with intended limb movement. They then decode the acquired neural signals into motor commands used to control the movement of robotic prosthetic limbs, computer interfaces, or other devices.




Figure 1-1: **Conceptual Design of a BrainMachine Interface System.** An *Electrode Array* implanted in the brain detects neural signals that are transmitted via a *Flexible Cable* to an *Implanted Unit*, which amplifies, digitizes, and decodes the neural signals into control parameters capable of operating an *External Device*. Decoded signals are transmitted wirelessly through the skin from the *Internal Unit* to the *External Unit*, which then broadcasts the control signals across a *Wireless Link* to the target device. Copyright © 2007 by Benjamin I. Rapoport.

This thesis describes core contributions to a collaboration to design, build, and test a micropower, microelectronic system that could form the technological core of the first fully implantable—and hence potentially clinically viable—neuromotor prosthesis. The overall system, a microchip-based prosthesis with brain-implanted electrodes, neural signal conditioning and decoding circuitry, a wireless transceiver, and wireless recharging circuitry, should run for the lifetime of a patient.

1.2 Neural Decoding in General

Every brain is unique, and so the meanings of neural signals differ from person to person. Neural prosthetics therefore cannot be programmed in advance with all of the information required to decode signals from a particular brain. Rather, they must incorporate flexible algorithms capable of learning personalized neural codes after implantation in the brain. Neural decoding algorithms for clinically viable neural prostheses must also be implementable in circuitry that consumes so little power that implantation in a living brain will not harm surrounding brain tissue. This thesis describes two such decoding systems. The first, an analog system, consists of a set of biomimetic algorithms and a microelectronic architecture for adaptively and automatically learning to interpret patterns of electrical activity in neuronal populations—in real time. The system is *biomimetic* in that it implements a real-time artificial neural network with analog synaptic dynamics resembling those of real neurons. The second, a digital system of extremely low computational complexity and suited for highly scalable implantable systems, is an explicit implementation of a Turing-type computing machine. It requires no arithmetic operations, and decodes neural signals using only computationally inexpensive logic operations.

1.3 Micropower Neural Decoders for Implantable Brain–Machine Interfaces

Brain–machine interfaces have proven capable of decoding neuronal population activity in real time to derive instantaneous control signals for robotic prosthetic limbs and other devices, and clinically viable neural prosthetics are an eagerly anticipated advance in the field of rehabilitation medicine. Algorithmically and energetically efficient computational architectures that operate in real time are essential for clinically useful neural prosthetic devices. Such devices not only decode raw neural data to generate control signals for external devices, but also perform data compression and dramatically reduce the bandwidth and consequently the power expended in wireless transmission of raw data from implantable brain–machine interfaces. The first section of this thesis describes a biomimetic algorithm and micropower analog circuit architecture for decoding neural cell ensemble signals. The decoding algorithm implements a continuous-time artificial neural network, using a bank of adaptive linear filters with kernels that emulate synaptic dynamics. The filters transform neural signal inputs into control-parameter outputs, and can be tuned automatically in an on-line learning process.

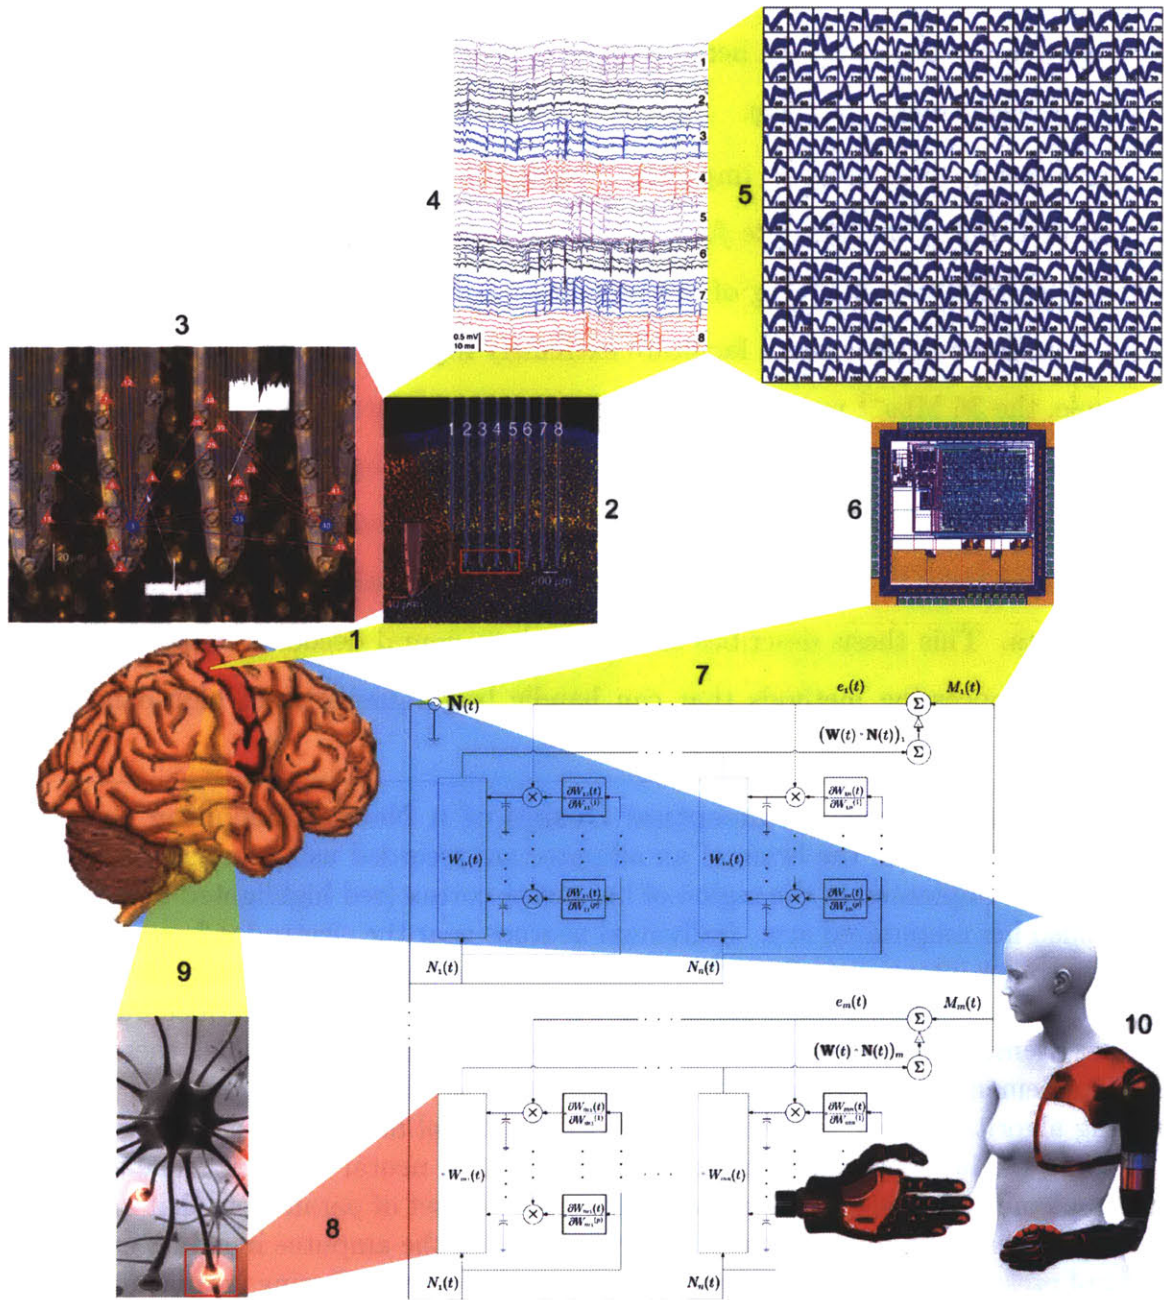
The analog system presented here represents a departure from previous work in the field of neural decoding, as all of the decoding systems demonstrated to date have operated by analyzing digitized neural data [15, 94, 88, 55, 29, 73, 34].

Development of brain–machine interfaces that wirelessly transmit neural data to external devices will represent an important step toward clinical viability. The general model for such devices has two components, as shown in Figure 1-2: a brain-implanted unit directly connected to a multielectrode array collecting raw neural data; and a unit outside the body for data processing, decoding, and control. Data transmission between the two units is wireless. A 100-channel, 12-bit-precise digitization of raw neural waveforms sampled at 30 kHz generates 36 Mbs^{-1} of data; the power costs in digitization, wireless communication, and population signal decoding all scale with this high data rate. Consequences of this scaling, as seen for example in cochlear-

implant systems, include unwanted heat dissipation in the brain, decreased longevity of batteries, and increased size of the implanted unit. Recent designs for system components have addressed these issues in several ways. However, with the exception of [4] and [64], which describe work done by the author and colleagues in pursuit of the research agenda described here, *essentially no work has been done in the area of power-efficient neural decoding.*

In addition to playing an important functional role in neuroprosthetic devices, *neural decoding is the ultimate form of compression of neural data.* Direct, power-efficient analysis and decoding of analog neural data within the implanted unit of a brain-machine interface could facilitate extremely high data compression ratios. For example, the 36 Mbs^{-1} required to transmit raw neural data from 100 channels could be compressed more than 100,000-fold to 300 bs^{-1} of 3-channel motor output information updated with 10-bit precision at 10 Hz. Such dramatic compression brings concomitant reductions in the power required for communication and digitization of neural data. This thesis describes an approach to neural decoding using low-power analog preprocessing methods that can handle large quantities of high-bandwidth

Figure 1-2 (*facing page*): **Conceptual Design of a Neural Decoding System.** 1 Neural signals from the brain of an amputee are recorded using an array of microelectrodes 2 implanted in the region of her motor cortex (red highlighted region) that controlled her amputated arm. Individual neurons near the electrodes 3 respond electrically to her intention to move her amputated arm, generating extracellular voltage waveforms 4 containing sharp spikes, known as ‘action potentials,’ with characteristic shapes identifying different neurons 5. My neural decoder 6 examines these spikes, and implements a decoding algorithm illustrated in the block diagram 7. The decoding algorithm itself simulates the properties of real neural networks 9, with each decoding module corresponding to a synapse 8 in a neural network. The synapses are implemented as signal filters, and each one has a set of parameters that are tuned during the learning phase of the algorithm, in which the amputee imagines executing a set of prescribed movements designed to train the system. During normal operation, the output of the neural decoder is a stream of control parameters used to operate a prosthetic limb 10. *Image Credits for Figure Components: Neuron illustration by Shutterstock Images, Prosthetic Arm illustration by the Johns Hopkins University Applied Physics Laboratory and Bryan Christie Design, Multichannel Neural Recording Grid by Miguel A. L. Nicolelis (PNAS 2003), Microelectrode images and Voltage Tracings by G. Buzsáki (Nature Neuroscience 2004).*



analog data, processing neural input signals slowly and in parallel to generate low-bandwidth control outputs.

Multiple approaches to neural signal decoding have been demonstrated by a number of research groups employing highly programmable, discrete-time, digital algorithms, implemented in software by processors located outside the brain. I am unaware of any work on continuous-time analog decoders or analog circuit architectures for neural decoding. The first neural signal decoder presented here is designed to complement and integrate with existing approaches. Optimized for implementation in micropower analog circuitry, it sacrifices some algorithmic programmability to reduce the power consumption and physical size of the neural decoder, facilitating use as a component of a unit implanted within the brain. Trading off the flexibility of a general-purpose digital system for the efficiency of a special-purpose analog system may be undesirable in some neural prosthetic devices. Therefore, the analog decoder is meant to be used not as a substitute for digital signal processors but rather as an adjunct to digital hardware, in ways that combine the efficiency of embedded analog preprocessing options with the flexibility of a general-purpose external digital processor. The micropower digital decoder developed in the second part of this thesis explicitly manipulates the trade-off between computationally and energetically efficient implanted system components, on the one hand, and complementary, more complex decoding elements located outside the implant, on the other.

For clinical neural prosthetic devices, the necessity of highly sophisticated decoding algorithms remains an open question, since both animal [88, 13, 55, 91] and human [29] users of even first-generation neural prosthetic systems have proven capable of rapidly adapting to the particular rules governing the control of their brain-machine interfaces. In the present work we therefore focus on implantable architectures that implement simple, continuous-time decoding algorithms. The approaches we present here can be generalized to implement circuit architectures for generalized linear models and Bayesian algorithms; examples of related systems include analog probabilistic decoding circuit architectures used in speech recognition and error correcting codes [38, 43].

1.4 Energy Harvesting for Bioimplantable Electronics

Biologically implantable devices are increasing both in prevalence and in importance as treatment modalities in contemporary medicine and as tools in basic scientific research. In particular, four programmable, implantable electronic devices, have emerged as standards of care in the treatment of some electrophysiologic disorders: implantable cardiac pacemakers and implantable cardioverter-defibrillators have extended the healthy lifespans of millions of people with heart arrhythmias [23], cochlear implants have afforded functional hearing to hundreds of thousands of children who would otherwise have grown up deaf [54], and deep brain stimulators have been used to treat tens of thousands of people with Parkinson's disease and related movement disorders [87]. These devices all use implanted power sources, as will many implantable electronic devices of the future.

Various solutions to the problem of providing power to biologically implanted devices have been proposed, prototyped, or implemented. Two principal solutions are currently in widespread use: single-use batteries, such as those used in implantable pulse generators for cardiac pacing, defibrillation, and deep brain stimulation, which are designed to have finite lifetimes and to be replaced surgically at intervals of several years [23]; and inductive power transfer systems, typically functioning transcutaneously at radio frequencies, as in cochlear implants [76]. Recent advances in battery technology and related fields, leading to increased energy and power densities in small devices such as supercapacitors [11, 35] as well as thin film lithium and thin film lithium ion batteries [20], will facilitate improvements in systems based on these two solutions, particularly by shrinking battery sizes and extending battery lifetimes.

An alternative approach to powering biologically implanted devices is to employ schemes for harvesting power from the biological environment. Important objectives of such *energy scavenging* systems are to provide stable, long-lasting sources of power, and to obviate the need for surgical replacement of implantable batteries. Systems for transducing light, thermal energy, mechanical vibration, low-frequency mechanical

stresses or strains, and near- or far-field electromagnetic radiation, into electrical energy have been described and implemented, and complementary electronic design techniques required to make use of the power generated by such systems have been described and reviewed [14].

One approach to harvesting energy from a physiologic environment is to extract it from physiologically available biological fuel substrates such as glucose. Three major design paradigms for glucose-based fuel cells have been described, differing principally in the materials used to catalyze electrode reactions: In one paradigm the catalysts are abiotic; in the second paradigm the catalysts are isolated enzymes fixed to electrode substrates; and in the third paradigm oxidation is performed by exoelectrogenic bacteria colonizing a fuel cell anode. Numerous designs representing each of these fuel cell paradigms have been described and reviewed in the scientific and patent literatures. An extensive review of the scientific and patent literatures on abiotic implantable glucose fuel cells is provided in [36]. Bioimplantable fuel cells based on enzymatic catalysis are reviewed in [27]. Microbial fuel cells are reviewed in Appendix A. While several proposals for biologically implantable fuel cells of each type have been described, no design has truly been reduced to practice *in vivo*.

1.5 A Glucose-Based Bioimplantable Fuel Cell

As implantable electronic devices become increasingly prevalent in scientific research and in the diagnosis, management, and treatment of human disease, there is a correspondingly increasing demand for devices (including pacemakers, implantable cardiac defibrillators, cochlear implants, and deep brain stimulators) with extremely long functional lifetimes. The first chapters of this thesis focus on developing implantable electronic devices under microwatt power constraints, in order to avoid damaging sensitive brain tissue when such devices are used in the context of brain-machine interfaces. Such micropower implantable electronic devices, which consume power at rates comparable to those of their host biological systems, beg the question of whether such electronics can be powered in the same manner as are their surrounding tissues.

A biologically implantable system that could harvest all of the power it required from its surroundings would represent a major step toward meeting this demand by escaping an important constraint on system lifetime, the need to recharge or replace an implanted power supply, in favor of a power supply with an effectively unlimited lifetime. In developing an implantable biofuel cell designed for a micropower electronic system, the second major project presented in this thesis seeks to demonstrate the feasibility of such a paradigm.

This thesis describes a biologically implantable biofuel cell that converts extracellular glucose into electrical power. As described in the previous section, there are three principal approaches to developing a glucose-based bioimplantable fuel cell. The work described here focuses on the first of those approaches. The fuel cell uses solid-state catalysts to oxidize glucose at the anode, liberating electrons to generate an electronic current at a potential defined by the electrochemical reactions at the anode and cathode.

In general, a fuel cell generates power by catalyzing complementary electrochemical reactions (oxidation and reduction) at a pair of corresponding electrodes (the anode and cathode, respectively), as a fuel flows across the anode and an oxidant flows across the cathode. A fuel substrate is electrooxidized at the anode, which collects electrons from the oxidation current and conducts them through an external load to the cathode. Typically an ion-selective membrane partitions the anode and cathode into separate chambers, facilitating a unidirectional flow of protons from anode to cathode that leads to overall charge balance as electrons arriving at the cathode are used to reduce the oxidant in the cathode chamber.

In a glucose-based biofuel cell, glucose is oxidized at the anode, while oxygen is reduced to water at the cathode. This scheme is illustrated in Figure 1-3. The nature of the catalyst residing at the anode determines the extent of glucose oxidation and the associated oxidation products. Traditional solid-state anode catalysts predominantly oxidize glucose to a single product such as gluconic acid, liberating a single pair of electrons. By contrast, using living microorganisms, such as exoelectrogenic bacteria, to catalyze the anodic reaction can result in complete oxidation of glucose, liberating

twenty-four electrons per molecule of glucose consumed. This much higher catalytic efficiency, and correspondingly greater power-generating capabilities, have led to a considerable body of work on microbial fuel cells, including some proposals for using such fuel cells as components of implantable systems. These ideas are reviewed in greater detail in the literature review of Section A.3 of Appendix A.

The availability of fuel cell reactants, oxygen and glucose, only as a mixture in the physiologic environment, has traditionally posed a design challenge, as net current production requires oxidation and reduction to occur separately and selectively at the anode and cathode, respectively; glucose oxidation at the cathode and oxygen reduction at the anode generate electrochemical short circuits by abolishing the potential difference that induces electron flow through the external circuit. The fuel cell developed here is configured in a half-open geometry that shields the anode while exposing the cathode, resulting in an oxygen gradient that strongly favors oxygen reduction at the cathode. Glucose reaches the shielded anode by diffusing through the porous cathode, which does not catalyze glucose oxidation, and through the ion-selective membrane, which is permeable to small neutral and cationic species.

One important innovation of the work described here is the use of cerebrospinal fluid as a physiologic niche for an implantable power source. The thesis describes an intracranial configuration of the biofuel cell that enables it to harvest glucose from the cerebrospinal fluid surrounding the brain in order to power brain-implanted microelectronic devices. The cerebrospinal fluid represents a promising and unexplored environment for a bioimplantable fuel cell: it is virtually acellular, it is under minimal immune surveillance, it has a lower protein content than blood and other tissues and is therefore less prone to induce biofouling of implanted devices, and its glucose levels under normal conditions are approximately half those of blood plasma.

The bioavailabilities of glucose and oxygen to a fuel cell residing in the subarachnoid space are modeled in detail in Chapter 4, with a view toward estimating the amount of power that can safely be drawn from cerebrospinal fluid glucose and oxygen without appreciably altering cerebrospinal fluid physiology.

Another important innovation in the present work relates to its manufacturing


process, which facilitates true integration with microelectronic circuits at the most basic level. The fuel cell is constructed using conventional semiconductor fabrication processes, and so is well suited for manufacture together with traditional integrated circuits on a single silicon wafer to yield implantable microchips with long-lifetime, integrated power sources.

The fuel cell is used to power a neural signal amplifier, the fundamental front-end used to generate input signals for a neural decoder, in a prototype of the first *brain-powered brain-machine interface*.

1.6 Literature Reviews

Detailed discussions in Appendix A review the foundational and most relevant scientific literature published to date in fields spanned by this thesis. Section A.1 considers prior work on brain-machine interfaces and neural decoding, while Section A.2 surveys the field of bioimplantable device engineering, with a focus on implantable power sources and areas of emerging research related to the work described in this thesis.

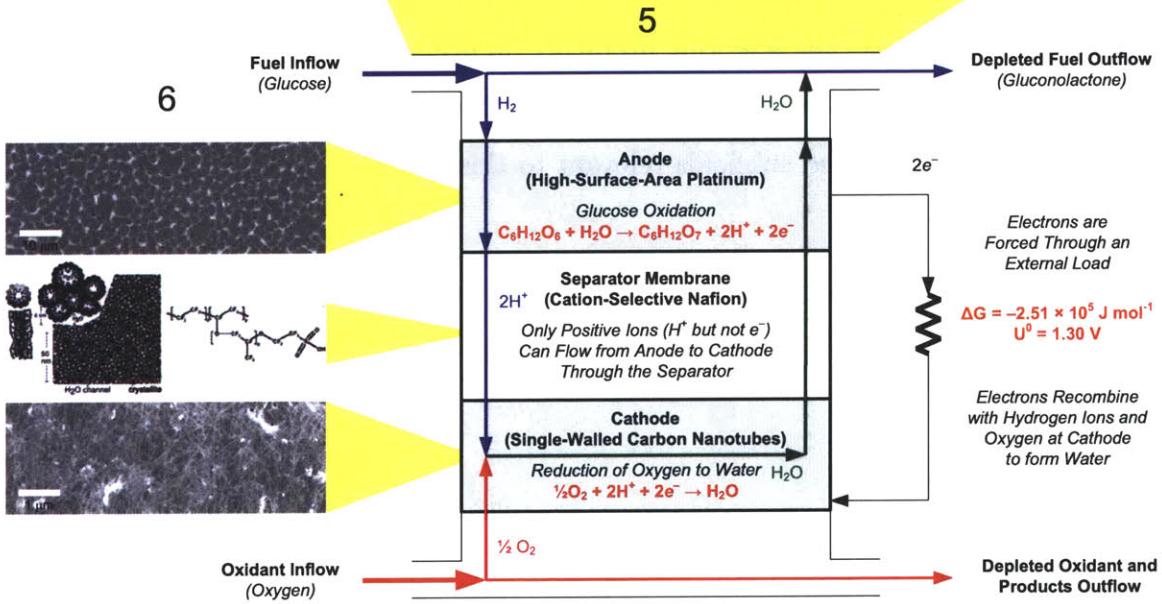
Figure 1-3 (*facing page*): **Conceptual Design of a Brain-Implantable Glucose-Based Fuel Cell.** **1** Cerebrospinal fluid flows within the subarachnoid space, illustrated in the first inset, surrounding the brain. A biofuel cell placed within the subarachnoid space **2** could harvest glucose from the cerebrospinal fluid to generate electrical power. A prototype of one such fuel cell, manufactured on a silicon dioxide substrate using semiconductor microfabrication techniques, is shown in micrograph **3**. Its principles of operation are illustrated in schematics **4** and **5**, while special materials properties of its catalytic anode, cation-selective membrane, and current-collecting cathode are highlighted in the micrographs and diagrams of **6**. *Image Credits for Figure Components: Skull and Subarachnoid Space Inset illustration by A.D.A.M. Medical Images, Meninges and Human Head illustration by Karolinska Institute 3D Brain Project, Nafion Chemical Structure from Wikimedia Commons, Nafion Water Channel Model illustration by Schmidt-Rohr and Chen (Nature Materials 2008).* All other figure components are original micrographs and diagrams generated by the author.



Oxygen Gradient
 Oxygen Reduction at Cathode Prevents
 Oxygen from Reaching Anode



Glucose Gradient
 Carbon Does Not Catalyze Glucose Oxidation
 Porous Nanotube Mesh and Permeable Nafion Pass
 Glucose Freely to Anode



1.7 Structure of This Thesis

This thesis is structured as follows:

The present chapter, Chapter 1, places the work presented here in context, introducing background material relevant to the scientific and technological visions to which this work has contributed, and outlining a conceptual framework developed in detail in subsequent chapters. Chapter 2 presents work on biomimetic algorithms for neural decoding, together with micropower circuit architectures for implementing such algorithms in the context of fully brain-implantable brain-machine interfaces for neural prosthetic devices and neurorehabilitation in victims of traumatic limb loss, amputation, and paralysis; and in patients suffering from locked-in syndrome or severe and degenerative neuromotor diseases such as amyotrophic lateral sclerosis (ALS or Lou Gehrig’s disease). This work is developed further in Chapter 3, which discusses digital approaches to neural decoding using arithmetic-free, purely logical primitives, and establishes universal computing architectures for neural decoding. Chapter 4 provides a detailed discussion of the design and performance of a biologically implantable glucose-powered fuel cell. Finally, Chapter 5 presents conclusions, summarizing the achievements of the work described here.

Appendix A reviews the scientific literature on brain-machine interfaces and neural decoding, as well as the scientific and patent literatures on bioimplantable power sources and several related subfields relevant to this thesis.

Chapter 2

Biomimetic Algorithms and Micropower Circuit Architectures for Real-Time Neural Decoding

2.1 Overview

This chapter presents a corpus of work on biomimetic algorithms for neural decoding, together with micropower circuit architectures for implementing such algorithms in the context of fully brain-implantable brain-machine interfaces. The primary applications of this work are in the field of neurorehabilitation, designing brain-controlled prosthetic devices to restore natural movement and communication to victims of traumatic limb loss, amputation, and paralysis; as well as to patients suffering from locked-in syndrome, or severe and degenerative neuromotor diseases such as amyotrophic lateral sclerosis (ALS or Lou Gehrig’s disease). As discussed in detail in Chapter 1, the Introduction to this thesis, *neuroprosthetic* devices decode ‘thoughts’ encoded in electrical signals from neuronal populations to control mechanical or computer interfaces. The electronic systems that implement the neural decoding algorithms described here are *biomimetic* in that they compute in ways similar to biological neurons and neuronal networks.

At the time this chapter was composed, the work it describes had already been published as a conference paper [4] and a book chapter [64].

2.2 A Biomimetic Adaptive Algorithm and Low-Power Architecture for Implantable Neural Decoders

This section was previously published as a paper in the *Proceedings of the 31st International Meeting of the IEEE Engineering in Medicine and Biology Society* [4], Copyright 2009 IEEE. It is reprinted here with the permission of the publisher.

A Biomimetic Adaptive Algorithm and Low-Power Architecture for Implantable Neural Decoders

Benjamin I. Rapoport, *Student Member, IEEE*, Woradorn Wattanapanitch, *Student Member, IEEE*, Hector L. Penagos, Sam Musallam, Richard A. Andersen, and Rahul Sarpeshkar, *Senior Member, IEEE*

Abstract—Algorithmically and energetically efficient computational architectures that operate in real time are essential for clinically useful neural prosthetic devices. Such devices decode raw neural data to obtain direct control signals for external devices. They can also perform data compression and vastly reduce the bandwidth and consequently power expended in wireless transmission of raw data from implantable brain-machine interfaces. We describe a biomimetic algorithm and micropower analog circuit architecture for decoding neural cell ensemble signals. The decoding algorithm implements a continuous-time artificial neural network, using a bank of adaptive linear filters with kernels that emulate synaptic dynamics. The filters transform neural signal inputs into control-parameter outputs, and can be tuned automatically in an on-line learning process. We provide experimental validation of our system using neural data from thalamic head-direction cells in an awake behaving rat.

Index Terms—Brain-machine interface, Neural decoding, Biomimetic, Adaptive algorithms, Analog, Low-power

I. INTRODUCTION

BRAIN-MACHINE interfaces have proven capable of decoding neuronal population activity in real-time to derive instantaneous control signals for prosthetics and other devices. All of the decoding systems demonstrated to date have operated by analyzing digitized neural data [1]–[7]. Clinically viable neural prosthetics are an eagerly anticipated advance in the field of rehabilitation medicine, and development of brain-machine interfaces that wirelessly transmit neural data to external devices will represent an important step toward clinical viability. The general model for such devices has two components: a brain-implanted unit directly connected to a multielectrode array collecting raw neural data; and a unit outside the body for data processing, decoding, and control. Data transmission between the two units is wireless. A 100-channel, 12-bit-precise digitization of raw neural waveforms sampled at 30 kHz generates 36 Mbs^{-1} of data; the power costs in digitization, wireless communication, and population signal decoding all scale with this high data rate. Consequences of

this scaling, as seen for example in cochlear-implant systems, include unwanted heat dissipation in the brain, decreased longevity of batteries, and increased size of the implanted unit. Recent designs for system components have addressed these issues in several ways. However, *almost no work has been done in the area of power-efficient neural decoding.*

In this work we describe an approach to neural decoding using low-power analog preprocessing methods that can handle large quantities of high-bandwidth analog data, processing neural input signals in a slow-and-parallel fashion to generate low-bandwidth control outputs.

Multiple approaches to neural signal decoding have been demonstrated by a number of research groups employing highly programmable, discrete-time, digital algorithms, implemented in software or microprocessors located outside the brain. We are unaware of any work on continuous-time analog decoders or analog circuit architectures for neural decoding. The neural signal decoder we present here is designed to complement and integrate with existing approaches. Optimized for implementation in micropower analog circuitry, it sacrifices some algorithmic programmability to reduce the power consumption and physical size of the neural decoder, facilitating use as a component of a unit implanted within the brain. Trading off the flexibility of a general-purpose digital system for the efficiency of a special-purpose analog system may be undesirable in some neural prosthetic devices. Therefore, our proposed decoder is meant to be used not as a substitute for digital signal processors but rather as an adjunct to digital hardware, in ways that combine the efficiency of embedded analog preprocessing options with the flexibility of a general-purpose external digital processor.

For clinical neural prosthetic devices, the necessity of highly sophisticated decoding algorithms remains an open question, since both animal [3], [4], [8], [9] and human [5] users of even first-generation neural prosthetic systems have proven capable of rapidly adapting to the particular rules governing the control of their brain-machine interfaces. In the present work we focus on an architecture to implement a simple, continuous-time analog linear (convolutional) decoding algorithm. The approach we present here can be generalized to implement analog-circuit architectures of general Bayesian algorithms; examples of related systems include analog probabilistic decoding circuit architectures used in speech recognition and error correcting codes [10], [11]. Such architectures can be extended through our mathematical approach to design circuit architectures for Bayesian decoding.

B. I. Rapoport, W. Wattanapanitch, and R. Sarpeshkar are with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology (MIT), Cambridge, Massachusetts 02139 USA; B. I. Rapoport is also with the Harvard-MIT Division of Health Sciences and Technology, Cambridge, Massachusetts 02139 and Harvard Medical School, Boston, Massachusetts 02115; E-mail: rahuls@mit.edu.

H. L. Penagos is with the MIT Department of Brain and Cognitive Sciences and the Harvard-MIT Division of Health Sciences and Technology.

S. Musallam is with the Department of Electrical and Computer Engineering, McGill University, Montreal, Quebec, Canada.

R. A. Andersen is with the Division of Biology, California Institute of Technology, Pasadena, California 91125.

Manuscript received April 2009.

II. A BIOMIMETIC ADAPTIVE ALGORITHM FOR DECODING NEURAL CELL ENSEMBLE SIGNALS

In convolutional decoding of neural cell ensemble signals, the decoding operation takes the form

$$\vec{M}(t) = \mathbf{W}(t) \circ \vec{N}(t) \quad (1)$$

$$M_i(t) = \sum_{j=1}^n W_{ij}(t) \circ N_j(t); \quad i \in \{1, \dots, m\}, \quad (2)$$

where $\vec{N}(t)$ is an n -dimensional vector containing the neural signal (n input channels of neuronal firing rates, analog signal values, or local field potentials, for example) at time t ; $\vec{M}(t)$ is a corresponding m -dimensional vector containing the decoder output signal (which in the examples presented here corresponds to motor control parameters, but could correspond as well to limb or joint kinematic parameters or to characteristics or states of nonmotor cognitive processes); \mathbf{W} is a matrix of convolution kernels $W_{ij}(t)$ (formally analogous to a matrix of dynamic synaptic weights), each of which depends on a set of p modifiable parameters, W_{ij}^k , $k \in \{1, \dots, p\}$; and \circ indicates convolution. Accurate decoding requires first choosing an appropriate functional form for the kernels and then optimizing the kernel parameters to achieve maximal decoding accuracy. Since the optimization process is generalizable to any choice of kernels that are differentiable functions of the tuning parameters, we discuss the general process first. We then explain our biophysical motivations for selecting particular functional forms for the decoding kernels; appropriately chosen kernels enable the neural decoder to emulate the real-time encoding and decoding processes performed by biological neurons.

Our algorithm for optimizing the decoding kernels uses a gradient-descent approach to minimize decoding error in a least-squares sense during a learning phase of decoder operation. During this phase the correct output $\hat{M}(t)$, and hence the decoder error $\vec{e}(t) = \vec{M}(t) - \hat{M}(t)$, is available to the decoder for feedback-based learning. We design the optimization algorithm to evolve $\mathbf{W}(t)$ in a manner that reduces the squared decoder error on a timescale set by the parameter τ , where the squared error is defined as

$$E(\mathbf{W}(t), \tau) = \int_{t-\tau}^t |\vec{e}(u)|^2 du \quad (3)$$

$$= \sum_{i=1}^m \int_{t-\tau}^t |\vec{e}_i(u)|^2 du \equiv \sum_{i=1}^m E_i, \quad (4)$$

and the independence of each of the m terms in Equation 4 is due to the independence of the m sets of np parameters W_{ij}^k , $j \in \{1, \dots, n\}$ $k \in \{1, \dots, p\}$ associated with generating each component $M_i(t)$ of the output. Our strategy for optimizing the matrix of decoder kernels is to modify each of the kernel parameters W_{ij}^k continuously and in parallel, on a timescale set by τ , in proportion to the negative gradient of $E(\mathbf{W}(t), \tau)$ with respect to that parameter:

$$-\vec{\nabla}_{ij}^k E(\mathbf{W}(t), \tau) \equiv -\frac{\partial E}{\partial W_{ij}^k} \quad (5)$$

$$= -\sum_{l=1}^m \int_{t-\tau}^t du \left\{ 2 \left(M_l(u) - \sum_{j=1}^n W_{lj}(u) \circ N_j(u) \right) \times \left(-\frac{\partial W_{ij}(u)}{\partial W_{ij}^k} \circ N_j(u) \right) \right\} \quad (6)$$

$$= 2 \sum_{l=1}^m \int_{t-\tau}^t e_l(u) \left(\frac{\partial W_{ij}(u)}{\partial W_{ij}^k} \circ N_j(u) \right) du. \quad (7)$$

The learning algorithm refines \mathbf{W} in a continuous-time fashion, using $-\vec{\nabla}E(t)$ as an error feedback signal to modify $\mathbf{W}(t)$, and incrementing each of the parameters $W_{ij}^k(t)$ in continuous time by a term proportional to $-\vec{\nabla}_{ij}^k E(\mathbf{W}(t))$ (the proportionality constant, ϵ , must be large enough to ensure quick learning but small enough to ensure learning stability). If $\mathbf{W}(t)$ is viewed as an array of linear filters operating on the neural input signal, the quantity $-\vec{\nabla}_{ij}^k E(\mathbf{W}(t), \tau)$ used to increment each filter parameter can be described as the product, averaged over a time interval of length τ , of the error in the filter output and a secondarily filtered version of the filter input. The error term is identical for the parameters of all filters contributing to a given component of the output, $M_i(t)$. The secondarily filtered version of the input is generated by a secondary convolution kernel, $\frac{\partial W_{ij}(u)}{\partial W_{ij}^k}$, which depends on the functional form of each primary filter kernel and in general differs for each filter parameter. Figure 1 shows a block diagram for an analog circuit architecture that implements our decoding and optimization algorithm.




Fig. 1: Block diagram of a computational architecture for linear convolutional decoding and learning.

Many functional forms for the convolution kernels are both theoretically possible and practical to implement using low-power analog circuitry. Our approach has been to emulate

biological neural systems by choosing a biophysically inspired kernel whose impulse response approximates the postsynaptic currents biological neurons integrate when encoding and decoding neural signals *in vivo* [12]. Combining our decoding architecture with the choice of a first-order low-pass decoder kernel enables our low-power neural decoder to implement a biomimetic, continuous-time artificial neural network. Numerical experiments have also indicated that decoding using such biomimetic kernels can yield results comparable to those obtained using optimal linear decoders [13]. But in contrast with our on-line optimization scheme, optimal linear decoders are computed off-line after all training data have been collected. We have found that this simple choice of kernel offers effective performance in practice, and so we confine the present analysis to that kernel.

Two-parameter first-order low-pass filter kernels account for trajectory continuity by exponentially weighting the history of neural inputs:

$$W_{ij} = \frac{A_{ij}}{\tau_{ij}} e^{-\frac{t}{\tau_{ij}}}, \quad (8)$$

where the two tunable kernel parameters are $W_{ij}^{k=1} = A_{ij}$, the low-pass filter gain, and $W_{ij}^{k=2} = \tau_{ij}$, the decay time over which past inputs $\vec{N}(t')$, $t' < t$, influence the present output estimate $\vec{M}(t) = \mathbf{W} \circ \vec{N}(t)$. The filters used to tune the low-pass filter kernel parameters can be implemented using simple and compact analog circuitry. The gain parameters are tuned using low-pass filter kernels of the form

$$\frac{\partial W_{ij}(t)}{\partial W_{ij}^{k=1}} = \frac{1}{\tau_{ij}} e^{-\frac{t}{\tau_{ij}}}, \quad (9)$$

while the time-constant parameters are tuned using band-pass filter kernels:

$$\frac{\partial W_{ij}(t)}{\partial W_{ij}^{k=2}} = \frac{A_{ij}}{\tau_{ij}^2} e^{-\frac{t}{\tau_{ij}}} \left(\frac{t}{\tau_{ij}} - 1 \right). \quad (10)$$

When decoding discontinuous trajectories, such as sequences of discrete decisions, we can set the τ_{ij} to zero, yielding

$$W_{ij}(t) = W_{ij}^{k=1} \delta(t) = A_{ij} \delta(t). \quad (11)$$

Such a decoding system, in which each kernel is a zeroth-order filter characterized by a single tunable constant, performs instantaneous linear decoding, which has successfully been used by others to decode neuronal population signals in the context of neural prosthetics [5], [14]. With kernels of this form, $\mathbf{W}(t)$ is analogous to matrices of synaptic weights encountered in artificial neural networks, and our optimization algorithm resembles a ‘delta-rule’ learning procedure [15].

III. RESULTS

Head direction was decoded from the activity of $n = 6$ isolated thalamic neurons according to the method described in [16]. The adaptive filter parameters $W_{ij}^{(p)} \in \{A_{ij}, \tau_{ij}\}$ were implemented as micropower analog circuits and simulated in

SPICE; they were optimized through gradient descent over training intervals of length T during which the decoder error, $e_i(t) = M_i(t) - \hat{M}_i(t)$ (where $\vec{M}(t) = (\cos(\theta(t)), \sin(\theta(t)))$ and θ denotes the head direction angle), was made available to the adaptive filter in the feedback configuration described in Section II for $t \in [0, T]$. Following these training intervals feedback was discontinued and the performance of the decoder was assessed by comparing the decoder output $\vec{M}(t)$ with $\hat{\vec{M}}(t)$ for $t > T$.

Figure 2 compares the output of the decoder to the measured head direction over a 240 s interval. The filter parameters were trained over the interval $t \in [0, T = 120]$ s. The figure shows $\vec{M}(t)$ (gray) tracking $\hat{\vec{M}}(t)$ (black) with increasing accuracy as training progresses, illustrating that while initial predictions are poor, they improve with feedback over the course of the training interval. Feedback is discontinued at $t = 120$ s. Qualitatively, the plots on the interval $t \in [120, 240]$ s illustrate that the output of the neural decoder reproduces the shape of the correct waveform, predicting head direction on the basis of neuronal spike rates.

IV. DISCUSSION

Simulations using basic circuit building-blocks for the modules shown in Figure 1 indicate that a single decoding module (corresponding to an adaptive kernel W_{ij} and associated optimization circuitry, as diagrammed in Figure 1) should consume approximately 54 nW from a 1 V supply in 0.18 μm CMOS technology and require less than 3000 μm^2 . Low power consumption is achieved through the use of subthreshold bias currents for transistors in the analog filters and other components. Analog preprocessing of raw neural input waveforms is accomplished by dual thresholding to detect action potentials on each input channel and then smoothing the resulting spike trains to generate mean firing rate input signals. SPICE simulations indicate that each analog preprocessing module should consume approximately 241 nW from a 1 V supply in 0.18 μm CMOS technology. A full-scale system with $n = 100$ neuronal inputs comprising $\vec{N}(t)$ and $m = 3$ control parameters comprising $\vec{M}(t)$ would require $m \times n = 300$ decoding modules and consume less than 17 μW in the decoder and less than 25 μW in the preprocessing stages.

Direct and power-efficient analysis and decoding of analog neural data within the implanted unit of a brain-machine interface could also facilitate extremely high data compression ratios. For example, the 36 Mbs^{-1} required to transmit raw neural data from 100 channels could be compressed more than 100,000-fold to 300 bs^{-1} of 3-channel motor-output information updated with 10-bit precision at 10 Hz. Such dramatic compression brings concomitant reductions in the power required for communication and digitization of neural data. Ultra-low-power analog preprocessing prior to digitization of neural signals could thus be beneficial in some applications.

V. CONCLUSIONS

The algorithm and architecture presented here offer a practical approach to computationally efficient neural signal decoding, independent of the hardware used for their

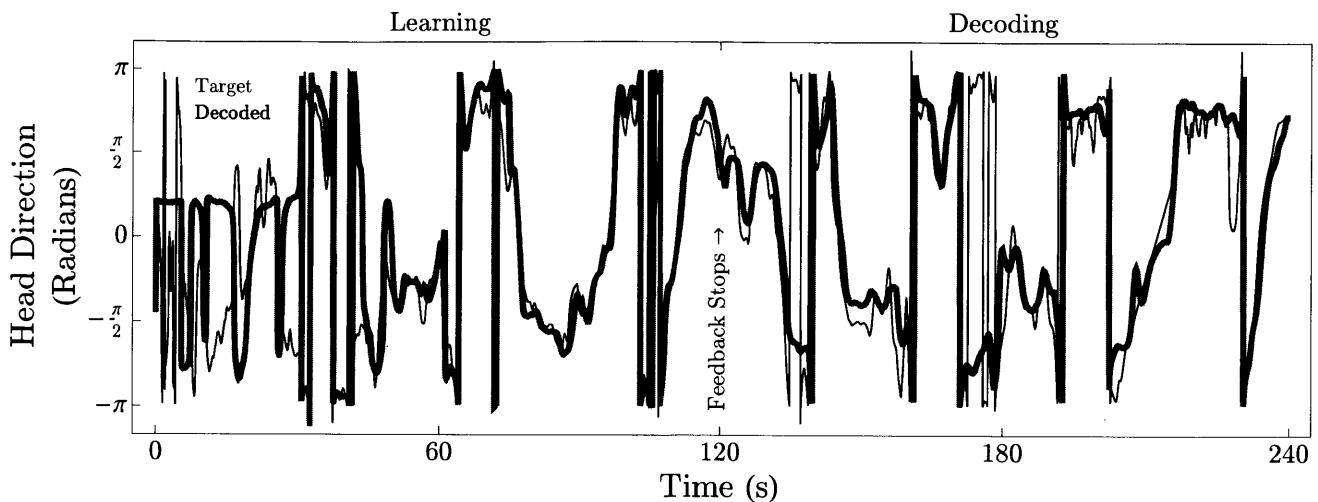


Fig. 2: Continuous decoding of head direction from neuronal spiking activity.

implementation. While the system is suitable for analog or digital implementation, we suggest that a micropower analog implementation trades some algorithmic programmability for reductions in power consumption that could facilitate implantation of a neural decoder within the brain. In particular, circuit simulations of our analog architecture indicate that a 100-channel, 3-motor-output neural decoder can be built with a total power budget of approximately $43 \mu\text{W}$. Our work could also enable a 100,000-fold reduction in the bandwidth needed for wireless transmission of neural data, thereby reducing to nanowatt levels the power potentially required for wireless data telemetry from a brain implant. Our work suggests that highly power-efficient and area-efficient analog neural decoders that operate in real time can be useful components of brain-implantable neural prostheses, with potential applications in neural rehabilitation and experimental neuroscience. Through front-end preprocessing to perform neural decoding and data compression, algorithms and architectures such as those presented here can complement digital signal processing and wireless data transmission systems, offering significant increases in power and area efficiency at little cost.

ACKNOWLEDGMENTS

This work was funded in part by National Institutes of Health grants R01-NS056140 and R01-EY15545, the McGovern Institute Neurotechnology Program at MIT, and National Eye Institute grant R01-EY13337. Rapoport received support from a CIMIT-MIT Medical Engineering Fellowship.

REFERENCES

- [1] J. K. Chapin, K. A. Moxon, R. S. Markowitz, and M. L. Nicolelis. Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex. *Nature Neuroscience*, 2:664–670, 1999.
- [2] J. Wessberg, C. R. Stambaugh, J. D. Kralik, P. D. Beck, M. Laubach, J. K. Chapin, J. Kim, S. J. Biggs, M. A. Srinivasan, and M. A. L. Nicolelis. Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature*, 408:361–365, November 2000.
- [3] D. M. Taylor, S. I. H. Tillery, and A. B. Schwartz. Direct cortical control of 3d neuroprosthetic devices. *Science*, 296:1829–1832, June 2002.

- [4] S. Musallam, B. D. Corneil, B. Greger, H. Scherberger, and R. A. Andersen. Cognitive control signals for neural prosthetics. *Science*, 305:258–262, July 2004.
- [5] L. R. Hochberg, M. D. Serruya, G. M. Friehs, J. A. Mukand, M. Saleh, A. H. Caplan, A. Branner, D. Chen, R. D. Penn, and J. P. Donoghue. Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature*, 442:164–171, July 2006.
- [6] G. Santhanam, S. I. Ryu, B. M. Yu, A. Afshar, and K. V. Shenoy. A high-performance brain-computer interface. *Nature*, 442:195–198, July 2006.
- [7] A. Jackson, J. Mavoori, and E. E. Fetz. Long-term motor cortex plasticity induced by an electronic neural implant. *Nature*, 444:55–60, November 2006.
- [8] J. M. Carmena, M. A. Lebedev, R. E. Crist, J. E. O'Doherty, D. M. Santucci, D. F. Dimitrov, P. G. Patil, C. S. Henriquez, and M. A. L. Nicolelis. Learning to control a brain-machine interface for reaching and grasping by primates. *Public Library of Science Biology*, 1(2):1–16, October 2003.
- [9] M. Velliste, S. Perel, M. C. Spalding, A. S. Whitford, and A. B. Schwartz. Cortical control of a prosthetic arm for self-feeding. *Nature*, 453(7198):1098–1101, June 2008.
- [10] John Lazzaro, John Wawrzynek, and Richard P. Lippmann. A micropower analog circuit implementation of hidden markov model state decoding. *IEEE Journal of Solid-State Circuits*, 32(8):1200–1209, August 1997.
- [11] Hans-Andrea Loeliger, Felix Tarköy, Felix Lustenberger, and Markus Helfenstein. Decoding in Analog VLSI. *IEEE Communications Magazine*, pages 99–101, April 1999.
- [12] A. Arenz, R. A. Silver, A. T. Schaefer, and T. W. Margrie. The contribution of single synapses to sensory representation in vivo. *Science*, 321:977–980, August 2008.
- [13] C. Eliasmith and C. H. Anderson. *Neural Engineering*, chapter 4, pages 112–115. MIT Press, 2003.
- [14] J. Wessberg and M. A. L. Nicolelis. Optimizing a linear algorithm for real-time robotic control using chronic cortical ensemble recordings in monkeys. *Journal of Cognitive Neuroscience*, 16(6):1022–1035, 2004.
- [15] S. Haykin. *Neural Networks: A Comprehensive Foundation*. Prentice Hall, Upper Saddle River, New Jersey, 1999.
- [16] R. Barbieri, L. M. Frank, M. C. Quirk, M. A. Wilson, and E. N. Brown. A Time-Dependent Analysis of Spatial Information Encoding in the Rat Hippocampus. *Neurocomputing*, 32-33:629–635, 2000.

2.3 A Biomimetic Adaptive Algorithm and Micropower Circuit Architecture for Implantable Neural Decoders

This section has been published as a chapter in *System and Circuit Design for Biologically-Inspired Learning* [64], Copyright 2011 IGI Global, www.igi-global.com.

It is reprinted here with the permission of the publisher.

System and Circuit Design for Biologically-Inspired Intelligent Learning

Turgay Temel
Bahcesehir University, Istanbul, Turkey



MEDICAL INFORMATION SCIENCE REFERENCE

Hershey • New York

Director of Editorial Content: Kristin Klinger
Director of Book Publications: Julia Mosemann
Acquisitions Editor: Lindsay Johnston
Development Editor: David DeRicco
Typesetter: Michael Brehm
Production Editor: Jamie Snavely
Cover Design: Lisa Tosheff

Published in the United States of America by
Medical Information Science Reference (an imprint of IGI Global)
701 E. Chocolate Avenue
Hershey PA 17033
Tel: 717-533-8845
Fax: 717-533-8661
E-mail: cust@igi-global.com
Web site: <http://www.igi-global.com>

Copyright © 2011 by IGI Global. All rights reserved. No part of this publication may be reproduced, stored or distributed in any form or by any means, electronic or mechanical, including photocopying, without written permission from the publisher. Product or company names used in this set are for identification purposes only. Inclusion of the names of the products or companies does not indicate a claim of ownership by IGI Global of the trademark or registered trademark.

Library of Congress Cataloging-in-Publication Data

System and circuit design for biologically-inspired learning / Turgay Temel,
editor.

p. cm.

Includes bibliographical references and index.

Summary: "The objective of the book is to introduce and bring together well-known circuit design aspects, as well as to cover up-to-date outcomes of theoretical studies in decision-making, biologically-inspired, and artificial intelligent learning techniques"--Provided by publisher.

ISBN 978-1-60960-018-1 (hardcover) -- ISBN 978-1-60960-020-4 (ebook) 1.

Neural networks (Computer science) 2. Biologically-inspired computing. 3.

Logic circuits. I. Temel, Turgay, 1970-

QA76.87.S97 2010

006.3'2--dc22

2010017198

British Cataloguing in Publication Data

A Cataloguing in Publication record for this book is available from the British Library.

All work contributed to this book is new, previously-unpublished material. The views expressed in this book are those of the authors, but not necessarily of the publisher.

Chapter 10

A Biomimetic Adaptive Algorithm and Micropower Circuit Architecture for Implantable Neural Decoders

Benjamin I. Rapoport

Massachusetts Institute of Technology, USA & Harvard Medical School, USA

Rahul Sarpeshkar

Massachusetts Institute of Technology, USA

ABSTRACT

Algorithmically and energetically efficient computational architectures that operate in real time are essential for clinically useful neural prosthetic devices. Such architectures decode raw neural data to obtain direct motor control signals for external devices. They can also perform data compression and vastly reduce the bandwidth and consequently power expended in wireless transmission of raw data from implantable brain-machine interfaces. We describe a biomimetic algorithm and micropower analog circuit architecture for decoding neural cell ensemble signals. The decoding algorithm implements a continuous-time artificial neural network, using a bank of adaptive linear filters with kernels that emulate synaptic dynamics. The filters transform neural signal inputs into control-parameter outputs, and can be tuned automatically in an on-line learning process. We demonstrate that the algorithm is suitable for decoding both local field potentials and mean spike rates. We also provide experimental validation of our system, decoding discrete reaching decisions from neuronal activity in the macaque parietal cortex, and decoding continuous head direction trajectories from cell ensemble activity in the rat thalamus. We further describe a method of mapping the algorithm to a highly parallel circuit architecture capable of continuous learning and real-time operation. Circuit simulations of a subthreshold analog CMOS instantiation of the architecture reveal that its performance is comparable to the predicted performance of our decoding algorithm for a system decoding three control parameters from 100 neural input channels at microwatt levels of power consumption. While the algorithm and decoding architecture are suitable for analog or digital implementation, we indicate how a micropower analog system trades some algorithmic programmability for reductions in power and area consumption that could facilitate implantation of a

DOI: 10.4018/978-1-60960-018-1.ch010

neural decoder within the brain. We also indicate how our system can compress neural data more than 100,000-fold, greatly reducing the power needed for wireless telemetry of neural data.

INTRODUCTION

Brain–machine interfaces have proven capable of decoding neuronal population activity in real time to derive instantaneous control signals for prosthetics and other devices. All of the decoding systems demonstrated to date have operated by analyzing digitized neural data (Chapin, Moxon, Markowitz, & Nicolelis, 1999; Hochberg et al., 2006; Jackson, Mavoori, & Fetz, 2006; Musallam, Corneil, Greger, Scherberger, & Andersen, 2004; Santhanam, Ryu, Yu, Afshar, & Shenoy, 2006; Taylor, Tillery, & Schwartz, 2002; Wessberg et al., 2000). Clinically viable neural prosthetics are an eagerly anticipated advance in the field of rehabilitation medicine, and development of brain–machine interfaces that wirelessly transmit neural data to external devices will represent an important step toward clinical viability. The general model for such devices has two components: a brain–implanted unit directly connected to a multielectrode array collecting raw neural data; and a unit outside the body for data processing, decoding, and control. Data transmission between the two units is wireless. A 100-channel, 12-bit-precise digitization of raw neural waveforms sampled at 30 kHz generates 36 Mbs⁻¹ of data; the power costs in digitization, wireless communication, and population signal decoding all scale with this high data rate. Consequences of this scaling, as seen for example in cochlear-implant systems, include unwanted heat dissipation in the brain, decreased longevity of batteries, and increased size of the implanted unit. Recent designs for system components have addressed these issues in several ways, including micropower neural amplification (Holleman & Otis, 2007; Wattanapanitch, Fee, & Sarpeshkar, 2007); adaptive power biasing to reduce recording power in multielectrode arrays (Sarpeshkar et al., 2008); low-power data telemetry (Ghovanloo &

Atluri, 2007; Mandal & Sarpeshkar, 2007, 2008; Mohseni, Najafi, Eliades, & Wang, 2005); ultra-low-power analog-to-digital conversion (Yang & Sarpeshkar, 2006); low-power neural stimulation (Theogarajan et al., 2004); energy-efficient wireless recharging (Baker & Sarpeshkar, 2007); and low-power circuits and system designs for brain–machine interfaces (Sarpeshkar et al., 2008; Sarpeshkar et al., 2007). Power-conserving schemes for compressing neural data before transmission have also been proposed (Olsson III & Wise, 2005). However, almost no work has been done in the area of power-efficient neural decoding (Rapoport et al., 2009).

Direct and power-efficient analysis and decoding of analog neural data within the implanted unit of a brain–machine interface could facilitate extremely high data compression ratios. For example, the 36 Mbs⁻¹ required to transmit raw neural data from 100 channels could be compressed more than 100,000-fold to 300 bs⁻¹ of 3-channel motor-output information updated with 10-bit precision at 10 Hz. Such dramatic compression brings concomitant reductions in the power required for communication and digitization of neural data. Ultra-low-power analog preprocessing prior to digitization of neural signals could thus be beneficial in some applications. Related considerations arise in the design of cochlear implants, and prior work in that field has demonstrated that power-efficient, analog preprocessing before digitization can be used to achieve high data compression ratios, leading to order-of-magnitude reductions in power consumption relative to fully digital data-processing schemes (Sarpeshkar, Baker et al., 2005; Sarpeshkar, Salhouse et al., 2005). Such power savings have been achieved while preserving programmability, as well as robustness to multiple sources of noise and transistor mismatch. Importantly, a processor based on this low-power

design paradigm functioned successfully in a deaf patient, enabling her to understand speech on her very first attempt (Sarpeshkar, 2006).

In this chapter we describe an approach to neural decoding using low-power analog preprocessing methods that can handle large quantities of high-bandwidth analog data, processing neural input signals in a slow-and-parallel fashion to generate low-bandwidth control outputs. Parallel architectures for data compression constitute an area in which analog systems perform especially well relative to digital systems, as has been demonstrated analytically (Sarpeshkar, 1998), and as exemplified by biological systems such as the retina and the cochlea (Mead, 1989; Sarpeshkar, 2006).

Multiple approaches to neural signal decoding have been implemented successfully by a number of research groups, as mentioned in the section entitled “Implementation of the Decoding Algorithm in Analog Circuitry.” All of these have employed highly programmable, discrete-time, digital algorithms, implemented in software or microprocessors located outside the brain. We are unaware of any work on continuous-time analog decoders or analog circuit architectures for neural decoding (Rapoport et al., 2009). The neural signal decoder we present here is designed to complement and integrate with existing approaches. Optimized for implementation in micropower analog circuitry, it sacrifices some algorithmic programmability to reduce the power consumption and physical size of the neural decoder, facilitating use as a component of a unit implanted within the brain.

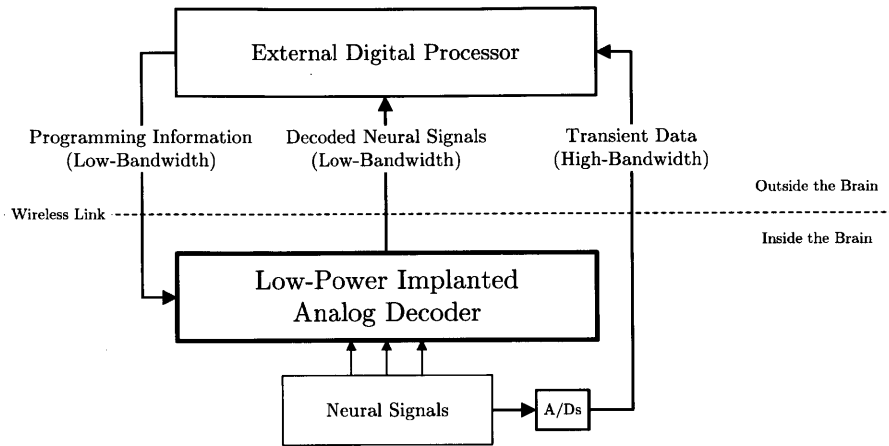
We demonstrate a method of designing highly power-efficient adaptive circuit architectures for neural decoding in two stages. We first generalize a discrete-time algorithm for adaptive-filter decoding to a continuous-time analog decoder. We then provide an approach to translating this decoder into a circuit architecture.

In this chapter we show that our algorithm and analog architecture offer good performance when learning to interpret real neural cell ensemble codes

in the rat thalamus, from which we decode continuous head direction trajectories, and the macaque monkey parietal cortex, from which we decode discrete reach decisions. We further show that the system is sufficiently flexible to accept local field potentials or mean spike rates as inputs. Circuit simulations of a subthreshold analog CMOS implementation of our circuit architecture suggest that a spike processor and decoder with 100 input channels and 3 output control parameters could be built with a total power budget of approximately 42 μW ; this would reduce the power expended in wireless telemetry to approximately 300 nW and the total power consumption of the implanted unit would be under 43 μW .

Our decoding scheme sacrifices the flexibility of a general-purpose digital system for the efficiency of a special-purpose analog system. This tradeoff may be undesirable in some neural prosthetic devices. Therefore, our proposed decoder is meant to be used not as a substitute for digital signal processors but rather as an adjunct to digital hardware, in ways that synergistically fuse the efficiency of embedded analog preprocessing options with the flexibility of a general-purpose external digital processor. The schematic for such a combined system is shown in Figure 1. First, an external digital processor transiently analyzes raw, high-bandwidth neural data to determine spike thresholds and DAC-programmable analog parameters for the implanted analog decoder. Such parameters could control which input channels are turned on or off, set filter time constants for converting spikes to mean firing rates, and establish amplitude thresholds that ensure optimal spike sorting. Such analysis requires a great deal of flexible programmability and so is best done digitally, by an external processor, taking into account all available raw neural data. The analysis generates configuration settings that can be downloaded over a low-bandwidth link to reconfigure the implanted unit, including a low-power analog decoder, digitally. The analog system then learns, decodes, and outputs low-bandwidth informa-

Figure 1. Schematic diagram of a brain–machine interface system incorporating an implantable analog neural decoder



tion in a low-power fashion for most of the time that the brain–machine interface is operational. Digital analysis and recalibration of the analog system may be done infrequently (perhaps once a day, as in early reports of a clinically useful brain–machine interface (Hochberg et al., 2006)) so that the average power consumption is always low even though transient power consumption of the external digital system may be relatively high during recalibration. The architecture of Figure 1 permits exploration of other decoding algorithms in the external unit, allowing high-power operation when necessary to achieve sufficient programmability.

For clinical neural prosthetic devices, the necessity of highly sophisticated decoding algorithms remains open to question, since both animal (Carmena et al., 2003; Musallam et al., 2004; Taylor et al., 2002; Velliste, Perel, Spalding, Whitford, & Schwartz, 2008) and human (Hochberg et al., 2006) users of even first-generation neural prosthetic systems have proven capable of rapidly adapting to the particular rules governing the control of their brain–machine interfaces. In the present work we focus on an architecture to implement a simple, continuous-time analog linear (convolutional) decoding algorithm. The

approach we present here can be generalized to implement analog-circuit architectures of more general Bayesian algorithms; examples of related systems include analog probabilistic decoding circuit architectures used in speech recognition and error correcting codes (Lazzaro, Wawrzyniek, & Lippmann, 1997; Loeliger, Tarköy, Lustenberger, & Helfenstein, 1999). Related architectures can be adapted through our mathematical approach to design circuits for Bayesian neural decoding.

This chapter is organized as follows: The section entitled “Methods of Designing and Testing a Micropower Neural Decoding System” begins by presenting a derivation of our adaptive-kernel decoder; a corresponding description of an example instantiation in subthreshold micropower analog VLSI is presented in the subsection “Implementation of the Decoding Algorithm in Analog Circuitry.” “Methods of Testing the Neural Signal Decoding System” describes methods for testing the algorithm and corresponding circuit architecture in several decoding tasks: continuous trajectory decoding from simulated local field potential (LFP) input signals, continuous head-direction trajectory decoding from spike-train inputs recorded from the thalamus of a freely moving rat, and discrete arm-reach intentions

decoded from spike-train inputs recorded from the parietal cortex of a macaque monkey. The results of these tests are subsequently presented in the section entitled “Evaluating the Performance of the Neural Decoding System.” The following section, “Evaluating Design Tradeoffs in the Construction of an Implantable Neural Decoder,” discusses the simulated performance characteristics of our decoding system and alternative digital implementations from the perspectives of reducing power consumption for implantable brain–machine interfaces. The chapter concludes with a summary of our work and anticipated directions for future research in the field of neural decoding in the context of brain–machine interfaces.

METHODS OF DESIGNING AND TESTING A MICROPOWER NEURAL DECODING SYSTEM

An Algorithm for Adaptive Convolutional Decoding of Neural Cell Ensemble Signals

The function of neuronal population decoding is to map neural signals onto higher-level cognitive processes to which they correspond. This section presents the mathematical foundations of an adaptive convolutional decoder whose kernels can be modified according to a gradient-descent–based learning algorithm that optimizes decoding performance in an on-line, real-time fashion.

In convolutional decoding of neural cell ensemble signals, the decoding operation takes the form

$$\begin{aligned} \vec{M}(t) &= \mathbf{W}(t) \circ \vec{N}(t) \\ M_i(t) &= \sum_{j=1}^n W_{ij}(t) \circ N_j(t) \end{aligned} \quad (1.1)$$

where $i \in \{1, \dots, m\}$; $\vec{N}(t)$ is an n -dimensional vector containing the neural signal (n input channels of neuronal firing rates, analog signal values, or local field potentials, for example) at time t ; $\vec{M}(t)$ is a corresponding m -dimensional vector containing the decoder output signal (which in the examples presented here corresponds to motor control parameters, but could correspond as well to limb or joint kinematic parameters or to characteristics or states of nonmotor cognitive processes); \mathbf{W} is a matrix of convolution kernels $W_{ij}(t)$ (formally analogous to a matrix of dynamic synaptic weights), each of which depends on a set of p modifiable parameters, $W_{ij}^k, k \in \{1, \dots, p\}$; and \circ indicates convolution. Accurate decoding requires first choosing an appropriate functional form for the kernels and then optimizing the kernel parameters to achieve maximal decoding accuracy. Since the optimization process is generalizable to any choice of kernels that are differentiable functions of the tuning parameters, we discuss the general process first. We then explain our biophysical motivations for selecting particular functional forms for the decoding kernels; appropriately chosen kernels enable the neural decoder to emulate the real-time encoding and decoding processes performed by biological neurons.

Our algorithm for optimizing the decoding kernels uses a gradient-descent approach to minimize decoding error in a least-squares sense during a learning phase of decoder operation. During this phase the correct output $\hat{\vec{M}}(t)$, and hence the decoder error $\vec{e}(t) = \vec{M}(t) - \hat{\vec{M}}(t)$, is available to the decoder for feedback-based learning. We design the optimization algorithm to evolve $\mathbf{W}(t)$ in a manner that reduces the squared decoder error on a timescale set by the parameter τ , where the squared error is defined as

$$\begin{aligned}
 E(\mathbf{W}(t), \tau) &= \int_{t-\tau}^t |\vec{e}(u)|^2 du \\
 &= \sum_{i=1}^m \int_{t-\tau}^t |e_i(u)|^2 du \quad (1.2) \\
 &\equiv \sum_{i=1}^m E_i,
 \end{aligned}$$

and the independence of each of the m terms in Equation (1.2) is due to the independence of the m sets of $n \times p$ parameters W_{ij}^k , $j \in \{1, \dots, n\}$, $k \in \{1, \dots, p\}$ associated with generating each component $M_i(t)$ of the output. Our strategy for optimizing the matrix of decoder kernels is to modify each of the kernel parameters W_{ij}^k continuously and in parallel, on a timescale set by τ , in proportion to the negative gradient of $E(\mathbf{W}(t), \tau)$ with respect to each parameter:

$$\begin{aligned}
 -\vec{\nabla}_{ij}^k E(\mathbf{W}(t), \tau) &\equiv -\frac{\partial E}{\partial W_{ij}^k} \\
 &= -\int_{t-\tau}^t du \left[2 \left(M_i(u) - \sum_{j=1}^n W_{ij}^k(u) \circ N_j(u) \right) \times \left(-\frac{\partial W_{ij}^k(u)}{\partial W_{ij}^k} \circ N_j(u) \right) \right] \\
 &= 2 \int_{t-\tau}^t e_i(u) \left(\frac{\partial W_{ij}^k(u)}{\partial W_{ij}^k} \circ N_j(u) \right) du. \quad (1.3)
 \end{aligned}$$

The learning algorithm refines \mathbf{W} in a continuous-time fashion, using $-\vec{\nabla} E(t)$ as an error feedback signal to modify $\mathbf{W}(t)$, and incrementing each of the parameters $W_{ij}^k(t)$ in continuous time by a term proportional to $-\vec{\nabla}_{ij}^k E(\mathbf{W}(t))$ (the proportionality constant, ε , must be large enough to ensure quick learning but small enough to ensure learning stability). If $\mathbf{W}(t)$ is viewed as an array of linear filters operating on the neural input signal, the quantity $-\vec{\nabla}_{ij}^k E(\mathbf{W}(t), \tau)$ used to increment each filter parameter can be described as the product, averaged over a time interval of length τ , of the error in the filter output and a secondarily filtered version of the filter input. The error term is identical for the parameters of all

filters contributing to a given component of the output, $M_i(t)$. The secondarily filtered version of the input is generated by a secondary convolution

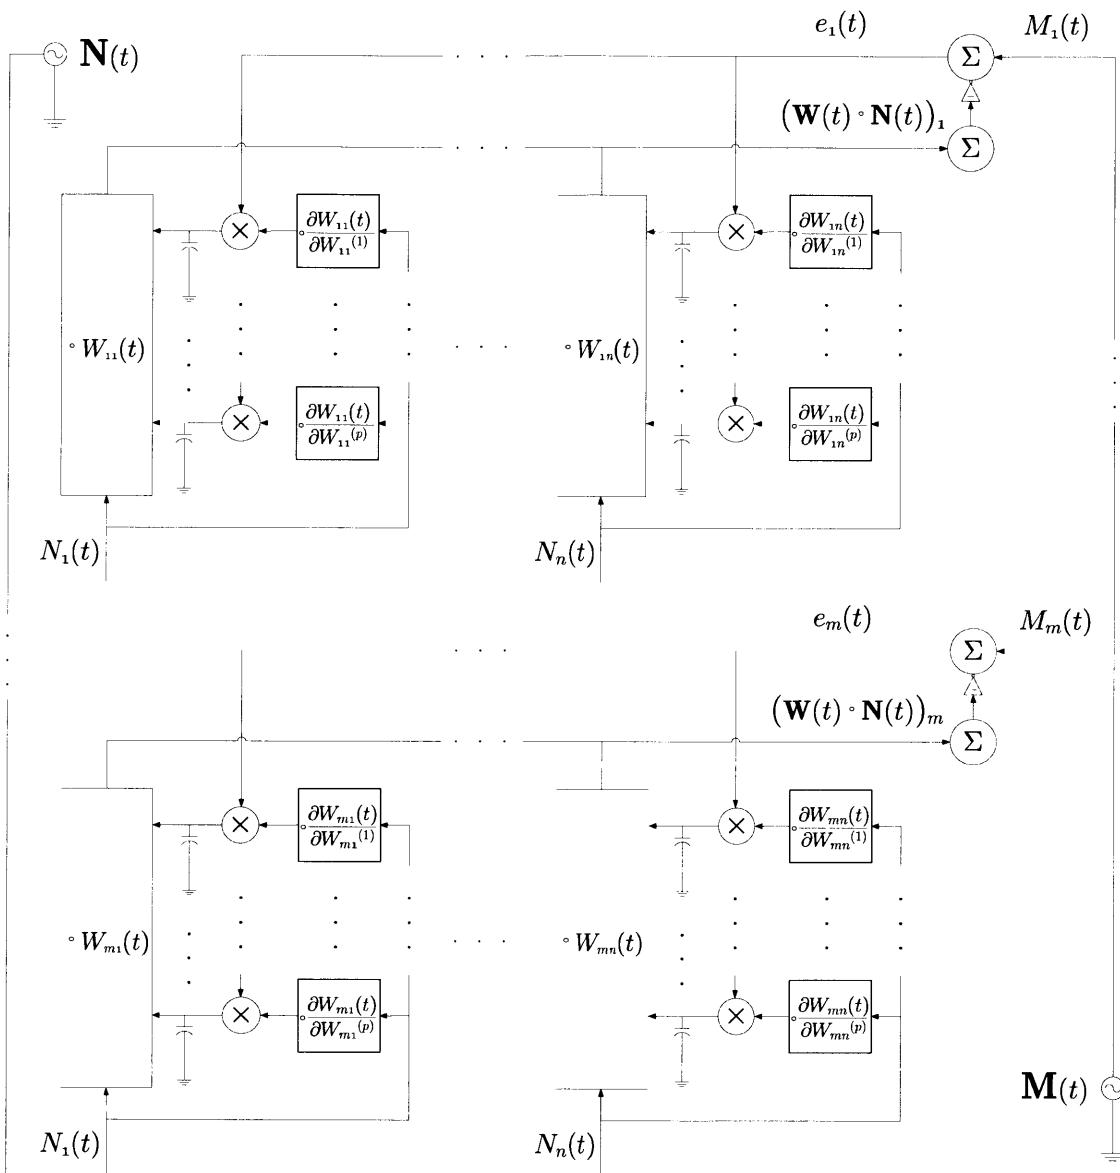
kernel, $-\frac{\partial W_{ij}^k(u)}{\partial W_{ij}^k}$, which depends on the func-

tional form of each primary filter kernel and in general differs for each filter parameter. Figure 2 shows a block diagram for an analog circuit architecture that implements our decoding and optimization algorithm.

Many functional forms for the convolution kernels are both theoretically possible and practical to implement using low-power analog circuitry. Our approach has been to emulate biological neural systems by choosing a biophysically inspired kernel whose impulse response approximates the postsynaptic currents biological neurons integrate when encoding and decoding neural signals *in vivo* (Arenz, Silver, Schaefer, & Margrie, 2008). Combining our decoding architecture with the choice of a first-order low-pass decoder kernel enables our low-power neural decoder to implement a biomimetic, continuous-time artificial neural network. Numerical experiments have also indicated that decoding using such biomimetic kernels can yield results comparable to those obtained using optimal linear decoders (Eliasmith & Anderson, 2003). But in contrast with our on-line optimization scheme, optimal linear decoders are computed off-line after all training data have been collected. We have found that this simple choice of kernel offers effective performance in practice, and so we confine the present analysis to that kernel.

We also offer a heuristic justification for using first-order low-pass kernels. Kernels for decoding continuous trajectories should be designed to anticipate low-frequency variation in the input and output signals. Two-parameter first-order low-pass filter kernels account for trajectory continuity by exponentially weighting the history of neural inputs:

Figure 2. Block diagram of an analog circuit architecture for linear convolutional decoding and learning



$$W_{ij} = \frac{A_{ij}}{\tau_{ij}} e^{-\frac{t}{\tau_{ij}}}, \quad (1.4)$$

where the two tunable kernel parameters are $W_{ij}^{k=1} = A_{ij}$, the low-pass filter gain, and $W_{ij}^{k=2} = \tau_{ij}$, the decay time over which past inputs

$\vec{N}(t')$, $t' < t$, influence the present output estimate $\vec{M}(t) = \mathbf{W} \circ \vec{N}(t)$. The filters used to tune the low-pass filter kernel parameters can be implemented using simple and compact analog circuitry. The gain parameters are tuned using low-pass filter kernels of the form

$$\frac{\partial W_{ij}(t)}{\partial W_{ij}^{k=1}} = \frac{1}{\tau_{ij}} e^{-\frac{t}{\tau_{ij}}}, \quad (1.5)$$

while the time-constant parameters are tuned using band-pass filter kernels:

$$\frac{\partial W_{ij}(t)}{\partial W_{ij}^{k=2}} = \frac{A_{ij}}{\tau_{ij}^2} e^{-\frac{t}{\tau_{ij}}} \left(\frac{t}{\tau_{ij}} - 1 \right). \quad (1.6)$$

When decoding discontinuous trajectories, such as sequences of discrete decisions, we use the limiting case of this kernel as $\tau_{ij} \rightarrow 0$:

$$W_{ij}(t) = W_{ij}^{k=1} \delta(t) = A_{ij} \delta(t). \quad (1.7)$$

Such a decoding system, in which each kernel is a zeroth-order filter characterized by a single tunable constant, performs instantaneous linear decoding, which has successfully been used by others to decode neuronal population signals in the context of neural prosthetics (Hochberg et al., 2006; Wessberg & Nicolelis, 2004). With kernels of this form, $\mathbf{W}(t)$ is analogous to matrices of synaptic weights encountered in artificial neural networks, and our optimization algorithm resembles a ‘delta-rule’ learning procedure (Haykin, 1999).

Implementation of the Decoding Algorithm in Analog Circuitry


In this section we describe the implementation of our adaptive decoder in an analog circuit architecture. This approach represents a divergence from conventional approaches to neural signal decoding. The ease of implementing sophisticated algorithms in digital systems has led to a proliferation of effective approaches to decoding and learning. Some of the most popular such algorithms in the context of neural decoding involve construction of optimal linear filters (Serruya, Hatsopoulos, Fellows, Paninski, & Donoghue, 2003; Warland,

Reinagel, & Meister, 1997); training of artificial neural networks (Sanchez, Erdogmus, Principe, Wessberg, & Nicolelis, 2005); use of Kalman filters (Wu, Black et al., 2004; Wu, Gao, Bienenstock, Donoghue, & Black, 2006; Wu, Shaikhouni, Donoghue, & Black, 2004) and adaptive Kalman filters (Wu & Hatsopoulos, 2008); estimation based on Bayesian inference techniques (Brockwell, Rojas, & Kass, 2004; Srinivasan, Eden, Mitter, & Brown, 2007) and point-process models (Eden, Frank, Barbieri, Solo, & Brown, 2004); and decoding on the basis of frequency-domain data (such as the spectral content of local field potentials) (Musallam et al., 2004; Pesaran, Pezaris, Sahani, Mitra, & Andersen, 2002; Shenoy et al., 2003) or wavelet decompositions of neural signals (Musallam et al., 2004). However, to our knowledge the present work represents the first description of an analog circuit architecture for neural signal decoding.

We separate the present discussion into three parts. The first part treats the preprocessing of the two principal classes of neural input signals, local field potentials (LFPs) and action potentials (‘spikes’). We then address the decoding architecture itself. Finally, we describe our estimates of power consumption by each module of the decoder architecture. Figure 3 is a diagram of the circuit modules required to implement a single adaptive kernel of the decoder. These modules fall into six functional classes: (1) Adaptive filters corresponding to the kernels $W_{ij}(t)$; (2) Parameter-learning filters to tune the $\{W_{ij}^{(k)}\} = \{A_{ij}, \tau_{ij}\}$; (3) Biasing circuits for the parameter-learning filters; (4) Multipliers; (5) Adders and subtractors; and (6) Memory units for storing learned parameter values. The second part of this section presents the design for each functional subunit of our decoding architecture in turn.

Note that the descriptions given here correspond to an unoptimized, proof-of-concept circuit implementation. Current-mode techniques, circuit optimizations not described here, and noise-robust

Figure 3. Block diagram indicating the functional component circuits required to implement a single adaptive kernel of the convolutional decoder



analog biasing techniques such as those described in (Sarpeshkar, Salthouse et al., 2005) will be necessary to ensure robust, programmable, and efficient operation in a practical implementation.

Input Signals for the Neural Decoder

Local-Field-Potential-based decoding

Local field potentials encode information about well defined cognitive states and can also be used as control signals for neural prostheses. Furthermore, it has been found that particular classes of information are encoded in distinct frequency bands of the LFP power spectrum (Pesaran et al., 2002). Information encoded in this manner can be extracted from the LFP by passing the raw LFP signal through a band-pass amplifier tuned to the spectral band of interest, rectifying the output of the band-pass filter, and passing the result through a peak-detection circuit to generate an envelope waveform. Prior work solved this design problem in the context of ultra-low-power bionic-ear (cochlear implant) processors (Sarpeshkar, Baker et al., 2005; Sarpeshkar, Salthouse et al., 2005) for which micropower band-pass amplifier (Salthouse

& Sarpeshkar, 2003) and envelope detector circuits (Zhak, Baker, & Sarpeshkar, 2003) were built that can be employed to process neural signals at lower frequencies.

Spike-Based Decoding

Neuronal action potential voltage spikes typically have widths on the order of 1 ms, corresponding to frequencies in the kilohertz range. Spike-based inputs are transformed to lower-frequency signals (of order 1–10 Hz) through time-domain averaging and the resulting spike rates are used as input signals for the decoder. Such averaging can be implemented by low-pass interpolation filters. The simplest such filter is a first-order low-pass filter with frequency-domain transfer function $H_1(s) = \frac{1}{1 + \tau_1 s}$, and cutoff frequency $f_c = (2\pi\tau_1)^{-1}$ Hz. Smoother interpolation can be obtained by cascading first-order filters. The analog implementation of such filters can be achieved using a G_m -C design: G_m refers to the transconductance of an operational transconductance amplifier (OTA) component,

while C denotes a filter capacitance. In such a filter $f_c = \frac{1}{2\pi} \frac{G_m}{C}$, so a low cutoff frequency requires C to be large or G_m to be small. Circuit layout area restrictions will constrain the maximum value of C to approximately 4 pF, so a low f_c requires G_m to be small. Wide-linear-range transconductors with subthreshold bias currents (Sarpeshkar, Lyon, & Mead, 1997) allow G_m to be small enough to achieve such low corner frequencies in a reliable fashion. We convert raw neural input waveforms into smooth time-averaged spike rates (Figure 5) using the circuitry shown in Figure 4, which uses dual thresholding to detect action potentials on each input channel and then smooths the resulting spike trains to generate mean firing rate input signals. Circuit simulations indicate that the charge pump and smoothing filter modules consume approximately 1.1 nW and 40 pW of power from a 1 V supply in 0.18 μm CMOS technology, respectively. As a result, power consumption of the preprocessing stages will be dominated by the requirements of the comparator; a comparator of the kind described in (Yang & Sarpeshkar, 2006), operating at 30 kHz, requires approximately 240 nW.

Functional Circuit Subunits of the Decoder Architecture

Primary Adaptive Filters

Each of the $m \times n$ tunable kernels used to implement our decoding architecture can be understood as an adaptive filter whose frequency-domain transfer function is obtained from the Laplace transform of the time-domain kernel in Equation (1.4):

$$\overline{W}_{ij}(s) = \frac{A_{ij}}{1 + \tau_{ij}s} \quad (1.8)$$

where the gain A_{ij} can be positive or negative. This transfer function can be obtained from a filter having the topology shown in Figure 6, which contains four standard, nine-transistor wide-range operational transconductance amplifiers (OTAs) of the form described in (Mead, 1989).

Every OTA is operated subthreshold such that its transconductance is linear in its bias current. The gain of the filter is determined by the three OTAs, $G_{m+}^{(A_j)}$, $G_{m-}^{(A_j)}$, and $G_m^{(R_j)}$, which have transconductances $\frac{\kappa I_+^{(A_j)}}{V_T}$, $\frac{\kappa I_-^{(A_j)}}{V_T}$, and $\frac{\kappa I^{(R_j)}}{V_T}$, respectively, where κ denotes the gate-coupling coefficient of the MOS transistor, and in this analysis $\kappa = 0.7$ is assumed for all transistors; and $V_T = kT/q$, where k denotes the Boltzmann constant, q denotes the electron charge, and T denotes the Kelvin temperature. By applying Kirchoff's Current Law (KCL) at node V_x , we obtain the voltage gain from V_{in} to V_x as

$$\begin{aligned} \frac{V_x}{V_{in}} &= \frac{G_{m+}^{(A_j)} - G_{m-}^{(A_j)}}{G_m^{(R_j)}} \\ &= \frac{I_+^{(A_j)} - I_-^{(A_j)}}{I^{(R_j)}}. \end{aligned} \quad (1.9)$$

If we set the sum of $I_+^{(A_j)}$ and $I_-^{(A_j)}$ equal to a constant current $I_{tot}^{(A_j)}$, the expression in Equation (1.9) can be reduced to

$$\frac{V_x}{V_{in}} = \frac{2I_+^{(A_j)} - I_{tot}^{(A_j)}}{I^{(R_j)}}. \quad (1.10)$$

The transfer function from V_x to V_{out} can be expressed as

Figure 4. Analog preprocessing circuitry for converting raw neural signals into smooth analog inputs for an analog decoder

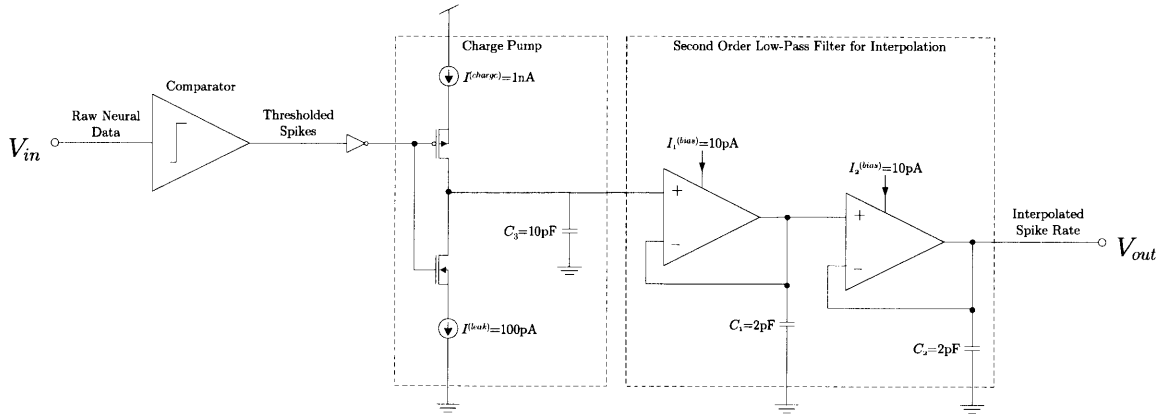


Figure 5. Analog preprocessing of raw neural signals to generate smooth analog inputs to an implantable analog decoder. The raw neural signal (top trace) is thresholded to generate a spike train (middle trace), which is converted to a smooth time-averaged spike rate (bottom trace) using a second-order low-pass filter

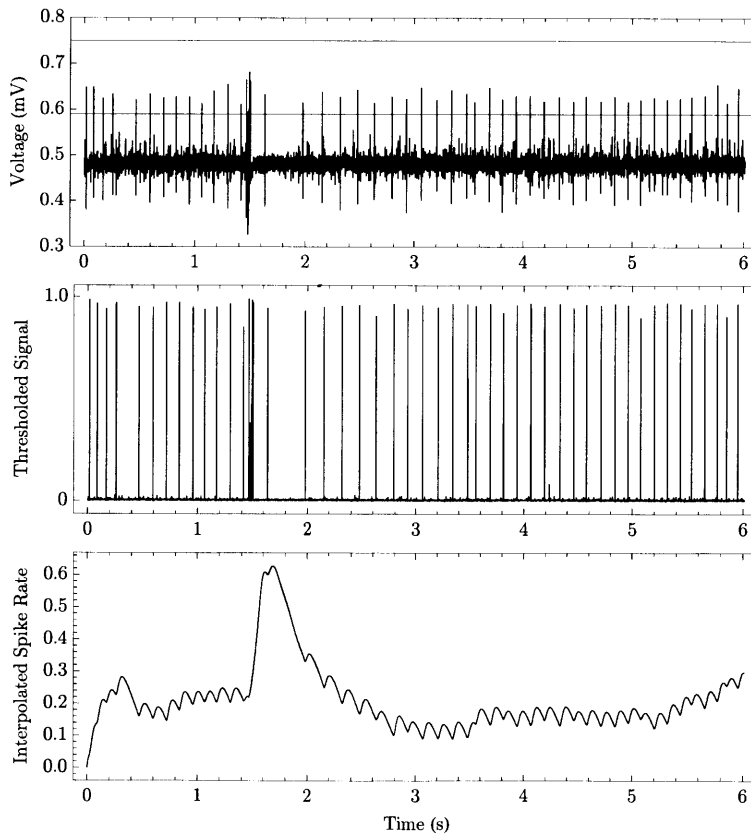
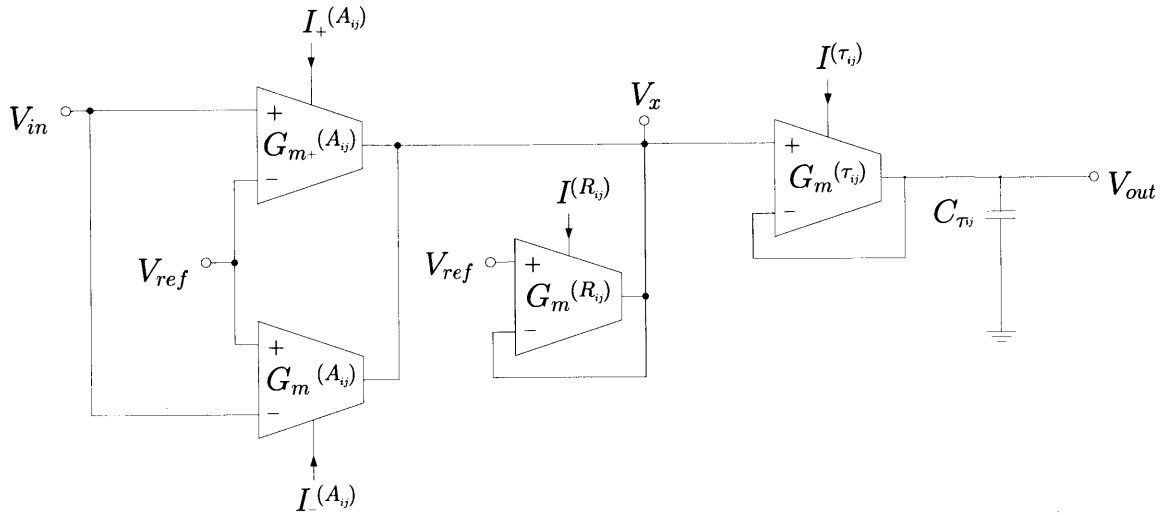


Figure 6. Adaptive filter with tunable parameters for learning the optimal convolution kernels for neural signal decoding



$$\frac{V_{out}}{V_x}(s) = \frac{1}{1 + s \frac{C_{\tau_{ij}}}{G_m^{\tau_{ij}}}}. \quad (1.11)$$

As a result, we can express the overall transfer function from V_{in} to V_{out} as

$$\frac{V_{out}}{V_{in}}(s) = \frac{2I_+^{(A_{ij})} - I_{tot}^{(A_{ij})}}{I^{(R_{ij})}} \frac{1}{1 + s \frac{C_{\tau_{ij}}}{G_m^{\tau_{ij}}}}. \quad (1.12)$$

Comparing Equation (1.12) to Equation (1.8), we obtain expressions for the gain and time constant parameters:

$$A_{ij} = \frac{2I_+^{(A_{ij})} - I_{tot}^{(A_{ij})}}{I^{(R_{ij})}} \quad (1.13)$$

$$\tau_{ij} = \frac{C_{\tau_{ij}}}{G_m^{\tau_{ij}}}.$$

Since $I_+^{(A_{ij})}$ can be adjusted from 0 to $I_{tot}^{(A_{ij})}$, A_{ij} can vary from $-I_{tot}^{(A_{ij})} / I^{(R_{ij})}$ to $I_{tot}^{(A_{ij})} / I^{(R_{ij})}$. The time constant τ_{ij} can be adjusted by tuning the bias current in $G_m^{\tau_{ij}}$.

Secondary Parameter Tuning Filters

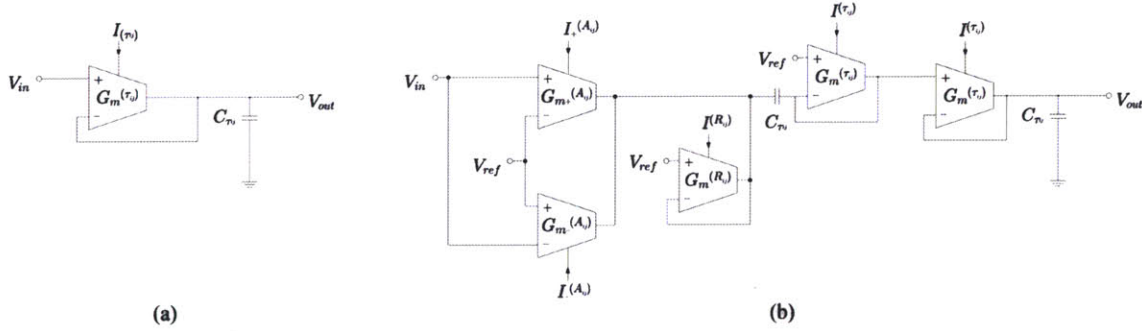
Real-time gradient-descent-based optimization of the convolution kernels can be achieved through tuning the filter parameters $\{W_{ij}^{(k)}\} = \{A_{ij}, \tau_{ij}\}$ using signals proportional to $-\vec{\nabla}_{ij}^{(k)} E$. Construction of such signals requires convolution kernels proportional to $\frac{\partial W_{ij}(t)}{\partial W_{ij}^{(k)}}$. These convolution ker-

nels can be implemented by ‘parameter-learning filters’ of the kind shown in Figure 7. Figure 7(a) shows a first-order low-pass G_m - C filter with transfer function $W_{ij}^{A_{ij}}(s) = \frac{1}{1 + \tau_{ij}s}$ to be used

as a ‘gain-learning filter.’ Figure 7(b) shows a second-order band-pass G_m - C filter with transfer

Figure 7. Parameter-learning filters for tuning adaptive filter parameters based on error-signal feedback.

(a) A first-order low-pass G_m -C filter with transfer function $W_{ij}^{A_{ij}}(s) = \frac{1}{1 + \tau_{ij}s}$ to be used as a ‘gain-learning filter.’ (b) A second-order band-pass G_m -C filter with transfer function $W_{ij}^{\tau_{ij}}(s) = \frac{\tau_{ij}s}{(1 + \tau_{ij}s)^2}$ to be used as a ‘time-constant-learning filter’



function $W_{ij}^{\tau_{ij}}(s) = \frac{\tau_{ij}s}{(1 + \tau_{ij}s)^2}$ to be used as a

‘time-constant-learning filter.’ The time constant

$$\tau_{ij} = \frac{C_{\tau_{ij}}}{G_m(\tau_{ij})}$$

for the two parameter-learning filters

is identical to that of the adaptive filter, as described in the section entitled “Primary adaptive filters.” Correspondingly, the bias currents and therefore

the transconductances $G_m(\tau_{ij})$ in all three types of filter are identical, so the time constants of all filters in the learning architecture are updated simultaneously. Note that the actual transfer function of the time-constant-learning filter need only

be proportional to $W_{ij}^{\tau_{ij}}(s)$, so the factor of τ_{ij}/A_{ij} between the transfer function of Equation (1.8) and the filter shown in Figure 7(b) is acceptable. Implementation of the negation required to adapt the time constants is addressed in the section entitled “Multipliers.”

Multipliers

The multipliers that perform the operations

$$e_i(t) \times \frac{\partial W_{ij}(t)}{\partial W_{ij}^{(k)}}$$

required by the gradient-descent algorithm, denoted by the symbol \times in Figure 2 and Figure 3, can be implemented using wide-range four-quadrant Gilbert multipliers of the kind described in (Mead, 1989), which accept four voltage inputs $V_i, i \in \{1, 2, 3, 4\}$ and a bias current I_b , and generate an output current; low-power performance can be obtained by operating the multiplier circuit with all transistors in the sub-threshold regime. The input-output characteristic of the Gilbert multiplier is given by

$$I_{out} = I_b \tanh \frac{\kappa(V_1 - V_2)}{2V_T} \tanh \frac{\kappa(V_3 - V_4)}{2V_T} \approx I_b \left(\frac{\kappa}{2V_T} \right)^2 (V_1 - V_2)(V_3 - V_4), \quad (1.14)$$

where the approximation of Equation (1.14) is valid in the intended operating region, where V_1

$\approx V_2$ and $V_3 \approx V_4$. Noninverting multiplication, as required for adapting the A_{ij} , can be implemented by feeding e_i into V_1 , $\frac{\partial W_{ij}(t)}{\partial W_{ij}^{(A_{ij})}}$ into V_3 , and setting V_2 and V_4 to a constant reference voltage. On the other hand, inverting multiplication, as required for adapting the τ_{ij} , can be implemented by interchanging the roles of V_3 and V_4 , feeding $\frac{\partial W_{ij}(t)}{\partial W_{ij}^{(\tau_{ij})}}$ into V_4 , and setting V_3 to a constant reference voltage.

The output of each multiplier is a current, I_{out} , so the integrations required in Equation (1.3) for updating the parameters $\{W_{ij}^{(k)}\} = \{A_{ij}, \tau_{ij}\}$ are conveniently implemented by linear capacitors, as indicated in Figure 3. The voltages $V_C^{(W_{ij}^{(k)})}$ on the capacitors $C^{(W_{ij}^{(k)})}$ and used in adapting the parameter values $W_{ij}^{(k)}$ are therefore given by

$$V_C^{(W_{ij}^{(k)})} = \frac{I_b}{C^{(W_{ij}^{(k)})}} \frac{\kappa^2}{4V_T^2} \int_{u=0}^t e_i(u) \left(\frac{\partial W_{ij}(u)}{\partial W_{ij}^{(k)}} \circ N_j(u) \right) du \quad (1.15)$$

which has the form required by Equation (1.3). The filter parameters therefore vary continuously in time, and the time variation of the control voltages can be obtained by differentiating Equation (1.15).

Biasing Circuits

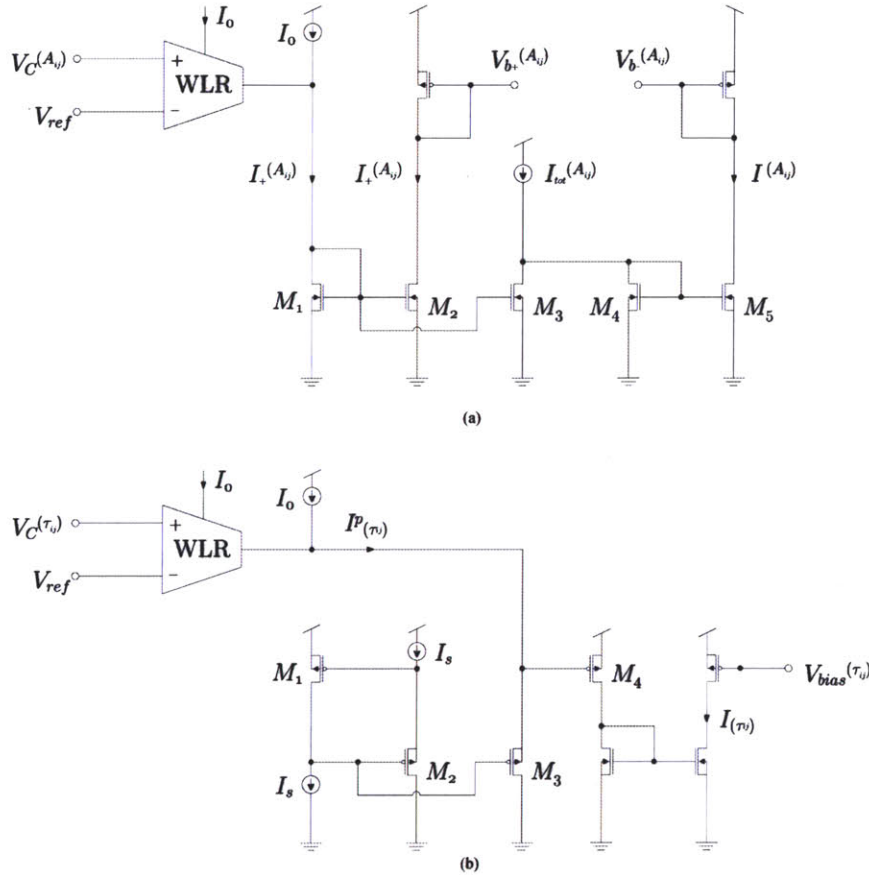
As discussed in the section entitled ‘‘Multipliers,’’ the filter parameters A_{ij} and τ_{ij} defining the transfer function of adaptive filter W_{ij} are stored on the capacitors $C^{(A_{ij})}$ and $C^{(\tau_{ij})}$, respectively. Furthermore, as indicated in the section entitled ‘‘Primary adaptive filters,’’ the values of the filter parameters can be tuned by adjusting the bias currents that determine $G_m^{(A_{ij})}$ and $G_m^{(\tau_{ij})}$. Since the

gain A_{ij} depends on $\frac{I_{(A_{ij})}}{I_{(R_{ij})}} \propto I_{(A_{ij})}$, while the time constant τ_{ij} depends on $\frac{C_{\tau_{ij}}}{G_m^{(\tau_{ij})}} \propto \frac{1}{I_{(\tau_{ij})}}$, real-time adaptive parameter tuning requires a scheme for modifying $I_{(A_{ij})}$ in proportion to $V_C^{(A_{ij})}$ and $I_{(\tau_{ij})}$ in inverse proportion to $V_C^{(\tau_{ij})}$.

Tuning $I_{(A_{ij})}$ in proportion to variations in the capacitor voltage $V_C^{(A_{ij})}$ can be accomplished by converting $V_C^{(A_{ij})}$ into a current proportional to $V_C^{(A_{ij})}$ and then using a current mirror to generate a copy of that current that is in turn used to set the transconductance $G_m^{(A_{ij})}$ of the adaptive filter. The conversion of $V_C^{(A_{ij})}$ into a current proportional to $V_C^{(A_{ij})}$ can be performed by a wide-linear-range transconductance amplifier (WLR) of the form described in (Sarpeshkar et al., 1997). Figure 8(a) shows a schematic of the gain-biasing circuit used to generate $I_{(A_{ij})} \propto V_C^{(A_{ij})}$. This biasing circuit is intended to make $I_+^{(A_{ij})} + I_-^{(A_{ij})} = I_{tot}^{(A_{ij})}$. Assuming that all of the NMOS transistors are well matched, the current in M_3 must be equal to that in M_2 , which is $I_+^{(A_{ij})}$. Using KCL at the drain of M_3 , the current in M_4 is thus $I_{tot}^{(A_{ij})} - I_+^{(A_{ij})}$, which is equal to the current in M_5 , so $I_-^{(A_{ij})} = I_{tot}^{(A_{ij})} - I_+^{(A_{ij})}$ as required.

Tuning $I_{(\tau_{ij})}$ in inverse proportion to variations in the capacitor voltage $V_C^{(\tau_{ij})}$ can be accomplished using the circuit shown in Figure 8(b), which operates as follows. First, $V_C^{(\tau_{ij})}$ is converted into a proportional current $I_{(\tau_{ij})}^P$ as in the gain-biasing circuit. A translinear circuit, formed by the four

Figure 8. Circuits for setting the bias currents and transconductances that determine the adaptive filter parameters: (a) Bias-current-setting circuit for filter gains, (b) Bias-current-setting circuit for filter time constants



well matched MOS transistors M_1 – M_4 , is then used to invert $I_{(\tau_y)}^p$, producing $I_{(\tau_y)} = \frac{I_s^2}{I_{(\tau_y)}^p}$, where I_s is a current reference that scales the inversion. A mirror copy of $I_{(\tau_y)}$ is then used as the bias current that sets transconductance $G_m^{(\tau_y)}$.

Adders and Subtractors

Each adder, denoted by the symbol Σ in Figure 2 and Figure 3, sums the n outputs of each set $\{W_{ij}\}$, $j \in \{1, \dots, n\}$ of adaptive filters contributing to

$(\mathbf{W}(t) \circ \mathbf{N}(t))_i = \hat{M}_i(t)$. The adders can be implemented using a follower-aggregation circuit of the kind described in (Mead, 1989). The corresponding error signal, $e_i(t)$, is generated by performing the subtraction $M_i(t) - \hat{M}_i(t)$. This operation can be implemented by another adder, with a unity-gain inverting amplifier negating the adder input from $\hat{M}_i(t)$.

Parameter Memory

A completely analog implementation of our decoder could include analog memory units for

storing parameter values when learning ends. Such units could consist of analog memory elements operating in a switched sample-and-hold scheme to permit memoryless adaptation during the learning phase and parameter storage as soon as learning terminates. Analog memory circuits with 8-bit hold times of 3.9 hours and 12-bit hold times of 14.5 minutes have been developed (O’Halloran & Sarpeshkar, 2004, 2006) and could be used in this context. Output from the memory and biasing circuits could be multiplexed onto the adaptive filter nodes whose voltages correspond to the adaptive filter parameters, using a CMOS transmission gate. This scheme is indicated in Figure 3. However, even using digital memory elements does not increase total power consumption significantly, since the termination of a learning phase is a rare event and therefore writing to memory, with its associated power cost, occurs only infrequently.

Power Consumption of the Decoder

Table 1 shows an estimate of the power consumed by each circuit block needed to implement our neural decoder, based on a SPICE simulation of the decoder, the performance of which is shown in Figure 11. The supply voltage of the entire system is 1 V. The bias current of each module is chosen so that the circuit has enough bandwidth to process input signals band-limited to 1 kHz (note that this choice of input signal bandwidth is more than sufficient because the bandwidth of mean-firing-rate changes is typically much smaller than the reciprocal of a refractory period. As indicated in Table 1, the total power consumption of one decoding module is approximately 54 nW.

In order to ensure robust, stable decoding, the learned (optimized) filter parameters for the neural decoder will be programmed using digital-to-analog converters (DACs) as indicated in the “Introduction” section. These DACs do not appreciably increase system power consumption. One DAC is required per filter parameter to set the bias current for each filter constant, so a sys-

Table 1. Power consumption in decoder circuit modules

Portion of Channel	Power Consumption (nW)
Tunable Decoding Filter	16
Gain-Learning Filter	2
Time-Constant-Learning Filter	2
Multipliers (2)	2
Gain Biasing Circuit	10
Time Constant Biasing Circuit	10
Analog Memories (2)	5
Total (One Decoder Module)	54 nW

tem decoding 3 motor parameters from 100 neural input channels would require 600 current DACs. The static current in a current DAC is approximately equal to that of a transistor whose gate is connected to the associated capacitors in Figure 3, which is approximately 100 pA. The total power consumed by all 600 DACs, when operated from a 1.8-V supply, is therefore only approximately 110 nW.

Methods of Testing the Neural Signal Decoding System

In this section we describe three approaches to testing the ability of our system to decode neural signals. The first experiment is based on a simulation of local field potential input signals, which are used to decode continuous-time motor trajectories. The second experiment involves continuous-time decoding of head direction from spike train data recorded from a small set of neurons in the thalamus of an awake, behaving rat. In the third experiment we decode discrete arm-reach intentions from spike train data recorded using a multielectrode array in the posterior parietal cortex of a macaque monkey.

Testing the Ability to Decode Trajectory Information from Multichannel Local Field Potential Neural Input Signals

A recent set of experiments has shown that parietal cortex neurons tend to exhibit electrical activity predictive of arm or eye movement in a single preferred direction. Increases in γ -band (25–90 Hz) spectral activity of such tuned parietal neurons anticipate movements in the preferred directions of those neurons, so a potentially useful signal for decoding intended movement from neural activity is therefore an envelope curve describing the modulated amplitude of the power transmitted in the gamma band (Andersen, Musallam, & Pesaran, 2004; Musallam et al., 2004; Pesaran et al., 2002). As a preliminary demonstration of the ability of our convolutional decoder to interpret LFP-type neural input signals, we therefore generated simulated γ -band power envelopes in order to model the local field potentials recorded by a set of n neural recording electrodes. We modeled γ -band power envelopes using a set of sinusoids with randomized amplitudes and phases and a constant offset term, and stored the corresponding waveforms in the vector $\vec{N}(t)$, which was used as the input to the decoder. We then randomly generated an $m \times n$ matrix, \mathbf{W}^* , and used it to construct a vector of m motor control parameters, $\vec{M}(t)$. The decoder was permitted to observe $\vec{M}(t)$ during a learning period of variable length, over which it sought to optimize its $m \times n$ -dimensional convolution kernel $\mathbf{W}(t)$. The parameter ε was set to 0.1 during these simulations, and both the $N_i(t)$ and the $M_i(t)$ were transformed using a hyperbolic-tangent normalization to constrain them to the interval $[-1, 1]$. We evaluated decoder performance using a scale-invariant and dimensionless conventional figure of merit, the normalized mean squared error, η , of the estimated trajectory, defined as

$$\eta \equiv \frac{1}{T} \int_{t_1=T}^{t_2=2T} dt \sum_{i=1}^m \left(\frac{M_i(t) - \hat{M}_i(t)}{L_i} \right)^2, \quad (1.16)$$

where L_i denotes the maximum extent of excursions permitted to $M_i(t)$ and the time T denotes the length of the training interval; $\eta^{(1)}$ is used to denote the average value of η for a single output dimension. It is possible to consider other figures of merit, including ones based on correlation rather than absolute error between $\vec{M}(t)$ and $\hat{\vec{M}}(t)$, but other authors have agreed that η -like figures of merit tend to reflect decoding system performance most reasonably (Wu et al., 2006). The results of these simulations are provided in the section entitled “Model Motor Trajectory Decoding from Simulated Multichannel Local Field Potential Inputs.”

Testing the Decoding of Continuous Trajectories Using Neuronal Spike Recordings from the Thalamus of an Awake Behaving Rat

Head direction cells of the rat thalamus are neurons known to exhibit receptive fields tuned to specific orientations of the head relative to the environment (Taube, 1995). We explored the ability of our system to decode the temporal firing patterns of such cells in real time. The spike trains used as input signals to the decoder in this set of experiments were derived from tetrode recordings made in the thalamus of a laboratory rat that rat ran back and forth between two points on a circular maze for food reward over a 30-minute period. The position and head direction of the animal were continuously tracked using a pair of head-mounted light-emitting diode (LED) arrays imaged at a sampling frequency of 30 Hz using a 300×300 -pixel charge-coupled device (CCD) array. Data from the imager were time-stamped in order to ensure synchronization with the neural

recordings, and were used to generate a target output signal $\hat{M}(t)$ for the decoder. In order to avoid learning the discontinuity associated with a mod- 2π -based definition of the head direction angle $\theta(t)$, $\hat{M}(t)$ was constructed as a two-dimensional vector having components $M_1(t) \equiv \cos\theta(t)$ and $M_2(t) \equiv \sin\theta(t)$. Spike-sorting analysis of the tetrode-derived waveforms isolated $n = 6$ single units, and the activity of each of these units was converted to a normalized spike rate $N_i(t)$, $i \in \{1, \dots, n = 6\}$ so that $\vec{N}(t)$ could be used as an n -channel analog input to the decoder. As indicated in the section entitled “Input Signals for the Neural Decoder,” the conversion of spike trains to analog input signals was achieved by treating each spike train as a train of pulses having uniform amplitude and pulse width, then passing the pulse trains through a third-order interpolation filter with transfer function

$$H_3(s) = \frac{1}{(1 + \tau_1 s)^3} \quad (1.17)$$

with $\tau_1 = 500$ ms. The filter output was then rescaled and recentered about a zero-offset in order to generate the normalized spike rates $N_i(t)$. The performance of the convolutional decoding algorithm in learning to map $\vec{N}(t)$ to $\vec{M}(t)$ was studied using a software implementation of the decoder, as well as a SPICE analog-circuit simulation of the convolutional decoder, the details of which are described in the section entitled “Functional Circuit Subunits of the Decoder Architecture.” We discuss the results of these simulations in the section entitled “Continuous Real-Time Decoding of Head Direction from Neuronal Spike Activity in the Rat Thalamus.”

Testing the Decoding of Discrete Decisions Using Neuronal Spike Recordings from the Posterior Parietal Cortex of a Macaque Monkey Engaged in an Arm-Reaching Task

Neurons in the posterior parietal cortex have been shown to encode intention to execute limb movements in both humans (Connolly, Andersen, & Goodale, 2003) and nonhuman primates (Snyder, Batista, & Andersen, 1997). The ability to decode signals from this region has provided proof of principle for neural prosthetic devices based on cognitive (as opposed to explicitly motor) control signals (Musallam et al., 2004). Using neural spike trains recorded during these proof-of-principle experiments, we investigated the performance of our system in the real-time decoding of discrete arm-reaching decisions.

The spike trains used as input signals to the decoder in this set of experiments were derived from recordings by a multielectrode array chronically implanted in the medial intraparietal area (within the ‘parietal reach region’) of a macaque monkey. The animal had previously been trained to perform a standard stimulus-response task involving center-out arm-reaching movements between visual targets. This task, which has been described in detail elsewhere (Musallam et al., 2004), was designed to isolate the neural correlates of motor intention from those of actual movement. The monkey initiated each iteration of the task by touching a central cue target and looking at a nearby visual fixation point at $t = -800$ ms (its gaze was monitored using an eye-tracking device, and cues were presented on a touch-sensitive screen). After a delay of 500 ms a peripheral cue target was flashed from $t = -300$ to $t = 0$ ms at one of four locations displaced up, right, down, or left from the starting point. The animal was rewarded if it touched the indicated target at the end of a memory period of 1500 ± 300 ms. Neural activity as monitored by the implanted electrode array was recorded continuously during each trial. A spike-sorting algorithm isolated 54 units from the recorded signals, and spike trains from

each unit were smoothed to spike rate waveforms using a filter of the form provided in Equation (1.17), with $\tau_1 = 50$ ms. Filter output was then rescaled and recentered about a zero-offset in order to generate the normalized spike rate input signals $N_i(t)$ for each of the isolated units indexed by $i \in \{1, \dots, n - 1 = 54\}$; the $(n = 55)$ -dimensional neural input signal $\vec{N}(t)$ contained an additional constant-offset component $N_n(t) = 1$.

The task of the decoder was to predict imminent arm movement on the basis of intention-related neuronal activity. Therefore, in order to ensure the absence from $\vec{N}(t)$ of residual neuronal activity corresponding to actual arm movement (such artifact signals are sometimes present at the beginning and end of a memory interval), only the segment of $\vec{N}(t)$ from $t \in [200, 1100]$ ms was fed into the decoder for each reach. The motor output $\vec{M}(t)$ was defined as a two-dimensional position vector corresponding to the target to which the monkey reached at the end of a corresponding memory period. Reach target positions were encoded as follows:

Analog outputs generated by the decoder were thresholded according to

$$M_i \rightarrow \text{sgn } M_i, \quad (1.18)$$

so that positive and negative outputs were interpreted as +1 and -1, respectively.

We studied the performance of a software implementation of our convolutional decoding system in learning to use neural signals from the parietal reach region to predict the direction of subsequent reaching movements. We discuss the results of these simulations in the section entitled “Decoding Discrete Arm Reaches from Neuronal Spike Activity in the Macaque Monkey Posterior Parietal Cortex.”

Table 2. Encoding vectors for reach targets

Direction	M_1	M_2
Up	+1	+1
Right	+1	-1
Down	-1	-1
Left	-1	+1


EVALUATING THE PERFORMANCE OF THE NEURAL DECODING SYSTEM

Model Motor Trajectory Decoding from Simulated Multichannel Local Field Potential Inputs

Parts (a) and (b) of Figure 9 respectively illustrate the training and post-training phases of decoder operation as the system learns to trace a three-dimensional trajectory in real time during a simulation of the kind described in the section entitled “Testing the Ability to Decode Trajectory Information from Multichannel Local Field Potential Neural Input Signals,” with $(n, m) = (10, 3)$. The figure illustrates qualitatively that $\vec{M}(t)$ converges toward $\hat{\vec{M}}(t)$ reasonably quickly on the timescale set by full-scale variations in the trajectory.

Figure 10 presents the results of a set of computations of $\eta^{(1)}$ for the performance of the neural decoding system in simulations of the kind described in the section entitled “Testing the Ability to Decode Trajectory Information from Multichannel Local Field Potential Neural Input Signals.” The system was trained for intervals of varying length up to one minute, $T \in [0, 60]$ s, and the value of $\eta^{(1)}$ was computed for each of 50 trials at each value of T . In a fraction f of trials at each value of T , the decoding performance as reflected by $\eta^{(1)}(T)$ was significantly worse than in the remaining fraction, $1-f$, of cases. In such cases markedly improved decoding, comparable

Figure 9. Trajectory decoding from simulated local field potential input signals. (a) The correct trajectory, $\hat{\vec{M}}(t)$, is plotted in light gray over the initial segment $t \in [0, 10]$ s of a 40-second training interval, while the decoded trajectory, $\vec{M}(t) \equiv \mathbf{W} \circ \vec{N}(t)$, is plotted as a darker line whose shade lightens as time progresses; $\vec{M}(t)$ converges toward $\hat{\vec{M}}(t)$. (b) For $t \in [40, 80]$ s following training, the correct trajectory is plotted in light gray while the decoded trajectory evolves with time from dark to lighter shades




to that achieved in the $(1-f)$ -majority of cases, could be achieved by randomly reinitializing the parameters $W_{ij}^{(k)}$ and decoding again. The data presented in Figure 10 were obtained by setting $f = 0.1$. Error in trajectory estimation by the decoder decreases rapidly as the training interval increases. At $T = 30$ s, for example, $\langle \eta^{(l)} \rangle \approx 0.008 \pm 0.008$, as compared with a baseline value of $\langle \eta^{(l)} \rangle \approx 0.118 \pm 0.106$ computed for an untrained system ($T = 0$) over 1000 trials.

Continuous Real-Time Decoding of Head Direction from Neuronal Spike Activity in the Rat Thalamus

Head direction was decoded from the activity of the $n = 6$ isolated thalamic neurons according to the method described in the section entitled “Testing the Decoding of Continuous Trajectories Using Neuronal Spike Recordings from the Thalamus of an Awake Behaving Rat.” The adaptive filter parameters $\{W_{ij}^{(p)}\} = \{A_{ij}, \tau_{ij}\}$ were optimized through gradient descent over training intervals of length T during which the decoder error,

Figure 10. Mean squared trajectory prediction error as a function of training time for the convolutional decoding algorithm



$$\begin{aligned} e_i(t) &= M_i(t) - \hat{M}_i(t) \\ \vec{M}(t) &= (\cos \theta(t), \sin \theta(t)), \end{aligned} \quad (1.19)$$

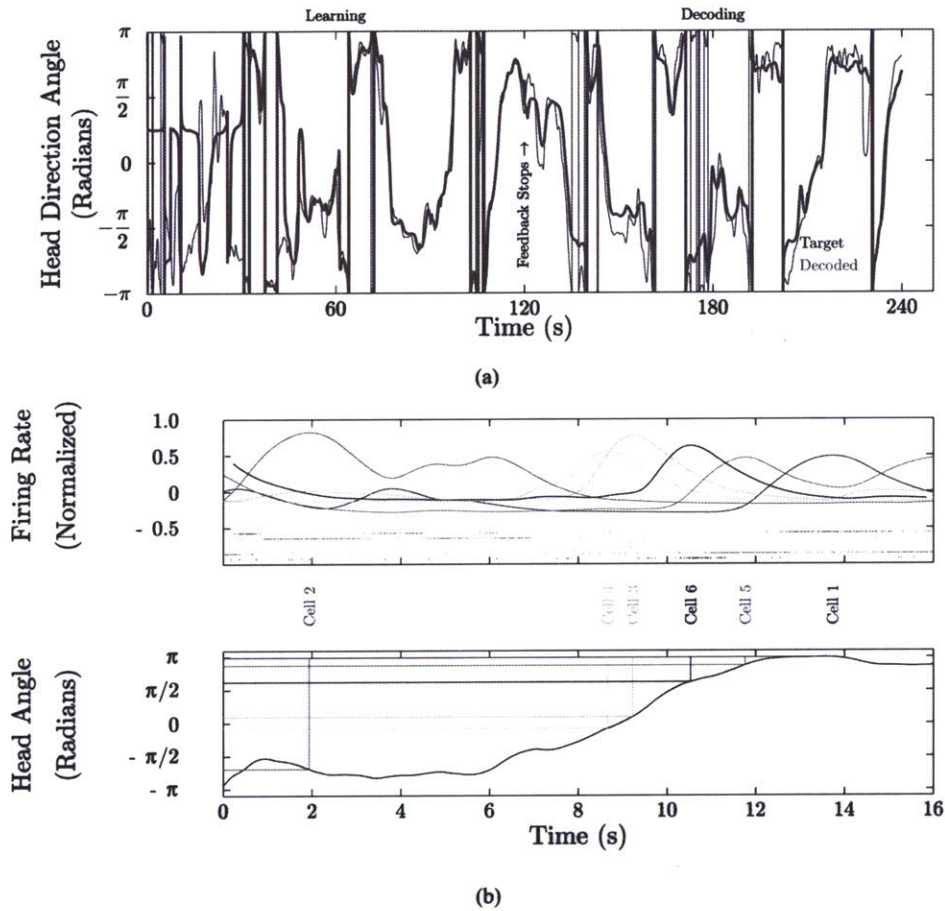
(where θ denotes the head direction angle) was made available to the adaptive filter in the feedback configuration described in the section entitled “An Algorithm for Adaptive Convolutional Decoding of Neural Cell Ensemble Signals” for $t \in [0, T]$. Following these training intervals feedback was discontinued and the performance of the decoder was assessed by comparing the decoder output $\vec{M}(t)$ with $\hat{\vec{M}}(t)$ for $t > T$.

Figure 11(a) compares the output of the decoder to the measured head direction over a 240-s interval. The filter parameters were trained over the interval $t \in [0, T = 120]$ s. The figure shows $\vec{M}(t)$ (thick, light line) tracking $\hat{\vec{M}}(t)$ (thin, dark line) with increasing accuracy as training progresses, illustrating that while initial predictions are poor, they improve with feedback over the course of the training interval. Feedback is discontinued at $t = 120$ s. Qualitatively, the plots

on the interval $t \in [120, 240]$ s illustrate that the output of the neural decoder reproduces the shape of the correct waveform, predicting head direction on the basis of neuronal spike rates. Figure 11(b) displays output over a brief interval from the neurons in the ensemble whose activity was decoded, showing individual action potentials as raster plots below the corresponding smoothed firing rate curve for each cell, as well as a plot of head direction over the same time interval; individual neurons evidently have distinct receptive fields in head-direction space, and their spatial selectivity makes decoding both possible and, in this case, intuitive.

The performance of the decoder in predicting head direction was assessed quantitatively using the normalized mean-squared error measure $\eta^{(1)}$ (in this context $L_i = 2$, $t_1 = T$, and $t_2 = T + 60$ s). In order to quantify the accuracy of head direction decoding as a function of training time T , $\eta^{(1)}$ was computed for a set of training and decoding trials with increasingly long training periods, averaging over randomized initial settings of the filter parameters and different choices of training interval.

Figure 11. (a) Continuous decoding of head direction from neuronal spiking activity. (b) Spiking activity in head direction cells and corresponding head direction plotted as functions of time. The paired plots illustrate neuronal receptive fields and the distribution of their peaks over the range of possible head direction angles




The results of this computation are displayed in Figure 12, which shows improving accuracy of head direction decoding with increased training.

Decoding Discrete Arm Reaches from Neuronal Spike Activity in the Macaque Monkey Posterior Parietal Cortex

The input signals to the neural decoder when decoding discrete reaches were piecewise-constant, as described in the section entitled “Testing

the Decoding of Discrete Decisions Using Neuronal Spike Recordings from the Posterior Parietal Cortex of a Macaque Monkey Engaged in an Arm-Reaching Task.” This form of input reflects a qualitative difference between the decoding problem in this experiment, which requires the decoder to make a series of decisions from among a finite set of options, and the decoding problems framed in the sections entitled “Testing the Ability to Decode Trajectory Information from Multichannel Local Field Potential Neural Input Signals” and “Testing the Decoding of Continuous

Figure 12. Performance improvements in head direction decoding with increasing training times



Trajectories Using Neuronal Spike Recordings from the Thalamus of an Awake Behaving Rat,” and simulated in the sections entitled “Model Motor Trajectory Decoding from Simulated Multichannel Local Field Potential Inputs” and “Continuous Real-Time Decoding of Head Direction from Neuronal Spike Activity in the Rat Thalamus,” which require the decoder to estimate smooth trajectories as functions of time. While the gradient-descent least-squares approach is applicable to both kinds of problem, the convolution kernel chosen to implement the neural decoder,

$$W_{ij} = \frac{A_{ij}}{\tau_{ij}} e^{-\frac{t}{\tau_{ij}}},$$

is designed to exploit the predictive value of past input signals. The degree to which past inputs $\vec{N}(t' < t)$ have predictive value is reflected by the value of the time constant τ_{ij} , and as $\tau_{ij} \rightarrow 0$ the time interval over which $\vec{N}(t' < t)$ contributes significantly to the present-time output $\vec{M}(t)$ correspondingly vanishes. In this experiment the signal to be decoded corresponds to a time series of discrete decisions made every $\Delta t = w$; consecutive reach cues were guar-

anteed to be independent through experimental design. Consequently, $\vec{N}(t')$ is completely uncorrelated from $\vec{N}(t)$ and $\vec{M}(t)$ for $|t - t'| \geq w$. (In concrete terms, since successive reaches are independent, neural activity preceding one reach contains no predictive information concerning the direction of the next reach.) As a result, meaningful decoding requires $\tau_{ij} = 0$, which we enforce by taking τ_{ij} to be a small, unmodifiable constant. Therefore, our decoding scheme as applied to the reach-intention neural data reduces to an instantaneous linear decoder, analogous to a single-layer artificial neural network implemented in continuous time and trained with continuous-time feedback.

The neural data used in the experiments reported here were first obtained and analyzed in connection with a previously reported set of experiments (Musallam et al., 2004). The decoding method used in those experiments involved an analysis of variance to preselect a subset of neurons exhibiting the greatest directional tuning, followed by Bayesian inference on the mean firing rates and higher-order Haar wavelet coefficients

of the signals obtained from the selected neurons. We consequently had the opportunity to compare the performance of the adaptive-filter decoder to that of the Bayesian decoder. The principal performance measure reported using the Bayesian decoder was a 64.4% success rate in predicting the correct one of four allowed reach directions (Musallam et al., 2004). Under corresponding training conditions the neural decoding system described in the present work generated accurate predictions in $65\% \pm 9\%$ of trials (the uncertainty figure preceded by the \pm symbol indicates the magnitude of one standard deviation). Improved decoding performances were demonstrated in (Musallam et al., 2004) by considering higher-order coefficients in the Haar wavelet decomposition of the neural input signals, but in the present study we considered only the zeroth-order coefficients, corresponding to mean firing rates.

The decoding scheme originally used to decode the data analyzed in this section was based in part on the known tendency of direction-sensitive neurons to ‘tune’ to preferred directions in the sense that only movement in certain preferred directions induces such neurons to modulate their firing rates away from a baseline (Cohen & Andersen, 2002). In preparing the Bayesian decoder, an off-line analysis of variance on the training set of spike trains (corresponding to arm reaches in each direction for each isolated neuron) was required to rank the isolated neurons by degree of directional sensitivity. The computational intensity of this decoding scheme was sufficiently high that inputs from only a subset of isolated neurons were used in decoding after the learning period ended, and this ranking provided a means of prioritizing neurons for use as decoder inputs. By contrast, the adaptive-filter decoder described here easily handles all 54 neuronal inputs in computer simulations of real-time decoding; in a certain sense the adaptive-filter decoder automatically learns which neural inputs are the most highly directionally tuned (Rapoport, 2007). Moreover, the analog-circuit-based implementation of the

decoder described here processes all neuronal inputs in parallel, so the computational intensity of the decoding task does not constrain the number of neuronal inputs the system can handle. Consequently, the adaptive-filter approach to decoding scales favorably with the number of neuronal inputs to the system. This is an important virtue of adaptive-filter decoding, as decoding accuracy typically improves and more complex decoding tasks can be performed without sacrificing accuracy as more neuronal inputs are used (Chapin, 2004). Furthermore, improvements in multielectrode neural recording methodologies and technologies continue to facilitate recording from increasing numbers of neurons (Harrison et al., 2007; M. A. L. Nicolelis et al., 2003; Suner, Fellows, Vargas-Irwin, Nakata, & Donoghue, 2005; Wise, Anderson, Hetke, Kipke, & Najafi, 2004).

The ability of the adaptive-filter decoder to handle large numbers of neurons provides an opportunity to explore a decoding regime less amenable to the Bayesian analysis of (Musallam et al., 2004). Figure 13 plots decoder performance as a function of the number of neuronal inputs (the horizontal line across the plot indicates the 0.25 threshold corresponding to unbiased guessing). While the corresponding performance curves for the Bayesian decoding algorithm were originally analyzed for up to sixteen neurons (Musallam et al., 2004), the computational efficiency of the adaptive filter enables performance to be evaluated for considerably larger numbers of neuronal inputs; the computation illustrated in Figure 13 is limited only by the total number of neuronal inputs available. The computation was performed under the standard training condition of 30 trials, and the error bars indicate the magnitude of a standard deviation after averaging over sets of randomized initial conditions and training inputs. As expected, decoder performance increases from just above the 25% chance threshold to the maximum of approximately 65% reported earlier in this section. The lower curve corresponds to random selections of the input neurons, while the upper

curve corresponds to preselection of neurons in order of decreasing variance in the mean firing rates over the four reach directions (higher variance indicates greater directional selectivity). The latter curve suggests that decoding input signals from the subset of neurons transmitting the greatest amount of directional information results in performance nearly equivalent to that obtained from using the full set of available signals.

The 65% success rate of the decoding system, while comparable to that of other decoding algorithms tested on the same data (Musallam et al., 2004), indicates that there is considerable room for improvement. An outstanding question in the field of neural prosthetics concerns the degree to which intelligent users can compensate for imperfect decoding through biofeedback. Marked improvements in performance along these lines have been observed over time in both monkeys and humans (Carmena et al., 2003; Hochberg et al., 2006; Musallam et al., 2004; Taylor et al., 2002), but such contributions from biological learning are evidently insufficient. The work of (Musallam et al., 2004) indicates that performance can be improved by considering the temporal structure of the input neural signals at higher than zeroth order, and this might be achieved through changing the form of the filter kernels W_{ij} .

EVALUATING DESIGN TRADEOFFS IN THE CONSTRUCTION OF AN IMPLANTABLE NEURAL DECODER

Power Efficiency

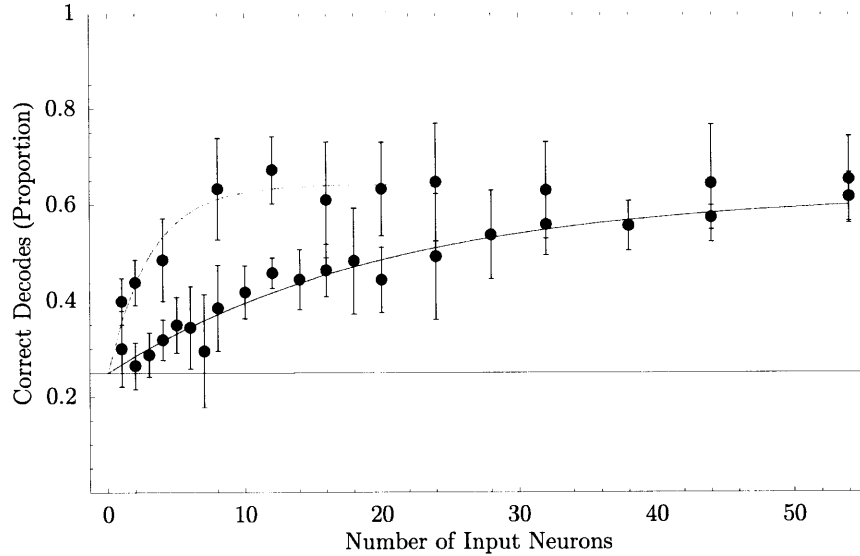
Simulations using the circuit designs presented in the section entitled “Implementation of the Decoding Algorithm in Analog Circuitry” indicate that a single decoding module (corresponding to an adaptive kernel W_{ij} and associated optimization circuitry, as diagrammed in Figure 3) should consume approximately 54 nW from a 1-V supply

in 0.18 μm CMOS technology (and require less than 3000 μm^2 in area). This low power consumption is achieved through the use of subthreshold bias currents for transistors in the analog filters and other components. Analog preprocessing of raw neural input waveforms can be accomplished by dual thresholding to detect action potentials on each input channel and then smoothing the resulting spike trains to generate mean firing rate input signals. Simulations of the circuits presented in the section entitled “Spike-based decoding” indicate that each analog preprocessing module should consume approximately 241 nW from a 1-V supply in 0.18 μm CMOS technology. A full-scale system with $n = 100$ neuronal inputs comprising $\vec{N}(t)$ and $m = 3$ control parameters comprising $\vec{M}(t)$ would require $m \times n = 300$ decoding modules and consume less than 17 μW in the decoder and less than 25 μW in the preprocessing stages.

Savings in Telemetry Power

An advantage of our system is that its suitability for power- and area-efficient analog implementation can enable decoding in the implanted unit of a brain-machine interface, saving power by obviating the need for analog-to-digital conversion of neural signals before decoding, and by compressing data before wireless transmission. In this section we compare the power costs of internal analog decoding to those of external digital decoding. System power will depend on the bandwidth required to transmit digitized neural data to an external unit, so we consider three alternatives for digitization and telemetry, reflecting three approaches to trading algorithmic flexibility in decoding for total system power: (1) Transmission of digitized raw neural waveforms, preserving the greatest amount of neural information and requiring the greatest bandwidth; (2) Transmission of threshold-crossing events from neural waveforms,

Figure 13. Decoding performance as a function of neuron number for randomly selected neurons (dark) and neurons selected on the basis of directional selectivity (light)



preserving only spike timing information and requiring intermediate bandwidth; (3) Transmission of control-parameter outputs from an implantable analog decoder, requiring minimal bandwidth and internal power consumption but sacrificing algorithmic flexibility.

In cases (1) and (2) initial 8-bit digitization of the input data at 30 kHz would require approximately 1 μ W per channel (Yang & Sarpeshkar, 2006), contributing 100 μ W in total. In these cases we consider implementing the external decoding system using a highly power-efficient digital signal processor (DSP) (“TMS320C55x Technical Overview (Literature Number SPRU393),” 2000) and analog-to-digital converters (A/Ds) (Yang & Sarpeshkar, 2006). (1) The first alternative requires transmitting all of the digitized data, corresponding to a bandwidth of 24 Mbs⁻¹ for all $n = 100$ channels; at a rate of 1 mW per Mbs⁻¹ (assuming typical link geometries and efficient topologies) (Mandal & Sarpeshkar, 2007, 2008) wireless data telemetry would consume 24 mW. (2) The second

alternative requires dual thresholding for spike detection, followed by transmission of one bit per channel indicating the presence of a detected spike. Two comparison operations per period at 30 kHz across $n = 100$ input channels constitutes 6×10^6 operations per second. In practice (due to chip and board level parasitics) the most power-efficient DSPs operate at efficiencies of several hundred microwatts per MIPS (Millions of Instructions Per Second) (Verret, 2003); assuming an efficiency of 250 μ W/MIPS, the power associated with these thresholding operations would be 1.5 mW. The bandwidth required to transmit the thresholded waveforms, assuming a maximal spike rate of 1 Hz per channel, would be 0.1 Mbs⁻¹ for all $n = 100$ channels, corresponding to a transmission power of 0.1 mW. So the total power consumed by a system adopting this alternative for data digitization, compression and transmission would be approximately 1.7 mW.

In estimating the power consumed in cases (1) and (2) we have not included the overhead of a

DSP in excess of the power consumed in computations. In practice such overhead costs can be significant, contributing several milliwatts or more, depending on the processor used and its operating settings (Verret, 2003). However, it is important to note that all DSPs incur static power costs due to CMOS leakage currents, independent of device activity and operating frequency. Leakage power depends primarily on operating voltage and temperature, and the processing core of one highly power-efficient DSP dissipates approximately 180 μW at 1.6 V and room temperature (25°C); this leakage power increases to approximately 277 μW at body temperature (37°C) (Verret, 2003). Therefore, our power estimates for DSP-based processing represent lower bounds on what would be consumed in a practical implementation.

When our decoding algorithm is executed by a power-efficient DSP, the power expended in performing computations is negligible relative to the leakage power of the processor core. In a digital implementation of our decoder the adaptive filter kernels are transformed into discrete-time filters. A discrete-time implementation of the filter kernel $Ae^{-\frac{t}{\tau}}$ has a two-parameter z -transform, so each update cycle requires two multiplications. An update frequency of 10 Hz requires 20 multiplications per second, and scaling by $n \times m = 300$ yields the computational rate required to execute the algorithm, 0.006×10^6 multiplications per second. At an efficiency of 250 $\mu\text{W}/\text{MIPS}$, the power cost of the algorithm itself would therefore be only approximately 1.5 μW .

(3) As discussed in the section entitled “Power Efficiency,” our simulations indicate that a low-power-analog implementation of our first-order adaptive kernel decoder with $(n, m) = (100, 3)$ could be built with a power budget of approximately 17 μW (54 nW per decoding module) and that an additional 25 μW is required for the preprocessing stages. Transmission of $m = 3$ decoded parameters for real-time control of an external device with

10-bit precision and an update frequency of 10 Hz requires a transmission rate of 300 bs^{-1} and an associated power of 300 nW. (Here we have assumed an impedance modulation telemetry scheme similar to the one reported in (Mandal & Sarpeshkar, 2007) and (Mandal & Sarpeshkar, 2008), which operates at an efficiency of < 1 nJ per bit even for transmission rates greater than 1 Mbs^{-1} .) The total power consumption associated with internal decoding would therefore be approximately 43 μW .

So an implantable analog decoding system is capable of operating with 40- to 500-fold less power than would be required by digital preprocessing and telemetry in external decoding systems. Moreover, the 43 μW required by an implantable decoding system is 4–6 times less than even the leakage power of state-of-the-art power-efficient DSPs; this implies that when the implantable analog decoder is operating, it should consume at least 4 times less power than an external digital decoder that is doing nothing. The efficiency of our analog architecture arises from its exploitation of subthreshold currents (which are wasted as ‘leakage’ power in traditional digital implementations) to perform computations.

A further point of comparison concerns the sizes of the devices required to perform the decoding. In our simulations, a single decoder module implemented in 0.18 μm technology would require less than 3000 μm^2 , so a system decoding 3 control parameters from 100 input channels would consume less than 1 mm^2 of chip area. By contrast, industry-standard low-power DSPs are typically 100 times larger, consuming approximately 1 cm^2 of board area (“TMS320C54x BGA Mechanical Data (Document MPBG021C),” 2002). Space limitations and skull curvature restrict the planar surface area of devices implanted within the cranial cavity to approximately 2×2 cm^2 , so individual components that consume 1 cm^2 of area are likely impractical as elements of implantable units.

FUTURE RESEARCH DIRECTIONS

Neural Decoding in the Context of Implantable Brain–Machine Interfaces

In this section we describe five important emerging areas in the field of neural decoding for brain–machine interfaces: (1) *Biocompatibility and Stability*: Ideal neuroprosthetics must be capable of implantation in the brain for indefinite periods without damaging or eliciting unintended reactions from surrounding neural tissue, macro- and microanatomically optimized for the appropriate neural targets, and capable of stably recording neural signals of the selected types for indefinite periods of time; (2) *Biological versus Artificial Learning*: Achieving optimal long-term performance from brain–machine interfaces requires understanding the interplay between artificial and biological learning; (3) *Closed-Loop Neuroprosthetics*: Neuroprosthetics must be able to adapt in real time based on feedback from sensory input, neural populations, and other environmental variables; (4) *Integrated Electronic Systems*: General-purpose devices suitable for clinical intervention or neuroscientific research must have all functional subcomponents integrated within a single modularized package; (5) *Scalable Multi-Region Brain–Machine Interfaces*: Brain–machine interfaces should be capable of parallel or networked operation, in which multiple independent modules can operate either independently from neighboring or disparate neural populations without sacrificing performance quality, or while sharing information in a networked fashion in real time.

Biocompatibility and Stability

The difficulty of maintaining stable single-unit neural recordings for extended periods has long been acknowledged by electrophysiologists, and while some notable examples have been reported in the literature, such reports tend to be

the exception rather than the rule. The field of brain–machine interfaces has responded to this problem in two principal ways, at once seeking to optimize electrode design and electrode-tissue interactions for stable long-term recording, while simultaneously exploring a range of neural signals, including local field potentials and unsorted multiunit activity as alternatives to single-unit action potentials. Biocompatibility as a concern for implantable brain–machine interface systems extends to multiple aspects of the implanted system, including the design and materials chosen for the packaging, but we restrict our comments here to neural electrodes.

Decoding neural signals presupposes the availability of such signals, and the quality of neural decoding depends not only on the particular algorithms used for the decoding, but also on the quality of the neural signals used as inputs to the neural decoder. An ideally biocompatible neural recording electrode must satisfy two sets of constraints, which respectively concern the static and dynamic properties of the tissue-electrode interface. Ideal static properties of the electrode include the following: (1) Suitability for insertion in a manner that does minimal damage to surrounding neural and vascular tissue; (2) Mechanical stiffness equal to that of the surrounding tissue, in order to avoid tissue-electrode shearing and electrode displacement during head acceleration; (3) Minimal electrode volume to minimize neural tissue displacement after insertion; (4) Shank and recording site geometry optimized for stereotactic placement relative to desired macro- and microanatomic targets; and (5) Tip or recording site geometry optimized for the neural signal of choice. Ideal dynamic properties of a neural recording electrode include the following: (1) Optimized electrode impedance given the local tissue properties and neural signals of interest; and (2) Ability to regulate the response of neural tissue to the presence of the electrode over time.

Almost all of these static and dynamic properties represent active areas of current research and

potentially fruitful future work. In particular, microfabrication techniques (Kipke et al., 2008; Wise et al., 2004) and nontraditional materials (such as conducting polymers) (Abidian & Martin, 2008; Ludwig, Uram, Yang, Martin, & Kipke, 2006; Richardson-Burns, Hendricks, & Martin, 2007) present large parameter spaces within which to optimize electrode designs. Additionally, it is well known that the microanatomic environment of a recording electrode changes over time, generally in ways that degrade electrode performance, as neighboring neurons tend to die or recede from the electrode, for example, and astrocytes and microglia ensheath the electrode and change the impedance at the tissue-electrode interface. Consequently, understanding and controlling the dynamic response of neural and neurovascular tissue to neural recording electrodes is also an extremely important area for future research (Bjornsson et al., 2006; Shain et al., 2003).

Some differences of opinion exist as to the most appropriate signals to use as inputs for neural prosthetic devices. The difficulty of first reliably obtaining and subsequently maintaining stable single-unit recordings for indefinite periods has motivated the use of other forms of input signal, which in a number of studies have yielded equivalent decoding performance while exhibiting greater stability over time. Action potentials, multiunit activity, and local field potentials associated with movement, imagined movement, and motor planning have all been shown capable of affording neural control of external devices in conjunction with a variety of neural decoding techniques (Andersen, Burdick, Musallam, Pesaran, & Cham, 2004; Mehring et al., 2003; Pesaran et al., 2002; Shenoy et al., 2003; Stark & Abeles, 2007).

A related question concerns the ways in which the brain represents motor control information: Precisely what kinds of state space variables do neural populations encode, and how stable are the representations (Rokni, Richardson, Bizzi, & Seung, 2007)? In particular, in the context of neuro-

motor prosthetics, under what circumstances does neural activity reflect internal signals to muscles, external parameters such as those relating to limb kinematics or dynamics, or higher-level control signals such as those related to movement planning (Kakei, Hoffman, & Strick, 1999)? Questions such as these are particularly relevant to systems such as those in which functional electrical stimulation (FES) of muscles in a paralyzed limb is placed under real-time neural control to restore lost function (Moritz, Perlmutter, & Fetz, 2008).

Biological vs. Artificial Learning

A major question in the area of neural decoding to achieve neural control of prosthetic devices concerns the extent to which the brain can learn to decode processed neural data on its own. Given a particular prosthetic interface controlled by neural signals from a subject that receives feedback on decoder performance, what is the minimal decoder required to operate the device reliably? Under what circumstances, and with what degree of reliability, can biological learning compensate for deficiencies in an artificial neural decoder?

As indicated earlier in this chapter, experiments in both animals (Musallam et al., 2004; Taylor et al., 2002) and humans (Hochberg et al., 2006) have provided early evidence that users of neural prosthetic systems are capable of adapting to the particular rules governing the control of their brain-machine interfaces, demonstrating improved performance over time even after termination of the learning phases for their software-based decoders. Nevertheless, some of the same experiments have also demonstrated the apparent insufficiency of such contributions from biological learning, and so the degree to which a neural prosthetic system should rely on artificial or biological learning remains unclear. More recent work (Ganguly & Carmena, 2009) has helped clarify some of these issues by underscoring the importance of stable long-term neural recordings for reliable neural decoding and improvements

through biological learning. In particular, it appears that the brain is capable not only of learning complex decoders, but also of correcting errors in corrupt decoders, provided the decoders are based on stable neural recordings. Much of the existing work on neural decoding has been limited by the inability to maintain stable recordings from multiple units over long time periods, and by the consequent need to retrain decoders at each training session. As biological learning can occur over longer timescales than those set by the instability of typical single- or multiunit neural recordings, optimal conditions for biological learning in the context of neural prosthesis control have rarely been achieved. Recent evidence suggests that a reliably high level of performance can be obtained as the brain forms motor memories in response to learning a decoder based on stable, long-term recordings. The implications of these findings for the design of neural decoders, neural electrodes, and neural prosthetic systems in general remain to be explored and will be an important area for future research.

Relatedly, decoder optimality is likely to be application dependent, as specific contexts will determine the constraints and performance criteria to be optimized by a particular neural decoder. Nevertheless, it may be important to ask under what circumstances is it possible to define and construct optimal neural decoders, and to understand how to make tradeoffs between biological and artificial learning that optimize system performance.

Closed-Loop Neuroprosthetics

Translating advances in neural decoding into clinical practice will require prosthetic devices, including robotic prosthetic limbs, that confer natural functionality (Aaron et al., 2006; Adee, 2009; Kuiken et al., 2009; Kuniholm, 2009). The form and quality of sensory and direct neural feedback signals are important in enabling devices to achieve such lifelike performances. Therefore, in

the context of motor prosthetics, realistic sensory feedback from proprioceptive, haptic, and other modalities encoded by physiologic signals will improve motor prosthetic performance (Miguel A. L. Nicolelis, Chapin, & Wessberg, 2007). Implementation of such a system will require simultaneous neural recording and neural stimulation of somatosensory and motor brain regions.

An important class of future neural prosthetic systems will involve closed-loop neural decoding and neural stimulation, in which decoded neural firing patterns can be used to control neural stimulation in a feedback loop. This paradigm is applicable to deep brain stimulation (DBS) in which the stimuli delivered depend on sensory data and the state of local neural networks; current DBS systems use fixed, nonadaptive stimulation paradigms. Neural decoding for seizure detection, coupled with stimulation to preempt the seizure in a feedback fashion, represents another potential clinical application for closed-loop neural decoding and stimulation. The scope of clinical problems amenable to solutions of this kind is likely to expand over time.

Integrated Electronic Systems

We make reference in the Introduction to recent progress developing modular electronic components of implantable brain-machine interfaces, particularly in the areas of micropower neural amplification, low-power data telemetry, ultra-low-power analog-to-digital conversion, low-power and power-conserving schemes for neural stimulation, adaptive power biasing of multi-electrode arrays, and energy-efficient wireless recharging circuits (Sarpeshkar et al., 2008). Significant progress has also been made toward integrating and packaging these individual modules into functioning systems, in several cases together with neural recording electrodes (Harrison, 2008; Song et al., 2007; Wise et al., 2004). Work in this area is an important avenue for continuing research toward

brain–machine interface systems that can be used in both clinical and experimental neuroscience.

Power efficiency and power management are extremely important concerns in small, biologically implantable systems. The development of high-performance biologically implantable power sources that are safe, compact, and easily rechargeable or schemes for scavenging energy from the biological environment could lift a variety of design constraints on the implantable electronics, as implantable batteries are typically the largest and most massive components of biologically implantable electronic devices.

Scalable Multi-Region Brain–Machine Interfaces

Scalability represents an important trend in contemporary neuroscience (Narasimhan, 2004). Neural computations underlying motor and other functions are distributed across a variety of brain regions, so it is likely that optimal control of neural prosthetic devices will be achieved using neural decoders that coordinate input signals from multiple brain regions (Andersen, Burdick et al., 2004; Chapin, 2004; M. A. L. Nicolelis et al., 2003). Coordinating signal processing over multiple input streams derived from distinct regions requires that brain–machine interface systems be designed in a modular fashion that permits individual modules not only to report data independently, but also to share information in a networked fashion in real time. Such scalable, networked systems represent an extremely interesting direction for future research, particularly in light of the observation that even rudimentary artificial networking of neural regions can induce long-term motor plasticity (Jackson et al., 2006). Systems of this kind might be used not only to establish functional connections among multiple regions within an individual brain, but also to connect individual brains to one another.

CONCLUSION

We have described a continuous-time, adaptive, biomimetic algorithm and a corresponding micropower analog circuit architecture for decoding neural signals (Figure 1 and Figure 2). The system is suitable for decoding both local field potentials and mean spike rates; in experimental trials it has successfully decoded both discrete decisions from the parietal cortex of a behaving primate (Figure 13) and continuous trajectories from the thalamus of a behaving rodent (Figure 11). The algorithm and architecture presented here offer a practical approach to computationally efficient neural signal decoding, independent of the hardware used for their implementation. While the system is suitable for analog or digital implementation, we discuss how a micropower analog implementation trades some algorithmic programmability for reductions in power consumption that could facilitate implantation of a neural decoder within the brain. In particular, circuit simulations of our analog architecture (comprised of standard building-block circuits described in the section entitled “Functional Circuit Subunits of the Decoder Architecture”) indicate that a 100-channel, 3-motor-output neural decoder can be built with a total power budget of approximately 43 μ W. Our work could also enable a 100,000-fold reduction in the bandwidth needed for wireless transmission of neural data, thereby reducing to nanowatt levels the power potentially required for wireless data telemetry from a brain implant.

Our work suggests that highly power-efficient and area-efficient analog neural decoders that operate in real time can be useful components of brain–implantable neural prostheses, with potential applications in neural rehabilitation and experimental neuroscience. Through front-end preprocessing to perform neural decoding and data compression, algorithms and architectures such as those presented here can complement digital signal processing and wireless data transmission systems (Figure 1), offering significant increases in power and area efficiency at little cost.

ACKNOWLEDGMENT

This work was supported in part by a grant from the National Institutes of Health, NS-056140, and in part by an MIT–CIMIT Medical Engineering Fellowship. Neural data were provided by Sam Musallam of the Andersen Laboratory at the California Institute of Technology, and by Hector Penagos of the Wilson Laboratory at the Massachusetts Institute of Technology.

REFERENCES

- Aaron, R. K., Herr, H. M., Ciombor, D. M., Hochberg, L. R., Donoghue, J. P., & Briant, C. L. (2006). Horizons in Prosthesis Development for the Restoration of Limb Function. *The Journal of the American Academy of Orthopaedic Surgeons*, *14*(10), S198–S204.
- Abidian, M. R., & Martin, D. C. (2008). Experimental and theoretical characterization of implantable neural microelectrodes modified with conducting polymer nanotubes. *Biomaterials*, *29*, 1273–1283. doi:10.1016/j.biomaterials.2007.11.022
- Adee, S. (2009). The revolution will be prosthetized. *IEEE Spectrum*, *46*(1), 44–48. doi:10.1109/MSPEC.2009.4734314
- Andersen, R. A., Burdick, J. W., Musallam, S., Pesaran, B., & Cham, J. G. (2004). Cognitive neural prosthetics. *Trends in Cognitive Sciences*, *8*(11), 486–493. doi:10.1016/j.tics.2004.09.009
- Andersen, R. A., Musallam, S., & Pesaran, B. (2004). Selecting the signals for a brain-machine interface. *Current Opinion in Neurobiology*, *14*, 1–7. doi:10.1016/j.conb.2004.10.005
- Arenz, A., Silver, R. A., Schaefer, A. T., & Margrie, T. W. (2008). The Contribution of Single Synapses to Sensory Representation in Vivo. *Science*, *321*, 977–980. doi:10.1126/science.1158391
- Baker, M. W., & Sarpeshkar, R. (2007). Feedback Analysis and Design of RF Power Links for Low-Power Bionic Systems. *IEEE Transactions on Biomedical Circuits and Systems*, *1*(1), 28–38. doi:10.1109/TBCAS.2007.893180
- Bjornsson, C. S., Oh, S. J., Al-Kofahi, Y. A., Lim, Y. J., Smith, K. L., & Turner, J. N. (2006). Effects of insertion conditions on tissue strain and vascular damage during neuroprosthetic device insertion. *Journal of Neural Engineering*, *3*, 196–207. doi:10.1088/1741-2560/3/3/002
- Brockwell, A. E., Rojas, A. L., & Kass, R. E. (2004). Recursive Bayesian decoding of motor cortical signals by particle filtering. *Journal of Neurophysiology*, *91*, 1899–1907. doi:10.1152/jn.00438.2003
- Carmena, J. M., Lebedev, M. A., Crist, R. E., O’Doherty, J. E., Santucci, D. M., & Dimitrov, D. F. (2003). Learning to Control a Brain-Machine Interface for Reaching and Grasping by Primates. *Public Library of Science Biology*, *1*(2), 1–16.
- Chapin, J. K. (2004). Using multi-neuron population recordings for neural prosthetics. *Nature Neuroscience*, *7*(5), 452–455. doi:10.1038/nn1234
- Chapin, J. K., Moxon, K. A., Markowitz, R. S., & Nicolelis, M. L. (1999). Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex. *Nature Neuroscience*, *2*, 664–670. doi:10.1038/10223
- Cohen, Y. E., & Andersen, R. A. (2002). A Common Reference Frame for Movement Plans in the Posterior Parietal Cortex. *Nature Reviews Neuroscience*, *3*, 553–562. doi:10.1038/nrn873
- Connolly, J. D., Andersen, R. A., & Goodale, M. A. (2003). fMRI evidence for a ‘parietal reach region’ in the human brain. *Experimental Brain Research*, *153*, 140–145. doi:10.1007/s00221-003-1587-1

- Eden, U. T., Frank, L. M., Barbieri, R., Solo, V., & Brown, E. N. (2004). Dynamic analyses of neural encoding by point process adaptive filtering. *Neural Computation*, *16*(5), 971–998. doi:10.1162/089976604773135069
- Eliasmith, C., & Anderson, C. H. (2003). *Neural Engineering*.
- Ganguly, K., & Carmena, J. M. (2009). Emergence of a Stable Cortical Map for Neuroprosthetic Control. *Public Library of Science Biology*, *7*(7).
- Ghovanloo, M., & Atluri, S. (2007). A wideband power-efficient inductive wireless link for implantable microelectronic devices using multiple carriers. *IEEE Transactions on Circuits and Systems. I, Fundamental Theory and Applications*, *54*(10), 71–77. doi:10.1109/TCSI.2007.905187
- Harrison, R. R. (2008). The design of integrated circuits to observe brain activity. *Proceedings of the IEEE*, *96*, 1203–1216. doi:10.1109/JPROC.2008.922581
- Harrison, R. R., Watkins, P. T., Kier, R. J., Lovejoy, R. O., Black, D. J., & Greger, B. (2007). A Low-Power Integrated Circuit for a Wireless 100-Electrode Neural Recording System. *IEEE Journal of Solid-state Circuits*, *42*(1), 123–133. doi:10.1109/JSSC.2006.886567
- Haykin, S. (1999). *Neural Networks: A Comprehensive Foundation*. Upper Saddle River, New Jersey: Prentice Hall.
- Hochberg, L. R., Serruya, M. D., Friehs, G. M., Mukand, J. A., Saleh, M., & Caplan, A. H. (2006). Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature*, *442*, 164–171. doi:10.1038/nature04970
- Holleman, J., & Otis, B. (2007). 29th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, *A sub-microwatt low-noise amplifier for neural recording* (pp. 3930–3933).
- Jackson, A., Mavoori, J., & Fetz, E. E. (2006). Long-term motor cortex plasticity induced by an electronic neural implant. *Nature*, *444*, 55–60. doi:10.1038/nature05226
- Takei, S., Hoffman, D. S., & Strick, P. L. (1999). Muscle and Movement Representations in the Primary Motor Cortex. *Science*, *285*, 2136–2139. doi:10.1126/science.285.5436.2136
- Kipke, D. R., Shain, W., Buzsáki, G., Fetz, E. E., Henderson, J. M., & Hetke, J. F. (2008). Advanced Neurotechnologies for Chronic Neural Interfaces: New Horizons and Clinical Opportunities. *The Journal of Neuroscience*, *28*(46), 11830–11838. doi:10.1523/JNEUROSCI.3879-08.2008
- Kuiken, T. A., Li, G., Lock, B. A., Lipschultz, R. D., Miller, L. A., & Stubblefield, K. A. (2009). Targeted Muscle Reinnervation for Real-time Myoelectric Control of Multifunction Artificial Arms. *Journal of the American Medical Association*, *301*(6), 619–628. doi:10.1001/jama.2009.116
- Kuniholm, J. (2009). Open Arms. *IEEE Spectrum*, *46*(3), 36–41. doi:10.1109/MSPEC.2009.4795454
- Lazzaro, J., Wawrzynek, J., & Lippmann, R. P. (1997). A Micropower Analog Circuit Implementation of Hidden Markov Model State Decoding. *IEEE Journal of Solid-state Circuits*, *32*(8), 1200–1209. doi:10.1109/4.604076
- Loeliger, H.-A., Tarköy, F., Lustenberger, F., & Helfenstein, M. (1999). Decoding in Analog VLSI. *IEEE Communications Magazine*, 99–101. doi:10.1109/35.755457
- Ludwig, K. A., Uram, J. D., Yang, J., Martin, D. C., & Kipke, D. R. (2006). Chronic neural recordings using silicon microelectrode arrays electrochemically deposited with a poly(3,4-ethylenedioxythiophene) (PEDOT) film. *Journal of Neural Engineering*, *3*, 59–70. doi:10.1088/1741-2560/3/1/007

- Mandal, S., & Sarpeshkar, R. (2007). A Bidirectional Wireless Link for Neural Prostheses that Minimizes Implanted Power Consumption. *Proceedings of the IEEE Biomedical Circuits and Systems Conference (BioCAS)*, 45-48.
- Mandal, S., & Sarpeshkar, R. (2008). Power-Efficient Impedance-Modulation Wireless Data Links for Biomedical Implants. *IEEE Transactions on Biomedical Circuits and Systems*, 2(4), 301–315. doi:10.1109/TBCAS.2008.2005295
- Mead, C. (1989). *Analog VLSI and Neural Systems*. Addison Wesley.
- Mehring, C., Rickert, J., Vaadia, E., de Oliveira, S. C., Aertsen, A., & Rotter, S. (2003). Inference of hand movements from local field potentials in monkey motor cortex. *Nature Neuroscience*, 6(12), 1253–1254. doi:10.1038/nn1158
- Mohseni, P., Najafi, K., Eliades, S. J., & Wang, X. (2005). Wireless Multichannel Biopotential Recording Using an Integrated FM Telemetry Circuit. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 13(3), 263–271. doi:10.1109/TNSRE.2005.853625
- Moritz, C. T., Perlmutter, S. I., & Fetz, E. E. (2008). Direct control of paralysed muscles by cortical neurons. *Nature*, 456, 639–642. doi:10.1038/nature07418
- Musallam, S., Corneil, B. D., Greger, B., Scherberger, H., & Andersen, R. A. (2004). Cognitive Control Signals for Neural Prosthetics. *Science*, 305, 258–262. doi:10.1126/science.1097938
- Narasimhan, K. (2004). Scaling up neuroscience. *Nature Neuroscience*, 7(5), 425. doi:10.1038/nn0504-425
- Nicolelis, M. A. L., Chapin, J. K., & Wessberg, J. (2007). United States of America Patent No. Official Gazette of the United States Patent and Trademark Office Patents: U. S. P. a. T. Office.
- Nicolelis, M. A. L., Dimitrov, D., Carmena, J. M., Crist, R., Lehew, G., & Kralik, J. D. (2003). Chronic, multisite, multielectrode recordings in macaque monkeys. *Proceedings of the National Academy of Sciences of the United States of America*, 200(19), 11041–11046. doi:10.1073/pnas.1934665100
- O’Halloran, M., & Sarpeshkar, R. (2004). A 10-nW 12-bit Accurate Analog Storage Cell with 10-aA Leakage. *IEEE Journal of Solid-state Circuits*, 39(11), 557–560. doi:10.1109/JSSC.2004.835817
- O’Halloran, M., & Sarpeshkar, R. (2006). An Analog Storage Cell with 5e/sec Leakage. *Proceedings of the IEEE International Symposium on Circuits and Systems (ISCAS)*, 557-560.
- Olsson, R. H. III, & Wise, K. D. (2005). A Three-Dimensional Neural Recording Microsystem With Implantable Data Compression Circuitry. *IEEE Journal of Solid-state Circuits*, 40(12), 2796–2804. doi:10.1109/JSSC.2005.858479
- Pesaran, B., Pezaris, J. S., Sahani, M., Mitra, P. P., & Andersen, R. A. (2002). Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nature Neuroscience*, 5(8), 805–811. doi:10.1038/nn890
- Rapoport, B. I. (2007). *Neural Prosthetics for Paralysis: Algorithms and Low-Power Analog Architectures for Decoding Neural Signals*. Cambridge, Massachusetts: Massachusetts Institute of Technology.
- Rapoport, B. I., Wattanapanitch, W., Penagos, H. L., Musallam, S., Andersen, R. A., & Sarpeshkar, R. (2009). *A Biomimetic Adaptive Algorithm and Low-Power Architecture for Implantable Neural Decoders*. Paper presented at the 31st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC 2009).

- Richardson-Burns, S. M., Hendricks, J. L., & Martin, D. C. (2007). Electrochemical polymerization of conducting polymers in living neural tissue. *Journal of Neural Engineering*, 4, L6–L13. doi:10.1088/1741-2560/4/2/L02
- Rokni, U., Richardson, A. G., Bizzi, E., & Seung, H. S. (2007). Motor Learning with Unstable Neural Representations. *Neuron*, 54, 653–666. doi:10.1016/j.neuron.2007.04.030
- Salthouse, C., & Sarpeshkar, R. (2003). A Practical Micropower Programmable Bandpass Filter for use in Bionic Ears. *IEEE Journal of Solid-state Circuits*, 38(1), 63–70. doi:10.1109/JSSC.2002.806286
- Sanchez, J. C., Erdogmus, D., Principe, J. C., Wessberg, J., & Nicolelis, M. A. L. (2005). Interpreting spatial and temporal neural activity through a recurrent neural network brain machine interface. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 13(2), 213–219. doi:10.1109/TNSRE.2005.847382
- Santhanam, G., Ryu, S. I., Yu, B. M., Afshar, A., & Shenoy, K. V. (2006). A high-performance brain-computer interface. *Nature*, 442, 195–198. doi:10.1038/nature04968
- Sarpeshkar, R. (1998). Analog Versus Digital: Extrapolating from Electronics to Neurobiology. *Neural Computation*, 10, 1601–1638. doi:10.1162/089976698300017052
- Sarpeshkar, R. (2006). Brain Power: Borrowing from Biology Makes for Low-Power Computing. *IEEE Spectrum*, 43(5), 24–29. doi:10.1109/MSPEC.2006.1628504
- Sarpeshkar, R., Baker, M., Salthouse, C., Sit, J.-J., Turicchia, L., & Zhak, S. (2005). An Analog Bionic Ear Processor with Zero-Crossing Detection. *Proceedings of the IEEE International Solid State Circuits Conference (ISSCC)*, 78-79.
- Sarpeshkar, R., Lyon, R. F., & Mead, C. (1997). A low-power wide-linear-range transconductance amplifier. *Analog Integrated Circuits and Signal Processing*, 13(1-2), 123–151. doi:10.1023/A:1008292213687
- Sarpeshkar, R., Salthouse, C., Sit, J.-J., Baker, M., Zhak, S., & Lu, T. (2005). An Ultra-Low-Power Programmable Analog Bionic Ear Processor. *IEEE Transactions on Bio-Medical Engineering*, 52(4), 711–727. doi:10.1109/TBME.2005.844043
- Sarpeshkar, R., Wattanapanitch, W., Arfin, S. K., Rapoport, B. I., Mandal, S., & Baker, M. W. (2008). Low-Power Circuits for Brain-Machine Interfaces. *IEEE Transactions on Biomedical Circuits and Systems*, 2(3), 173–183. doi:10.1109/TBCAS.2008.2003198
- Sarpeshkar, R., Wattanapanitch, W., Rapoport, B., Arfin, S. K., Baker, M. W., Mandal, S., et al. (2007). Low-Power Circuits for Brain-Machine Interfaces. *Proceedings of the IEEE Symposium on Circuits and Systems*, 2068-2071.
- Serruya, M. D., Hatsopoulos, N. G., Fellows, M. R., Paninski, L., & Donoghue, J. P. (2003). Robustness of neuroprosthetic decoding algorithms. *Biological Cybernetics*, 88(3), 219–228. doi:10.1007/s00422-002-0374-6
- Shain, W., Spataro, L., Dilgen, J., Haverstick, K., Retterer, S., & Isaacson, M. (2003). Controlling cellular reactive responses around neural prosthetic devices using peripheral and local intervention strategies. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 11, 186–188. doi:10.1109/TNSRE.2003.814800
- Shenoy, K. V., Meeker, D., Cao, S. Y., Kureshi, S. A., Pesaran, B., & Buneo, C. A. (2003). Neural prosthetic control signals from plan activity. *Neuroreport*, 14(4), 591–596. doi:10.1097/00001756-200303240-00013

- Snyder, L. H., Batista, A. P., & Andersen, R. A. (1997). Coding of intention in the posterior parietal cortex. *Nature*, *386*, 167–170. doi:10.1038/386167a0
- Song, Y.-K., Patterson, W. R., Bull, C. W., Borton, D. A., Li, Y., Nurmikko, A. V., et al. (2007). *A Brain Implantable Microsystem with Hybrid RF/IR Telemetry for Advanced Neuroengineering Applications*. Paper presented at the 29th Annual International Conference of the IEEE Engineering in Medicine and Biology Society.
- Srinivasan, L., Eden, U. T., Mitter, S. K., & Brown, E. N. (2007). General-purpose filter design for neural prosthetic devices. *Journal of Neurophysiology*, *98*, 2456–2475. doi:10.1152/jn.01118.2006
- Stark, E., & Abeles, M. (2007). Predicting movement from multiunit activity. *The Journal of Neuroscience*, *27*(31), 8387–8394. doi:10.1523/JNEUROSCI.1321-07.2007
- Suner, S., Fellows, M. R., Vargas-Irwin, C., Nakata, G. K., & Donoghue, J. P. (2005). Reliability of signals from a chronically implanted, silicon-based electrode array in non-human primate primary motor cortex. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, *13*(4), 524–541. doi:10.1109/TNSRE.2005.857687
- Taube, J. S. (1995). Head Direction Cells Recorded in the Anterior Thalamic Nuclei of Freely Moving Rats. *The Journal of Neuroscience*, *15*(1), 70–86.
- Taylor, D. M., Tillery, S. I. H., & Schwartz, A. B. (2002). Direct Cortical Control of 3D Neuroprosthetic Devices. *Science*, *296*, 1829–1832. doi:10.1126/science.1070291
- Theogarajan, L., Wyatt, J., Rizzo, J., Drohan, B., Markova, M., Kelly, S., et al. (2004). *Low power neural stimulator for a retinal prosthesis*. Paper presented at the Annual Meeting of the Association for Research in Vision and Ophthalmology.
- TMS320C54x BGA Mechanical Data (Document MPBG021C). (2002). Texas Instruments Incorporated.
- TMS320C55x Technical Overview (Literature Number SPRU393). (2000). Texas Instruments Incorporated.
- Velliste, M., Perel, S., Spalding, M. C., Whitford, A. S., & Schwartz, A. B. (2008). Cortical control of a prosthetic arm for self-feeding. *Nature*, *453*(7198), 1098–1101. doi:10.1038/nature06996
- Verret, R. (2003). *TMS320VC5510 Power Consumption Summary (Application Report)*. Texas Instruments Incorporated. Document Number.
- Warland, D., Reinagel, P., & Meister, M. (1997). Decoding visual information from a population of retinal ganglion cells. *Journal of Neurophysiology*, *78*(5), 2336–2350.
- Wattanapanitch, W., Fee, M., & Sarpeshkar, R. (2007). An Energy-Efficient Micropower Neural Recording Amplifier. *IEEE Transactions on Biomedical Circuits and Systems*, *1*(2), 136–147. doi:10.1109/TBCAS.2007.907868
- Wessberg, J., & Nicolelis, M. A. L. (2004). Optimizing a Linear Algorithm for Real-Time Robotic Control using Chronic Cortical Ensemble Recordings in Monkeys. *Journal of Cognitive Neuroscience*, *16*(6), 1022–1035. doi:10.1162/0898929041502652
- Wessberg, J., Stambaugh, C. R., Kralik, J. D., Beck, P. D., Laubach, M., & Chapin, J. K. (2000). Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature*, *408*, 361–365. doi:10.1038/35042582
- Wise, K. D., Anderson, D. J., Hetke, J. F., Kipke, D. R., & Najafi, K. (2004). Wireless implantable microsystems: High-density electronic interfaces to the nervous system. *Proceedings of the IEEE*, *92*(1), 76–97. doi:10.1109/JPROC.2003.820544

- Wu, W., Black, M. J., Mumford, D., Gao, Y., Bienenstock, E., & Donoghue, J. P. (2004). Modeling and decoding motor cortical activity using a switching Kalman filter. *IEEE Transactions on Bio-Medical Engineering*, 51(6), 933–942. doi:10.1109/TBME.2004.826666
- Wu, W., Gao, Y., Bienenstock, E., Donoghue, J. P., & Black, M. J. (2006). Bayesian population decoding of motor cortical activity using a Kalman filter. *Neural Computation*, 18(1), 80–118. doi:10.1162/089976606774841585
- Wu, W., & Hatsopoulos, N. G. (2008). Real-Time Decoding of Nonstationary Neural Activity in Motor Cortex. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 16(3), 213–222. doi:10.1109/TNSRE.2008.922679
- Wu, W., Shaikhouni, A., Donoghue, J. P., & Black, M. J. (2004). Closed-loop neural control of cursor motion using a Kalman filter. *Proceedings of the IEEE Engineering in Medicine and Biology Society*, 4126–4129.
- Yang, H., & Sarpeshkar, R. (2006). A Bio-inspired Ultra-Energy-Efficient Analog-to-Digital Converter for Biomedical Applications. [Special Issue on Life Sciences and System Applications]. *IEEE Transactions on Circuits and Systems I*, 53(11), 2349–2356. doi:10.1109/TCSI.2006.884463
- Zhak, S., Baker, M., & Sarpeshkar, R. (2003). A Low Power Wide Dynamic Range Envelope Detector. *IEEE Journal of Solid-state Circuits*, 38(10), 1750–1753. doi:10.1109/JSSC.2003.817599
- Brown, E. N. (2004). Theory of Point Processes for Neural Systems. In B. G. Carson Chow, David Hansel, Claude Meunier, Jean Dalibard (Ed.), *Methods and Models in Neurophysics, Session LXXX: Lecture Notes of the Les Houches Summer School 2003* (pp. 691–727): Elsevier.
- Brown, E. N., Barbieri, R., Eden, U. T., & Frank, L. M. (2004). Likelihood Methods for Neural Spike Train Data Analysis. In Feng, J. (Ed.), *Computational Neuroscience: A Comprehensive Approach* (pp. 253–286). London: CRC Press.
- Bullock, T. H. (1997). Signals and signs in the nervous system: The dynamic anatomy of electrical activity is probably information-rich. *Proceedings of the National Academy of Sciences of the United States of America*, 94(1-6), 1-6.
- Buzsáki, G. (2004). Massively Parallel Recording of Unit and Local Field Potentials With Silicon-Based Electrodes. *Nature Neuroscience*, 7(5), 446–451.
- Chorost, M. (2006). *Rebuilt: My Journey Back to the Hearing World*. Boston: Houghton Mifflin.
- Csicsvari, J., Henze, D. A., Jamieson, B., Harris, K. D., Sirota, A., & Barthó, P. (2003). Massively Parallel Recording of Unit and Local Field Potentials With Silicon-Based Electrodes. *Journal of Neurophysiology*, 90, 1314–1323. doi:10.1152/jn.00116.2003
- Eden, U. T. (2005). *Point Process Filters in the Analysis of Neural Spiking Models*. Cambridge: Unpublished Doctor of Philosophy, Harvard University.
- Eliasmith, C., & Anderson, C. H. (2004). *Neural Engineering: Computation, Representation, and Dynamics in Neurobiological Systems*. Cambridge: MIT Press.

ADDITIONAL READING

Brockwell, A. E., Rojas, A. L., & Kass, R. E. (2004). Recursive Bayesian Decoding of Motor Cortical Signals by Particle Filtering. *Journal of Neurophysiology*, 91, 1899–1907. doi:10.1152/jn.00438.2003

- Georgopoulos, A. P., Kalaska, J. F., Caminiti, R., & Massey, J. T. (1982). On the Relations Between the Direction of Two-Dimensional Arm Movements and Cell Discharge in Primate Motor Cortex. *The Journal of Neuroscience*, 2(11), 1527–1537.
- Georgopoulos, A. P., Schwartz, A. B., & Kettner, R. E. (1986). Neural Population Coding of Movement Direction. *Science*, 233, 1416–1419. doi:10.1126/science.3749885
- Haykin, S. (2008). *Neural Networks and Learning Machines*. New York: Prentice Hall.
- He, J., Ma, C., & Herman, R. (2008). Engineering Neural Interfaces for Rehabilitation of Lower Limb Function in Spinal Cord Injured. *Proceedings of the IEEE*, 96(7), 1152–1166. doi:10.1109/JPROC.2008.922593
- Kandel, E., Schwartz, J., & Jessell, T. (2000). *Principles of Neural Science* (4th ed.). New York: McGraw-Hill.
- Lewicki, M. S. (1998). A review of methods for spike sorting: the detection and classification of neural action potentials. *Network (Bristol, England)*, 9, R53–R78. doi:10.1088/0954-898X/9/4/001
- MacKay, D. J. C. (2006). *Information Theory, Inference, and Learning Algorithms*. Cambridge: Cambridge University Press.
- Martin, R. (2005). Mind Control. *Wired*, 13(3).
- Maynard, E. M., Nordhausen, C. T., & Normann, R. A. (1997). The Utah Intracortical Electrode Array: a recording structure for potential brain-computer interfaces. *Electroencephalography and Clinical Neurophysiology*, 102, 228–239. doi:10.1016/S0013-4694(96)95176-0
- Mead, C. (1989). *Analog VLSI and Neural Systems*. Reading: Addison-Wesley.
- Mitra, P., & Bokil, H. (2007). *Observed Brain Dynamics*. Oxford University Press. doi:10.1093/acprof:oso/9780195178081.001.0001
- Nicolelis, M. A. L. (Ed.). (1999). *Methods for Neural Ensemble Recordings*. New York: CRC Press.
- Nordhausen, C. T., Rousche, P. J., & Normann, R. A. (1994). Optimizing recording capabilities of the Utah Intracortical Electrode Array. *Brain Research*, 637, 27–36. doi:10.1016/0006-8993(94)91213-0
- Paradiso, J., Borriello, G., & Bonato, P. (2008). Implantable Electronics. *IEEE Pervasive Computing / IEEE Computer Society [and] IEEE Communications Society*, 7(1), 12–13. doi:10.1109/MPRV.2008.11
- Ventura, V. (2008). Spike train decoding without spike sorting. *Neural Computation*, 20(4), 923–963. doi:10.1162/neco.2008.02-07-478
- Zhang, K., Ginzburg, I., McNaughton, B. L., & Sejnowski, T. J. (1998). Interpreting Neuronal Population Activity by Reconstruction: Unified Framework With Application to Hippocampal Place Cells. *Journal of Neurophysiology*, 79(2), 1017–1044.

KEY TERMS AND DEFINITIONS

Adaptive Algorithm: In the present work we use the term “adaptive algorithm” to denote an optimization process in which a filter adjust its own transfer function. The associated filter is referred to as an “adaptive filter,” and the adaptation process is referred to as “tuning.” In other scientific literature, the term “adaptive algorithm” is sometimes used to refer to an algorithm that modifies its behavior in accordance with the computational resources available.

Biomimetic: Designed according to principles used by biological systems.

Brain–Machine Interface (BMI): A system enabling direct transfer of information between a brain and a machine; when the machine in question is a computer, such a system is sometimes called a “brain–computer interface” (BCI). The transfer of information is typically “direct” in one of two senses: Either information obtained from the brain remains encoded in physiologic signals (such as intracortical or electroencephalographic [EEG] voltage traces) that are recorded from the brain and transmitted to an external system for decoding, or information is sent to the brain by electrochemical stimulation of neural tissue. The present chapter focuses on brain-machine interfaces of the former kind, which decode information from electrophysiologically recorded neural signals.

Motor Control: “Motor” in the neurophysiologic sense refers to muscle activity and associated body movements coordinated by the nervous system; such neural coordination is referred to as “motor control.”

Neural Cell Ensemble (Neuronal Ensemble): A group of neurons whose activity is considered simultaneously. The cells in the ensemble may or may not share functional or anatomic connections.

Neural Decoding: A process by which physiologic neural signals are translated into objectively defined state variables (which may correspond to parameters such as position, speed, or discrete values from a set of possibilities).

Neural Prosthetic (Neural Prosthesis): The term “prosthesis” refers to an artificial replacement for a body part, usually one whose function has been lost or impaired. The term “neural prosthesis” may refer literally to a system designed to replace or augment part of a brain or nervous system. The term may also refer to a prosthetic device, such as an artificial limb, that is controlled by a brain–machine interface.

Parietal Cortex: An anatomic region of the cerebral cortex whose functions relate primarily to spatial orientation and proprioception (the perception of body positioning), based on integration of information from multiple sensory information pathways. Experiments described in the present chapter involve neural data from the macaque monkey posterior parietal cortex, which computes transformations between eye-based and head-based coordinate systems, and which has been correspondingly implicated in the coordination of limb and eye movements for reaching and grasping.

Receptive Field: The “receptive field” of a neuron consists of all locations from which a stimulus is capable of altering the firing pattern of the neuron in question. When the cell of interest belongs to a sensory pathway, its receptive field may be spatial, corresponding to regions of the visual field for neurons of the visual system or to regions on the body surface for touch-sensitive neurons; the receptive field may also exist in a more abstract parameter space, corresponding, for example, to ranges in pitch for neurons of the auditory system. Functionally connected neurons may also be considered to comprise receptive fields for one another: One neuron whose firing constitutes a stimulus for a second neuron can be considered part of the receptive field of the second cell.

Thalamus: A structure in the brain whose functions include transmitting sensory information to the cerebral cortex, and relaying motor information from the cerebral cortex to the peripheral nervous system. The present chapter describes a set of experiments involving neural decoding from “head direction cells” of the rat thalamus, which are neurons within this region of the rat brain that are known to exhibit receptive fields tuned to specific orientations of the head relative to the environment.

Chapter 3

Efficient Universal Computing Architectures for Decoding Neural Activity

3.1 Overview

The ability to decode neural activity into meaningful control signals for prosthetic devices is critical to the development of clinically useful brainmachine interfaces. This chapter presents a computationally efficient, low-power, implantable system for real-time neural decoding. In particular, it describes a Turing-machine-like digital architecture that uses virtually arithmetic-free, purely logical primitives to establish a universal computing architecture for neural decoding. It then describes in detail a concrete realization of the neural decoding architecture in a low-power, field-programmable gate array (FPGA) that dissipates 0.5 mW, and validates system performance by decoding electrophysiologic data from a behaving rodent.

3.2 Efficient Universal Computing Architectures for Decoding Neural Activity

This section consists of a manuscript in preparation for submission to *PLoS ONE* [63].

Efficient Universal Computing Architectures for Decoding Neural Activity

Benjamin I. Rapoport^{1,2,*}, Lorenzo Turicchia³, Woradorn Wattanapanitch³, Rahul Sarpeshkar³

¹ M.D.–Ph.D. Program, Harvard Medical School, Boston, Massachusetts, United States of America

² Department of Electrical Engineering and Computer Science and Division of Health Sciences and Technology, Massachusetts Institute of Technology, Cambridge, Massachusetts, United States of America

³ Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, Massachusetts, United States of America

* E-mail: brapoport@post.harvard.edu

Abstract

The ability to decode neural activity into meaningful control signals for prosthetic devices is critical to the development of clinically useful brain–machine interfaces (BMIs). Such systems require input from tens to hundreds of brain-implanted recording electrodes in order to deliver robust and accurate performance; they must also minimize power dissipation in order to avoid damaging neural tissue; and they must transmit data wirelessly in order to minimize the risk of infection associated with chronic, transcutaneous implants. Electronic architectures for brain–machine interfaces must therefore minimize size and power consumption, while maximizing the ability to compress data to be transmitted over limited-bandwidth wireless channels. Here we present a system of extremely low computational complexity, designed for real-time decoding of neural signals, and suited for highly scalable implantable systems. Our programmable architecture is an explicit implementation of a universal computing machine emulating the dynamics of a network of integrate-and-fire neurons; it requires no arithmetic operations except for counting, and decodes neural signals using only computationally inexpensive logic operations. The simplicity of this architecture does not compromise its ability to compress raw neural data by factors greater than 10^5 . We describe a set of decoding algorithms based on this computational architecture, one designed to operate within an implanted system, minimizing its power consumption and data transmission bandwidth; and a complementary set of algorithms for learning, programming the decoder, and post-processing the decoded output, designed to operate in an external, nonimplanted unit. We demonstrate their implementation in a field-programmable gate array (FPGA) decoding system with 32 neural input channels that dissipates 0.5 mW of power, occupies a $6 \times 6 \times 1\text{-mm}^3$ footprint, and requires fewer than 5000 operations per second. We validate the performance of the system by decoding electrophysiologic data from a behaving rodent.

Author Summary

Programmable, implantable medical devices such as cardiac pacemakers, cochlear implants, and deep-brain stimulators, have transformed hundreds of thousands of lives. Brain–machine interfaces, already established as tools for basic neuroscience, are poised to emerge with similarly transformative impact in medicine, repairing and augmenting the human nervous system in conditions ranging from paralysis and amputation to Parkinsonian dyskinesia, blindness, and epilepsy. Clinically useful brain–machine interfaces must decode neural signals in real time into control signals for prosthetics and other external devices. The clinical viability of brain–machine interfaces hinges on their ability to operate in miniature, implantable electronic systems with low power budgets. Yet state-of-the-art neural decoding systems currently neglect power and size constraints, relying on computationally intensive, software-based algorithms. Here we present a universal computing architecture for neural decoding, and show how it can be programmed to decode neural signals by emulating the dynamics of a biological neural network.

In contrast with existing approaches to neural decoding it requires no arithmetic operations, and uses only computationally inexpensive logic operations. We also demonstrate a field-programmable gate array (FPGA) prototype of the implantable component of the decoding system, which dissipates 0.5 mW of power, occupies only $6 \times 6 \times 1\text{-mm}^3$, and operates with fewer than 5000 operations per second. We validate the performance of the system by decoding neural signals from a behaving rodent.

Introduction

Implantable Neural Decoding Systems for Brain–Machine Interfaces

Recent years have seen dramatic progress in the field of brain–machine interfaces, with implications for rehabilitation medicine and basic neuroscience [1–3]. One emerging goal is the development of an implantable system capable of recording and decoding neural signals, and wirelessly transmitting raw and processed neural data to external devices. Early versions of such systems have shown promise in developing prosthetic devices for paralyzed patients [4], retinal implants to restore sight to the blind [5,6], deep brain stimulators for treating Parkinson’s disease and related disorders [7], and systems for predicting and preventing seizures [8]. Neural decoding has been essential to these systems, conferring the adaptive ability to learn to extract from neural data meaningful signals for controlling external devices in real time.

Electronics implanted in the brain must be sufficiently energy-efficient to dissipate very little power while operating, so as to avoid damaging neural tissue; conserving power also extends device lifetimes. Yet experimental and clinical neuroscience demand ever-increasing bandwidth from such systems [9]: sampling rates on the order of 30 kbps per recording channel are commonly used in applications requiring discrimination of action potentials generated by individual cells, and while contemporary systems rarely record from more than hundreds of neurons simultaneously, much more extensive sampling will be required to probe the state of an entire human brain containing on the order of 10^{11} neurons. In this context, neural decoding can be viewed not only as a computational approach to extracting meaning from vast quantities of data [10], but also as a means of compressing such data. In previous work [11], we have shown that an implanted neural decoder can compress neural data by a factor of 100,000.

Considerable attention has been devoted to meeting the low-power operation constraint for brain implantation in the context of signal amplification [12,13], analog-to-digital conversion [14], power and data telemetry [15–18], neural stimulation [19,20], and overall low-power circuit and system architecture [21,22]. A small amount of work has also been conducted on power-efficient neural data compression [23]. However, almost no systematic effort has been devoted to the problem of power-efficient neural decoding [11,24].

Multiple approaches to neural decoding have been implemented by several research groups. As we have discussed in [11], nearly all of these have employed highly programmable algorithms, using software or microprocessors located outside the body [25–40]. An implantable, low-power decoder, designed to complement and integrate with existing approaches, would add the efficiency of embedded preprocessing options to the flexibility of a general-purpose external processor. As illustrated in Figure 1, our decoding architecture is designed to couple a power-efficient, bandwidth-reducing, implanted decoder, with an external unit that is less power-constrained and can therefore bear a computational load of greater complexity when postprocessing the decoded neural data. Being optimized for low power consumption, it sacrifices a small amount of algorithmic programmability—posing algorithmic challenges with which we deal in this paper—to reduce power consumption and physical size, facilitating inclusion of the decoder within an implanted unit.

As we have described in previous work [21,22], such an implanted unit consists of circuits for neural signal amplification, digitization, decoding, and near-field power and data telemetry. In association with an external unit that manages wireless power transfer and far-field data telemetry, such an implanted unit forms the electronic core of a brain–machine interface.

Biological and Universal Computing Primitives for Neural Decoding

Our computational architecture for neural decoding operates explicitly as a Turing-type universal computing machine, in which the decoding operation is programmed by selecting the rule array of the machine, which can also reprogram itself, resulting in an overall system that emulates the dynamics of a network of integrate-and-fire neurons. In contrast with existing approaches to neural decoding, this framework facilitates extreme power efficiency, *requiring no arithmetic operations except for counting*.

Our architecture decomposes the operation of neural signal decoding, allocating the computational load across two processing units: one implanted within the body, and therefore power-constrained, and the other located outside the body, and therefore less power-constrained. The overall architecture strategically imbalances the computational load of decoding in a way that leverages the relatively high computational power of the external unit to minimize power consumption in the implanted unit. Simultaneously, the system minimizes data throughput between the internal and external units in order to reduce the power costs of wireless communication between the two units.

Figure 1 shows the overall architecture of our neural decoding system. The architecture is decomposed into a set of operations implemented by Turing-type computing machines, shown as a collection of heads (data processing units) reading from and writing to a set of corresponding tapes (programs and data streams). Amplification and digitization of raw neural data, and decoding of that data, are performed by heads N and I , respectively, in the implanted unit. The computations of these two system components are streamed across a wireless data channel to an external unit, which performs more power-intensive external computations to postprocess the decoded output. In particular, further processing of the decoded data is performed externally by head E , and the final output of the system is reported by head O .

The core decoding function executed by the internal unit is an evaluation of the likelihood that the system is in each of its available states. At each time step, t , the internal unit reports a one-bit binary score $d_i(t)$, $i \in \{1 \dots n_s\}$ for each of the n_s possible states, based on neural data observed at each time step. The binary vector of likelihood scores, $\vec{d}(t)$, is processed by the external unit, which decides, on the basis of system history and other *a priori* information, which single state is most probable. It then broadcasts its decision, for example to be used in controlling external devices.

The detailed operation of the internal unit is diagrammed in Figure 2, in which functional blocks are color-coded in accord with the scheme used in Figure 1. Neural inputs from an ($n = 32$)-channel array are amplified and digitized, and the resulting digital bits are copied to the high-throughput neural data tape. As indicated by the red rectangle in Figure 2, these operations correspond to the function of the N head in Figure 1. The internal decoding computations implemented by the I head in Figure 2 are shown in detail within the green box in Figure 1. Digital circuits in this subsystem monitor each input channel during successive time windows of length t_w , counting the number of spikes whose amplitudes exceed channel-specific, programmable levels. The resulting spike counts are evaluated by a program stored in memory, which constitutes the core of the internal decoder. The program defines a set of rules, one or more for each possible state, that are configured during a learning period and then used to evaluate the likelihood scores d_i on the basis of the spike counts observed at each time step. Each rule identifies the n_i channels most informative in decoding its corresponding state, along with spike count thresholds that discriminate the state from its neighbors with statistically defined sensitivity and specificity.

The components of $\vec{d}_i(t)$ are written to the decoder output tape. The external unit examines the decoder output tape in a noncausal manner, postprocessing the decoded output generated by the internal unit to find a single most probable state. Our architecture permits a wide variety of postprocessing schemes, consistent with comparative studies of neural decoding algorithms and their underlying assumptions, which have formally and systematically demonstrated that movement smoothing is the most significant algorithmic factor influencing decoder performance [41]. We therefore implement a general-purpose decoding algorithm in the implanted system, while permitting application-specific choices in the external unit. Here we implement the postprocessing using a Viterbi algorithm.

Pattern-matching algorithms conceptually related to the one implemented in our internal unit have

previously been used to decode neuronal activity, notably in the context of memory replay during dreams [42, 43], but until now the computational complexity of such approaches has limited their applicability to off-line, software-based implementations. Pattern matching systems have the useful property of being able to emulate receptive field structures—essential computational primitives of biological neurons—in a direct and intuitive way: they learn and store a set of *templates*, patterns corresponding to the activity of a given ensemble of neurons in response to a particular set of external states. Classical implementations of decoding by pattern matching function by comparing observed neuronal activity against stored templates (the system must store at least one template for each state to be decoded) and choosing a best match. This approach is typically computationally expensive for two reasons. First, the ability to quantify the degree to which observed neuronal activity matches a given template requires a defined metric, the value of which must be computed for every stored template at every time step of the decoder. And second, useful metrics themselves typically require computationally expensive operations, such as multiplication, root extraction, and division (or normalization). Computation of continuous-valued metrics in a digital context can also be accomplished only to a specified limit of precision.

The efficiency of application-specific digital microcontrollers and digital signal processors (DSPs) arises in large part from their ability to identify and prioritize the computational primitives, such as Fourier transformation or specific kinds of filtering, that are of greatest importance in particular applications [44]. In seeking a minimal digital decoding system whose operation is consistent with the computing primitives of biological neural networks, we have retained pattern matching as an approach to embedding neuronal receptive fields within the decoding architecture. However, we have reduced the template-matching metric to a set of rules in programmable logic. The structure of these rules as implemented in the example system described here results in a decoding architecture that behaves like a network of integrate-and-fire neurons. However, the programmability of the system and its explicitly rule-based architecture ensure that its scope encompasses even complex, multimodal receptive fields, but is not limited to such emulations [45]; the decoding architecture presented here is an example of a universal computing machine customized for neural decoding.

This paper is structured as follows: In this Introduction Section and in the Discussion Section, we address the implications of this work in the context of implantable brain-machine interfaces for clinical applications and basic neuroscience. We present results illustrating the performance of our neural decoding architecture in an initial Results Section, which includes a subsection discussing techniques for noise reduction. In the Methods Section we describe the acquisition and format of our input signals, and develop the decoding and smoothing algorithms themselves. A Methods subsection describes in detail a concrete, hardware implementation of the neural decoding architecture in a low-power field-programmable gate array (FPGA).

Results

Neural Decoding

We applied our neural decoding system to decode whole-body position trajectories from place cell ensemble activity in the hippocampus of a behaving rat. Place cells in rat hippocampus exhibit receptive fields tuned to specific locations in the environment [46–48]. Our system was able to decode temporal firing patterns of ensembles of such cells in real-time simulations using recorded neural data. Spike train inputs were derived from spike-sorted tetrode recordings from the hippocampus of a rat bidirectionally traversing a maze for food reward, as described in [49].

In the context of our place-cell-based position decoding problem, the states to be decoded, s_i , $i \in 1 \dots n_s = m = 32$, are 32 equally sized, discrete sections of a one-dimensional track maze, constituting an arbitrary discretization of the continuous, linear, 10-meter track.

Figure 3 illustrates the encoding of position in our ensemble of $n = 32$ place cells. The columns of

the color-coded array are normalized representations of spike activity for the place cells in the ensemble, with bins (columns) corresponding to discretized positions in the one-dimensional track maze. The rows have been sorted based on the locations associated with maximal spike activity. Figure 3 shows that the receptive fields within this ensemble of neurons are distributed over the available one-dimensional space, forming a basis for effective decoding.

Figure 4 graphically displays the data structure, g , in which the decoding templates have been stored. We used $n_t = 2$, $t_s = 0.5$, and $t_p = 0.25$ to compute the decoding templates, as described in the Methods Section. In Figure 4, $g_{s_i,j}$ is displayed as an $(n_s = m = 32) \times (n = 32)$ array, in which row s_i contains the template used to establish the likelihood threshold for state s_i . Hence, each row contains $n_t = 2$ maximally informative elements, color-coded in one of $2^{b_e=4}$ shades of gray, with remaining row elements colored white. (In the case shown here, every element of g is smaller than 2^3 , so black elements in Figure 4 correspond to threshold values of 7.)

Figure 5 displays the performance of our decoding algorithms at both the spike-to-state and state-to-state stages. At each time step, the output of the decoder is displayed vertically, with black pixels representing the ones in the binary vector $\vec{d}(t)$. The locations of those ones are observed to cluster along a trajectory reflecting the position of the rat in time, with stray ones due to noise and decoder error. The effects of noise and spike-to-state decoder errors are reduced by the state-to-state smoothing algorithm (in this case a Viterbi algorithm implemented by the external unit), whose output is also shown. Using a window length of $t_w = 360$ ms, our decoded output matches the correct trajectory with a Pearson correlation coefficient of 0.70; the correlation rises to 0.91 when the time window for histogram generation and input spike counting is widened to $t_w = 1440$ ms. Our performance is comparable to those of other implementations [28, 50, 51].

Compression Factor

Real-time decoding compresses neural data as it is acquired, by extracting only meaningful information of interest from a high-bandwidth, multichannel signal. For data acquired from an n -channel array and digitized to b_p bits of precision at s samples per second, our algorithm yields a compression factor of

$$c = \frac{nsb_p}{m/t_w}. \quad (1)$$

In our work, data are digitized to at least $b_p = 8$ bits of precision, typically at a rate of 31,250 Hz; time windows are rarely shorter than $t_w = 90$ ms. In an $(n = 32)$ -channel system decoding $m = 32$ states, decoding therefore yields a compression factor of $c \geq 23,130$. In this example we have used $m = n$ to generate a conservative value for c ; in practice, the number of states decoded is often fewer than the number of input channels, resulting in larger compression ratios.

Computational Efficiency

We explicitly calculate the computational efficiency of our decoding architecture in the Methods Section, following a detailed description of its operation. We find that the total computational load, L , associated with neural decoding, scales as

$$L = \beta mn_t f_w, \quad (2)$$

where $f_w \equiv \frac{1}{t_w}$, and $\beta = \frac{13}{2}$ for the implementation employed here, as explained in detail in the Methods Section. Thus, for $m = n_s = 32$, $n_t = 2$, and $f_w = \frac{1}{90 \text{ ms}}$, $L \approx 4623$ operations per second, or approximately 4.7×10^{-3} MIPS (millions of instructions per second).

Noise Reduction

Input noise degrades the performance of the decoder, but our system has a number of mechanisms for mitigating the effects of noise. In considering the impact of noise on system performance, it is helpful to distinguish between correlated and uncorrelated noise, where the correlation is with reference to the collection of neural input signals, $\vec{h}(t)$, and $h_j(t)$, $j \in \{1 \dots n\}$, refers to the signal obtained from input channel j . More precisely, the covariance matrix for $\vec{h}(t)$, computed over a designated time interval, reflects the degree of correlation across input channels. In the context of neural signal recordings, cross-channel correlations (for electrodes spaced tens of micrometers apart) can arise from low-frequency components of the electroencephalogram (EEG) or from motion artifacts (movement of the recording array with respect to the brain, as may occur with head acceleration). Uncorrelated noise may be attributed to intrinsic properties of the recording system or to the biological signal itself, as discussed extensively in [22].

A useful feature of our internal decoding algorithm is its ability explicitly to suppress output noise and tune performance by adjusting t_w , the duration of the window over which neural data are collected at each time step before making a prediction. In particular, t_w can be scaled in proportion to $\frac{1}{f_c}$, where f_c denotes a low-frequency cutoff in the noise spectrum. Lengthening t_w sacrifices system response speed for improved performance accuracy by integrating over more input data. Longer t_w intervals require more memory to store decoding templates of correspondingly higher resolution, so the total memory available to the implanted system (in our case, 18 kbits of RAM) places an upper bound on decoding performance.

The effects of correlated noise can also be suppressed in postprocessing, by the nonimplanted component of the system designated ‘External Computations’ in Figure 1. By virtue of its ability to use *a priori* constraints, such as those imposed by continuity as well as the physical properties and relationship of the states being decoded, the performance of the external system is not entirely limited by the statistical properties of the input signals and their associated noise levels.

Uncorrelated noise restricted to individual channels, associated with low channel-specific signal-to-noise ratios (SNRs), can be handled in the context of our decoding architecture by tuning the value of n_t on a channel-by-channel basis.

Generalizations

Our system for neural decoding, implemented as described here, can be generalized in several ways.

First, consider the combinatorial logic we use to interpret spike train data. We designed the decoding architecture to facilitate applying primitive operations derived from the structure of neuronal receptive fields. We model the receptive field of a neuron with respect to a set of states as the probability distribution of action potential firing as a function of state index. Discretizing this model to a finite level of precision transforms the frequency distribution to a histogram over states. The general problem of evaluating the degree to which input spike counts represent evidence of particular states has often been cast in terms of Bayesian analysis [52]. In the context of a discrete-time digital implementation, however, the most general analysis of a set $\vec{h}(t)$ of input spike counts can be conducted by comparing each element $h_j(t)$ to a finite set of thresholds. As the set of all possible composite results of all such comparisons (made over all channels and all states) is enumerable and finite, discrete-time digital decoding can be implemented by a finite set of threshold-comparison rules. In particular, although the template-matching algorithm we describe here consists of logical conjunction operators applied to single, unidirectional threshold crossings on designated sets of channels, our decoding architecture is capable of using much more general combinatorial logic in analyzing $\vec{h}(t)$.

The combinatorial logic function we implement explicitly in this work, described by Equation 3, is appropriate when decoding from neuronal populations that encode information in receptive fields structured like those of the hippocampal place cells from which we demonstrate effective decoding here. The prototypical place cell has an approximately Gaussian, unimodal receptive field, which can be coarsely

approximated by an impulse or threshold function that assumes a constant, nonzero value for states in which the cell is most active, and vanishes over all other states.

While such receptive-field encoding patterns are common, a more precise and more general model based on the one we describe here will have broader applicability. In a system with more memory or more logic circuits, a decoding scheme could use more logical functions with more conditions, operating on multiple upper- and lower-bound thresholds per channel. Such an approach could more smoothly model the state dependence of neuronal activity present in most receptive fields, through level-dependent conditions, where spike count levels are defined by pairs of upper- and lower-bound thresholds. This approach should be especially useful in decoding from multimodal receptive fields, or from input channels carrying unsorted multiunit activity, as in [53]. As indicated by the dashed ‘Operation Selection’ line in Figure 2, the most general decoder can employ different logical functions for each state.

Second, consider the information content of channel silence. Our decoding scheme takes action on the basis of observed spikes. Yet the absence of spikes also conveys information. A more general version of the system we describe here could exploit the information content of channel silence, constructing histograms and templates for time-windowed spike absence in analogy with those described here for observed spikes.

Third, consider cross-channel correlations in neural activity. The logical operations and probabilistic assumptions we have described here treat input channels, and indeed channel activity in each position, as independent. While it is convenient to assume such independence, neural activity across channels and positions may in general exhibit nonzero correlations. In a more elaborate version of our decoding system, such correlations could be exploited, for example by implementing combinational logic functions depending simultaneously on activity across multiple channels.

Discussion

We have designed a processor architecture for decoding neural signals. Its core decoding functions require no arithmetic operations except for simple counting: by relying entirely and explicitly on computationally inexpensive combinational logic, the architecture can be implemented in devices that dissipate power at rates low enough to permit safe implantation in the human body [54]. Our system applies a pattern recognition algorithm to windowed spike trains to decode and thereby compress large quantities of neural data. Pattern-matching templates are constructed in a way that permits approximate quantification of decoding reliability, and data integration times can be tuned to control decoder output noise. We also describe ways in which more powerful and computationally expensive routines, executed outside the body, can apply physical constraints and other prior knowledge about the system being decoded, to smooth and correct the output of the implanted processor. The resulting low-power, implantable neural decoder is designed for integration into a compact unit of implanted electronics providing neural signal amplification, analog-to-digital conversion, and wireless telemetry for brain-machine interfaces of the kind we have described in previous work [21, 22]. Processors designed using our architecture could be used to monitor cognitive states and to provide natural, brain-based control of prosthetic limbs and other devices for paralyzed patients and amputees.

Advanced neural prosthetic systems that control multiple degrees of freedom currently rely on neural input signals from hundreds of parallel recording channels [55–57], and the trend over the past decade has been toward systems with increasing numbers of channels [9]. The power required to transmit recorded neural data using wireless telemetry increases with the bandwidth required for transmission, which in turn scales linearly with system channel count. Because the total power consumption of fully implanted systems is constrained below the threshold for heat-induced damage of biological tissue, effective strategies for compressing neural data in real time, such as the one we describe here, will be important for high-channel-count implantable—and hence clinically viable—brain-machine interfaces.

Our decoding system, programmed as we have described here, illustrates how a simple universal computing architecture can implement an effective, biomimetic neural decoding. In particular, the rule-

based decoding program we describe can be understood as implementing a two-layer network of digitized integrate-and-fire neurons. The first layer of this network consists of n input neurons, corresponding to the n input channels of the system. The second layer consists of m neurons, each connected to the n_t neurons in the first layer indicated by the decoding rules, and each synaptic weight set to the reciprocal of the corresponding threshold. Each second-layer neuron integrates its inputs for the duration of each time window, resetting to zero every t_w , and firing when it accumulates a value of n_t . The output of the neurons in this second layer constitute the decoder output we have designated $\vec{d}(t)$.

The particular scheme we describe here for programming our decoding architecture has an intuitive interpretation as an emulation of a network of integrate-and-fire neurons. However, the rule-based programming structure of the system is extremely versatile, and by no means limited to intuitive, biomimetic computations. Indeed, the computational universality of such rule-based systems has been explored and described in detail by Wolfram [45].

Methods

Spike Detection

Our neural decoder is designed to process full-bandwidth, multichannel neural signals in real time, by interfacing directly with the digitized output stream from an array of amplifiers and neural recording electrodes. Such a direct connection of the logic to the amplifier array facilitates great flexibility in the spike detection and decoding algorithms that the associated system can implement. In the present work we detect neuronal action potentials (‘spikes’) using a single-threshold method; individual thresholds are programmable, and can be fine-tuned channel by channel.

More advanced spike-detection methods are also possible. One such approach would be to implement dual-threshold detection, with or without a refractory period. Another would be to use more computationally intensive detection criteria, such as nonlinear energy operators [58]. However, we have found that a basic single-threshold method works well in practice.

Importantly, we can also estimate the noise level on each channel in real time. Estimating the noise level is necessary in high-channel-count systems because in such systems it is typically desirable to set spike detection thresholds automatically. A common approach is to set the detection threshold to a multiple of the noise level, isolating spike events with approximate statistical confidence bounds.

Neural Decoding Algorithm

Overall Scheme of Operation

The objective in neural decoding is to infer an aggregate neural state, such as the intention to move a limb in a particular way, from signals sampled from a population of neurons encoding that state. The computational task of a neural decoding system operating in discrete time can therefore be described in the language of universal computing architectures, and our neural decoding architecture can be understood in terms of a Turing-machine-like architectures with tapes and heads, as shown in Figure 1. At each time step, the system accumulates a sample of activity from a targeted neuronal population. This sample is interpreted as a symbol, or set of symbols, on the ‘tape’ read by our processor. The state of the system at that time step is inferred from the symbols by a discrete set of rules, stored in the memory of the system and used in reading the ‘tape.’

The rules for symbol interpretation are derived using a statistical procedure we describe here in detail. Intuitively, the procedure treats neural spike counts from each input channel as test statistics, used to evaluate the likelihood that the system is in each of the possible states. The distributions of these test statistics, conditioned on being in or out of each of the possible states, are approximated as histograms

collected during a learning phase. Maximally informative spike-count values (thresholds) can therefore be derived and converted to logical rules involving only comparison operations.

Our decoding architecture operates as follows:

Neural spikes are detected on each of n input channels using digitally programmable thresholds, as described in the Methods Section under ‘Spike Detection.’ Spikes detected from each channel during an observation window of duration t_w (typically tens to hundreds of milliseconds) are registered in b_c -bit counters. At the end of each time window, the set of counter values is stored in an n -dimensional spike-count vector, $\vec{h}(t)$, describing neural activity across all channels. This vector is compared, component by component, to a set of m stored templates embodying the decoding rules. Each template g_i , $i \in \{1 \dots m\}$ is also an n -dimensional vector, and its components g_{ij} , $j \in \{1 \dots n\}$ constitute spike-count thresholds for corresponding channels $h_j(t)$ in $\vec{h}(t)$. The data structure g that contains the templates is therefore an $m \times n$ array, mnb_c bits in size; one such template array is illustrated in Figure 4. The set of templates contains at least one member for each of the n_s states to be decoded.

At the end of every time window, $\vec{h}(t)$ is compared component-wise to each of the m templates, g_{s_k} , $k \in \{1 \dots m \geq n_s\}$, $s_k \in \{1 \dots n_s\}$; we index the templates in this way to make clear that we allow states to be encoded by multiple templates. The decoder output vector, $\vec{d}(t)$, is then defined as follows:

$$d_{s_k} = \bigvee_{j=1}^n h_{s_k j}(t) > g_{s_k j}, \quad (3)$$

where the condition specified in Equation 3 is the logical *and* of all the component-wise comparisons between the spike counts in the input at time t , and the corresponding stored thresholds in the templates encoding a given state, s_k . Thus, at every time step, the decoder independently evaluates the strength of the evidence that the state s_k is encoded by the neural population over the time window $[t - t_w, t)$. The decoder output, $\vec{d}(t)$, is an m -dimensional binary vector whose components coarsely encode the corresponding likelihood estimates.

In practice, as discussed in the Methods Section under ‘FPGA Implementation,’ the decoder output can be computed by implementing a reduced version of the full m -input logical *and*. This is accomplished by storing the templates g_{s_k} not as n -component vectors, but as $(n_t < \frac{n}{2})$ -component vectors containing the n_t most informative threshold values for each state, paired with pointers to their corresponding input channels. As we discuss in the Results Section under ‘Generalizations,’ and as indicated by the dashed ‘Operation Selection’ line in Figure 2, a more general form of the decoder architecture can also store a pointer to a logical operation to apply to the indicated channel thresholds. This pointer would select among alternatives to the operation described in Equation 3, which we apply uniformly across all channels. In a more general decoder, a more flexible scheme along these lines might be appropriate, in which, for example, individualized sum-of-product operators are selectively applied to each channel.

The decoding algorithm itself makes no *a priori* assumptions about the nature of the encoded states. As a result, it is equally well suited to decoding tasks involving continuous trajectories (controlling a computer mouse, for example, or a prosthetic limb) and to those involving discrete and discontinuous decisions (such as typing on a computer keyboard, controlling the click-state of a mouse, or selecting from among preprogrammed grip states of a prosthetic hand). Decoder output predictions of multiple simultaneous states, null predictions, and state-to-state transitions that are implausible due to physical or other constraints, are all handled by the smoothing algorithm described in the Methods Section under ‘Trajectory Smoothing Algorithm.’

Computational Efficiency

In order to determine the total computational load associated with our decoding architecture, we account for the number of operations required in each computational frame, defined as a single cycle through the

input channels, in which a single spike-count vector, $\vec{h}(t)$, is acquired. Each spike-count vector must be compared against n_t thresholds in each of m stored templates. As illustrated in Figure 2, each comparison is associated with the set of

$$\beta = 5 + \beta_L + \frac{1}{n_t} \quad (4)$$

operations enumerated in Table 1, where the parameter β denotes the total number of basic operations per frame, and the parameter β_L denotes the number of basic operations associated with the combinational logic used to compare incoming spike-count vectors to each stored template. In the implementation we demonstrate here, that logic consists of a single comparison, and so $\beta_L = 1$. In general, however, when more intricate logic is used to evaluate $\vec{h}(t)$, such as when a sum-of-products scheme is used to account for receptive fields with multimodal activity patterns (or, equivalently, input channels carrying multiunit spike activity), β_L may be greater than 1.

Spike-count vectors $\vec{h}(t)$ are acquired at frequency f_w , so the total computational load, as indicated in Equation 2, is $\beta mn_t f_w$.

Computation of Templates

Decoding templates are constructed so as to be, in a statistical sense, both sensitive and specific to their associated states (which may correspond to discretized positions, discretized velocity vectors, or discrete choices, among many other possibilities). By *sensitive* and *specific* we mean, respectively, that each template should detect its corresponding state with high probability when that state occurs, and that templates should discriminate well between states.

The decoding templates are computed based on data gathered during a training period, as shown in Figure 6, and are based on the approximation that the system state is stationary over the chosen time window, t_w . From the training data the decoder learns the tuning and firing rate properties of each input channel with respect to the defined set of states, in the following way:

The template construction algorithm begins by computing two sets, H^+ and H^- , of joint frequency histograms for neural spike activity over all pairs of channels and all states. The sets H^+ and H^- are organized by state and by channel. More precisely, $H_{s_i}^+$, $s_i \in \{1 \dots m\}$, is a data structure containing n joint frequency histograms, $H_{s_i,j}^+$, $j \in \{1 \dots n\}$, of the number of spikes detected on channel j in a time interval of length t_w during which the system encodes state s_i . Similarly, $H_{s_i}^-$ stores corresponding histograms of the number of spikes detected on channel j in a time interval of length t_w during which the system encodes a state *other than* s_i . As described in the Methods Section under ‘Overall Scheme of Operation,’ a template g_{s_k} consist of a set of n thresholds, corresponding to the minimum number of spikes that must be detected on each channel during a single time window in order to meet the likelihood criterion for state s_i .

We use the following heuristic for setting these thresholds, designed to yield decoding performance meeting at least a tunable minimum quality. We begin by setting two parameters, $0 \leq t_s, t_p \leq 1$, representing a global minimum sensitivity and a global minimum positive predictive value, respectively. We then consider each pair of histograms, $H_{s_i,j}^\pm$, and for every possible threshold value from 0 to b_s we compute the sensitivity and positive predictive value with which the threshold discriminates between s_i and \bar{s}_i , where the overbar denotes the set complement (\bar{s}_i signifies *not- s_i* , any state other than s_i). The template value $g_{s_i,j}$ is set to the lowest threshold value yielding a discrimination sensitivity greater than or equal to t_s and simultaneously a positive predictive value greater than or equal to t_p . In the reduced implementations described in the Methods Section under ‘Overall Scheme of Operation,’ t_s and t_p can be tuned either globally or on a state-by-state basis until no template contains more than n_t predictive thresholds.

FPGA Implementation

Figure 2 shows our architecture for implementing the neural decoding algorithm described in the Methods Section under ‘Overall Scheme of Operation.’ As a practical demonstration of our computational architecture for neural decoding, we implemented our system in a low-power, $6 \times 6 \times 1\text{-mm}^3$, AGL060 FPGA of the Actel IGLOO Low-Power Flash family, which is tightly constrained by a number of parameters [59]. The measured power consumption of the decoding architecture, obtained as the difference between total power consumption of the FPGA (the product of operating current and the 1.2-V supply voltage) when it is programmed to contain or exclude the neural decoding architecture, is $537 \mu\text{W}$. We note that the same architecture, implemented in an advanced-process application-specific integrated circuit (ASIC), should consume significantly less power, as even the lowest-power FPGAs are power-inefficient in comparison with custom ASICs.

Our system accepts input from $n = 32$ channels carrying digitized neural data from an array of recording electrodes and neural amplifiers [21]. The signal on each channel is thresholded to detect spikes, as described in the Methods Section under ‘Spike Detection,’ using a programmable comparator. Threshold-crossing events on each channel are registered by a ($b_c = 4$)-bit windowed counter, and the counters on all channels reset synchronously every $t_w = 100$ ms. The set of counter values registered at the end of each time window constitute the input spike-count vector $\vec{h}(t)$ described in the Methods Section under ‘Overall Scheme of Operation.’

As described in the Methods Sections under ‘Overall Scheme of Operation’ and ‘Computation of Templates,’ the decoding algorithm functions by evaluating a set of m logical expressions at each time step. The elements of these expressions are drawn from $\vec{h}(t)$ and the m templates, g_{s_i} ; the templates thus implement programmable rules in combinational logic that are stored in the memory of the FPGA. In order to compress the decoding algorithm into the limited memory of our FPGA, we implement a reduced version of the template matching scheme, as described in the Methods Sections under ‘Overall Scheme of Operation,’ with $n_t = 2$ rules per state; the thresholds and corresponding pointers are stored in RAM.

In our implementation we use $m = 32$ reduced templates to decode $n_s = 32$ states. At each time step, for each template, the $n_t = 2$ counters identified by the template pointers are compared to the corresponding thresholds stored in the template. The logical results of these comparisons are stored in 1-bit shift registers, and each component of the decoder output, $d_{s_k}(t)$, is formed from the logical *and* of the associated $n_t = 2$ registers. The m components of $\vec{d}(t)$ are saved in an ($n = 32$)-bit shift register, and $\vec{d}(t)$ is transmitted every t_w , after all m templates have been applied.

Trajectory Smoothing Algorithm

Our decoding architecture separately implements routines based purely on *spike-to-state* correlation statistics, and routines based on physical or other constraints on *state-to-state* transitions, independent of neuronal activity. This separation permits us to prioritize the former, the decoding computations most closely associated with raw neural activity and most critical for data compression, for execution within the implanted component of the system. The latter routines, described in this section, can then operate on the compressed data, outside the body. In the context of an implanted system with a small power budget, this scheme conserves power that would otherwise be spent in telemetry, by greatly reducing the bandwidth required for data transmission; and in computation, by offloading some of the decoding to the external components of the system. It also preserves algorithmic flexibility, as spike-to-state decoding, based on a generalized notion of neuronal receptive fields, is far less context-dependent than state-to-state decoding, which may employ a very broad range of constraints and context-specific statistical priors. For example, the priors employed to constrain letter-to-letter and word-to-word transitions in a typing task are very different from those used to smooth the trajectory of a prosthetic limb, yet both decoding problems can be addressed at the spike-to-state level by a pattern matching algorithm of the kind we describe

here. The external unit, more accessible and less subject to power and size constraints, is better suited to tasks requiring more intense computation and frequent reprogramming.

We applied our decoding system to a problem involving the decoding of continuous movement, as described in the Results Section under ‘Neural Decoding,’ and so the state-to-state component of our decoding algorithm involves reconstruction of smooth trajectories from the spike-to-state predictions of the decoding algorithm described in the Methods Section under ‘Neural Decoding Algorithm’ and ‘FPGA Implementation.’

The classic Viterbi algorithm [60] is an efficient, optimal, recursive approach to estimating the sequence of hidden states in a Markov process observed in noise. In casting our decoding problem in terms of a discrete-time, finite-state Markov process, we can construct trajectories by treating discretized positions as states in a hidden Markov model, and applying a modified Viterbi algorithm to convert the decoded output into a maximum-likelihood trajectory.

In the context of our decoder, the observed states are the values generated at the output of the decoder, $\vec{d}(t)$, at each time step, t . The actual positions $s(t) \in \{s_i\}$ at the corresponding times are the associated hidden states. Implementing the Viterbi algorithm requires recursively computing

$$\hat{s}(t) = \max_i P\left(\vec{d}(t)|s_i(t)\right) P(s_i(t)|s(t-1)), \quad (5)$$

where $\hat{s}(t)$ is the optimal estimate for the trajectory position at time step t . In the following paragraphs we explain the other terms in Equation 5, as well as our approaches to computing them.

The first term in Equation 5 denotes the probability of obtaining the decoder output $\vec{d}(t)$ when the corresponding, true position is $s_i(t)$. We use the following method to estimate these conditional probabilities. As the decoder output $\vec{d}(t)$ consists of m single-bit indicators reflecting the likelihood that $s(t) = s_i$, we treat each component $d_i(t)$ of $\vec{d}(t)$ as being independent of the other $m - 1$. Under this assumption of independence,

$$P\left(\vec{d}(t)|s_i(t)\right) \approx \prod_i P\left(d_j(t)|s_i(t)\right), \quad (6)$$

where the index j in Equation 6 runs over all components i of $\vec{d}(t)$ for which $d_i(t) = 1$. We estimate the probabilities forming the individual terms in the product of Equation 6 empirically, by computing a confusion matrix, C , describing the performance of the decoder during its learning period. The element c_{ij} of this confusion matrix is the ratio of the number of time windows during which the decoder generated $d_i(t) = 1$ when the correct state was s_j , to the total number of time windows during which the decoder generated $d_i(t) = 1$; hence, c_{ij} approximates $P\left(d_j(t)|s_i(t)\right)$.

The second term in Equation 5 represents a probability of transition from one state to the next; in our system, these transitions correspond to movements from one position to another. Such movements must obey physical constraints, and so the most probable transitions at adjacent time steps are those between each position and itself, and between a given position and its nearest neighbors. We therefore set the prior probability of transition between positions s_i and s_j according to the diffusion-like expression in Equation 7:

$$P(s_i(t)|s_j(t-1)) \equiv \frac{1}{Z_j} e^{-\alpha \frac{(\Delta(s_i, s_j))^2}{\Delta t}} \quad (7)$$

$$Z_j \equiv \sum_{i=1}^{n_s} e^{-\alpha \frac{(\Delta(s_i, s_j))^2}{\Delta t}}, \quad (8)$$

where $\Delta(s_i, s_j)$ denotes the physical distance between the indicated states, and Z_i is the normalization constant for transitions from s_j to s_i . The inverse proportionality to Δt , the time interval since the decoder generated at least one nonzero bit, permits transitions between more distant states with increased probability as time elapses, by broadening the probability distribution with elapsed time between informative observations; α is a constant, related in our system to mean speed of movement, that tunes the rate of this spreading. We set $\alpha = 0.85$ to obtain the results shown in Figure 5.

Acknowledgments

The authors would like to thank Thomas J. Davidson, formerly of the Wilson Laboratory, Picower Institute for Learning and Memory, MIT Department of Brain and Cognitive Sciences, who provided behavioral neurophysiologic data used in the present work.

References

1. Patil PG (2009) Introduction: Advances in brain-machine interfaces. *Neurosurgical Focus* 27: E1.
2. Nicolelis MAL (2003) Brain-machine interfaces to restore motor function and probe neural circuits. *Nature Reviews Neuroscience* 4: 417-422.
3. Schwartz A, Cui XT, Weber D, Moran D (2006) Brain-controlled interfaces: Movement restoration with neural prosthetics. *Neuron* 56: 205-220.
4. Hochberg LR, Serruya MD, Friehs GM, Mukand JA, Saleh M, et al. (2006) Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature* 442: 164-171.
5. Chader GJ, Weiland J, Humayun MS (2009) Artificial vision: needs, functioning, and testing of a retinal electronic prosthesis. *Neurotherapy: Progress in Restorative Neuroscience and Neurology* 175: 317-332.
6. Shire DB, Kelly SK, Chen JH, Doyle DP, Gingerich MD, et al. (2009) Development and implantation of a minimally invasive wireless subretinal neurostimulator. *IEEE Transactions on Biomedical Engineering* 56: 2502-2511.
7. Tarsy D, Vitek J, Starr P, Okun M, editors (2008) *Deep brain stimulation in neurological and psychiatric disorders*. Current Clinical Neurology. Totowa: Humana Press.
8. Salanova V, Worth R (2007) Neurostimulators in epilepsy. *Current Neurology and Neuroscience Reports* 7: 315-319.
9. Narasimhan K (2004) Scaling up neuroscience. *Nature Neuroscience* 7: 425.
10. Eliasmith C, Anderson CH (2003) *Neural Engineering*, MIT Press, chapter 4. pp. 112-115.
11. Rapoport B, Sarpeshkar R (2011) A Biomimetic Adaptive Algorithm and Micropower Circuit Architecture for Implantable Neural Decoders, IGI Global, chapter 10. pp. 216-254. In *System and Circuit Design for Biologically-Inspired Learning*, Edited by T. Temel.
12. Holleman J, Otis B (2007) A sub-microwatt low-noise amplifier for neural recording. 29th Annual International Conference of the IEEE Engineering in Medicine and Biology Society : 3930-3933.
13. Wattanapanitch W, Fee M, Sarpeshkar R (2007) An energy-efficient micropower neural recording amplifier. *IEEE Transactions on Biomedical Circuits and Systems* 1: 136-147.

14. Yang H, Sarpeshkar R (2006) A bio-inspired ultra-energy-efficient analog-to-digital converter for biomedical applications. *IEEE Transactions on Circuits and Systems I (Special Issue on Life Sciences and System Applications)* 53: 2349-2356.
15. Ghovanloo M, Atluri S (2007) A wideband power-efficient inductive wireless link for implantable microelectronic devices using multiple carriers. *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications* 54: 71-77.
16. Baker MW, Sarpeshkar R (2007) Feedback analysis and design of RF power links for low-power bionic systems. *IEEE Transactions on Biomedical Circuits and Systems* 1: 28-38.
17. Mandal S, Sarpeshkar R (2008) Power-efficient impedance-modulation wireless data links for biomedical implants. *IEEE Transactions on Biomedical Circuits and Systems* 2: 301-315.
18. Mohseni P, Najafi K, Eliades SJ, Wang X (2005) Wireless multichannel biopotential recording using an integrated fm telemetry circuit. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 13: 263-271.
19. Theogarajan LS (2008) A Low-Power Fully Implantable 15-Channel Retinal Stimulator Chip. *IEEE Journal of Solid-State Circuits* 43: 2322-2337.
20. Kelly SK, Wyatt JL (2011) A power-efficient neural tissue stimulator with energy recovery. *IEEE Transactions on Biomedical Circuits and Systems* 5: 20-29.
21. Sarpeshkar R, Wattanapanitch W, Arfin SK, Rapoport BI, Mandal S, et al. (2008) Low-power circuits for brain-machine interfaces. *IEEE Transactions on Biomedical Circuits and Systems* 2: 173-183.
22. Sarpeshkar R (2010) *Ultra Low Power Bioelectronics: Fundamentals, Biomedical Applications, and Bio-Inspired Systems*. Cambridge University Press.
23. III RHO, Wise KD (2005) A three-dimensional neural recording microsystem with implantable data compression circuitry. *IEEE Journal of Solid-State Circuits* 40: 2796-2804.
24. B I Rapoport RS, Wattanapanitch W, Arfin SK, Rapoport BI, Mandal S, et al. (2009) A biomimetic adaptive algorithm and low-power architecture for implantable neural decoders. *Proceedings of the 31st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC09)* .
25. Warland D, Reinagel P, Meister M (1997) Decoding visual information from a population of retinal ganglion cells. *Journal of Neurophysiology* 78: 2336-2350.
26. Serruya M, Hatsopoulos NG, Fellows M, Paninski L, Donoghue J (2003) Robustness of neuroprosthetic decoding algorithms. *Biological Cybernetics* 88: 219-228.
27. Sanchez JC, Erdogmus D, Principe JC, Wessberg J, Nicolelis MAL (2005) Interpreting spatial and temporal neural activity through a recurrent neural network brain machine interface. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 13: 213-219.
28. Wu W, Black MJ, Mumford D, Gao Y, Bienenstock E, et al. (2004) Modeling and decoding motor cortical activity using a switching Kalman filter. *IEEE Transactions on Biomedical Engineering* 51: 933-942.
29. Wu W, Gao Y, Bienenstock E, Donoghue JP, Black MJ (2006) Bayesian population decoding of motor cortical activity using a Kalman filter. *Neural Computation* 18: 80-118.

30. Wu W, Shaikhouni A, Donoghue JP, Black MJ (2004) Closed-loop neural control of cursor motion using a Kalman filter. *Proceedings of the IEEE Engineering in Medicine and Biology Society* : 4126-4129.
31. Wu W, Hatsopoulos NG (2008) Real-time decoding of nonstationary neural activity in motor cortex. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 16: 213-222.
32. Brockwell AE, Rojas AL, Kass RE (2004) Recursive bayesian decoding of motor cortical signals by particle filtering. *Journal of Neurophysiology* 91: 1899-1907.
33. Brown EN, Frank LM, Tang D, Quirk MC, Wilson MA (1998) A statistical paradigm for neural spike train decoding applied to position prediction from ensemble firing patterns of rat hippocampal place cells. *Journal of Neuroscience* 18: 7411-7425.
34. Eden UT, Frank LM, Barbieri R, Solo V, Brown EN (2004) Dynamic analyses of neural encoding by point process adaptive filtering. *Neural Computation* 16: 971-998.
35. Srinivasan L, Eden UT, Mitter SK, Brown EN (2007) General-purpose filter design for neural prosthetic devices. *Journal of Neurophysiology* 98: 2456-2475.
36. Sarma SV, Eden UT, Cheng ML, Williams ZM, Hu R, et al. (2010) Using point process models to compare neural spiking activity in the subthalamic nucleus of parkinsons patients and a healthy primate. *IEEE Transactions on Biomedical Engineering* 57: 1297-1305.
37. Pesaran B, Pezaris JS, Sahani M, Mitra PP, Andersen RA (2002) Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nature Neuroscience* 5: 805-811.
38. Musallam S, Corneil BD, Greger B, Scherberger H, Andersen RA (2004) Cognitive control signals for neural prosthetics. *Science* 305: 258-262.
39. Santhanam G, Ryu SI, Yu BM, Afshar A, Shenoy KV (2006) A high-performance brain-computer interface. *Nature* 442: 195-198.
40. Shenoy KV, Meeker D, Cao SY, Kureshi SA, Pesaran B, et al. (2003) Neural prosthetic control signals from plan activity. *Neuroreport* 14: 591-596.
41. Koyama S, Chase SM, Whitford AS, Velliste M, Schwartz AB, et al. (2010) Comparison of brain-computer interface decoding algorithms in open-loop and closed-loop control. *Journal of Computational Neuroscience* 1-2: 73-87.
42. Louie K, Wilson MA (2001) Temporally Structured Replay of Awake Hippocampal Ensemble Activity during Rapid Eye Movement Sleep. *Neuron* 29: 145-156.
43. Lee AK, Wilson MA (2002) Memory of Sequential Experience in the Hippocampus during Slow Wave Sleep. *Neuron* 36: 1183-1194.
44. Madisetti V, editor (2009) *The Digital Signal Processing Handbook*. Electrical Engineering Handbook. Boca Raton, Florida: CRC Press, 2nd edition.
45. Wolfram S (2002) *A New Kind of Science*. Champaign, Illinois: Wolfram Media.
46. O'Keefe J, Nadel L (1978) *The Hippocampus as a Cognitive Map*. Oxford, England: Clarendon Press.
47. Best PJ, White AM, Minai A (2001) Spatial processing in the brain: The activity of hippocampal place cells. *Annual Review of Neuroscience* 24: 459-86.

48. McNaughton BL, Battaglia FP, Jensen O, Moser FI, Moser MB (2006) Path integration and the neural basis of the 'cognitive map'. *Nature Reviews Neuroscience* 7: 663-678.
49. Davidson TJ (2009) Replay of memories of extended behavior in the rat hippocampus. Doctoral Thesis, Massachusetts Institute of Technology. Department of Brain and Cognitive Sciences.
50. Hochberg LR, Serruya MD, Friehs GM, Mukand JA, Saleh M, et al. (2006) Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature* 442: 164-171.
51. Paninski L, Fellows MR, Hatsopoulos NG, Donoghue JP (2004) Spatiotemporal tuning of motor cortical neurons for hand position and velocity. *Journal of Neurophysiology* 91: 515-532.
52. Pouget A, Dayan P, Zemel RS (2003) Inference and computation with population codes. *Annual Reviews of Neuroscience* 26: 381-410.
53. Stark E, Abeles M (2007) Predicting movement from multiunit activity. *Journal of Neuroscience* 27: 8387-8394.
54. Wolf PD (2008) *Indwelling Neural Implants: Strategies for Contending with the In Vivo Environment*, CRC Press, chapter 3: Thermal Considerations for the Design of an Implanted Cortical BrainMachine Interface (BMI).
55. Nicolelis MAL, Dimitrov D, Carmena JM, Crist R, Lehw G, et al. (2003) Chronic, multisite, multielectrode recordings in macaque monkeys. *Proceedings of the National Academy of Sciences of the United States of America* 100: 11041-11046.
56. Velliste M, Perel S, Spalding MC, Whitford AS, Schwartz AB (2008) Cortical control of a prosthetic arm for self-feeding. *Nature* 453: 1098-1101.
57. Clanton ST, Zohny Z, Velliste M, Schwartz AB (2010) Simultaneous 7-dimensional cortical control of an arm and hand robot via direct brain interface. In: *Society for Neuroscience Annual Meeting (Neuroscience 2010) Abstracts*. San Diego, California: The Society for Neuroscience.
58. Rizk M, Wolf PD (2009) Optimizing the Automatic Selection of Spike Detection Thresholds Using a Multiple of the Noise Level. *Medical and Biological Engineering and Computing* 47: 955-966.
59. Corporation A. IGLOO Low Power Flash FPGAs with Flash*Freeze Technology. Data Sheet.
60. Forney GD (1973) The viterbi algorithm. *Proceedings of the IEEE* 61: 268-278.

Figure Legends

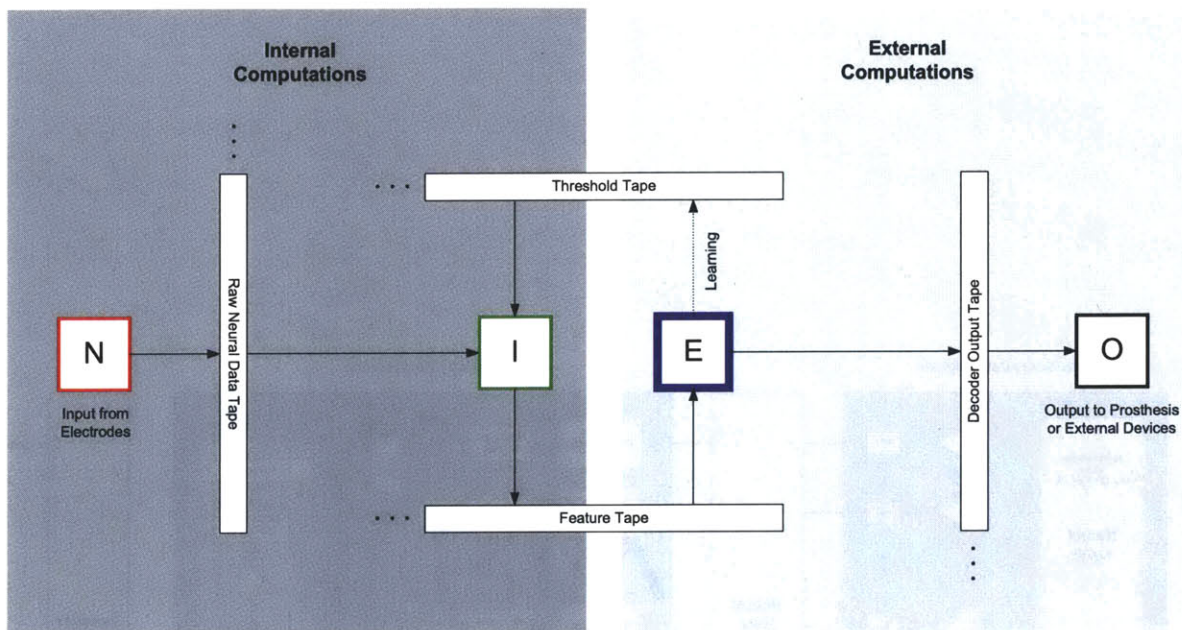


Figure 1. Universal Computing Architecture for Neural Decoding. The overall architecture of a neural decoding system is decomposed into a set of operations implemented by Turing-type computing machines, shown here as a collection of *heads* (data processing units) reading from and writing to a set of corresponding *tapes* (programs and data streams). Amplification and digitization of raw neural data, and decoding of that data, are performed by heads *N* and *I*, respectively, in a biologically implanted unit. The ‘Internal Computations’ of these two system components are streamed across a wireless data channel to an external unit, which performs more power-intensive ‘External Computations’ to post-process the decoded output. Further processing of the decoded data is performed externally by head *E*, and the final output of the system is reported by head *O*. The external system implements a learning algorithm that is used to write the program on the threshold tape, which is executed by the internal unit.

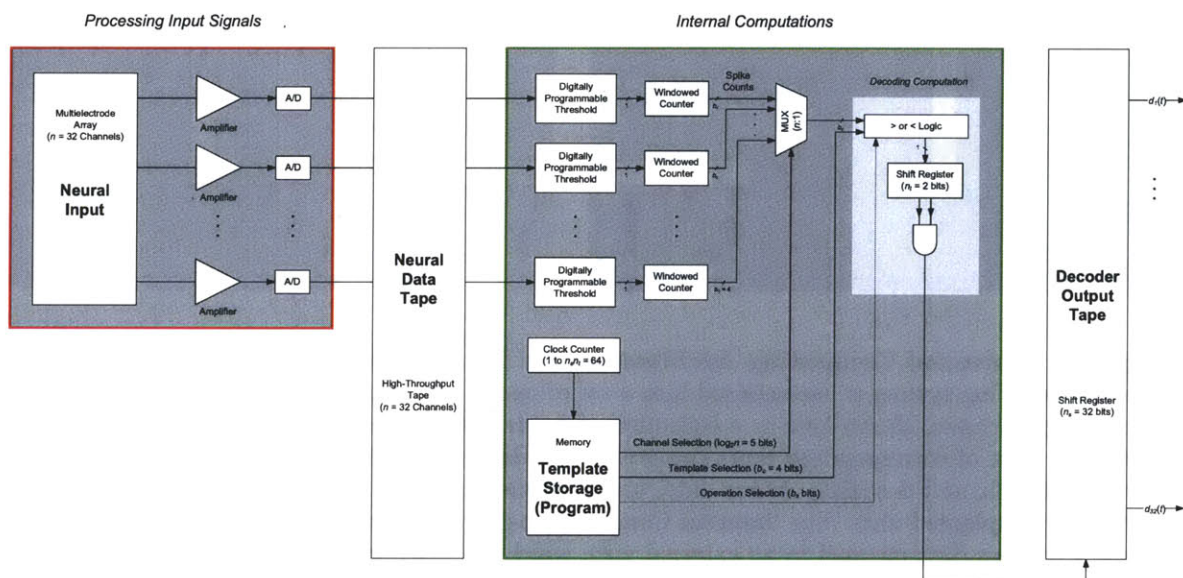


Figure 2. Decoding Architecture. Block diagram of the low-power processing system of the internal component of our neural decoder, as implemented in one instantiation of our architecture. Functional blocks are color-coded in accord with the scheme used in Figure 1.

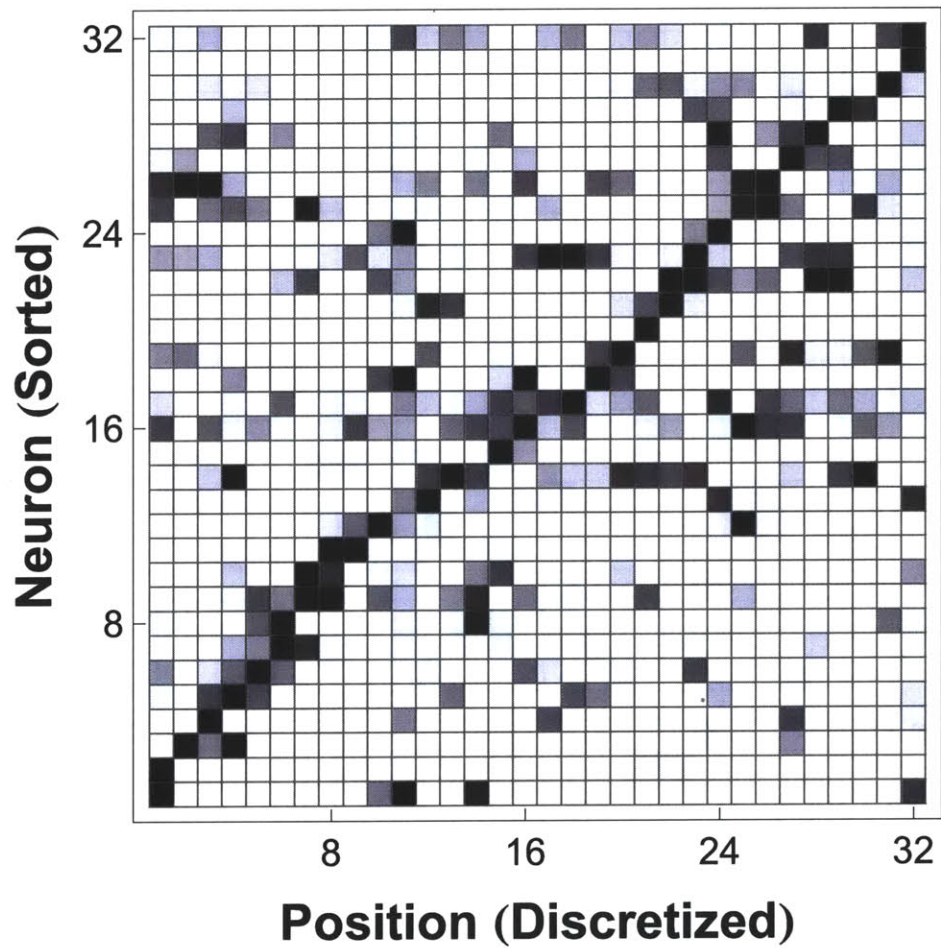


Figure 3. Encoding of Position by Place Cell Receptive Fields. Normalized spike rate for each of $n = 32$ neurons in $n_s = m = 32$ equal-length intervals along a one-dimensional track maze. Neurons (rows) have been sorted according to their positions of maximal activity to illustrate that the receptive fields of the place cells in this population cover the one-dimensional space of interest. (Black: Maximal Spike Rate, White: Zero Spike Rate, Gray: Intermediate Spike Rates.)




Figure 4. Decoder Logic Program: Finite-State Automaton Rules for Neural Decoding.

Decoding template array g stored in system memory, resulting in the output shown in Figure 5, using the $n_t = 2$ most informative threshold values for each position state. (White: Unused, Light Gray: 1 Spike per 1440-ms Window, Black: 7 Spikes per 1440-ms Window.) Intuitively, this set of templates can be understood as the tape-reading rules for a Turing machine, whose symbols are generated by the time-windowed spike counts on neural input channels, and whose states correspond to a discretized set of position states encoded by the underlying neuronal populations. At each time step, the neural decoder scans down each column in the array to determine the states, if any, whose rules have been satisfied; the decoded output elements $d_i(t)$ are set to 1 for those states, and to 0 otherwise. The rules displayed graphically in the rectangular array are encoded numerically in the table displayed above the array. The columns of the table are aligned with the states in the array for which they contain decoding data, comprising the indices of the two most informative channels, C_1 and C_2 , and the corresponding spike thresholds, θ_1 and θ_2 .

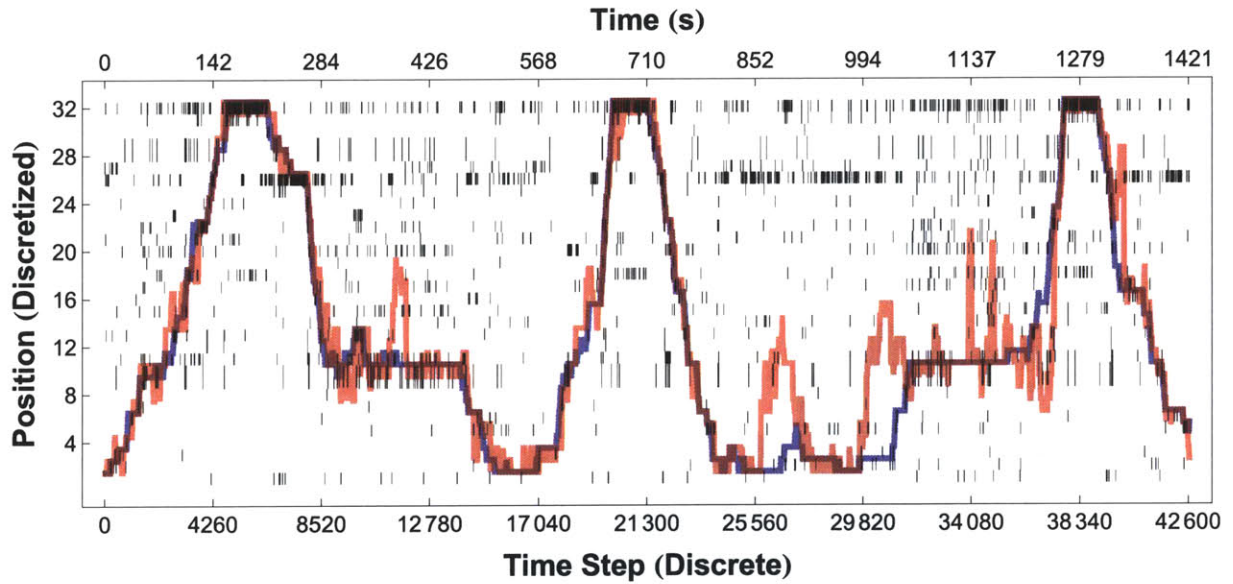


Figure 5. Decoder Output when Decoding Position from Hippocampal Place Cells. Our system decodes the location of a maze-roaming rat, from spike trains recorded from thirty-two hippocampal place cells. Raw output of the decoding algorithm, $\vec{d}(t)$, is shown as a raster array, with the output at each time step displayed as a vertical column of pixels (black pixels correspond to $\vec{d}_i(t) = 1$, white pixels to $\vec{d}_i(t) = 0$). Red lines show the trajectories obtained after applying our Viterbi algorithm to the raw decoder output, as described in the text. The actual trajectories of the rat are shown in blue. Decoding accuracy and decoder noise are affected by the length of the time window over which spikes are collected at each time step: $t_w = 2^{b_w=4} \times 90$ ms ($t_w = 1440$ ms). Here the decoded trajectory matches the actual trajectory with a correlation coefficient of 0.91.




Figure 6. Histograms and Threshold. Histograms collected during the training phase of the decoding algorithm facilitate computation of thresholds for windowed spike activity, which are stored as templates in memory and used to discriminate between states. This histogram of spike activity, collected from recording channel 20, demonstrates that a threshold of $g_{29,20} = 7$ spikes per 16×90 -ms window, on recording channel 20, is sensitive and specific for state s_{29} (Sensitivity: 0.63, Specificity: 0.85, Positive Predictive Value: 0.86). This threshold is written on the threshold tape used to program the internal unit of the decoder, and can be seen numerically in Figure 4 as the C_1 and θ_1 row entries of column 29, and graphically as the corresponding pixels in the rule array.

Tables

Table 1. Basic Operations per Computational Frame of the Decoding Architecture

Operation	Instances per Frame
Clock Counter ($\log_2(n_s n_t) = 6$ bits)	1
Memory Access	1
Multiplexer ($n = 32$): 1	1
Comparison (1 bit)	$\beta_L = 1$
Binary Logic (<i>and</i>)	$\frac{1}{n_t} = \frac{1}{2}$
Shift Register ($n_t = 2$ bits)	1
Shift Register ($n_t = 2$ bits)	1
Total	$\beta = 5 + \beta_L + \frac{1}{n_t} = \frac{13}{2}$

Detailed accounting of the basic operations required in each computational frame of neural decoding, corresponding to the system block diagram of Figure 2. The parameter β denotes the total number of basic operations per frame. The parameter β_L denotes the number of basic operations associated with the combinational logic used to compare incoming spike-count vectors to each stored template.

Chapter 4

Bioimplantable Glucose-Based Fuel Cells

4.1 Overview

This chapter presents work toward the development of a *brain-powered brain-machine interface*, with emphasis on the design, fabrication, and testing of a glucose-based implantable fuel cell.

4.2 An Implantable Glucose Fuel Cell for Brain–Machine Interfaces

This section consists of a manuscript in preparation for submission to *PLoS ONE* [61].

An Implantable Glucose Fuel Cell for Brain–Machine Interfaces

Benjamin I. Rapoport^{1,2,3,*}, Jakub T. Kedzierski^{2,**}, Rahul Sarpeshkar^{1,***}

1 Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, Massachusetts, United States of America

2 Advanced Silicon Technology Group, Lincoln Laboratory, Massachusetts Institute of Technology, Lexington, Massachusetts, United States of America

3 M.D.–Ph.D. Program, Harvard Medical School, Boston, Massachusetts, United States of America

* E-mail: brapoort@post.harvard.edu

** E-mail: jakub@ll.mit.edu

*** E-mail: rahuls@mit.edu

Abstract

As implantable electronic devices become increasingly prevalent in the diagnosis, management, and treatment of human disease, there is a correspondingly increasing demand for integrated electronic devices with very long functional lifetimes that integrate seamlessly into host biological systems, harvesting power from their biological environments.

To that end, we have developed an implantable fuel cell that generates power through glucose oxidation, producing $3.4 \mu\text{W cm}^{-2}$ steady-state power and up to $180 \mu\text{W cm}^{-2}$ peak power.

The fuel cell is constructed using conventional semiconductor fabrication processes, and so is well suited for manufacture together with traditional integrated circuits on a single silicon wafer, to yield implantable microelectronic systems with long-lifetime, integrated power sources.

The fuel reactions are mediated by robust, solid-state catalysts. Glucose is oxidized at the nanostructured surface of an activated platinum anode. Oxygen is reduced to water at the surface of a self-assembled network of single-walled carbon nanotubes, embedded in a Nafion film, which forms the cathode and is exposed to the biological environment. The catalytic electrodes are separated by a Nafion membrane.

The availability of fuel cell reactants, oxygen and glucose, only as a mixture in the physiologic environment, has traditionally posed a design challenge, as net current production requires oxidation and reduction to occur separately and selectively at the anode and cathode, respectively; glucose oxidation at the cathode and oxygen reduction at the anode generate electrochemical short circuits by abolishing the potential difference that induces electron flow through the external circuit. The fuel cell is configured in a half-open geometry that shields the anode while exposing the cathode, resulting in an oxygen gradient that strongly favors oxygen reduction at the cathode. Glucose reaches the shielded anode by diffusing through the nanotube mesh, which does not catalyze glucose oxidation, and the Nafion layers, which are permeable to small neutral and cationic species.

We propose that the cerebrospinal spinal fluid space surrounding the brain constitutes a physiologic niche particularly well suited for a glucose-based fuel cell, being virtually acellular and under minimal surveillance by the immune system, and yet containing glucose concentrations approximately half of those found in blood plasma. We then demonstrate computationally that the natural recirculation of cerebrospinal fluid around the human brain theoretically permits glucose energy harvesting at a rate on the order of at least 1 mW with no adverse physiologic effects. Low-power brain–machine interfaces can thus potentially benefit from having their implanted units powered or recharged by glucose fuel cells.

Introduction

As implantable electronic devices become increasingly prevalent in the diagnosis, management, and treatment of human disease, there is a correspondingly increasing demand for devices with unlimited functional lifetimes that integrate seamlessly into their host biological systems. Consequently, a holy grail of bio-

electronics is to engineer biologically implantable systems that can be embedded without disturbing their local environments while harvesting from their surroundings all of the power they require. In particular, micropower implantable electronics beg the question of whether such electronics can be powered from their surrounding tissues. Here we discuss how to construct an implantable glucose fuel cell suitable for such applications, and how it may potentially be powered from cerebrospinal fluid in the brain.

Bioimplantable Power Sources

Various solutions to the problem of providing power to biologically implanted devices have been proposed, prototyped, or implemented. Two principal solutions are currently in widespread use: single-use batteries, such as those used in implantable pulse generators for cardiac pacing, defibrillation, and deep brain stimulation, which are designed to have finite lifetimes and to be replaced surgically at intervals of several years [1]; and inductive power transfer, typically accomplished transcutaneously at radio frequencies, as in cochlear implants [2, 3]. Inductive schemes can be used either to supply power continuously or to recharge an implanted power source. Recent advances in battery technology and related fields, leading to increased energy and power densities in small devices such as supercapacitors [4, 5] as well as thin film lithium and thin film lithium ion batteries [6], will facilitate improvements in systems based on these two solutions, particularly by shrinking battery sizes and extending battery lifetimes.

Power Scavenging and Power Requirements for Implantable Electronics

Systems for transducing light [7], heat [8], mechanical vibration [9], as well as near- [3, 10] or far-field [11] electromagnetic radiation, into electrical energy, have been described and implemented. Several of these energy-harvesting techniques, as well as electronic design techniques required to make use of them, have been discussed in depth in [12, 13]. Here we focus on powering biologically implanted devices by harvesting energy from glucose in the biological environment.

The emergence of ultra-low-power bioelectronics as a field over the last decade [13] has led to the development of highly energy-efficient, implantable medical devices with power budgets in the microwatt regime. This new generation of low-power devices has driven interest in a range of sustainable power sources and energy scavenging systems that, while impractical for conventionally designed electronic devices, are entirely practical in the context of micropower electronics [13]. For example, in brain-machine interfaces, the combination of low-power circuit design [14] and adaptive power biasing [15] can be used to build sub-microwatt neural amplifiers for multi-electrode arrays. Impedance-modulation radio-frequency (RF) telemetry techniques can drastically reduce implanted-unit power consumption and operate at less than 1 nJ bit^{-1} even for transcutaneous data rates as high as 3 Mbps in brain-machine interfaces. Finally, ultra-low-power analog processing techniques [13, 16] can enable 100-channel neural decoding at micropower levels [17, 18] and dramatically reduce the data rates needed for communication, further reducing total power consumption. The combination of these advances in energy-efficient amplification, communication, and computation implies that implanted components in brain-machine interfaces that operate with tens of microwatts of total power consumption are feasible today. Therefore, implantable biofuel cells such as the one presented here, which generate power at densities on the order of $1\text{--}100 \mu\text{W cm}^2$, provide useful power sources for such ultra-low-power implantable medical devices.

Glucose Fuel Cells

A fuel cell generates power by catalyzing complementary electrochemical reactions (oxidation and reduction) at a pair of corresponding electrodes (the anode and cathode, respectively), as a reducing-agent fuel flows across the anode and an oxidant flows across the cathode. The fuel substrate is electrooxidized at the anode, which collects the liberated electrons and conducts them through an external load to the

cathode. Typically an ion-selective membrane partitions the anode and cathode into separate chambers, facilitating a unidirectional flow of protons, generated via oxidation at the anode, to the cathode. The protons arriving at the cathode through the solution, the electrons arriving at the cathode through the external circuit, and the oxidant at the cathode undergo a redox reaction that reestablishes charge neutrality in the overall cell.

One approach to harvesting energy from a physiologic environment is to extract it from physiologically available biological fuel substrates such as glucose. In a glucose-based biofuel cell, glucose is oxidized at the anode, while oxygen is reduced to water at the cathode. The nature of the catalyst residing at the anode determines the extent of glucose oxidation and the associated oxidation products. Three major design paradigms for glucose-based fuel cells have been described, differing principally in the materials used to catalyze electrode reactions: In one paradigm the catalysts are abiotic, solid-state materials; in the second paradigm the catalysts are isolated enzymes fixed to electrode substrates; and in the third paradigm oxidation is performed by exoelectrogenic bacterial biofilms colonizing a fuel cell anode. Numerous designs representing each of these fuel cell paradigms have been described and reviewed in the scientific and patent literatures. An extensive review of the scientific and patent literatures on abiotic implantable glucose fuel cells is provided in [19]. Bioimplantable fuel cells based on enzymatic catalysis are reviewed in [20]. Microbial fuel cells are reviewed in [21].

These three principal catalytic schemes yield systems that differ markedly from one another in efficiency and robustness.

Enzyme-based glucose fuel cells have high catalytic efficiency, which together with their small size results in high volumetric power density, yielding up to $8.3 \mu\text{A}$, and $4.3 \mu\text{W}$ of total power, in systems with footprints on the order of 1 mm^2 and volumes less than 10^{-2} mm^3 [20]. Such fuel cells are often constructed as tethered-enzyme systems, in which oxidation and reduction of fuel cell substrates are catalyzed *ex vivo* by enzymes molecularly wired to threads of conductive material. Enzyme-based glucose fuel cells described in the recent literature have typically generated on the order of $100 \mu\text{W cm}^{-2}$ [22,23]. Fuel cells of this kind may be capable of continuous operation for up to several weeks, but their lifetimes are limited by the tendency of enzymes to degrade and ultimately degenerate with time.

Using living microorganisms, such as exoelectrogenic bacteria, to catalyze the anodic reaction results in complete oxidation of glucose, liberating twenty-four electrons per molecule of glucose consumed. Microbial fuel cells are thus very efficient catalytically, and can produce more than $1300 \mu\text{A cm}^{-2}$ and $1900 \mu\text{W cm}^{-2}$ [24,25]. Moreover, in contrast with enzymatic systems, which have short lifetimes limited by the degradation of tethered enzymes *ex vivo*, microbial fuel cells are inherently self-regenerating: they use a fraction of input biomass to power and supply molecular substrates for maintenance functions such as resynthesis of degraded enzymes [21]. Microbial glucose fuel cells described in the recent literature have typically generated on the order of $1000 \mu\text{W cm}^{-2}$ [22,23,26]. However, microbial fuel cells of the present generation are not yet suitable for biologically implanted applications. The prospect of implanting even nonpathogenic bacteria raises concerns of safety and biocompatibility.

Solid-state anode catalysts such as those we use in the work described here are capable of oxidizing glucose to gluconic acid, liberating one pair of electrons, and yielding further oxidation products with reduced probability [19,27]. As a result, they represent the least catalytically efficient of the three design paradigms we consider. Yet while glucose fuel cells based on solid-state catalysts typically only generate on the order of $1\text{--}10 \mu\text{A cm}^{-2}$ and $1\text{--}10 \mu\text{W cm}^{-2}$ [19], they have proven reliable as implantable power sources for many months [28,29].

Cerebrospinal Fluid as a Physiologic Niche Environment for a Fuel Cell

An innovation in the work described here is the use of cerebrospinal fluid as a physiologic niche for an implantable power source. Implantable fuel cells have typically been designed for use in blood or interstitial fluid; to the best of our knowledge, the operation of a biofuel cell in the cerebrospinal fluid has not previously been described (we are aware of one incidental reference, made in [30]). The cerebrospinal

fluid represents a promising environment for an implantable fuel cell: It is virtually acellular, it is under minimal immune surveillance, it has a hundred-fold lower protein content than blood and other tissues and is therefore less prone to induce biofouling of implanted devices, and its glucose levels are comparable to those of blood and other tissues [31]. The bioavailability of glucose and oxygen to a fuel cell residing in the subarachnoid space is addressed in detail in the Methods subsection entitled ‘Brain and Cerebrospinal Fluid as Sites for an Energy-Harvesting Fuel Cell.’

Fuel Cell Design

A variety of mechanical designs for bioimplantable biofuel cells have been described in a literature spanning at least half a century [19]. A number of physical, electronic, and electrochemical factors also influence the voltage and current output of a biofuel cell. These include fuel cell and electrode geometry; electrode and membrane spacings; redox potentials of fuel cell components; internal and load impedances; and environmental conditions in which the cell operates, including fuel substrate concentration, environmental temperature and pH, and the presence of chemical species capable of driving parasitic side reactions. Many of these aspects of fuel cell design have been modeled in detail and measured empirically in real systems [32, 33]. As described in the Methods Section and as shown in Figures 1, 2, and 3, the design we employ in the work described here is a version of a classic half-open, two-chamber design [19, 34, 35], sized and shaped to fit a particular anatomic compartment.




Figure 1. Power Extraction from Cerebrospinal Fluid by an Implantable Glucose Fuel Cell. Conceptual schematic design for a system that harvests power from the cerebrospinal fluid, showing a plausible site of implantation within the subarachnoid space. The inset at right is a micrograph of one prototype, showing the metal layers of the anode (central electrode) and cathode contact (outer ring) patterned on a silicon wafer. *Image Credits: Skull and Subarachnoid Inset, A.D.A.M. Medical Images; Meninges and Vascular Anatomy, Karolinska Institute 3D Brain Project.*

Structure of This Paper

This paper is structured as follows. In the Results and Discussion Section we describe the performance of our glucose fuel cell in detail. In the Methods Section we address several topics in detail. First, we discuss solid-state catalysis of glucose oxidation from a theoretical perspective, and explain the operating principles of our fuel cell, including its mechanism of separating the oxidation and reduction reactions, even though their reactants, glucose and oxygen, naturally occur mixed in physiologic compartments.




Figure 2. General Operational Scheme for an Implantable Glucose Fuel Cell. This schematic conceptually illustrates the structure of an abiotically catalyzed glucose fuel cell, including the essential half-cell and overall reactions, the sites at which they occur within the system, and the flows of reactants and products.

Next, we describe our CMOS-compatible process for fabricating implantable glucose fuel cells. In that subsection we also describe our approach to characterizing the materials and electrochemical properties, as well as the power-generating performance of the fuel cells. We then discuss the power available from circulating glucose in human physiologic compartments, and the suitability of the cerebrospinal fluid as a physiologic niche for power harvesting. We describe a detailed model of glucose and oxygen consumption by a fuel cell implanted in the subarachnoid space surrounding the human brain, and analyze the impact of such a fuel cell on glucose and oxygen homeostasis.

Results and Discussion

Device Characterization

Anode Surface Roughness by Scanning Electron Microscopy

The efficiency of the fuel cell critically depends on its ability to catalyze the oxidation of glucose at the anode. Our device uses a solid-state platinum anode catalyst, whose catalytic capacity is directly related to the number of atomic sites it can provide on its surface. We describe a CMOS-compatible process for electrode surface roughening to increase effective electrode surface area, and hence catalytic capacity, in the Methods Section. Briefly, by alloying the platinum electrode with aluminum, and then reactively etching away all of the aluminum, we generate a high-surface-area anode with a nanostructure




Figure 3. Glucose Fuel Cell in Cross Section. This schematic cross-section of the glucose fuel cell illustrates the structure of the device, as well as the oxygen and glucose concentration gradients crucially associated with its cathode and anode half-cell reactions, and underlying their respective site specificity.

very different from that of atomically smooth platinum. Here, we illustrate the effects of that roughening procedure in a series of micrographs. Figure 4 shows the contrast between atomically smooth and roughened platinum metal layers as patterned on a silicon dioxide wafer substrate, as seen under optical microscopy. Figure 5 shows a series of scanning electron micrographs, taken at increasing magnifications (from single-micrometer to single-nanometer resolution), and showing the persistence of pore-like structures at every scale, as expected on removing one element of a bimetallic alloy. The scanning electron micrographs in Figure 6 demonstrates the contrast between atomically smooth platinum and roughened platinum at the nanometer scale.

Anode Surface Roughness by Atomic Force Microscopy

We used atomic force microscopy to quantify the surface area enhancements generated by the anode roughening procedure we employed, using a high-aspect-ratio silicon nitride probe with radius of curvature 5 nm. Figure 7 shows two pairs of surface scans, contrasting atomically smooth platinum with roughened platinum at $10\mu\text{m}$ and $1\mu\text{m}$ resolution. While this technique is limited both by the geometry of the scanning probe and by the single-axis nature of the measurements, variance in surface height nevertheless provides a quantitative means of assessing surface roughness. Sampling at 256 sites over $10\mu\text{m}$ square patches, we measured a variance of 10.3 nm in the roughened anodes, as compared with a variance of 1.2 nm in atomically smooth platinum (an 8.5-fold increase in z -roughness).

Power Output and Electrochemical Characterization

Power Output

We characterized the glucose fuel cell by obtaining a polarization curve, as shown in Figure 8. Using the fuel cell as a current source by operating a potentiostat in controlled-current mode, we drew current from the fuel cell at a rate increasing from zero by $10\mu\text{Acm}^{-2}\text{ s}^{-2}$, until the cell potential was abolished.




Figure 4. Anode Roughening. These optical micrographs illustrate the effect of the roughening technique on the anode surface, showing the atomically smooth platinum traces (narrow metallic strips) leading to the anode in contrast with the anode itself (large rectangular area). The roughness of the anode surface is detectable optically as an abrupt change in color and texture. The wire traces are $100\ \mu\text{m}$ wide. The image at right is an enlargement of the central region of the image at left, focusing on the boundary between the smooth and rough platinum surfaces (rotated with the wire trace set vertical).

Figure 8 shows the cell voltage and power output generated during this procedure. The fuel cell generated an open-circuit voltage of $192\ \text{mV}$, and generated peak power of $> 180\ \mu\text{W cm}^{-2}$ when sourcing $1.5\text{--}1.85\ \text{mA cm}^{-2}$. Deep discharges at currents greater than $1.85\ \text{mA cm}^{-2}$ were found to damage the fuel cells irreversibly.

Techniques that exploit resonant transformer action can be used to efficiently convert even $20\ \text{mV}$ energy-harvesting outputs to $1\text{--}5\ \text{V}$ levels needed for powering electronic chips [36]. In an actual brain-machine interface, such techniques will be needed to increase the voltage level of our glucose fuel cell as well.

Impedance Matching for Optimizing Power Output

We determined the steady-state power output of the glucose fuel cell when driving a range of resistive loads. The results are shown in Figure 9 for a $1\ \text{mm}^2$ device, which achieved maximum steady-state power output of $3.4\ \mu\text{W cm}^{-2}$ when driving a load of $550\ \text{k}\Omega$.

Power Output as a Function of Glucose Concentration

We tested our fuel cells in standard $10\ \text{mM}$ phosphate buffered saline at $\text{pH } 7.4$ to simulate the physiologic environment of the cerebrospinal fluid, loading the medium with glucose at various concentrations. The performance data provided earlier in this section was obtained in the context of $10\ \text{mM}$ glucose [37]. As expected, and as quantified in Figure 10, the anodic current generated by the fuel cell varies as a function of ambient glucose concentration. Figure 10 shows segments from a set of cyclic voltammograms, swept through oxidizing potentials at $0.1\ \text{V s}^{-1}$ while the anode was kept in media containing glucose at concentrations from 0 to $500\ \text{mM}$. The increases in oxidation current with concentration at physiologic glucose levels, reflected by the areas of the glucose oxidation peaks in the region of $700\ \text{mV}$, is sublinear, confirming that fuel cell performance is not primarily substrate-limited at these concentrations.





Figure 5. Anode Micro- and Nanostructure. This set of scanning electron micrographs, taken of a fuel cell anode at increasing levels of magnification (as indicated by the scale bars in each image), illustrates the effects of the roughening procedure on electrode surface structure over a hierarchy of length scales from nanometers to micrometers.

Methods

Solid-State Catalysis of Glucose Oxidation

The electrochemical reaction mechanisms of direct glucose fuel cells are discussed in detail by Kerzenmacher and colleagues in their thorough review of energy harvesting by implantable, abiotically catalyzed glucose fuel cells [19]. The complete oxidation of glucose to carbon dioxide and water is associated with the transfer of 24 electrons per molecule of glucose, as described by the reactions in Equation 1:



However, studies have confirmed that glucose is not completely oxidized when its oxidation is mediated by traditional, solid state catalysts; the theoretical maximum rate of electron transfer from glucose oxidation





Figure 6. Nanostructural Effects of Surface Roughening. This pair of scanning electron micrographs, taken at the same level of magnification, illustrates the effects of our surface roughening technique. The image at left is a high-magnification image of atomically smooth platinum deposited by evaporation on silicon dioxide. The image at right is taken from one of our roughened anodes, and shows the highly porous nanostructure of the electrode.

is not achieved in abiotically catalyzed glucose fuel cells. Instead, glucose is principally oxidized to gluconolactone in the reaction described by Equation 2, which transfers only a single pair of electrons:



Gluconolactone is typically then hydrolyzed to form gluconic acid. In principle, solid-state catalysts are capable of oxidizing gluconolactone further, and high-pressure liquid chromatography has detected tartaric and oxalic acids in such systems (resulting from 14- and 22-electron-transfer processes, respectively), among other oxidation products. In practice, the mean number of electrons transferred per molecule of glucose oxidized depends on the nature of the catalyst and on thermodynamic properties of the system, such as ambient pH. Gebhardt and colleagues have shown that the mean number of electrons transferred per molecule of glucose by Raney-type catalysts can be up to 17 [27]. As this family of high-surface-area catalysts has also been shown to generate current densities an order of magnitude greater than smooth noble metal catalysts, we chose to use Raney platinum anodes to catalyze glucose oxidation in our fuel cells.

Separation of Anode and Cathode Reactions

Maintenance of a net potential difference between the anode and the cathode of a fuel cell requires separation of the oxidation and reduction half-reactions in a way that restricts each to only one of




Figure 7. Atomic Force Microscopic Measurements of Anode Surface Roughness. Atomic force microscopic measurements of z -direction (plane-normal) surface roughness at the $10\mu\text{m}$ and $1\mu\text{m}$ scales (upper and lower images, respectively), comparing atomically smooth platinum (left images) with the roughened anodes we describe here (right images).

the fuel cell electrodes. As discussed in [19] and elsewhere, biologically implantable fuel cells pose a particular design challenge in that the fuel (glucose) and the oxidant (oxygen) must be extracted from same physiologic fluid, in which both are dissolved. By contrast, fuel cells are traditionally configured so as to isolate the anode from the cathode fluidically, in a compartmental arrangement that separates the two half-cell reactions.

The traditional configuration permits delivery of fuel to the anode and oxidant to the cathode via separate fluidic channels, an arrangement that can be effective because it physically prevents reduction (typically of oxygen) from occurring at the anode, and prevents oxidation of the fuel substrate from occurring at the cathode. These reverse reactions cause electrochemical short circuits by allowing both oxidation and reduction to occur at each electrode, eliminating the net potential difference across the fuel cell electrodes.

Preventing electrochemical short circuits in an implantable glucose fuel cell requires a different ap-




Figure 8. Fuel Cell Polarization Curve. The performance of the fuel cell is characterized through its output voltage (blue, left axis) and power density (red, right axis) as functions of output current density. A 2 mm^2 device exhibits an open-cell voltage of 192 mV and achieves maximum power output of $> 180 \mu\text{W cm}^{-2}$ when sourcing $1.5\text{--}1.85 \text{ mA cm}^{-2}$.

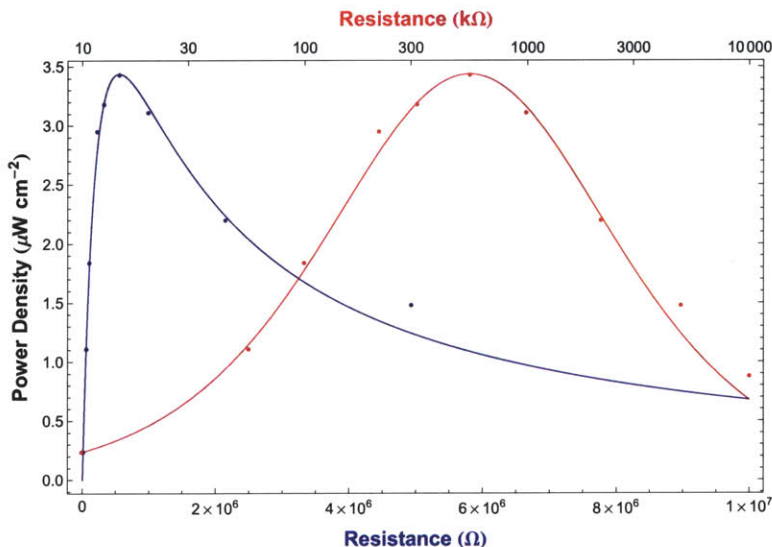


Figure 9. Impedance Matching to Maximize Output Power. Steady-state output power of the fuel cell exhibits characteristic second-order dependence on the magnitude of the resistive load. A 1 mm^2 device achieves maximum steady-state power output of $3.4 \mu\text{W cm}^{-2}$ when driving a load of $550 \text{ k}\Omega$. (Blue curve, lower horizontal axis, linear scale; red curve, upper horizontal axis, logarithmic scale.)

proach to restricting the fuel oxidation and oxygen reduction to the anode and cathode, respectively. Several designs are reviewed in [19]. We have adopted a modified version of a design first proposed by von Sturm and colleagues [34, 35], which uses an oxygen-selective cathode catalyst to shield a nonselec-

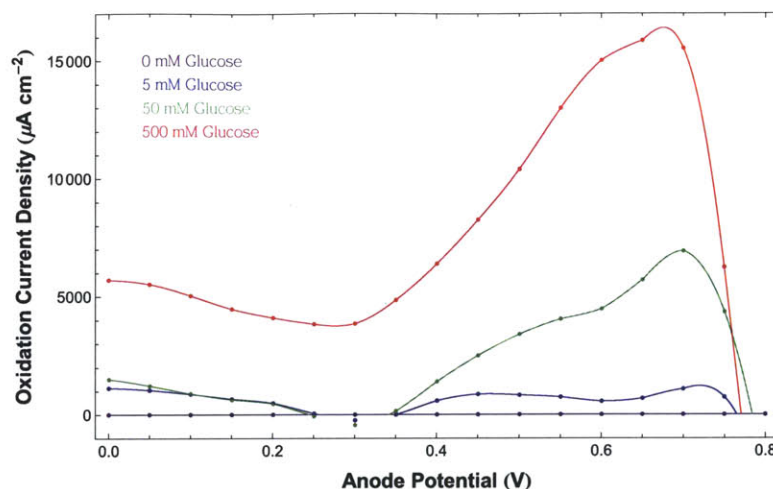


Figure 10. Oxidation Current Depends on Glucose Concentration. The oxidation current generated by the fuel cell anode, characterized here via cyclic voltammetry, varies with glucose concentration over several orders of magnitude.

tive anode from oxygen, while allowing glucose to reach the anode without reacting at the cathode. This scheme is illustrated in Figure 3.

As indicated in Figure 3, our fuel cell has a laminar structure: a porous mesh of single-walled carbon nanotubes, embedded in Nafion, comprises the cathode and forms the outermost layer; this layer is followed (from the outside in) by a Nafion separator membrane; a roughened (Raney-type) platinum anode; and finally the impermeable, silicon dioxide substrate. In this configuration, oxygen from the physiologic environment reacts at the carbon cathode, resulting in a gradient in the oxygen concentration that reaches its minimum, near zero, at the surface of the anode. The near absence of oxygen at the anode surface minimizes the rate of oxygen reduction at the anode, which in turn minimizes the electrochemical short-circuiting effects of such reactions. On the other hand, because carbon does not catalyze glucose oxidation, and because Nafion is permeable to glucose, physiologic glucose passes through the pores in the nanotube mesh unimpeded.

Microfabrication Methods

In contrast with most previous work in the area of glucose fuel cells [38], and in fulfillment of a recognized requirement for advancement of the field [39, 40], we manufactured our fuel cells entirely within a Class 10 (ISO 4) cleanroom, using microfabrication techniques and processing standards that are completely compatible with contemporary CMOS (complementary metaloxide semiconductor) integrated circuit manufacturing protocols. As a result, the fabrication methods we describe here can be used to enable wafer-level integration of fuel cell power sources with electronic circuits and microfluidics, facilitating the development of completely self-sufficient, embedded electronic and microfluidic systems.

Substrates, Masks, and Lithography

Electrodes, separator membranes, wire traces, and metal contacts were all patterned using conventional photolithography, with transparency masks printed to 2 μm feature-size tolerance (Infinite Graphics, Minneapolis, Minnesota). As substrates, we used 500 nm surface layers of silicon dioxide on conventional,

150-millimeter-diameter (6-inch) silicon wafers. To promote platinum adhesion, we deposited a 2 nm coat of titanium nitride on the oxide by evaporation.




Figure 11. Photolithography Masks. This image shows a set of glucose fuel cells of various sizes arranged for fabrication on a 150 mm (6-inch) silicon wafer. The schematic was constructed by overlaying the four process layers. Details of the individual masks are shown in the Appendix. Yellow, platinum; orange, roughened platinum anode (aluminum deposition for annealing); blue, Nafion; Green, cathode (single-walled carbon nanotubes in Nafion).

Raney Catalyst Anode

We fabricated Raney-type, activated platinum catalytic anodes using an approach similar to those described by several other groups [27, 41, 42]. This technique increases the catalytic capacity of a platinum electrode through roughening, converting an atomically smooth layer of platinum into a high-surface-area electrode. The overall approach involves patterning a platinum-aluminum alloy, then etching the aluminum out of the alloy to leave behind extremely porous platinum with a nanostructure similar to that described by Attard and colleagues [43, 44]. Our implementation proceeded as follows. We patterned all metal structures using a single mask, including the anode, cathode contact ring, wire traces, and electrical contacts, as in Figure 11 and the Appendix, depositing 100 nm of platinum by evaporation. We then deposited 100 nm of aluminum over the entire surface of the wafer. Using a second mask, we patterned photoresist over the regions designated for the anodes, protecting them from the etchant acting in the following step: using tetramethylammonium hydroxide (TMAH), we etched away the aluminum in all areas except for those designated as anodes. After stripping the remaining resist, we annealed the platinum and aluminum layers at 300°C for 60 minutes. This anneal step generated a platinum-aluminum alloy in the regions designated for the anodes. We then repeated the aluminum etch step, this time without protecting the anodes, in order to remove the aluminum from the alloy formed in the anode regions. This second etch produces high-surface-area Raney-type catalytic anodes, as shown in Figures 4, 5, 6, and 7, and as discussed in the associated Methods subsections.

Nafion Separator Membrane

Several groups have reported difficulty incorporating Nafion, a preferred ion-selective membrane in fuel cell applications, into microfabrication processes. Patterning Nafion lithographically does present some difficulties, as described by Gold and colleagues [45] and others. We used the following liftoff process for depositing Nafion on silicon dioxide and platinum.

We selected the areas to be coated with Nafion by first coating the wafers with 15-20 μm of photoresist (AZ 4620) and using a third mask to pattern and expose the regions to be coated with Nafion. We have found that using such thick layers of resist facilitates proper Nafion patterning and liftoff by enabling patterned regions of Nafion to be isolated in deep resist wells formed at the wafer surface, following resist development and evaporation of the Nafion dispersion solvent.

From a standard Nafion liquid dispersion (Nafion DE 521, DuPont), we formed a 0.83% Nafion solution via 1 : 5 dilution in 2-propanol. In a modified version of the protocol described in [46], we used a spin-on process to coat our wafers with the resulting dispersion, spinning the wafers at 750–1000 rpm for 3–10 s. In order to cure the Nafion and facilitate bonding to the substrate [47], we heated the wafers in a convection oven at 120° C for 20 minutes. This process generates Nafion layers approximately 60–420 nm in thickness, as measured by spectroscopic ellipsometry (Filmetrics, San Diego, California).

Carbon Nanotube Cathode

The cathode of our fuel cell comprises a conducting mesh of single-walled carbon nanotubes (swCNTs) embedded in Nafion, as shown in Figure 12, in electrical contact with a platinum ring (the cathode contact) on the wafer surface. This design is similar to the one described by Lee and colleagues [48]. The Nafion separator membrane electrically insulates the cathode from the anode, while permitting cationic exchange. Both the separator membrane and the cathode permit ambient glucose to reach the anode surface: Nafion is inherently permeable to glucose and the porosity of the swCNT mesh allows free transport of glucose.

We constructed the Nafion-coated swCNT mesh by allowing it to self-assemble. Using a procedure similar to the one described in [49], we suspended swCNTs at a concentration of 8 mg mL⁻¹ in an 0.83% Nafion dispersion of the kind described earlier in this Methods Section, under ultrasonic agitation for

60 minutes. We then patterned the Nafion-coated swCNT cathode using a fourth lithographic mask, and spin-on and liftoff processes identical to those described earlier in this Methods Section.




Figure 12. Fuel Cell Cathode. Scanning electron micrograph of the fuel cell cathode, showing the conducting mesh of carbon nanotubes encapsulated in Nafion ionomer. Scale Bar: 1 μm .

Biocompatibility and Anti-Fouling

Nafion is biocompatible [50], and in using it as an outermost layer encapsulating our fuel cell we enhance the biocompatibility of the entire device. Moreover, as a cation-selective ionomer, Nafion is impermeable to negatively charged proteins, small molecules, and ions, that are present in the physiologic environment and are known to cause fouling of catalytic electrodes. Nafion encapsulation has been used in the context of glucose sensors to increase the functional lifetimes of implanted electrodes [51].

Models of Energy Harvesting from Glucose in the Brain and Cerebrospinal Fluid

In this section we demonstrate that the cerebrospinal fluid is a physiologic niche particularly well suited for an implanted fuel cell, and we model the impact of an implanted glucose fuel cell on its biological environment. In particular, we derive a model used to determine the extent to which glucose and oxygen levels in the cerebrospinal fluid permit energy harvesting without adverse physiologic consequences.

Anatomy and Composition of the Cerebrospinal Fluid

The cerebrospinal fluid (CSF) comprises approximately 150 mL of fluid in the subarachnoid and intraventricular spaces, respectively surrounding and filling hollow structures within the human brain and spinal cord. The CSF is primarily produced by modified ependymal cells of the choroid plexus, located within the ventricles of the brain, as an ultrafiltrate of blood plasma whose composition is also regulated through active transport.

Under normal conditions the CSF is acellular. Its normal glucose content of 45–80 mg dL^{-1} (2.5 – 4.4 mM) [52] is comparable to, albeit systematically approximately 2-fold lower than that of plasma, which is normally regulated within the range of 70–120 mg dL^{-1} (3.9 – 6.7 mM) [53]. By contrast, normal

CSF protein content is 15 – 45 mg dL⁻¹ [54], which is at least 10²-fold lower than the normal range for plasma, 6000 – 8500 mg dL⁻¹ [55]. The ionic composition of CSF resembles that of interstitial fluid, and the two fluids are compared in Table 1.

Table 1. Ionic Composition of Cerebrospinal Fluid and Interstitial Fluid

Species	CSF Concentration	ISF Concentration
Na ⁺	154 mM	146 mM
K ⁺	3.0 mM	4.1 mM
Cl ⁻	128 mM	118 mM
HCO ₃ ⁻	23 mM	22 mM
H ⁺	pH ≈ 7.34	pH ≈ 7.44

Typical values for the principal ionic constituents of mammalian cerebrospinal and interstitial fluids [56].

With regard to the *in vitro* testing conditions for our glucose fuel cell, it is worth noting that a number of formulations for artificial cerebrospinal fluid (aCSF) have been described and developed in efforts to simulate the physiologic fluid (including trace components and gas partial pressures), both for experimental purposes [37] and in order to replace CSF lost during neurosurgical procedures [54]. In clinical practice, however, normal saline is routinely used during neurosurgery to replace CSF.

Availability and Use of Glucose in the Cerebrospinal Fluid

At physiologic glucose concentrations, high-efficiency biofuel cells such as the microbial biofuel cells developed by Rabaey and colleagues convert glucose to electricity with Coulombic efficiency exceeding $\eta_C = 80\%$ [57], and Lovley and colleagues have reported such efficiencies even at lower glucose concentrations [58]. At the opposite, low-efficiency extreme, abiotically catalyzed biofuel cells typically oxidize glucose incompletely to products such as gluconic acid, yielding only two electrons compared with the theoretical maximum of twenty-four electrons per molecule of glucose [19], corresponding to a maximum Coulombic efficiency of only approximately $\eta_C = 8\%$. The minimum glucose flux, J_g , required to fuel a glucose fuel cell generating power P , is

$$J_g = \frac{P}{\eta \Delta G_g}, \quad (3)$$

where η reflects the overall efficiency of the system (the Coulombic efficiency provides an upper bound on the overall efficiency), and $\Delta G_g = 2880 \text{ J mol}^{-1}$ denotes the heat of combustion of glucose. Using Equation 3, we can estimate the minimal glucose flux (amount of glucose per day) required to power a fuel cell generating $P = 1 \text{ mW}$:

$$J_g = \frac{(1 \text{ mW}) \times (86400 \text{ seconds per day})}{\frac{(0.08-0.80 \text{ efficiency}) \times (2880 \text{ kJ per completely oxidized mole glucose})}{180 \text{ grams per mole glucose}}} \quad (4)$$

$$= 6.75\text{--}67.5 \text{ mg (37.5--375 } \mu\text{mol) glucose per day.} \quad (5)$$

(The dashes in Equations 4 and 5 are used to denote ranges and should not be mistaken for minus signs.) Since cerebrospinal fluid in a typical adult human is produced at a rate of approximately $g_f = 550 \text{ mL}$ per day [59], the total flow of glucose through the subarachnoid space is approximately 250–440 mg per day. Thus, a biofuel cell in the configuration we propose consumes glucose at a rate of at most 2.8–28%

of the rate at which cerebrospinal fluid glucose is replenished; by comparison, physiologic fluctuations in cerebrospinal fluid glucose can exceed 25%. So availability of glucose is not likely to be a limiting factor and glucose usage by a biofuel cell of the kind we describe here should not interfere with normal physiologic processes. (At low glucose utilization efficiencies a more accurate model of CSF bulk flow might be required before this model can be considered sufficient, as the fraction of the total glucose flux available to a fuel cell as a result of CSF bulk flow will depend on the precise location of the fuel cell within the subarachnoid space.)

Availability and Use of Oxygen in the Cerebrospinal Fluid

Typical oxygen partial pressures in human cerebrospinal fluid are 25–50 mmHg [60]. These levels correspond to 35–70 $\mu\text{mol L}^{-1}$ or 1.1–2.2 mg L^{-1} if, as in [61], the solubility coefficient α of cerebrospinal fluid is considered equal to that of plasma ($\alpha = 0.003 \text{ mL}_{\text{O}_2} \text{ dL}^{-1} \text{ mmHg}^{-1} = 1.4 \mu\text{mol}_{\text{O}_2} \text{ L}_{\text{CSF}}^{-1} \text{ mmHg}^{-1}$). The bulk flow of oxygen through the subarachnoid space is therefore approximately $J_{\text{O}_2} = 19\text{--}38 \mu\text{mol per day} = 223\text{--}446 \text{ pmol s}^{-1}$. So physiologic oxygen levels are comparable to (slightly lower than) those under which biofuel cells such as those of Rabaey and colleagues have been tested. The electroreduction reactions of oxygen to water accompanying complete oxidation of glucose require $a_{g \rightarrow \text{O}_2} = 6$ moles of oxygen per mole of glucose, so the oxygen consumption rate, J_{O_2} , in the system we propose would be approximately

$$J_{\text{O}_2} = a_{g \rightarrow \text{O}_2} J_g \quad (6)$$

$$= \frac{(1 \text{ mW})}{\frac{(0.08\text{--}0.80 \text{ efficiency}) \times (2880 \text{ kJ per completely oxidized mole glucose})}{(6 \text{ moles O}_2 \text{ per completely oxidized mole glucose})}} \quad (7)$$

$$= 225\text{--}2250 \mu\text{mol O}_2 \text{ per day (2.6--26 nmol O}_2 \text{ s}^{-1}). \quad (8)$$

Thus, a biofuel cell in the configuration we propose consumes oxygen at a rate of 6–12 times the rate at which oxygen is replenished by cerebrospinal fluid bulk flow when operating at 80% efficiency, and ten times those rates when operating at 8% efficiency. Cerebrospinal fluid oxygen equilibrium under these conditions therefore depends on the ability of oxygen concentrations in cerebrospinal fluid and brain tissue interstitium to equilibrate through oxygen diffusion. A brain-implanted biofuel cell would exploit this equilibrating mechanism as a natural analog to continuous aeration of the cathode compartments of some laboratory-built biofuel cells. The diffusional transport characteristics of oxygen are such that equilibrium can be maintained with only a negligible perturbation to cerebrospinal fluid and interstitial oxygen concentrations, as the following calculations demonstrate.

As Lu and colleagues describe mathematically [61], the oxygen concentration C in the cerebrospinal fluid normally depends on four factors: (1) The oxygen partial pressure P_c in the choroid plexus capillaries, whose content is filtered and transported through the choroid endymal cells to become cerebrospinal fluid; (2) The oxygen partial pressure P_i within the brain tissue interstitial fluid; (3) The formation rate g_f of cerebrospinal fluid; and (4) The drainage rate g_d of cerebrospinal fluid. Under normal conditions $g_f = g_d \equiv g \approx 550 \text{ mL per day} = 6.4 \mu\text{L s}^{-1}$. Oxygen partial pressures in blood plasma, interstitial fluid, and cerebrospinal fluid are proportional to oxygen concentrations through solubility coefficients that have approximately the same value for all three fluids: $\alpha \approx 1.4 \mu\text{mol}_{\text{O}_2} \text{ L}^{-1} \text{ mmHg}^{-1}$. In the presence of a biofuel cell of the kind described here, the oxygen concentration also depends on a fifth factor: (5) The rate J_{O_2} at which the biofuel cell consumes oxygen. The rate of change of the concentration of oxygen in cerebrospinal fluid can therefore be expressed by the following equation:

$$\frac{dC}{dt} = \frac{1}{V_{\text{CSF}}} \left\{ g_f \alpha P_c - g_d C + D \left(P_i - \frac{C}{\alpha} \right) - J_{\text{O}_2} \right\}, \quad (9)$$

where $V_{\text{CSF}} \approx 150$ mL denotes the total volume of cerebrospinal fluid in the subarachnoid space, and $D = 17 \mu\text{mol mmHg}^{-1} \text{s}^{-1}$ is the diffusion capacity of oxygen between the interstitial fluid and the cerebrospinal fluid [61]. The five terms within the curly braces can be understood as follows: (1) The concentration C increases as oxygen flows into the cerebrospinal fluid from the choroid plexus capillaries at a rate proportional to g_f and concentration αP_c , hence the term $g_f \alpha P_c$; (2) The concentration C decreases as oxygen flows out of the cerebrospinal fluid at rate proportional to g_d and concentration C , hence the term $g_d C$; (3 & 4) Oxygen diffuses into the cerebrospinal fluid from the interstitial fluid in proportion to the partial pressure gradient between the interstitial fluid (at oxygen partial pressure P_i) and the cerebrospinal fluid (at oxygen partial pressure $P_c = \frac{C}{\alpha}$), with a proportionality constant of D , the diffusion capacity; (5) Oxygen is consumed by the biofuel cell at rate J_{O_2} . When the oxygen concentration C is at steady state, its rate of change vanishes, allowing us to solve for the equilibrium oxygen concentration, C_{eq} :

$$C_{eq} = \frac{g\alpha P_c + DP_i - J_{\text{O}_2}}{g + \frac{D}{\alpha}} \quad (10)$$

$$\approx \frac{DP_i - J_{\text{O}_2}}{\frac{D}{\alpha}}, \quad (11)$$

where the approximation holds because the diffusion transport is much faster than the bulk flow transport, allowing us to neglect the bulk flow terms containing g ($g = 6.4 \mu\text{L s}^{-1}$ is six orders of magnitude smaller than $\frac{D}{\alpha} = 12 \text{ L s}^{-1}$; and since P_c and P_i are on the same order of magnitude, $g\alpha P_c$ is six orders smaller than DP_i). Equation 10 permits us to compare the equilibrium concentrations of oxygen in the cerebrospinal fluid in the presence ($J_{\text{O}_2} = 2.6\text{--}26 \text{ nmol s}^{-1}$) and absence ($J_{\text{O}_2} = 0$) of the biofuel cell:

$$\frac{\Delta C_{eq}}{C_{eq}} = \frac{C_{eq}^{\text{present}} - C_{eq}^{\text{absent}}}{C_{eq}^{\text{absent}}} \quad (12)$$

$$= -\frac{J_{\text{O}_2}}{DP_i} \quad (13)$$

$$\approx \frac{(-2.6\text{--}26 \text{ nmol s}^{-1})}{(17 \mu\text{mol mmHg}^{-1} \text{s}^{-1}) \times (40 \text{ mmHg})} \quad (14)$$

$$= -3.8\text{--}38 \times 10^{-6}. \quad (15)$$

Thus, *the fractional change in oxygen concentration in the cerebrospinal fluid due to the presence of the biofuel cell will only be a few parts per million.* (The precise value of the interstitial pressure, approximated here as 40 mmHg, does not greatly impact this result). Note that although oxygen consumption by the biofuel cell will draw oxygen out of brain tissue into the cerebrospinal fluid, the corresponding decrease in interstitial fluid oxygen concentration will be even smaller in magnitude than the decrease in the cerebrospinal fluid, since the intracranial interstitial fluid volume is approximately twice that of the cerebrospinal fluid [61].

Oxygen Equilibration Time in Cerebrospinal Fluid

We can also estimate the time scale over which oxygen concentrations equilibrate in the cerebrospinal fluid. The system modeled by Equation 9 exhibits first-order kinetics with a time constant τ given by

$$\tau = \frac{V_{CSF}}{g + \frac{D}{\alpha}} \quad (16)$$

$$\approx \frac{V_{CSF}}{\frac{D}{\alpha}} \quad (17)$$

$$= \frac{0.150 \text{ L}}{12 \text{ L s}^{-1}} \quad (18)$$

$$= 12.5 \text{ ms.} \quad (19)$$

So the oxygen concentrations in this system equilibrate on the rapid timescale of tens of milliseconds.

Structure of Oxygen Concentration Gradients in Cerebrospinal Fluid due to Fuel Cell Oxygen Uptake

To demonstrate that the surface-area-to-volume characteristics of our proposed brain-implanted biofuel cell are consistent with a functioning system, we can model the approximate structure of the concentration gradient field around the cathode. In a first-order approximation, the concentration of oxygen will approximately vanish at the surface of the cathode, where oxygen is electroreduced to water; and at a certain distance d (which we will determine) from the surface the concentration will be equal to C , the average oxygen concentration of the cerebrospinal fluid. If d is much smaller than the characteristic dimension \sqrt{A} of the cathode (where A denotes the cathode area) and also much smaller than the width of the subarachnoid space at any location, the oxygen concentration $C(\vec{x})$ as a function of position will be approximately $C_{eq}^{present}$ everywhere except within a region close to the cathode surface; within that region the concentration satisfies $C(\vec{x}=0) = 0$ at the cathode surface, and rises to $C(\vec{x} = \hat{e}_\perp d) = C_{eq}^{present}$ at a distance d normal to the cathode surface.

We can demonstrate that the diffusion characteristics of oxygen in the cerebrospinal fluid set d to a value consistent with efficient fuel cell operation in a brain-implanted system. We have demonstrated that the average concentration of oxygen in the cerebrospinal fluid equilibrates rapidly to $C_{eq}^{present}$ when the fuel cell is operating, so the spatial distribution $C(\vec{x})$ is effectively constant at times more than several τ (several tens to hundreds of milliseconds) after the biofuel cell is turned on. We can therefore use Fick's First Law of Diffusion to determine d :

$$J = -D_{CSF} \nabla_\perp C(\vec{x}) \quad (20)$$

$$-\frac{J_{O_2}}{A} \approx -D_{CSF} \frac{C_{present}^{eq} - 0}{d}, \quad (21)$$

where the first equation is a statement of Fick's First Law for our system, and the second line expresses the first-order approximation we just described. The parameter $D_{CSF} \approx 3 \times 10^{-5} \text{ cm}^2 \text{ s}^{-1}$ is the diffusion coefficient of oxygen in cerebrospinal fluid (approximately equal to that in water), $J = \frac{J_{O_2}}{A}$ is the oxygen flux (here defined as an area-normalized flow) near the surface of the cathode, and $\nabla_\perp C(\vec{x}) \approx \frac{\Delta C}{d} \approx \frac{C_{present}^{eq} - 0}{d}$ is the oxygen concentration gradient near the cathode surface (linearized in this first-order approximation). Therefore,

$$d = D_{CSF} \frac{C_{eq}^{present}}{J_{O_2}} A, \quad (22)$$

so if the physiologic range of C is 35–70 $\mu\text{mol L}^{-1}$, d can be at most

$$d = \frac{(3 \times 10^{-5} \text{ cm}^2 \text{ s}^{-1}) \times (70 \text{ } \mu\text{mol L}^{-1}) \times (10^{-3} \text{ L cm}^{-3})}{(2.6 \text{ nmol s}^{-1})} \times (10 \text{ cm}^2) \quad (23)$$

$$= 80 \text{ } \mu\text{m} \quad (24)$$

for a device with surface area as great as $A = 10 \text{ cm}^2$. The value $d = 80 \text{ } \mu\text{m}$ satisfies our requirements that the concentration gradient vanish over a length scale much smaller than the typical dimension of the cathode and much smaller than the smallest width of the subarachnoid space, but it is also large enough to be physically achievable, as it is orders of magnitude larger than the mean free path of an oxygen molecule in water [62]. Thus, oxygen diffusion to the cathode in the system we propose will not be limited by the geometry of the system.

Acknowledgments

We thank Todd A. Thorsen and Shaun R. Berry of the Advanced Silicon Technology Group at MIT Lincoln Laboratory, and extend special thanks to Damiano Biasella of the Microelectronics Laboratory at MIT Lincoln Laboratory for extensive assistance with device fabrication. In addition, we thank Jongyoon Han and Yong-Ak Song of the MIT Research Laboratory of Electronics for collaborative discussions in the initial stages of this work. Finally, we thank Derek Lovley of the University of Massachusetts–Amherst Department of Microbiology, as well as Kelly Nevin and other members of the laboratory staff, for collaboration in the area of microbial fuel cells, and for their generosity in sharing time, expertise, and equipment during site visits.

References

1. Ellenbogen KA, Kay GN, Wilkoff BL, editors (2000) *Clinical Cardiac Pacing and Defibrillation*. Current Opinion in Pathology. Philadelphia: W. B. Saunders, 2nd edition.
2. Sarpeshkar R, Salthouse C, Sit JJ, Baker MW, Zhak SM, et al. (2005) An ultra-low-power programmable analog bionic ear processor. *IEEE Transactions on Biomedical Engineering* 52: 711-727.
3. Baker MW, Sarpeshkar R (2007) Feedback analysis and design of RF power links for low-power bionic systems. *IEEE Transactions on Biomedical Circuits and Systems* 1: 28-38.
4. Burke A (2000) Ultracapacitors: why, how, and where is the technology. *Journal of Power Sources* 91: 37-50.
5. Jayalakshmi M, Balasubramanian K (2008) Simple capacitors to supercapacitors—an overview. *International Journal of Electrochemical Science* 3: 1196-1217.
6. Dudney NJ, Neudecker BJ (1999) Solid state thin-film lithium battery systems. *Current Opinion in Solid State and Materials Science* 4: 479-482.
7. MacKay DJC (2009) *Sustainable Energy—Without the Hot Air*. Cambridge, England: Cambridge University Press.
8. Lenov V, Torfs T, Fiorini P, Hoof CV (2007) Thermoelectric converters of human warmth for self-powered wireless sensor nodes. *IEEE Sensors Journal* 7: 650-657.

9. Mitcheson PD, Yeatman EM, Rao GK, Holmes AS, Green TC (2008) Energy harvesting from human and machine motion for wireless electronic devices. *Proceedings of the IEEE* 96: 1457-1486.
10. Kurs A, Karalis A, Moffatt R, Joannopoulos JD, Fisher P, et al. (2007) Wireless power transfer via strongly coupled magnetic resonances. *Science* 317: 83-86.
11. Mandal S, Sarpeshkar R (2007) Low-power CMOS rectifier design for RFID applications. *IEEE Transactions on Circuits and Systems I: Regular Papers* 54: 1177-1188.
12. Rabaey J (2009) *Low Power Design Essentials. Integrated Circuits and Systems.* Springer.
13. Sarpeshkar R (2010) *Ultra Low Power Bioelectronics: Fundamentals, Biomedical Applications, and Bio-Inspired Systems.* Cambridge University Press.
14. Wattanapanitch W, Fee M, Sarpeshkar R (2007) An energy-efficient micropower neural recording amplifier. *IEEE Transactions on Biomedical Circuits and Systems* 1: 136-147.
15. Sarpeshkar R, Wattanapanitch W, Arfin SK, Rapoport BI, Mandal S, et al. (2008) Low-power circuits for brain-machine interfaces. *IEEE Transactions on Biomedical Circuits and Systems* 2: 173-183.
16. Sarpeshkar R, Salthouse C, Sit JJ, Baker M, Zhak S, et al. (2005) An ultra-low-power programmable analog bionic ear processor. *IEEE Transactions on Biomedical Engineering* 52: 711-727.
17. Rapoport B, Sarpeshkar R (2011) A Biomimetic Adaptive Algorithm and Micropower Circuit Architecture for Implantable Neural Decoders, IGI Global, chapter 10. pp. 216-254. In *System and Circuit Design for Biologically-Inspired Learning*, Edited by T. Temel.
18. B I Rapoport RS, Wattanapanitch W, Arfin SK, Rapoport BI, Mandal S, et al. (2009) A biomimetic adaptive algorithm and low-power architecture for implantable neural decoders. *Proceedings of the 31st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC09)* .
19. Kerzenmacher S, Ducreée J, Zengerle R, von Stetten F (2008) Energy harvesting by implantable abiotically catalyzed glucose fuel cells. *Journal of Power Sources* 182: 1-17.
20. Heller A (2004) Miniature biofuel cells. *Physical Chemistry Chemical Physics* 6: 209-216.
21. Logan BE (2008) *Microbial Fuel Cells.* Hoboken, New Jersey: John Wiley and Sons.
22. Bullen RA, Arnot TC, Lakeman JB, Walsh FC (2006) Biofuel cells and their development. *Biosensors and Bioelectronics* 21: 2015-2045.
23. Davis F, Higson PJ (2007) Biofuel cells—Recent advances and applications. *Biosensors and Bioelectronics* 22: 1224-1235.
24. Rabaey K, Boon N, Siciliano SD, and W Verstraete MV (2004) Biofuel cells select for microbial consortia that self-mediate electron transfer. *Applied and Environmental Microbiology* 70: 5373-5382.
25. J Niessen and Uwe Schr oder and F Scholz (2004) Exploiting complex carbohydrates for microbial electricity generation—a bacterial fuel cell operating on starch. *Electrochemistry Communications* 6: 955-958.

26. Osman MH, Shah AA, Walsh FC (2010) Recent progress and continuing challenges in bio-fuel cells. part ii: Microbial. *Biosensors and Bioelectronics* 26: 953-963.
27. Gebhardt U, Rao JR, Richter GJ (1976) A special type of raney-alloy catalyst used in compact biofuel cells. *Journal of Applied Electrochemistry* 6: 127-134.
28. Drake RF, Kusserow BK, Messinger S, Matsuda S (1970) A tissue implantable fuel cell power supply. *Transactions of the American Society for Artificial Internal Organs (ASAIO Journal)* 16: 199-205.
29. Weidlich E, Richter G, von Sturm F, Rao JR (1976) Animal experiments with biogalvanic and biofuel cells. *Biomaterials, Medical Devices, and Artificial Organs* 3-4: 227-306.
30. Heller A, Mano N, Kim HH, Zhang Y, Mao F, et al. (2008). Miniature biological fuel cell that is operational under physiological conditions, and associated devices and methods. U.S. Patent 7,368,190.
31. Davson H, Segal MB (1996) *Physiology of the CSF and Blood-Brain Barriers*. New York: CRC Press.
32. Bard AJ, Faulkner LR (2001) *Electrochemical Methods: Fundamentals and Applications*. Hoboken, New Jersey: John Wiley and Sons, 2nd edition.
33. Newman J, Thomas-Alyea KE (2004) *Electrochemical Systems*. Electrochemical Society Series. Hoboken, New Jersey: Wiley-Interscience.
34. von Sturm F, Richter G (1976). Pacemaker with biofuel cell. U.S. Patent 3,941,135. Assigned to Siemens Aktiengesellschaft.
35. Rao JR, Richter G, Weidlich E, von Sturm F (1972) Metal-oxygen and glucose-oxygen cells as power sources for implantable devices. *Physics in Medicine and Biology* 17: 738.
36. (2010) LT3108 Ultralow Voltage Step-Up Converter and Power Manager. Data sheet, Linear Technology Corporation, Milpitas, California.
37. An JH, Sua Y, Radmana T, Bikson M (2008) Effects of glucose and glutamine concentration in the formulation of the artificial cerebrospinal fluid (acsf). *Brain Research* 1218: 77-86.
38. Kerzenmacher S, Ducreé J, Zengerle R, von Stetten F (2008) An abiotically catalyzed glucose fuel cell for powering medical implants: Reconstructed manufacturing protocol and analysis of performance. *Journal of Power Sources* 182: 66-75.
39. von Stetten F, Kerzenmacher S, Ducreé, Zengerle R, Lozano P, et al. (2006) Implantable biofuel cells. In: *Healthy Aims Dissemination Day Proceedings*. p. WP 5b.
40. von Stetten F, Katakis I (2006) Direct glucose fuel cell as a power source for medical implants. In: *Healthy Aims Dissemination Day Proceedings*. p. WP 5b.
41. Kerzenmacher S, Schroeder M, Brämer R, Zengerle R, von Stetten F (2010) Raney-platinum film electrodes for potentially implantable glucose fuel cells. part 1: Nickel-free glucose oxidation anodes. *Journal of Power Sources* 195: 6516-6523.
42. Kerzenmacher S, Kräling U, Schroeder M, Brämer R, Zengerle R, et al. (2010) Raney-platinum film electrodes for potentially implantable glucose fuel cells. part 2: Glucose-tolerant oxygen reduction cathodes. *Journal of Power Sources* 195: 6524-6531.


43. Attard GS, Goltner CG, Corker JM, Henke S, Templer RH (1997) Liquid-crystal templates for nanostructured metals. *Angewandte Chemie International Edition* 36: 1315-1317.
44. Attard GS, Bartlett PN, Coleman NRB, Elliott JM, Owen JR, et al. (1997) Crystalline phases mesoporous platinum films from lyotropic liquid. *Science* 278: 838-840.
45. Gold S, Chua KL, Lua C, Shannon MA, Masel RI (2004) Acid loaded porous silicon as a proton exchange membrane for micro-fuel cells. *Journal of Power Sources* 135: 198-203.
46. Stetter JR, Maclay J (1989). Spin coating of electrolytes. U.S. Patent 4,795,543.
47. Ilic B, Neuzil P, Stanczyk T, Czaplowski D, Maclay GJ (1999) Low temperature nafion bonding of silicon wafers. *Electrochemical and Solid-State Letters* 2: 86-87.
48. Lee YJ, Park DJ, Park JY (2008) Fully packaged nonenzymatic glucose microsensors with nanoporous platinum electrodes for anti-fouling. *IEEE Sensors Journal* 8: 1922-1927.
49. Wu G, Xu BQ (2007) Carbon nanotube supported pt electrodes for methanol oxidation: A comparison between multi- and single-walled carbon nanotubes. *Journal of Power Sources* 174: 148-158.
50. Kim G, Kim H, Kim IJ, Kim JR, Lee JI, et al. (2009) Bacterial adhesion, cell adhesion and biocompatibility of nafion films. *Journal of Biomaterials Science* 20: 1687-1707.
51. Rishpon J, Gottesfeld S, Campbell C, Davey J, Jr TAZ (1994) Amperometric glucose sensors based on glucose oxidase immobilized in nafion. *Electroanalysis* 6: 17-21.
52. Ropper AH, Brown RH (2005) Adams and Victor's Principles of Neurology. New York: McGraw-Hill, 8th edition.
53. Kumar V, Abbas A, Fausto N (2005) Pathologic Basis of Disease. Philadelphia: Elsevier Saunders, 7th edition.
54. Greenberg MS (2010) Handbook of Neurosurgery. New York: Thieme Medical Publishers.
55. Chaudhry HJ, Grieco AJ, Macklis R, Mendelsohn M, Jr GHM (2004) Fundamentals of Clinical Medicine: An Introductory Manual. New York: Lippincott Williams and Wilkins.
56. Rosenberg GA (1990) Brain Fluids and Metabolism. Oxford: Oxford University Press.
57. Rabaey K, Lissens G, Siciliano SD, Verstraete W (2003) A microbial fuel cell capable of converting glucose to electricity at high rate and efficiency. *Biotechnology Letters* 25: 1531-1535.
58. Chaudhuri SK, Lovley DR (2003) Electricity generation by direct oxidation of glucose in mediatorless microbial fuel cells. *Nature Biotechnology* 21: 1229-1232.
59. Barrett KE, Barman SM, Boitano S, Brooks H (2010) Ganong's Review of Medical Physiology. Access Medicine. New York: McGraw-Hill, 23rd edition.
60. Zaharchuk G, Martin AJ, Rosenthal G, Manley GT, Dillon WP (2005) Measurement of cerebrospinal fluid oxygen partial pressure in humans using mri. *Magnetic Resonance in Medicine* 54: 113-121.
61. Lu K, J W Clark J, Ghorbel FH, Robertson CS, aand J B Zwischenberger DLW, et al. (2004) Cerebral autoregulation and gas exchange studied using a human cardiopulmonary model. *American Journal of Physiology: Heart and Circulatory Physiology* 286: H584-H601.
62. Denny MW (1993) Air and Water: The Biology and Physics of Life's Media. Princeton, New Jersey: Princeton University Press.


Appendix


This Appendix contains schematic and functional images of the four photolithographic masks referenced in the manuscript, and used to fabricate the fuel cell described here using conventional semiconductor manufacturing techniques.


The following descriptions apply to the images on the following pages:


1. Computer-aided design (CAD) image of the superimposed fabrication masks for an array of glucose fuel cells of varying dimensions on a 150 mm silicon wafer. Yellow, Platinum; Orange, Aluminum; Blue, Nafion; Green, Nafion containing Single-Walled Carbon Nanotubes (swCNTs).
2. CAD image of the layout of a single glucose fuel cell. Scale: The electrode contacts are 1 mm wide. Yellow, Platinum; Orange, Aluminum; Blue, Nafion; Green, Nafion containing Single-Walled Carbon Nanotubes.
3. CAD image detail showing the anode and cathode of a glucose fuel cell. The scale bar at left shows length in micrometers.
4. CAD image showing the platinum metal layer of a glucose fuel cell. Scale: The electrode contacts are 1 mm wide.
5. Transparency mask used to pattern the platinum metal layer for the array of glucose fuel cells of varying dimensions on a 150 mm silicon wafer, shown in the first CAD image.
6. Detail from the platinum-layer transparency mask.
7. Detail from the platinum-layer transparency mask.
8. CAD image showing the aluminum metal layer of a glucose fuel cell.
9. Transparency mask used to pattern the aluminum metal layer for the array of glucose fuel cells of varying dimensions on a 150 mm silicon wafer, shown in the first CAD image.
10. CAD image showing the Nafion layer of a glucose fuel cell.
11. Transparency mask used to pattern the Nafion layer for the array of glucose fuel cells of varying dimensions on a 150 mm silicon wafer, shown in the first CAD image.
12. Detail from the Nafion-layer transparency mask.
13. CAD image showing the Nafion-swCNT layer of a glucose fuel cell.
14. Transparency mask used to pattern the Nafion layer for the array of glucose fuel cells of varying dimensions on a 150 mm silicon wafer, shown in the first CAD image.



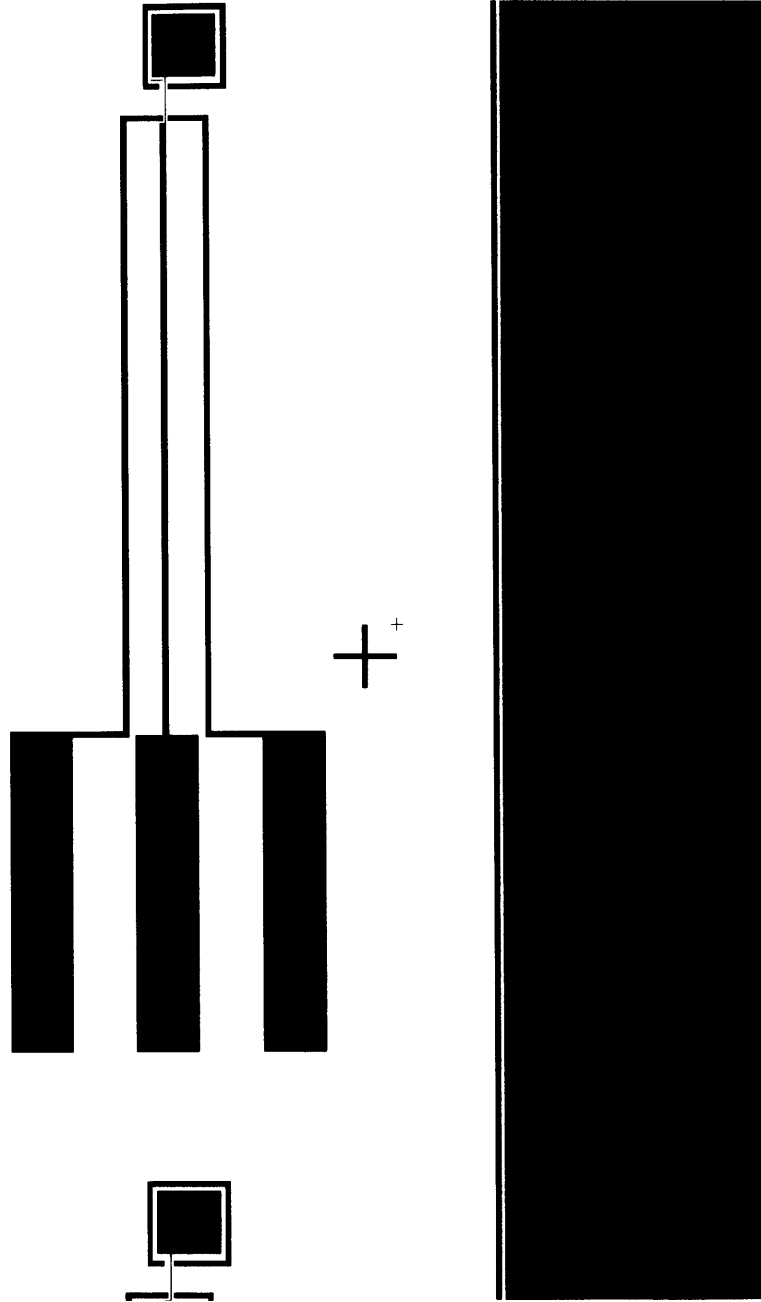










100-15750-1-000-0000-0000-0000








Al (250um by 100um) for Anneal to Roughen Pt Anode












Nafion + SWCNTs





Chapter 5

Conclusions

5.1 Motivations

Programmable, implantable medical devices such as cardiac pacemakers, implantable defibrillators, cochlear implants, and deep-brain stimulators, have already transformed the practice of medicine and the lives of hundreds of thousands of people. Brain–machine interfaces, already established as important research tools in basic neuroscience, are poised to emerge with similarly transformative impact in the medical world, repairing and augmenting the human nervous system in conditions ranging from paralysis and amputation to Parkinsonian dyskinesia, blindness, and intractable epilepsy. The viability of brain–machine interfaces as clinically useful devices hinges on their ability to operate in miniature, implantable electronic systems with microwatt power budgets. Major progress in that direction has been demonstrated by a number of groups in recent years, including our own, which have brought individual components of a future, fully implantable brain–machine interface system to an advanced level of development. Micropower multichannel neural signal amplifier arrays, power- and energy-efficient analog-to-digital converters, and low-power systems for wireless data and power telemetry all figure prominently among systems developed for neural electronics in the past decade, and some of these systems have achieved levels of performance approaching fundamental physical limits.

The work described in this thesis presents a set of solutions to two progress-

limiting problems that, in spite of the advances of the last decade, have until now remained unaddressed in the field of brain–machine interfaces:

1. *Implantable Neural Decoding Systems.* Clinically useful brain–machine interfaces must decode neural signals in real time into control signals for prosthetics and other external devices. The very high information content of neural signals essentially requires that a fully functional micropower system perform some or all of the neural decoding operation within the implanted device. Yet until now the state of the art in neural signal decoding has neglected power and size constraints, and the field as a whole has focused on computationally intensive, high-power neural decoding algorithms implemented in computer software. Neurons themselves constantly encode and decode neural signals on a nanowatt power budget, so biomimetics represents a natural approach to neural decoding. The work presented in this thesis demonstrates a set of micropower, scalable approaches to the decoding and compression of neural signals.
2. *Sustainable Long-Term Power Supplies for Implantable Electronic Systems.* Existing implantable medical devices have finite and not entirely predictable functional lifetimes, currently limited to several years by the performance of implantable batteries, which almost universally also comprise the largest components of implantable systems by both mass and volume. Yet the cells, tissues, and organs with which implantable devices interact harvest energy with high efficiency from fuel substrates that are typically overabundant in the physiologic environment. The work presented in this thesis includes the development of an implantable glucose fuel cell, representing a potentially ultralong-lifetime implantable power source for biomedical electronics, including brain–machine interfaces.

5.2 Summary of Work

5.2.1 Biomimetic Algorithms and Circuit Architectures for Micropower Analog Neural Decoding

Chapter 2, consisting of a published paper and a published book chapter, presents an algorithm and micropower analog circuit architecture for decoding neural signals in real time. The system architecture resembles that of a traditional a neural network, implemented in analog electronic hardware. The synapses of the neural network consist of linear filters, with first-order dynamics designed to resemble those of biological neurons. The entire system is capable of learning in real time to optimize performance in a self-tuning scheme inspired by the Hebbian learning paradigms observed in biological neural systems and used in training artificial neural networks.

Circuit simulations indicate that the system can perform at power consumption rates as low as 54 nW per decoding channel, and that a full-scale decoding system with 100 neural inputs and 3 motor outputs would consume less than 42 μ W, including power required for signal preprocessing. The same decoding architecture achieves compression factors on the order of 10^5 when used as a compression technique for reducing the bandwidth and hence the power required to transmit high-bandwidth neural signals wirelessly from an implanted system.

The decoding performance of this system, envisioned as a component of a neuro-motor prosthesis, is demonstrated using neural data from rodent head movement and primate limb movement.

5.2.2 Universal Computing Architectures for Micropower Digital Neural Decoding

Chapter 3, consisting of a manuscript recently submitted for publication, describes a digital approach to micropower neural decoding. This approach separates the neural decoding problem into two parts, to facilitate transmission of neural data across limited-bandwidth, wireless channels: (1) A set of preliminary decoding and compres-

sion operations to be performed within an implanted system, based on statistics and logical encoding primitives generic to all neural systems, and (2) A complementary decoding algorithm, implemented outside the implanted system, that applies physical and other implementation-specific constraints to generate accurate neural decoding.

In a departure from high-power, computationally expensive, software-based digital decoding, this system represents a highly scalable solution of extremely low computational complexity, suited for real-time decoding of neural signals in the context of an implantable system. The programmable architecture is an explicit implementation of a Turing-type computing machine. It achieves computational efficiency through an essentially arithmetic-free design based only on computationally inexpensive, binary logic operations. In particular, the implanted decoding system uses no multiplication.

A concrete realization of this neural decoding architecture is described, using a micropower, field-programmable gate array (FPGA) that dissipates less than $5 \mu\text{W}$. The performance of the system is validated by decoding 32-channel neural data from a behaving rodent.

5.2.3 Biologically Implantable Glucose Fuel Cells for Brain-Powered Brain–Machine Interfaces

Chapter 4, consisting of a manuscript recently submitted for publication, describes the design and performance of a biologically implantable fuel cell that generates power through glucose oxidation. The fuel cell produces $3.4 \mu\text{W cm}^{-2}$ of steady-state power and generates peak power of up to $180 \mu\text{W cm}^{-2}$ under physiologic conditions. In contrast with previous work, the fuel cell is constructed using conventional semiconductor fabrication processes, and relies on robust, solid-state catalysts, so is optimally suited for manufacture together with traditional integrated circuits on a single silicon wafer, to yield implantable microelectronic systems with long-lifetime, integrated power sources.

A modeling study accompanying the design of the glucose fuel cell proposes that the cerebrospinal spinal fluid surrounding the brain constitutes a physiologic niche

particularly well suited for a glucose-based fuel cell.

Finally, the chapter describes an energy harvesting circuit designed to couple the glucose fuel cell to a micropower neural amplifier, demonstrating the feasibility of a brain-powered brain-machine interface.

5.3 Future Outlook

I believe that electronic devices that function, compute, and power themselves like biological systems will change current paradigms of medical devices and surgical treatments, and will be fundamentally important to medicine of the coming decades. As I look forward to a career in neuroscience, medicine, and technology, I find this an irresistibly compelling vision, and feel certain that an emerging generation of such devices will transform our ability to care for patients and deepen our understanding of the brain.

Appendix A

Literature Reviews

A.1 Brain–Machine Interfaces and Neural Decoding

A.1.1 Brain–Machine Interfaces

Brain–machine interfaces have proven capable of decoding neuronal population activity in real-time to derive instantaneous control signals for prosthetics and other devices [15, 94, 88, 55, 29, 73, 34]. The field of brain–machine interface technology has been reviewed by several authors in recent years [39, 81]. All of the neural decoding systems demonstrated to date in connection with brain–machine-interfaces have operated by analyzing digitized neural data, and none have focused on power-efficient design; the power-efficient, hybrid analog-digital approach to neural decoding implemented by the biomimetic neural decoder presented in this thesis therefore represents a divergence from conventional approaches to neural signal decoding.

A.1.2 Microelectronics for Brain–Machine Interfaces

Clinically viable neural prosthetics are an eagerly anticipated advance in the field of rehabilitation medicine, and development of brain–machine interfaces that wirelessly transmit neural data to external devices will represent an important step toward clinical viability. The general model for such devices has two components: a brain-implanted unit directly connected to a multielectrode array collecting raw neural data; and a unit outside the body for data processing, decoding, and control. Data transmission between the two units is wireless. A brain–machine interface consisting of hundreds of channels, each sampled at tens of kilohertz and digitized to ten or more bits of precision, generates data at a rate of tens to hundreds of megabits per second; the power costs in digitization, wireless communication, and population signal decoding all scale with this high data rate. Consequences of this scaling include unwanted heat dissipation in the brain, decreased longevity of batteries, and increased size of the implanted unit. Recent designs for system components have addressed these issues in several ways, including micropower neural amplification [30, 93]; adaptive

power biasing to reduce recording power in multielectrode arrays [78, 77]; low-power data telemetry [26, 53, 49, 50]; ultra-low-power analog-to-digital conversion [99]; low-power neural stimulation [89]; energy-efficient wireless recharging [5]; and low-power circuits and system designs for brain-machine interfaces [78, 77]. Power-conserving schemes for compressing neural data before transmission have also been proposed [32]. However, *essentially no work has been done in the area of power-efficient neural decoding.*

A.1.3 Neural Decoding

The ease of implementing sophisticated algorithms in digital systems has led to a proliferation of effective approaches to decoding and learning. Some of the most popular such algorithms in the context of neural decoding involve construction of optimal linear filters [83, 92]; training of artificial neural networks [72]; use of Kalman filters [96, 97] and adaptive Kalman filters [98]; estimation based on Bayesian inference techniques [86, 10] and point-process models [21]; and decoding on the basis of frequency-domain data (such as the spectral content of local field potentials) [59, 85, 3] or wavelet decompositions of neural signals [55]. However, the biomimetic neural decoder developed in the thesis proposed here represents the first description of an analog circuit architecture for neural signal decoding.

A major question in the area of neural decoding to achieve neural control of prosthetic devices concerns the extent to which the brain can learn to decode processed neural data on its own, and with what degree of reliability biological learning can compensate for deficiencies in an artificial neural decoder. For clinical neural prosthetic devices, the necessity of highly sophisticated decoding algorithms remains an open question, since both animal [88, 13, 55, 91] and human [29] users of even first-generation neural prosthetic systems have proven capable of rapidly adapting to the particular rules governing the control of their brain-machine interfaces. Nevertheless, some of the same experiments have also demonstrated the apparent insufficiency of such contributions from biological learning, and so the degree to which a neural prosthetic system should rely on artificial or biological learning remains unclear.

More recent work [25] has helped clarify some of these issues by underscoring the importance of stable long-term neural recordings for reliable neural decoding and improvements through biological learning. In particular, it appears that the brain is capable not only of learning complex decoders, but also of correcting errors in corrupt decoders, provided the decoders are based on stable neural recordings. Much of the existing work on neural decoding has been limited by the inability to maintain stable recordings from multiple units over long time periods, and by the consequent need to retrain decoders at each training session. As biological learning can occur over longer timescales than those set by the instability of typical single- or multiunit neural recordings, optimal conditions for biological learning in the context of neural prosthesis control have rarely been achieved. Recent evidence suggests that a reliably high level of performance can be obtained as the brain forms motor memories in response to learning a decoder based on stable, long-term recordings.

A.2 Bioimplantable Device Engineering and Bioimplantable Power Sources

A.2.1 Bioimplantable Device Engineering

Implantable devices are increasing both in prevalence and in importance as treatment modalities and investigational tools in contemporary medicine. Four particular programmable devices, having emerged as standards of care in the treatment of some electrophysiologic disorders, are of particular relevance to the work pursued for this thesis as case studies: the implantable cardiac pacemaker and the implantable cardioverter-defibrillator [23], the cochlear implant [54], and the deep brain stimulator [87]. Emerging bioimplantable electronic technologies and issues facing the field of bioimplantable electronics have been the subject of recent reviews and special journal issues [58].

Designing bioimplantable electronic devices involves many of the same challenges faced in designing other implantable devices, such as those used in long-term drug

delivery [71]. Those challenges include designing with biocompatible materials [65, 7] as well as mechanically designing and packaging implantable devices in ways that facilitate access to the appropriate physiologic systems without compromising tissue viability or device integrity [9, 70].

Implantable electronic interfaces with the brain and nervous system also face a special set of design challenges related to biocompatibility and stability: Ideal neuroprosthetic devices must be capable of implantation in the brain for indefinite periods without damaging or eliciting unintended reactions from surrounding neural tissue, they must be macro- and microanatomically optimized for the appropriate neural targets, and they must be capable of stably recording neural signals of the selected types for indefinite periods of time.

The difficulty of maintaining stable single-unit neural recordings for extended periods has long been acknowledged by electrophysiologists, and while some notable examples have been reported in the literature, such reports tend to be the exception rather than the rule. The field of brain-machine interfaces has responded to this problem in two principal ways, at once seeking to optimize electrode design and electrode-tissue interactions for stable long-term recording; and simultaneously exploring a range of neural signals, including local field potentials and unsorted multiunit activity, as alternatives to single-unit action potentials. Biocompatibility as a concern for implantable brain-machine interface systems extends to multiple aspects of the implanted system, including the design and materials chosen for the packaging, but it is especially important in the context of neural recording electrodes.

Decoding neural signals presupposes the availability of such signals, and the quality of neural decoding depends not only on the particular algorithms used for the decoding, but also on the quality of the neural signals used as inputs to the neural decoder. An ideally biocompatible neural recording electrode must satisfy two sets of constraints, which respectively concern the static and dynamic properties of the tissue-electrode interface. Ideal static properties of the electrode include the following: (1) Suitability for insertion in a manner that does minimal damage to surrounding neural and vascular tissue; (2) Mechanical stiffness equal to that of the surrounding

tissue, in order to avoid tissue-electrode shearing and electrode displacement during head acceleration; (3) Minimal electrode volume to minimize neural tissue displacement after insertion; (4) Shank and recording site geometry optimized for stereotactic placement relative to desired macro- and microanatomic targets; and (5) Tip or recording site geometry optimized for the neural signal of choice. Ideal dynamic properties of a neural recording electrode include the following: (1) Optimized electrode impedance given the local tissue properties and neural signals of interest; and (2) Ability to regulate the response of neural tissue to the presence of the electrode over time.

Almost all of these static and dynamic properties represent active areas of current research. In particular, microfabrication techniques [37, 95] and nontraditional materials such as conducting polymers [1, 46, 68] present large parameter spaces within which to optimize electrode designs. Additionally, it is well known that the microanatomic environment of a recording electrode changes over time, generally in a way that tends to degrade electrode performance. For example, neighboring neurons may die or recede from the electrode, or astrocytes and microglia may ensheath the electrode and change the impedance at the tissue-electrode interface. Consequently, understanding and controlling the dynamic response of neural and neurovascular tissue to neural recording electrodes is also an important area of current research [6, 84].

A.2.2 Bioimplantable Power Sources

Various solutions to the problem of providing power to biologically implanted devices have been proposed, prototyped, or implemented. Two principal solutions are currently in widespread use: Single-use batteries, such as those used in implantable pulse generators for cardiac pacing, defibrillation, and deep brain stimulation, which are designed to have finite lifetimes and to be replaced surgically at intervals of several years [23]; and inductive power transfer, typically accomplished transcutaneously at radio frequencies as in cochlear implants [76].

Recent advances in battery technology and related fields, leading to increased energy and power densities in small devices such as supercapacitors [11, 35] as well as

thin film lithium and thin film lithium ion batteries [20], will facilitate improvements in systems based on these two solutions, particularly by shrinking battery sizes and extending battery lifetimes.

An alternative approach to powering biologically implanted devices is to employ schemes for harvesting power from the biological environment. Systems for transducing light, thermal energy, mechanical vibration, and low-frequency mechanical stresses or strains, as well as near- or far-field electromagnetic radiation into electrical energy have been described and implemented, and complementary electronic design techniques required to make use of the power generated by such systems have been reviewed [14].

One approach to harvesting energy from a physiologic environment is to extract it from physiologically available biological fuel substrates such as glucose. Three major design paradigms for glucose-based fuel cells have been described, differing principally in the materials used to catalyze electrode reactions: In one paradigm the catalysts are abiotic; in the second paradigm the catalysts are isolated enzymes fixed to electrode substrates; and in the third paradigm oxidation is performed by exoelectrogenic bacteria colonizing a fuel cell anode. Numerous designs representing each of these fuel cell paradigms have been described and reviewed in the scientific and patent literatures. An extensive review of the scientific and patent literatures on abiotic implantable glucose fuel cells is provided in [36]. Bioimplantable fuel cells based on enzymatic catalysis are reviewed in [27]. Microbial fuel cells are reviewed in the next section of this proposal. While several proposals for biologically implantable fuel cells of each type have been described, at the time this document is being written the author is not aware that any such design has been reduced to practice.

A.3 Microbial Fuel Cells

The concept of a microbial fuel cell was described as early as 1962 [18], but the field as it exists today originates from the discovery around the year 2000 that exoelectrogenic bacteria are capable of transferring electrons directly to extracellular substrates

without the need for membrane-soluble transport molecules. A recent text provides the most comprehensive review of this field to date [44] and includes an extensive bibliography on theoretical and applied aspects of microbial fuel cells.

Much of the literature on microbial fuel cells is concerned with heterogeneous, incompletely characterized microbial populations. By contrast, in an effort to engineer precisely characterized devices, this thesis focuses on single-species microbial fuel cells. While some species appear to have relatively greater capacities for current and power production (including *Geobacter sulfurreducens* and *Desulphovibrio desulfuricans*), it has also proven possible to adapt certain bacterial strains for increased current production, and there is no consensus as to the optimality of any particular microbial community or fuel cell design. This thesis therefore focuses on two promising exoelectrogenic species, *Geobacter sulfurreducens* and *Rhodoferax ferrireducens*: *Geobacter* because it is a well characterized, nonpathogenic electrogenic species that generates comparatively high current and power densities; and *Rhodoferax* because it is capable of exoelectrogenic activity fueled by a range of bioavailable substrates, including glucose, under physiologic conditions.

Geobacter is a genus of anaerobic, metal-reducing, environmental bacteria. This genus, and particularly the species *Geobacter sulfurreducens* [12], has been studied extensively in several contexts, of which the principal one relevant to this thesis is its role as an exoelectrogen capable of catalyzing the anodic reaction in microbial fuel cells. A small number of review articles focus specifically on *Geobacter sulfurreducens* in this context [8, 56], as does a recent book chapter on electricity production by electricigens [45], which also provides an extensive bibliography on this subject. In addition, the genomes of several *Geobacter* species, including *Geobacter sulfurreducens*, have been sequenced, analyzed, and annotated, facilitating rational genetic engineering to optimize metabolic function under conditions such as those found in microbial fuel cells [33, 31, 48].

The recently discovered exoelectrogenic bacterium *Rhodoferax ferrireducens*, is a facultative anaerobe that is capable of using glucose as a metabolic substrate under conditions similar to those found in mammalian cerebrospinal fluid. Only a small

number of primary research papers have been published to date on *Rhodoferax ferrireducens* [16, 24, 40, 41, 42]. Nevertheless, the sequencing and annotation of the *Rhodoferax ferrireducens* genome has facilitated *in silico* modeling of the metabolic reaction networks of this organism, with a view toward rationally optimizing its physiology for several applications, including microbial fuel cells [69].

A.4 Symbiosis

Symbiosis has been a somewhat neglected field over the past half-century, and a comparatively limited number of authors have treated the field as a whole in unified surveys during that time [2, 28, 57, 82, 19]. An intellectual history of symbiosis [74] explains this relative neglect while providing insight into approaches to studying and understanding the phenomenon, and into the development of field as a whole.

Pioneering work by Margulis in the area of symbiosis [51, 52], particularly concerning the endosymbiotic theory of the evolution of mitochondria and plastids, is relevant to a central idea of this thesis: Evolution on a microbiological scale of interactions between organisms and one another, and between organisms and their environments, can drive the acquisition of new functionality in ways that resemble, inspire, and facilitate the development of biologically implantable devices.

Reviews of standard perspectives on microbial symbiosis can be found in several major texts on the biology of microorganisms [47]. Such texts typically address classical topics in microbial metabolic diversity, including energy source tropisms, nitrogen fixation, and catabolism of organic compounds; microbial ecology, including microbial biogeochemistry, environmental cycling of nutrients, and bioremediation; as well as microbial disease and immunology.

There is also an emerging literature on the ecology, organismic and evolutionary biology, microbiology, and biochemistry, of symbiosis. *Nature Reviews Microbiology* recently produced a collection of reviews and a curated set of primary research articles on microbial symbiosis [66], and the field of symbiosis has a journal of its own, *Symbiosis* [67].

This very brief literature review neglects certain subfields of symbiosis that, while of great general importance, are of marginal relevance to the thesis proposed here. In particular, the fields of human microbial pathology and parasitology are associated with vast, extensive bodies of medical and scientific literature. However, some contemporary research initiatives have begun to emphasize the importance of understanding human biology from a symbiotic perspective, in the context of the vast number of interspecies associations human cells and organs naturally establish with parasitic as well as the mutualistic and commensal members of the human microbiome [90]. These approaches underscore the importance of cymbionics from the perspective of human biology: The human body, considered as an ecosystem, comprises an order of magnitude more cells of nonhuman than of human origin; its metagenome contains a multitude of microbial sequences, many of which are capable of evolution on time scales as short as days, and of driving functional adaptations to novel environments. *One claim of this thesis is that these rapidly evolvable elements of the human ecosystem, which some have termed the 'human supraorganism' [90], may be instrumental in enabling human adaptation to a new generation of bionic technologies.*

References

- [1] M. R. Abidian and D. C. Martin. Experimental and theoretical characterization of implantable neural microelectrodes modified with conducting polymer nanotubes. *Biomaterials*, 29:1273–1283, 2008.
- [2] V. Ahmadjian and S. Paracer. *Symbiosis: An Introduction to Biological Associations*. Oxford University Press, Oxford, 2nd edition, 2000.
- [3] R. A. Andersen, J. W. Burdick, S. Musallam, B. Pesaran, and J. G. Cham. Cognitive neural prosthetics. *Trends in Cognitive Sciences*, 8(11):486–493, 2004.
- [4] R. Sarpeshkar B. I. Rapoport, W. Wattanapanitch, S. K. Arfin, B. I. Rapoport, S. Mandal, M. W. Baker, M. S. Fee, S. Musallam, and R. A. Andersen. A biomimetic adaptive algorithm and low-power architecture for implantable neural decoders. *Proceedings of the 31st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC09)*, September 2009.
- [5] M. W. Baker and R. Sarpeshkar. Feedback analysis and design of RF power links for low-power bionic systems. *IEEE Transactions on Biomedical Circuits and Systems*, 1(1):28–38, April 2007.
- [6] C. S. Bjornsson, S. J. Oh, Y. A. Al-Kofahi, Y. J. Lim, K. L. Smith, J. N. Turner, S. De, B. Roysam, W. Shain, and S. J. Kim. Effects of insertion conditions on tissue strain and vascular damage during neuroprosthetic device insertion. *Journal of Neural Engineering*, 3:196–207, 2006.
- [7] J. Black. *Biological Performance of Materials: Fundamentals of Biocompatibility*. CRC Press, New York, 4th edition, 2005.
- [8] D. R. Bond and D. R. Lovley. Electricity production by geobacter sulfurreducens attached to electrodes. *Applied and Environmental Microbiology*, 69(3):1548–1555, 2003.
- [9] L. Bowman and J. D. Meindl. The packaging of implantable integrated sensors. *IEEE Transactions on Biomedical Engineering*, 33(2):248–255, February 1986.
- [10] A. E. Brockwell, A. L. Rojas, and R. E. Kass. Recursive Bayesian decoding of motor cortical signals by particle filtering. *Journal of Neurophysiology*, 91:1899–1907, 2004.
- [11] A. Burke. Ultracapacitors: why, how, and where is the technology. *Journal of Power Sources*, 91(1):37–50, November 2000.

- [12] F. Caccavo Jr., D.J. Lonergan, D.R. Lovley, M. Davis, J.F. Stolz, and M.J. McInerney. *Geobacter sulfurreducens* sp. nov., a hydrogen- and acetate-oxidizing dissimilatory metal-reducing microorganism. *Applied and Environmental Microbiology*, 60(10):3752–3759, 1994.
- [13] J. M. Carmenta, M. A. Lebedev, R. E. Crist, J. E. O’Doherty, D. M. Santucci, D. F. Dimitrov, P. G. Patil, C. S. Henriquez, and M. A. L. Nicolelis. Learning to control a brain-machine interface for reaching and grasping by primates. *Public Library of Science Biology*, 1(2):1–16, October 2003.
- [14] Anantha P. Chandrakasan, Naveen Verma, and Denis C. Daly. Ultralow-power electronics for biomedical applications. *Annual Reviews of Biomedical Engineering*, 10:247–274, 2008.
- [15] J. K. Chapin, K. A. Moxon, R. S. Markowitz, and M. L. Nicolelis. Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex. *Nature Neuroscience*, 2:664–670, 1999.
- [16] S. K. Chaudhuri and D. R. Lovley. Electricity generation by direct oxidation of glucose in mediatorless microbial fuel cells. *Nature Biotechnology*, 21(10):1229–1232, 2003.
- [17] L.L. Chen, R. Madhavan, B.I. Rapoport, and W.S. Anderson. Real-time brain oscillation detection and phase-locked stimulation using autoregressive spectral estimation and time-series forward prediction. *IEEE Transactions on Biomedical Engineering*. Accepted January 2011.
- [18] J. B. Davis. Preliminary experiments on a microbial fuel cell. *Science*, 137(3530):615, 1962.
- [19] A. E. Douglas. *Symbiotic Interactions*. Oxford University Press, Oxford, 1994.
- [20] N. J. Dudney and B. J. Neudecker. Solid state thin-film lithium battery systems. *Current Opinion in Solid State and Materials Science*, 4(5):479–482, October 1999.
- [21] U. T. Eden, L. M. Frank, R. Barbieri, V. Solo, and E. N. Brown. Dynamic analyses of neural encoding by point process adaptive filtering. *Neural Computation*, 16(5):971–998, 2004.
- [22] P.K. Eide, B.I. Rapoport, W.B. Gormley, and J.R. Madsen. A dynamic nonlinear relationship between the static and pulsatile components of intracranial pressure in patients with subarachnoid hemorrhage. *Journal of Neurosurgery*, 112(3):616–625, March 2010.
- [23] Kenneth A. Ellenbogen, G. Neal Kay, and Bruce L. Wilkoff, editors. *Clinical Cardiac Pacing and Defibrillation*. Current Opinion in Pathology. W. B. Saunders, Philadelphia, 2nd edition, 2000.
- [24] K. T. Finneran, C. V. Johnsen, and D. R. Lovley. *Rhodoferrax ferrireducens* sp nov., a psychrotolerant, facultatively anaerobic bacterium that oxidizes acetate with the reduction of Fe(III). *International Journal of Systematic and Evolutionary Microbiology*, 53:669–673, 2003.

- [25] K. Ganguly and Carmena J. M. Emergence of a stable cortical map for neuroprosthetic control. *Public Library of Science Biology*, 7(7), 2009.
- [26] M. Ghovanloo and S. Atluri. A wideband power-efficiency inductive wireless link for implantable microelectronic devices using multiple carriers. *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, 54(10):71–77, October 2007.
- [27] A. Heller. Miniature biofuel cells. *Physical Chemistry Chemical Physics*, 6:209–216, 2004.
- [28] S. Mark Henry, editor. *Symbiosis: Voume I (Associations of Microorganisms, Plants, and Marine Organisms) and II (Associations of Invertebrates, Birds, Ruminants, and Other Biota)*. Academic Press, New York, 1967.
- [29] L. R. Hochberg, M. D. Serruya, G. M. Friehs, J. A. Mukand, M. Saleh, A. H. Caplan, A. Branner, D. Chen, R. D. Penn, and J. P. Donoghue. Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature*, 442:164–171, July 2006.
- [30] J. Holleman and B. Otis. A sub-microwatt low-noise amplifier for neural recording. *29th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 3930–3933, August 2007.
- [31] D. E. Holmes, S. K. Chaudhuri, K. P. Nevin, T. Mehta, B. A. Methe, A. Liu, J. E. Ward, T. L. Woodard, J. Webster, and D. R. Lovley. Microarray and genetic analysis of electron transfer to electrodes in geobacter sulfurreducens. *Environmental Microbiology*, 8(10):1805–1815, 2006.
- [32] R. H. Olsson III and K. D. Wise. A three-dimensional neural recording microsystem with implantable data compression circuitry. *IEEE Journal of Solid-State Circuits*, 40(12):2796–2804, December 2005.
- [33] M. Izallen, R. Mahadevan, A. Burgard, B. Postier, R. Didonato, J. Sun, C. H. Schilling, and D. R. Lovley. Geobacter sulfurreducens strain engineered for increased rates of respiration. *Metabolic Engineering*, 10(5):267–275, September 2008.
- [34] A. Jackson, J. Mavoori, and E. E. Fetz. Long-term motor cortex plasticity induced by an electronic neural implant. *Nature*, 444:55–60, November 2006.
- [35] M. Jayalakshmi and K. Balasubramanian. Simple capacitors to supercapacitors—an overview. *International Journal of Electrochemical Science*, 3(11):1196–1217, November 2008.
- [36] S. Kerzenmachera, J. Ducreée, R. Zengerle, and F. von Stetten. Energy harvesting by implantable abiotically catalyzed glucose fuel cells. *Journal of Power Sources*, 182:1–17, March 2008.
- [37] Daryl R. Kipke, William Shain, György Buzsáki, Eberhard E. Fetz, Jaimie M. Henderson, Jamille F. Hetke, and Gerwin Schalk. Advanced neurotechnologies for chronic neural interfaces: New horizons and clinical opportunities. *Journal of Neuroscience*, 28(46):11830–11838, 2008.

- [38] John Lazzaro, John Wawrzynek, and Richard P. Lippmann. A micropower analog circuit implementation of hidden markov model state decoding. *IEEE Journal of Solid-State Circuits*, 32(8):1200–1209, August 1997.
- [39] M. A. Lebedev and M. A. Nicolelis. Brain-machine interfaces: Past, present and future. *Trends in Neurosciences*, 29(9):536–546, 2006.
- [40] Z. D. Liu, Z. W. Du, J. Lian, X. Y. Zhu, S. H. Li, and H. R. Li. Improving energy accumulation of microbial fuel cells by metabolism regulation using rhodospirillum rubrum as biocatalyst. *Letters in Applied Microbiology*, 44(4):393–398, 2007.
- [41] Z. D. Liu and H. R. Li. Effects of bio- and abio-factors on electricity production in a mediatorless microbial fuel cell. *Biochemical Engineering Journal*, 36(3):209–214, 2007.
- [42] Zhi-Dan Liu, Jing Lian, Zhu-Wei Du, and Hao-Ran Li. Construction of sugar-based microbial fuel cells by dissimilatory metal reduction bacteria. *Chinese Journal of Biotechnology*, 22(1):131–7, 2006.
- [43] Hans-Andrea Loeliger, Felix Tarköy, Felix Lustenberger, and Markus Helfenstein. Decoding in Analog VLSI. *IEEE Communications Magazine*, pages 99–101, April 1999.
- [44] B. E. Logan. *Microbial Fuel Cells*. John Wiley and Sons, Hoboken, New Jersey, 2008.
- [45] D.R. Lovley and K.P. Nevin. Electricity production with electricigens. In J. D. Wall, C. S. Harwood, and A. Demain, editors, *Bioenergy*, chapter 23, pages 295–306. ASM Press, Washington, D. C., 2008.
- [46] K. A. Ludwig, J. D. Uram, J. Yang, D.C. Martin, and D.R. Kipke. Chronic neural recordings using silicon microelectrode arrays electrochemically deposited with a poly(3,4-ethylenedioxythiophene) (pedot) film. *Journal of Neural Engineering*, 3:59–70, 2006.
- [47] M. T. Madigan, J. M. Martinko, P. V. Dunlap, and D. P. Clark. *Biology of Microorganisms*. Pearson Benjamin Cummings, New York, 12th edition, 2009.
- [48] R. Mahadevan, B. Yan, B. Postier, K.P. Nevin, T.L. Woodard, R. O’Neil, M.V. Coppi, B.A. Meth, and J. Krushkal. Characterizing regulation of metabolism in geobacter sulfurreducens through genome-wide expression data and sequence analysis. *OMICS*, 12(1):33–59, 2008.
- [49] S. Mandal and R. Sarpeshkar. A bidirectional wireless link for neural prostheses that minimizes implanted power consumption. In *Proceedings of the IEEE Biomedical Circuits and Systems Conference (BioCAS)*, pages 45–48, 2007.
- [50] S. Mandal and R. Sarpeshkar. Power-efficient impedance-modulation wireless data links for biomedical implants. *IEEE Transactions on Biomedical Circuits and Systems*, 2(4):301–315, December 2008.
- [51] Lynn Margulis. *Origin of Eukaryotic Cells: Evidence and Research Implications for a Theory of the Origin and Evolution of Microbial, Plant, and Animal Cells on the Precambrian Earth*. Yale University Press, New Haven, 1970.

- [52] Lynn Margulis. *Symbiosis in Cell Evolution: Life and its Environment on the Early Earth*. W. H. Freeman and Company, San Francisco, 1981.
- [53] P. Mohseni, K. Najafi, S. J. Eliades, and X. Wang. Wireless multichannel biopotential recording using an integrated fm telemetry circuit. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 13(3):263–271, September 2005.
- [54] A. R. Moller, editor. *Cochlear and Brainstem Implants*, volume 64 of *Advances in Oto-Rhino-Laryngology*. Karger, New York, 2006.
- [55] S. Musallam, B. D. Corneil, B. Greger, H. Scherberger, and R. A. Andersen. Cognitive control signals for neural prosthetics. *Science*, 305:258–262, July 2004.
- [56] K. P. Nevin, H. Richter, S. F. Covalla, J. P. Johnson, T. L. Woodard, A. L. Orloff, H. Jia, M. Zhang, and D. R. Lovley. Power output and columbic efficiencies from biofilms of geobacter sulfurreducens comparable to mixed community microbial fuel cells. *Environmental Microbiology*, 10(10):2505–2514, 2008.
- [57] P. S. Nutman and B. Mosse, editors. *Symbiotic Associations*, volume 13 of *Symbiosis of the Society for General Microbiology*. Cambridge University Press for the Society for General Microbiology, Cambridge, 1963.
- [58] J. Paradiso, G. Borriello, and P. Bonato. Implantable electronics. *IEEE Pervasive Computing*, 7(1):12–13, January 2008.
- [59] B. Pesaran, J. S. Pezaris, M. Sahani, P. P. Mitra, and R. A. Andersen. Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nature Neuroscience*, 5(8):805–811, August 2002.
- [60] B. I. Rapoport. Metabolic factors limiting performance in marathon runners. *Public Library of Science (PLoS) Computational Biology*, 6(10):e1000960, October 2010.
- [61] B. I. Rapoport, J. Kedzierski, and R. Sarpeshkar. An implantable glucose fuel cell for brain–machine interfaces. *Public Library of Science (PLoS) ONE*. Submitted April 2011.
- [62] B. I. Rapoport and R. Sarpeshkar. Cymbionics (symbionics): Rationally engineered symbioses for self-regenerative long-term sensing, actuation, information processing, and energy production. Technology Disclosure to Massachusetts Institute of Technology, Technology Licensing Office Case Number 14167, February 2010.
- [63] B. I. Rapoport, L. Turicchia, W. Wattanapanitch, and R. Sarpeshkar. Efficient universal computing architectures for decoding neural activity. *Public Library of Science (PLoS) ONE*. Submitted April 2011.
- [64] B.I. Rapoport and R. Sarpeshkar. *A Biomimetic Adaptive Algorithm and Micropower Circuit Architecture for Implantable Neural Decoders*, chapter 10, pages 216–254. IGI Global, 2011. In *System and Circuit Design for Biologically-Inspired Learning*, Edited by T. Temel.
- [65] Buddy Ratner, Allan Hoffman, Frederick Schoen, and Jack Lemons, editors. *Biomaterials Science : An Introduction to Materials in Medicine*. Academic Press, San Diego, 2nd edition, 2004.

- [66] David A. Relman. ‘til death do us part’: coming to terms with symbiotic relationships. *Nature Reviews Microbiology*, 6, October 2008.
- [67] David H.S. Richardson, Warwick B. Silvester, Margalith Galun, and Suzanne Trauffer, editors. *Symbiosis*. Balaban Publishers with the International Symbiosis Society, Rehovot, Since 1965.
- [68] Sarah M. Richardson-Burns, Jeffrey L. Hendricks, and David C. Martin. Electrochemical polymerization of conducting polymers in living neural tissue. *Journal of Neural Engineering*, 4:L6–L13, 2007.
- [69] C. Risso, B. Methe, J. Sun, R. Deboy, and D. R. Lovley. Complete genome sequence and genome-based in silico metabolic model of *rhodospirillum rubrum*. *Abstracts of the General Meeting of the American Society for Microbiology*, 108:576–577, 2008. 108th General Meeting of the American-Society-for-Microbiology June 01-05, 2008, Boston, MA, USA.
- [70] Steven S. Saliterman. *Fundamentals of BioMEMS and Medical Microdevices*. Wiley-Interscience. SPIE Press, Hoboken, 2006.
- [71] M. W. Saltzman. *Drug Delivery: Engineering Principles for Drug Therapy*. Topics in Chemical Engineering. Oxford University Press, Oxford, 2001.
- [72] J. C. Sanchez, D. Erdogmus, J. C. Principe, J. Wessberg, and M. A. L. Nicolelis. Interpreting spatial and temporal neural activity through a recurrent neural network brain machine interface. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 13(2):213–219, June 2005.
- [73] G. Santhanam, S. I. Ryu, B. M. Yu, A. Afshar, and K. V. Shenoy. A high-performance brain-computer interface. *Nature*, 442:195–198, July 2006.
- [74] Jan Sapp. *Evolution by Association: A History of Symbiosis*. Oxford University Press, Oxford, 1994.
- [75] R. Sarpeshkar, B.I. Rapoport, and W. Wattanapanitch. Low-power analog-circuit architecture for decoding neural signals. Patent Pending, May 2008. U.S. Patent Application Serial Number 127380.
- [76] R. Sarpeshkar, C. Salthouse, J.-J. Sit, M. W. Baker, S. M. Zhak, T. K.-T. Lu, L. Turicchia, and S. Balster. An ultra-low-power programmable analog bionic ear processor. *IEEE Transactions on Biomedical Engineering*, 52(4):711–727, April 2005.
- [77] R. Sarpeshkar, W. Wattanapanitch, S. K. Arfin, B. I. Rapoport, S. Mandal, M. W. Baker, M. S. Fee, S. Musallam, and R. A. Andersen. Low-power circuits for brain-machine interfaces. *IEEE Transactions on Biomedical Circuits and Systems*, 2(3):173–183, September 2008.
- [78] R. Sarpeshkar, W. Wattanapanitch, B. Rapoport, S. K. Arfin, M. W. Baker, S. Mandal, M. Fee, S. Musallam, and R. Andersen. Low-power circuits for brain-machine interfaces. *Proceedings of the IEEE Symposium on Circuits and Systems*, pages 2068–2071, May 2007.

- [79] R. Sarpeshkar, W. Wattanapanitch, and B.I. Rapoport. Micropower neural amplifier with adaptive input-referred noise. Patent Pending, May 2008. U.S. Patent Application Serial Number 127475.
- [80] R. Sarpeshkar, W. Wattanapanitch, B.I. Rapoport, and S. Mandal. Low-power analog architecture for brain-machine interfaces. Patent Pending, May 2008. U.S. Patent Application Serial Number 127380.
- [81] A. Schwartz, X. T. Cui, D. Weber, and D. Moran. Brain-controlled interfaces: Movement restoration with neural prosthetics. *Neuron*, 56:205–220, 2006.
- [82] J. Seckbach, editor. *Symbiosis: Mechanisms and Model Systems*, volume 4 of *Cellular Origin and Life in Extreme Habitats*. Kluwer Academic Publishers, Boston, 2002.
- [83] Mijail Serruya, Nicholas G. Hatsopoulos, Matthew Fellows, Liam Paninski, and John Donoghue. Robustness of neuroprosthetic decoding algorithms. *Biological Cybernetics*, 88(3):219–228, 2003.
- [84] W. Shain, L. Spataro, J. Dilgen, K. Haverstick, S. Retterer, M. Isaacson, M. Saltzman, and J. N. Turner. Controlling cellular reactive responses around neural prosthetic devices using peripheral and local intervention strategies. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 11:186–188, 2003.
- [85] K. V. Shenoy, D. Meeker, S. Y. Cao, S. A. Kureshi, B. Pesaran, C. A. Buneo, A. R. Batista, P. P. Mitra, J. W. Burdick, and R. A. Andersen. Neural prosthetic control signals from plan activity. *Neuroreport*, 14(4):591–596, March 2003.
- [86] L. Srinivasan, U. T. Eden, S. K. Mitter, and E. N. Brown. General-purpose filter design for neural prosthetic devices. *Journal of Neurophysiology*, 98:2456–2475, 2007.
- [87] D. Tarsy, J.L. Vitek, P.A. Starr, and M.S. Okun, editors. *Deep brain stimulation in neurological and psychiatric disorders*. Current Clinical Neurology. Humana Press, Totowa, 2008.
- [88] D. M. Taylor, S. I. H. Tillery, and A. B. Schwartz. Direct cortical control of 3d neuroprosthetic devices. *Science*, 296:1829–1832, June 2002.
- [89] L. Theogarajan, J. Wyatt, J. Rizzo, B. Drohan, M. Markova, S. Kelly, G. Swider, R. Milan, D. Shider, M. Gingrich, J. Lowenstein, and B. Yomtov. Minimally invasive retinal prosthesis. *Proceedings of the IEEE International Solid-State Circuits Conference*, pages 15–16, 2006.
- [90] Peter J. Turnbaugh, Ruth E. Ley, Micah Hamady, Claire M. Fraser-Liggett, Rob Knight, and Jeffrey I. Gordon. The human microbiome project. *Nature*, 449, October 2007.
- [91] M. Velliste, S. Perel, M. C. Spalding, A. S. Whitford, and A. B. Schwartz. Cortical control of a prosthetic arm for self-feeding. *Nature*, 453(7198):1098–1101, June 2008.
- [92] D. Warland, P. Reinagel, and M. Meister. Decoding visual information from a population of retinal ganglion cells. *Journal of Neurophysiology*, 78(5):2336–2350, 1997.

- [93] W. Wattanapanitch, M. Fee, and R. Sarpeshkar. An energy-efficient micropower neural recording amplifier. *IEEE Transactions on Biomedical Circuits and Systems*, 1(2):136–147, 2007.
- [94] J. Wessberg, C. R. Stambaugh, J. D. Kralik, P. D. Beck, M. Laubach, J. K. Chapin, J. Kim, S. J. Biggs, M. A. Srinivasan, and M. A. L. Nicolelis. Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature*, 408:361–365, November 2000.
- [95] K. D. Wise, D. J. Anderson, J. F. Hetke, D. R. Kipke, and K. Najafi. Wireless implantable microsystems: High-density electronic interfaces to the nervous system. *Proceedings of the IEEE*, 92(1):76–97, January 2004.
- [96] W. Wu, M. J. Black, D. Mumford, Y. Gao, E. Bienenstock, and J. P. Donoghue. Modeling and decoding motor cortical activity using a switching Kalman filter. *IEEE Transactions on Biomedical Engineering*, 51(6):933–942, June 2004.
- [97] W. Wu, Y. Gao, E. Bienenstock, J. P. Donoghue, and M. J. Black. Bayesian population decoding of motor cortical activity using a Kalman filter. *Neural Computation*, 18:80–118, 2006.
- [98] Wei Wu and Nicholas G. Hatsopoulos. Real-time decoding of nonstationary neural activity in motor cortex. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 16(3):213–222, June 2008.
- [99] H. Yang and R. Sarpeshkar. A bio-inspired ultra-energy-efficient analog-to-digital converter for biomedical applications. *IEEE Transactions on Circuits and Systems I (Special Issue on Life Sciences and System Applications)*, 53(11):2349–2356, November 2006.