

**DYNAMICS OF EMBODIED DISSOCIATED CORTICAL
CULTURES FOR THE CONTROL OF
HYBRID BIOLOGICAL ROBOTS**

A Dissertation
Presented to
The Academic Faculty

by

Douglas James Bakkum

In Partial Fulfillment
of the Requirements for the Degree
of Doctor of Philosophy in the
Interdisciplinary Bioengineering program
Department of Mechanical Engineering

Georgia Institute of Technology
Emory University
December 2007

Copyright © Douglas Bakkum 2007

**DYNAMICS OF EMBODIED DISSOCIATED CORTICAL
CULTURES FOR THE CONTROL OF
HYBRID BIOLOGICAL ROBOTS**

Approved by:

Dr. Steven M. Potter, Advisor
School of Biomedical Engineering
Georgia Institute of Technology

Dr. Robert J. Butera
School of Electrical and Computer
Engineering
Georgia Institute of Technology

Dr. Thomas D. DeMarse
School of Biomedical Engineering
University of Florida

Dr. Stephan P. DeWeerth
School of Biomedical Engineering
Georgia Institute of Technology

Dr. Eric Schumacher
School of Psychology
Georgia Institute of Technology

Date Approved: October 11, 2007

To my parents.

ACKNOWLEDGEMENTS

I often get asked how I came from studying Mechanical Engineering into researching basic Neuroscience, a turn I am very grateful I made and am grateful to those who have influenced me along the way.

Upon entering the University of Wisconsin Madison as an undergraduate, I deliberated between majoring in Psychology, Art, or Mechanical Engineering, ultimately choosing Mechanical Engineering. Learning about the brain and thought processes was enticing, but I determined not worthwhile if becoming a Psychologist meant listening to other peoples' problems all day (which turned out to be a limited viewpoint). Art, I figured, came from within, not from a lecture, and could be done anytime, although I found that time is not always available. Thus I chose Mechanical Engineering in order to learn about how the physical world works and how physics can be applied to create devices. Fortuitously, I was later able to work on a combination of each with the Meart project.

So first of all, I am grateful to the University of Wisconsin for providing excellent teachers, a diversity of classes to explore my interests, and an affordable education. I am grateful to the friends I met (Walt, Rick, Cory, Melissa, Rusty, Mike, Laura, and everyone else) who nurtured my curiosity with late night discussions. I am grateful to the department of Mechanical Engineering for providing a well respected degree and an analytic mind necessary to be successful as a researcher. In particular, undergraduate independent study in robotics with Dr. Ferrier opened my eyes to research in my senior year and motivated me to apply to graduate school.

I applied to research Mechanical Engineering and robotics in many schools, but was only accepted to those who did not read my essays, basing their decision on my grades and test scores. A large motivation to enter graduate school was also that I wasn't ready to have a boss every day. I didn't have a vision for my research or even a good reason to be in graduate school, which was reflected in the essays. However, I am grateful to Dr. Bill Wepfer, the graduate coordinator of Mechanical Engineering at Georgia Tech, for many reasons, but firstly for recruiting me. He is an alumni of Wisconsin and was keen to enroll his first Wisconsin recruit. I was not interested, but I did like the free trip to visit the campus, and was sold after discovering I could spend my first year at the campus in Metz, France. Again, not necessarily the right reasons. The graduate application at Georgia Tech consisted of mailing two parts: the essays and the academic record, and thankfully the essays did not arrive due to insufficient postage.

Nevertheless, I learned much about what constitutes research, and my education gave me direction. Mechanical Engineering was fun to learn, but not so fun to research. Consider a research field to be like a tree. Mechanical Engineering has been around for many years, since at least Newton. A discovery requires learning about all of the roots of the tree, the thick trunk, and following a branch to its end. Then the discovery is but a small bud at the outskirts, hence the phrase "designing a better bolt".

I was lost and ready for a change. While I was in the purgatory between finishing Mechanical Engineering early with a Masters degree and finding my current research in Bioengineering, Dr. Wepfer helped enormously by giving me advice, encouragement, and finding a Teaching Assistantship with Dr. Imme Ebert-Uphoff. I am also grateful to her for her advice and encouragement. She and Dr. Wepfer lead me to contact Dr. Ron

Arkin, a well respected roboticist in the Computer Science department. He was my first real research advisor, and I have gained valuable knowledge by watching how he ran his laboratory. My labmate Endo helped catch me up in the research and, along with Ron, gave a first view into alternative ways to study the brain (i.e. artificial intelligence and in particular the necessity of embodiment to cognition).

Upon returning to Atlanta, I was also fortunate enough to take classes to earn a certificate of cognitive science. In particular, I am grateful to Dr. Nancy Nercessian. Due to her teaching style, I learned more in my first class with her than in any other class I have taken. She opened my eyes to the many ways to study the brain (without the need to listen to other peoples' problems). The tree of research for brain studies is young. It has less developed roots, and just as many discoveries lead to the growth of whole branches as to the growth of small buds. This makes the research exciting, and I now had a direction. I applied to a variety of graduate programs in cognitive science to study artificial intelligence and brain or neuron models, and this time was accepted to each.

During the applications, Ron suggested I contact Dr. Steve Potter, a new professor in Bioengineering, who embodied living neurons with robotics. I figured if I wanted to study intelligence, why not study real intelligence. I emailed Steve and soon discovered we had very similar research philosophies. My path now led from Mechanical Engineering to Robotics/AI to Cognitive Science and finally to basic Neuroscience. I never thought I would find what I would love to do, but I have, and I owe that to Steve for following through on his vision and for giving me an opportunity.

Of course the path between beginning and finishing a Ph.D. can wind and twist even further. I am grateful to Steve, for being patient with my missing progress reports,

being open-minded about even the most illogical hypothesis (encouraging creativity and discovery), supporting forays into the art world with Meart, and being a friend, among many other reasons. I hope to carry out my future pursuits with the same enthusiasm and confidence he has. I am grateful to Dr. Tom DeMarse for much assistance and advice, at first as a post-doc in the lab and later as a committee member on my thesis. I am grateful to all my labmates, the post-docs, and the many undergraduates who have passed through the lab and become friends: Mark, Peter, John, Komal, Yixiao, Bhavesh, Ryan, Borris, Bobby, Blythe, Chu, Nathan, among others. I look forward to climbing the Himalayas with Radhika Madhavan and watching Zenas Chao's movies, two people who have become very close from sharing five years of the ups and downs of qualification exams, grants, and thesis writing. Our relationship reminds me of slime mold: we navigated our Ph.D.'s sometimes moving on our own and more often joined together as a symbiotic entity. John Ross, Edgar Brown, and Nakul Reddy have also been a joy to collaborate with on calcium imaging and neurostimulation experiments. I am grateful to have met and traveled the world with Guy Ben-Ary and Phil Gamblen (and his wife Dawn) through our work on Meart. I am also grateful that they forced me to eat sushi for the first time (there isn't much sushi in Wisconsin). Our always entertaining conversations put my research in perspective and gave me new ideas to explore. Working with Meart taught me how to be comfortable and speak in public. My wife, Rebekha, has made me a more complete and less isolated person through her contagious warmth. Her urging and support has helped me to complete my studies.

I am grateful to the other Neurolab professors for advice during my thesis and providing a diverse and warm work environment. I thank the Center for Behavioral

Neuroscience for acting like an extended family and providing financial support my last years of study, and also thank the National Institutes of Health, the National Science Foundation, and the Whitaker Foundation for supporting science. I would like to acknowledge former Neurolab manager Brian Williams and Potter lab managers Douglas Swelha, Pooja, and Dustin for making graduate life easier. Also, Kacy Cullen and the rest of Dr. Michelle Laplaca's laboratory have shared tissue, cell culture supplies, and their knowledge with our lab. I would like to thank Chris Ruffin, Amber Burris, Sandra Wilson, Rudy, and Penelope Pollard for getting me through many administrative hurdles.

Above all else, I am grateful to my parents, who have supported me in all my decisions and diversions with their love and attention.

Douglas J. Bakkum

Atlanta, November 13th, 2007

TABLE OF CONTENTS

	Page
ACKNOWLEDGEMENTS.....	IV
LIST OF FIGURES	XI
LIST OF ABBREVIATIONS.....	XIII
SUMMARY	XVI
CHAPTER 1. INTRODUCTION.....	1
CHAPTER 2. REMOVING SOME ‘A’ FROM AI.....	9
2.1 INTRODUCTION	9
2.2 EXAMPLES: THREE EMBODIED NEURAL SYSTEMS	13
2.2.1 Living neurons control a simulated animal	14
2.2.2 Living neurons control a mobile robot.....	16
2.2.3 Living neurons control a drawing arm.....	18
2.3 DISCUSSION.....	21
2.3.1 Embodying cultures: theory.....	21
2.3.2 Summary: integrating brain, body, and environment.....	26
CHAPTER 3. TETANUS INDUCED PLASTICITY AND REGION SPECIFIC CHANGES IN NETWORK ACTIVITY.....	28
3.1 INTRODUCTION	29
3.2 METHODS.....	32
3.2.1 Stimulation protocols.....	32
3.2.2 Neural activity statistics.....	33
3.2.3 Quantifying plasticity	34
3.3 RESULTS.....	35
3.3.1 Preliminary principal components analysis of plasticity.....	35
3.3.2 Incorporating spatial and temporal information: the center of activity trajectory (CAT).	37
3.4 DISCUSSION.....	41
3.4.1 Statistics of functional plasticity in extracellular multi-electrode recordings	41
3.4.2 Region-specific plasticity.....	42
3.4.3 Plasticity versus spontaneous bursting.....	43
3.4.4 Conclusion	44
CHAPTER 4. PLASTICITY INDUCED BY PATTERNED STIMULATION: CHANGES IN ACTION POTENTIAL PROPAGATION.....	45
4.1 INTRODUCTION	46
4.2 RESULTS.....	49
4.2.1 Propagation of direct electrically-evoked action potentials robustly detected using MEAs.	50
4.2.2 Dependence of signal propagation on neural activity and variation of stimulation pattern.	54
4.2.3 Non-synaptic expression.....	56
4.2.4 Theoretical computational capacity.	59
4.3 DISCUSSION.....	60
4.4 MATERIALS AND METHODS	61
4.4.1 Pharmacology.	62
4.4.2 Experiment parameters.	62
4.4.3 Automated detection of direct electrically-evoked action potentials (dAPs).	62
CHAPTER 5. SHAPING EMBODIED NEURAL NETWORKS FOR ADAPTIVE GOAL-DIRECTED BEHAVIOR.....	65

5.1 INTRODUCTION	66
5.2 METHODS:	69
5.2.1 Closed-loop training algorithm	69
5.2.2 Experiment design	74
5.3 RESULTS AND DISCUSSION	76
5.3.1 Training shifted neural activity towards the desired activity	76
5.3.2 Changes in motor output arose from neuronal plasticity, not an elastic dependency on stimulation history	80
5.3.3 Training required different PBS at different times.....	84
5.4 CONCLUSION.....	86
CHAPTER 6. FUTURE STUDIES	88
6.1 PROGRAMMABLE ARRAY MICROSCOPY (PAM) TO EVOKE ACTIVITY IN CHANNELRHODOPSIN-2 (CHR2) EXPRESSING NEURONS	89
6.2 FUTURE EMBODIMENTS.....	92
6.3 EXPLAINING DAPS.....	94
APPENDIX A. MEMOIRS OF A CYBORG ARTIST.....	96
A.1 INTRODUCTION	96
A.2 METHODS: MAKING THE SEMI-LIVING ARTIST.....	99
A.2.1 Preparing and caring for Meart's brain	100
A.2.2 Meart's body	103
A.2.3 Software development and Experimental design	106
A.3 RESULTS:	108
A.4 DISCUSSION	114
A.4.1 Art vs Science	114
A.4.2 Intelligence and embodiment	116
A.4.3 Nature of art and being an artist.....	118
A.4.4 Fear and the future	120
A.4.5 Conclusion	121
APPENDIX B. HOMEOSTASIS OF GLOBAL FIRING RATE	122
B.1 INTRODUCTION	122
B.2 METHODS	122
B.3 RESULTS	123
APPENDIX C. CLOSED-LOOP MOTH EXPERIMENT	125
APPENDIX D. USING THE NEUROMANCER STIMULATION BOARD FOR ARTIFACT SUPPRESSION	128
D.1 INTRODUCTION	128
D.2 ARTIFACT SUPPRESSION OF THE NEUROMANCER BOARD COMPARED TO THE RACS BOARD, ON MEAS PLATED WITH NEURONS.....	129
D.3 UPDATE SPEED AND TIMING ACCURACY OF THE NEUROMANCER STIMULATION BOARD FOR CLOSED-LOOP EXPERIMENTS.....	134
D.4 METHODS.....	137
APPENDIX E. HOW TO CODE CLOSED-LOOP EXPERIMENTS	138
APPENDIX F. FORMATING AND ANALYZING LARGE SPIKE DATA SETS.....	141
APPENDIX G. GENERAL METHODS	144
G.1 CELL CULTURING	144
G.2 ELECTROPHYSIOLOGY AND DATA ACQUISITION	144
G.3 CHANGE TO DRIFT RATIO, C/D	145
REFERENCES	146

LIST OF FIGURES

	Page
FIGURE 1.1 MICROSCOPE IMAGES OF GROWING AXONS AND DENDRITES.	3
FIGURE 2.1 CONNECTING NEURONS TO MULTI-ELECTRODE ARRAYS.....	11
FIGURE 2.2 HYBROT (HYBRID LIVING+ROBOTIC) SETUP.....	12
FIGURE 2.3 ANIMAT SETUP AND ACTIVITY.	15
FIGURE 2.4 LIVING NEURONS CONTROL A MOBILE ROBOT.	16
FIGURE 2.5 MEART–THE SEMI-LIVING ARTIST.....	19
FIGURE 3.1 WHOLE-INPUT–OUTPUT (WIO) VECTORS FOR ANALYZING PERFORMANCES OF DIFFERENT STATISTICS.....	34
FIGURE 3.2 SIMULTANEOUS TETANIZATION OF TWO ELECTRODES PRODUCED A LASTING CHANGE IN NETWORK RESPONSE PROPERTIES.	36
FIGURE 3.3 COMPARISON OF THE PERFORMANCE OF DIFFERENT STATISTICS ACROSS A TETANIZATION IN LIVING MEA CULTURES.	39
FIGURE 3.4 COMPARISON OF CAT AND CAT-ELS (ELECTRODE LOCATION SHUFFLED).....	40
FIGURE 4.1 MULTI-ELECTRODE ARRAYS (MEA) ROBUSTLY DETECTED DAPs.....	51
FIGURE 4.2 NEURAL ACTIVITY RECORDED ON ONE EXTRACELLULAR ELECTRODE IN RESPONSE TO STIMULATION AT ANOTHER CONSISTS OF AN EARLY DIRECTLY-EVOKED ACTION POTENTIAL (DAP) PHASE AND A LATER SYNAPTICALLY-EVOKED ACTION POTENTIAL (SAP) PHASE.....	52
FIGURE 4.3 DAPs WERE TIME-LOCKED TO BIPHASIC VOLTAGE STIMULI.	53
FIGURE 4.4 ACTION POTENTIAL PROPAGATION DEPENDED ON ONGOING NEURAL ACTIVITY AND STIMULATION PATTERN.....	55
FIGURE 4.5 THE CAUSE OF PLASTICITY WAS NEURAL ACTIVITY, AND THE SITE OF PLASTICITY WAS NOT SYNAPTIC.....	57
FIGURE 4.6 DAPs WERE AUTOMATICALLY DETECTED AND TRACKED FROM PEAKS AND VALLEYS IN CONSECUTIVE SMOOTHED POST-STIMULUS TIME HISTOGRAMS.	64
FIGURE 5.1 SCHEMATIC OF THE CLOSED-LOOP FEEDBACK AND ADAPTIVE TRAINING.....	70
FIGURE 5.2 STIMULATION SEQUENCE.....	71
FIGURE 5.3 NEURAL RESPONSE TO CLOSED-LOOP AND OPEN-LOOP TRAINING.....	78
FIGURE 5.4 LEARNING CURVES FOR ALL CLOSED-LOOP EXPERIMENTS.....	79
FIGURE 5.5 AVERAGE NORMALIZED LEARNING CURVES OF ALL CLOSED-LOOP AND OPEN-LOOP EXPERIMENTS.	80

FIGURE 5.6 LONG-TERM PLASTICITY OF MOVEMENT DIRECTION.....	82
FIGURE 5.7 PLASTICITY INDUCED BY CLOSED-LOOP TRAINING LASTED FOR 80 MINUTES.	83
FIGURE 5.8 TRAINING REQUIRED DIFFERENT PTS AT DIFFERENT TIMES.	85
FIGURE 6.6.1 SCHEMATIC OF PROGRAMMABLE ARRAY MICROSCOPY TO EVOKE CHANNELRHODOPSIN-2 EXPRESSING NEURONS.	89
FIGURE 6.6.2 RASTER PLOT OF EVOKED RESPONSES DETECTED ON ONE ELECTRODE IN RESPONSE TO REPEATED STIMULATION ON ANOTHER ELECTRODE.....	95
FIGURE A.1 MEART’S BODY.	97
FIGURE A.2 MEART’S BRAIN.....	98
FIGURE A.3 SCHEMATIC OF THE BIO-ROBOTIC SOFTWARE ALGORITHMS AND HARDWARE, I.E. MEART’S COMPONENTS.....	100
FIGURE A.4 LIFE-SUPPORT SYSTEM FOR MEART’S BRAIN.....	101
FIGURE A.5 THE BODY OF MEART AT THE MOSCOW BIENNALE AND DRAWINGS.....	102
FIGURE A.6 ACCURACY TEST OF THE ROBOTIC DRAWING MACHINE.	105
FIGURE A.7 PLASTIC CHANGES IN MEART AND ANIMAT BEHAVIOR.....	110
FIGURE A.8 NEURONAL PLASTICITY.	112
FIGURE A.9 TRAINING SERIES AND LEARNING CURVE FOR THE SIMULATED ANIMAT.	113
FIGURE A.10 DOES MEART CREATE VALUABLE ART? MEART DRAWING AND NOTES FROM AN EARLY ACCURACY TEST.	119
FIGURE B.1 NETWORK FIRING RATE (FR) AND DAP LATENCY ADAPTED TO STIMULATION PATTERN.	124
FIGURE C.1 CLOSED-LOOP MOTH EXPERIMENT.....	124

LIST OF ABBREVIATIONS

AI	Artificial intelligence
AMPA-R	Alpha-amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid receptor
APV	2-amino-5-phosphonovaleric acid
BMI	Bicuculline methiodide
Ca ⁺⁺	Calcium ion
CA	Center of activity
CAT	Center of activity trajectory
CAT-ELS	Center of activity trajectory with electrode locations shuffled
CCD	Charge coupled device (camera)
CCH	Cross-correlation histogram
C/D	Change / Drift ratio
ChR2	Channelrhodopsin-2
CNQX	6-cyano-7-nitroquinoxaline-2, 3-dione
CPS	Context-control probing sequence
CW	Center of weights
dAP	Direct electrically-evoked action potentials
DBS	Deep brain stimulation
DLP	Digital light processor
DST	Dynamic systems theory
FES	Functional electrical stimulation
fMRI	Functional magnetic resonance imaging
FR	Firing rate

FRH	Firing rate histogram
GABA-R	Gamma-aminobutyric acid receptor
GOF AI	Good old fashioned artificial intelligence
GUI	Graphical user interface
IPI	Inter-pulse-interval
ISI	Inter-stimulus-interval
JPSTH	Joint peri-stimulus time histogram
LTP	Long-term potentiation
MASC	Mean absolute synaptic change
MEA	Multi-electrode array
MEART	Multi-electrode array art
MCS	Multi-Channel Systems
Mg ⁺⁺	Magnesium ion
MI	Mutual information
NCA	Neurally-controlled animat
NMES	Neuromuscular electrical stimulation
NMDA-R	N-methyl-D-aspartic acid receptor
PAM	Programmable array microscopy
PCA	Principal Components Analysis
PET	Positron emission tomography
PID	Proportional-integral-derivative (controller)
PKC	Protein kinase C
PTS	Patterned training stimulation
RBS	Random background stimulation
RPS	Random probing sequence

sAP	Synaptically-evoked action potential
SBS	Shuffled background stimulation
SCCC	Shift-predictor corrected cross-correlogram
SEM	Standard error mean
STDP	Spike-timing-dependent plasticity
TCP/IP	Transmission control protocol / internet protocol
TTX	Tetrodotoxin
WIO vector	Whole-input-output vector

SUMMARY

One currency of the brain is action potentials. They relay sensations of the world into commands for muscles to move the body and produce speech, and their manipulation is responsible for adaptive behavior and cognition. However, much remains unknown about the fundamental rules governing such neural information processing in the brain. By growing networks of cortical neurons and glia over multi-electrode arrays (MEA), which can be used to both stimulate and record multiple neurons in parallel over durations up to months, a 2-way communication with neuronal network activity becomes feasible. In particular, I was interested in embodying these networks with robotics to study the importance of environmental interaction, or behavioral feedback, in neural processing. Here, the recorded activity of the neurons was transformed into movements in an artificial environment, and sensory feedback was transformed into electrical stimulation on multiple electrodes. Stimulation influences neural activity and in turn the subsequent movements, creating a closed-loop system we call a neurally-controlled animat.

My ultimate goal was to develop animats that could learn something about the environment and/or body given to them. Upon entering the lab, the technology to culture neurons for long durations on MEAs, robustly record neural activity, and stimulate an MEA's electrodes was recently achieved. However, the crucial ability to induce and detect neural plasticity was missing: methods were needed to determine appropriate sensory-motor mappings and training algorithms in order to produce any kind of adaptive behavior. I took a step back to first determine, in open-loop experiments, what types of

stimulation could induce plasticity and what kinds of activity statistics could identify plasticity. This knowledge was then applied to construct embodied systems.

To paraphrase the results, most any stimulation could induce neural plasticity, while the inclusion of temporal and/or spatial information about neural activity was needed to identify plasticity. Following a tangent from one of the open-loop experiments, the plasticity of action potential propagation was observed. This is a notion counter to the dominant theories of neuronal plasticity that focus on synaptic efficacies and is suggestive of a vast and novel computational mechanism for learning and memory in the brain.

The results from the open-loop experiments were next used to develop animats that achieved adaptive goal-directed behavior. The feedback of patterned training stimuli, contingent on behavioral performance, was found able to sculpt the network activity into desired states. Network plasticity was not just induced, but could be customized, suggesting a potential role for the rehabilitation of neural pathologies. Furthermore, in collaboration with artists from SymbioticA at the University of Western Australia, neurons were embodied with a robotic drawing machine and exhibited at galleries throughout the world. This provided a platform to educate the public and initiate critical discussions of biotechnology.

CHAPTER 1

INTRODUCTION

The brain is arguably the most complex system studied in science. A human brain contains about one hundred billion neurons, each of which communicate electrically and chemically with tens of thousands of other neurons to make up about a quadrillion (10^{15}) synaptic connections. The patterns and strengths of connections constantly change with experience and no two brains are alike. Neural action potentials filter through the connections and give rise to perceptions, interpretation, memories, consciousness, imagination, language, and ultimately, actions. Quantifying how humans, animals, and recently algorithms (artificial intelligence) can make use of such complexity to successfully live in the world comprises a vast array of studies on intelligence, and in particular, our lab's own niche in studies of embodied cortical networks.

Traditional neuroscience has studied the brain at many levels. At a global level, functional imaging techniques such as fMRI and PET are used to correlate behaviors with representational brain areas, consisting of tens of thousands of neurons. Experiments at the cellular level investigate the molecular mechanism and functional rules of cellular plasticity. However, the flow of information through the brain's interconnected neural networks could not be accessed in detail until the relatively recent adoption of multi-unit functional stimulation and recording technology, which remain fringe techniques.

We directly investigate the neural population dynamics that give rise to the brain's unique computational abilities and intelligence by growing networks of cortical neurons and glia over multi-electrode arrays (MEA) (see Appendix G for general methods). We embody the networks with a simulated or real body situated in an artificial environment, creating an "animat", to observe learning and behavior in concert with detailed electrophysiology (DeMarse, Wagenaar et al. 2001; Shkolnik 2003; Bakkum,

Shkolnik et al. 2004a; Chao, Bakkum et al. 2005). The MEA allows for both the stimulation and recording of multiple neurons in parallel over durations of months (Pine 1980; Gross, Williams et al. 1982; Potter and DeMarse 2001a). Cortical neurons are attractive because the cortex can be considered to be a type of general processing engine that is responsible for learning and memory in the brain. For example, altering sensory input of thalamic relays to cortical areas caused the cortex to develop structure appropriate to the new type of the sensory input (Sur, Garraghty et al. 1988). Intracellular electrodes, while allowing the direct measurement of a neuron's membrane voltage, were not desired because their recordings are short-term (a few hours at most) and interfere with the neuron's physiology due to dialysis of its cytosol and a rupture in its membrane, which eventually leads to cell death. An *in vitro* preparation was preferred over those *in vivo* because of better pharmacological access and control of variables: the confounding issues of competing goals, such as hunger or attention, must be accounted for during *in vivo* experiments but are absent here. Research in AI assesses network computation by creating algorithms for artificial networks, and has been a complimentary tool in our lab (Chao, Bakkum et al. 2005; Chao, Bakkum et al. 2007). The use of biological networks is conceptually simpler in one sense because our "control algorithm" has already been greatly developed through evolution. A better understanding of the processes leading to biological intelligence can, in turn, help create fundamentally different computational systems, through new AI algorithms and their hardware implementations or by using biology itself in control systems.

The fields of cognitive science and artificial intelligence began inquiries into the nature of intelligence in the middle of the last century without a concern for its substrate: intelligent thought was considered the manipulation of abstract concepts, analogous to how digital computations could be run on any manner of Turing machine. However, intelligence has not yet been attributed to digital computers or the robots they have been used to control despite the fact that the units (neurons) of an "intelligent" brain relay their

signals (action potentials) at least 6 orders of magnitude slower, on the order of milliseconds. Tasks trivial to humans have proven difficult for computers, such as adaptation, multitasking, real-time performance, fault tolerance, handling incomplete or occluded sensation, etc. This is likely due to significant differences in computational implementation, with brains using massively parallel processing, feedback loops on many scales, and relay switches (neurons) that learn and change function, among other differences.




Figure 0.1 Microscope images of growing axons and dendrites.

The images were taken on three consecutive days beginning the second day after plating the cells. The black circles are the electrodes. Even left to themselves without external input other than cell culture media, neurons re-establish connections with their neighbors and begin communicating electrically and chemically within days, demonstrating an inherent drive to network; the essence of neural network processing remains without the confounding behavioral drives present *in vivo*. (Pictures from P. Passaro)

Now becoming more accepted is the hypothesis that intelligence is not disembodied, but intimately entwined with the mechanics of the body and an interaction with the environment (Varela, Thompson et al. 1993; Pfeifer and Bongard 2007). The act of walking combines roles for neural signaling, proprioceptive feedback, the spring tension of muscles, the friction of shoes contacting pavement, and gravity to assist leg swing. Both our brains and bodies co-evolved to take advantage of the physics in the world, and the body is thought to be not just a medium possessed by the brain to behave and gather data, but instead another layer to filter behaviorally relevant information (Pfeifer and Bongard 2007). Neuroscientists consider the brain as the seat of intelligence. The ancient philosopher Aristotle imagined it to lie in the heart, and ancient Egyptians

thought so much of the brain that they threw it away prior to mummification. Maybe future generations will no longer throw away the body and its interaction with the environment in considering where lies intelligence.

Our ultimate goal was to follow the embodiment perspective and develop animats that exhibit adaptive goal-directed behavior, a trait at least one group considers to be the defining feature of intelligence (Sternberg and Salter 1988). When entering the lab, the technology to culture neurons for long durations on MEAs (Potter and DeMarse 2001b), robustly record activity (Wagenaar, DeMarse et al. 2005), and stimulate all electrodes (Wagenaar and Potter 2004a) was recently achieved. However, the crucial ability to induce and detect neural plasticity was missing: methods were needed to determine appropriate sensory-motor mappings and training algorithms in order to produce any kind of adaptive behavior. This required taking a step back and determining first, in open-loop experiments, what types of stimulation could induce plasticity and what kinds of activity statistics could identify plasticity (Chapters 3 and 4). We then applied this knowledge to construct embodied systems (Chapters 5 and Appendix A, with an overview in Chapter 2). To paraphrase the results, most any stimulation could induce neural plasticity, while the inclusion of temporal and/or spatial information about neural activity was needed to identify plasticity. In Chapter 5, adaptive goal-directed behavior in animats was achieved by using patterned training stimuli, contingent on animat performance, to sculpt the network into behaviorally appropriate functional states: network plasticity was not just induced, but could be customized. Appendix A discusses progress (mainly engineering) towards moving the embodiment from an artificial world to the real world using robots. Chapter 2 provides an introduction of our philosophy and history of using animats. Chapter 6 concludes with a discussion of future directions, written as a guide to future students who wish to build upon the work begun here.

Chapter summaries

Chapter 2, *Removing some 'A' from AI*, overviews earlier attempts to embody the cortical networks and discusses in further detail our scientific motivations.

Experiments in Chapter 3, *Tetanus-Induced plasticity and Region-Specific Changes in Network Activity*, used tetanic stimulation to induce neural plasticity. To observe the plasticity, we developed statistics that included spatial and temporal information about neural activity, called the center of activity trajectory (CAT). While comparisons of firing rates show plasticity in intracellular recordings, more detailed statistics incorporating spatiotemporal population activity patterns were needed to reveal plasticity in extracellular multi-electrode recordings. Synaptic noise across a chain of neurons (Kandel, Schwartz et al. 2000), convergent and divergent pathways (Abeles 1991), and homeostatic mechanisms that re-adjust firing rates in response to plasticity (Turrigiano and Nelson 2000; Spitzer, Borodinsky et al. 2005) all obscure firing rate measures of plasticity detected by extracellular MEAs. CAT was found to better detect tetanus-induced plasticity in both simulated and living networks than traditional statistics.

Experiments in Chapter 4, *Plasticity Induced by Patterned Stimulation and Changes in Action Potential Propagation*, were designed to test the ability of patterned stimulation to induce neural plasticity. To observe the plasticity, we quantified changes in the propagation delays of action potentials as they traveled through axons. From the viewpoint of a neuron, the specific timings of afferent synaptic potentials determines whether or not and when to fire an action potential. Tuning such input would provide a powerful mechanism to adjust the neuron's behavior and perhaps play a large role in learning and memory in general. However, axonal plasticity of action potential timing is counter to conventional notions of stable propagation and to the dominant theories of plasticity focusing on synaptic efficacies. Here we show activity-dependent plasticity of action potential propagation could indeed occur. We used a multi-electrode array to

induce, detect, and track changes in propagation in multiple neurons while they adapted to different patterned stimuli in controlled in vitro neocortical networks. Changes occurred in delays (up to 4ms or 40% after minutes and 13ms or 74% after hours) and amplitudes (up to 87%). The changes did not occur when the same stimulation was repeated while blocking synaptic activity. Even though plasticity depended on synaptically evoked action potentials, its expression was non-synaptic: action potential propagation. We concluded that propagation plasticity is a cellular mechanism underlying information processing in neuronal networks and potentially learning and memory in the brain.

Experiments in Chapter 5 were designed for *Shaping Embodied Neural Networks for Adaptive Goal-Directed Behavior*. Spatiotemporal patterns inspired by Chapter 4 were used as training signals, while the center of activity (CA instead of CAT; Chapter 3) was used to command motor output. We developed an adaptive training algorithm, whereby cortical networks learned to modulate their dynamics and achieve user-defined activity states within tens of minutes. *A priori* knowledge of functional connectivity was not necessary. Instead effective training sequences were continuously discovered and refined based on real-time feedback of animat performance. The short-term dynamics in response to training became engraved in the network, requiring fewer training stimuli later in time to achieve the same results. After 2 hours of training, plasticity was significantly greater than baseline for 80 min ($P < 0.05$). Interestingly, a given sequence of stimuli did not induce detectable plasticity, let alone desired activity, when replayed to the network, demonstrating the importance of context dependent feedback. The chapter discusses the potential for targeted electrical stimulation of the brain, contingent on the activity of the body or even of the brain itself, to treat neurological movement disorders.

Chapter 6 concludes with future directions, written as a guide to students who wish to build upon the work begun here. These include using light to evoke channelrhodopsin-2 expressing neurons with high spatiotemporal resolution,

investigating the molecular mechanisms and functional rules of changes in action potential propagation, and finding new horizon's for embodiment experiments.

Appendix A, *Memoirs of a Cyborg Artist*, describes the engineering and philosophy behind building a robotic drawing machine, named Meart, as part of an art-science collaboration. Experiments originally designed for Meart founded the experimental design in Chapters 4 and 5. In Chapter 5, a simulated body was used while improving the sensory-motor mappings and training algorithms. However in order to engineer biologically-based control systems, their performance must be tested in the real world where noise and non-linearity are commonplace. In the case of Meart, we applied patterned training stimuli and tested the neuronal network's ability to learn the dynamics of its body to achieve a goal-directed behavior of drawing a geometrical shape. The transformation from visual sensation into the training stimuli was fixed, and while neural plasticity occurred, successful behavior did not. Either prior knowledge of network connectivity or adaptive training was needed. A future step is to apply the algorithms developed with the simulated animat to control the drawings of Meart, creating a real-world biologically-based agent exhibiting goal-directed behavior.

Appendices B and C discuss two unpublished preliminary experiments, *Homeostasis of global firing rate*, and a *Closed-loop 'moth' experiment*. Appendix D describes testing of a custom built MEA stimulation board with artifact suppression (Neuromancer), developed by the DeWeerth laboratory. Appendices E and F overview code needed to conduct embodiment experiments and code to format and analyze spike data. Appendix G describes general methods for cell culturing, electrophysiology, and the common change-to-drift ratio (C/D) used to quantify plasticity. For further details, refer to the theses by Daniel Wagenaar and Radhika Madhavan for protocols to mix all solutions needed for cell culturing, using Meabench to record neural activity, and electrically stimulating the neural networks.

Knowledge gained from our studies could provide insight for many fields. Clinically, understanding the relationships between electrical stimulation and neural activity and the functional expression of neural plasticity would help direct the development of neuroprosthetics. Considering our research is of a basic science nature, unforeseen applications may arise: for example, a recent collaboration was formed to apply a stimulation paradigm developed in the lab (Wagenaar, Madhavan et al. 2005) to treat epileptic patients. Likewise, benefit may come to other fields including psychology and educational instruction (Geake and Cooper 2003), the development of artificial intelligence (Bakkum, Shkolnik et al. 2004a), telecommunications (Clarkson 1999), the evolution of mathematics and statistics such as in dynamical systems and emergence (Beer 2000) and multivariate statistical analysis (Lee and Seung 1999), philosophy of the mind (Manson 2004a), and art (Bakkum, Shkolnik et al. 2004a).

CHAPTER 2

REMOVING SOME ‘A’ FROM AI¹

We embodied networks of cultured biological neurons in simulation and in robotics. This is a new research paradigm to study learning, memory, and information processing in real time: the Neurally-Controlled Animat. Neural activity was subject to detailed electrical and optical observation using multi-electrode arrays and microscopy in order to access the neural correlates of animat behavior. Neurobiology has given inspiration to AI since the advent of the perceptron and consequent artificial neural networks, developed using local properties of individual neurons. We wish to continue this trend by studying the network processing of ensembles of living neurons that lead to higher-level cognition and intelligent behavior.

2.1 Introduction

We present a new paradigm for studying the importance of interactions between an organism and its environment using a combination of biology and technology: embodying cultured living neurons via robotics. From this platform, explanations of the emergent neural network properties leading to cognition are sought through detailed electrical observation of neural activity. A better understanding of the processes leading to biological cognition can, in turn, facilitate progress in understanding neural pathologies, designing neural prosthetics, and creating fundamentally different types of artificial intelligence. The Potter group is one of seven in the Laboratory for Neuroengineering (Neurolab) at the Georgia Institute of Technology, all working at the interface between neural tissue and engineered systems. We envision a future in which mechanisms employed by brains to achieve intelligent behavior are also used in artificial systems; we overview three preliminary examples of the Neurally-Controlled Animats approach below. By using biology directly, we hope to remove some of the 'A' from AI.

¹ Adapted from:
Bakkum DJ, Shkolnik AC, Ben-Ary G, Gamblen P, DeMarse TD, Potter SM, “Removing some of the ‘A’ from AI: embodied cultured networks”. In Proceedings of the Dagstuhl Conf on Embodied Artificial Intelligence, eds. Luc Steels and Rolf Pfeiffer, Springer, 2004

No one would argue that environmental interaction, or embodiment, is unimportant in the wiring of the brain; no one is born with the innate ability to ride a bicycle or solve algebraic equations. Practice is needed. An individual's unique environmental interactions lead to a continuous 'experience-dependent' wiring of the brain (Weiler, Hawrylak et al. 1995). This makes evolutionary sense as it is helpful to learn new abilities throughout life: if there are some advantageous features of an organism that can be attained through learning, then the ability to learn such features can be established through evolution (the Baldwin effect) (Dennett 1991). Thus, the ability to learn is innate (learning usually being defined as the acquisition of novel behavior through experience (Morris 1973)). We suggest that environmental interaction is needed to expose the underlying mechanisms for learning and intelligent behavior. Many researchers use in vitro models (brain slices or dissociated neural cell cultures) to study the basic mechanisms of neural plasticity underlying learning. We argue that because these systems are not embodied or situated, their applicability to learning in vivo is severely limited. We are developing systems to re-embodiment in vitro networks, and allow them to interact with an environment, so that we can watch the processes contributing to learning at the cellular and network levels *while they happen*.

We study networks of tens of thousands of brain cells in vitro (neurons and glia) on a scale of a few square millimeters. The cells in cortical tissue are separated using enzymes, and then cultured on a Petri dish with 59 electrodes embedded in the substrate, a multi-electrode array (MEA; from MultiChannel Systems) (Fig. 2.1) (Potter 2001; Potter and DeMarse 2001a). The neurons in these cultures spontaneously branch out (Fig. 1.1). Even left to themselves without external input other than nutrients (cell culture media), they re-establish connections with their neighbors and begin communicating electrically and chemically within days, demonstrating an inherent goal to network. Electrical and morphological observations suggest these cultures mature in about four weeks (Gross, Rhoades et al. 1993; Kamioka, Maeda et al. 1996; Watanabe, Jimbo et al.

1996; Wagenaar, Pine et al. 2006a). The neurons and supporting glia form a monolayer culture over the clear MEA substrate, amenable to optical imaging with conventional and two-photon microscopy (Potter 1996; Potter 2000; Potter, Lukina et al. 2001). With sub-micron resolution optical microscopy, we can observe learning-related changes in vitro with greater detail than is possible in living animals. The networks are also accessible to chemical or physical manipulation. We developed techniques to maintain neural cultures for up to two years, allowing for long-term continuous observation. For detailed methods, refer to (Potter and DeMarse 2001a) or Appendix G.




Figure 0.1 Connecting neurons to multi-electrode arrays

Left: Cells are plated inside a glass multi-electrode array culture dish such as this (Photo by Steve Potter). Right: recorded voltage traces in the lighter boxes overlay a microscope image of the neuronal network growing on a 60-electrode array (electrode diameter, 30 μm). The thick lines are the electrode leads. The voltage spikes are neural signals (Figure by Daniel Wagenaar).

A multi-electrode array records extracellular neural signals fast enough to detect the firing of nearby neurons as voltage spikes (Fig. 2.1, right). Thus, the activity of multiple neurons can be observed in parallel and network phenomena can be studied. In addition to the expression of spontaneous activity, supplying electrical stimulation through the multiple electrodes induces neural activity; we have built custom circuitry to continuously stimulate the 60 electrodes (Wagenaar and Potter 2004b). The MEA forms a long-term non-destructive two-way interface to cultured neural tissue. The recorded

signals can be used as motor commands, while the stimuli represent sensory inputs, in our embodied system. These techniques allow high resolution, long term, and continuous studies on the role of embodiment throughout the life of a cultured neural network.




Figure 0.2 Hybrot (Hybrid living+robotic) setup

Optical and electrical data from neurons on an MEA are analyzed and used to control various robotic devices, while time-lapse imaging is carried out to make movies of neuronal plasticity.

Wilson (Meyer and Wilson 1991) coined the term 'animat' (a computer simulated or robotic animal behaving in an environment) in his studies of intelligence in the interactions of artificial animals. Our interfacing of cultures to a simulated environment (described below) was the first Neurally-Controlled Animat (Fig. 2.2) (Potter, Fraser et al. 1997; DeMarse, Wagenaar et al. 2001; DeMarse, Wagenaar et al. 2002). For cultures interfaced to physical robots, we introduce the term 'hybrot' for hybrid biological robot. Mussa-Ivaldi's group created the first closed-loop hybrot by controlling a Khepera robot with a brain stem slice from a sea lamprey (Reger, Fleming et al. 2000). In a related approach, our Neurolab colleague Robert Butera studies detailed neural dynamics by coupling simulated neurons to real neurons using an artificial conductance circuit (Sharp, Abbott et al. 1992; Butera, Wilson et al. 2001). Stephen DeWeerth's group in the Neurolab develops and studies, among others things, silicon model neurons interfaced with living mollusk and leech neurons (Simoni, Cymbaluyk et al. 2001).

Using simulated environments is a good first step and provides easier control and repeatability compared to robotics. However, a 'real' environment's great complexity provides two advantages. First, many seemingly complex behaviors of animals are emergent: simple behavioral rules applied in a complex environment produce complex and productive behavior (Braitenberg 1984; Arkin 1999; Brooks 1999). Second, a complex environment produces a robust brain to take advantage of it: among other examples, this is evident in tool use (Clark 1997) and in exploiting properties such as the biomechanics of muscle tissue in repositioning an arm without excessive vibrations. It is difficult to simulate a complex environment with realistic physics. If physics plays an important role in the complex behavior of intelligent systems, then by using robots in the real world, the researcher gets the physics "for free." We believe that this merging of artificial intelligence concepts (including robotics) into neurobiological experiments can inform future AI approaches, making AI a bit less artificial.

2.2 Examples: three embodied neural systems

Creating a neurally controlled robot that handles a specific task begins with a hypothesis of how information is encoded in the brain. Much remains to be determined, but numerous schemes have been proposed, most based on the quantity and/or relative timing of the firing of neural signals. A neural network may be considered as a type of processing unit with an input (synaptic or electrical stimulation patterns), and an output (neural firing patterns), which can perform interesting mappings to produce behavior. Later work has found that both low frequency patterns of stimuli (Chapter 4, (Eytan, Brenner et al. 2003b)) and high frequency tetanic stimuli (Chapter 3, (Chao, Bakkum et al. 2007) (Reich, Victor et al. 1997; Jimbo, Robinson et al. 1998; Maeda, Kuroda et al. 1998; Jimbo, Tateno et al. 1999)) can induce plasticity. In addition behavioral feedback can be used to direct the plasticity (Chapter 5 and Appendix A).

Below are overviews of three such systems. These examples could have been conducted with artificial neural networks. We use biological neural networks not as substitutes to artificial neural networks, but to tease out the intricacies of *biological* processing to inform future development of *artificial* processing. In particular, we analyzed how the properties of neurons lead to real-time control and adaptation to novel environments.

2.2.1 Living neurons control a simulated animal

The first Neurally-Controlled Animat (DeMarse, Wagenaar et al. 2001) comprised a system for detecting spatio-temporal patterns of neural activity, which directed exploratory movement of a simulated animal in real time (Fig. 2.3). Neural firings were integrated over time to produce an activity vector every 200 ms, representing the current activity pattern, and recurring patterns were clustered in activity space. Each cluster was assigned a direction of movement (left, right, forward, backward). Proprioceptive and exteroceptive feedback via electrical stimulation was provided to the neural culture for each movement and for collisions with walls and barriers. The stimulation induced neural activity that, in turn, was detected through the activity vectors and used as commands for subsequent movements. We created the software and hardware necessary to enable a 15-ms sensory-motor feedback latency, since we feel it is important that a tight connection between the neural system and its environment is likely to be crucial to adaptive control and learning.

Within this real-time feedback loop, both spontaneous and stimulated neural activity patterns were observed. These patterns emerged over the course of the experiment, sometimes assembling into a recurrent sequence of patterns over several seconds, or the development of new patterns, as the system evolved. The overall effect of the feedback loop on neural activity was observed from the path of the animat's movement throughout its environment (Fig. 2.3). As the neural network moved its

artificial body, it received feedback and in turn produced more movement. The behavioral output was a direct result of both spontaneous activity within the network as well as activity produced by feedback due to the networks interaction with its virtual environment. Hence the path of the animat was indicative of current activity as well as the effects of feedback. Analyzing the change in behavior of the neurally-controlled animat provided a simple behavioral tool to study shifts in the states of neural activity. However, this first Neurally-Controlled Animat did not demonstrate noticeable goal-directed behavior, which the next example addresses explicitly.




Figure 0.3 Animat setup and activity.

Above: neural signals are used to control the movement of an animat, whose 'brain' is exposed to microscopic imaging; feedback from the environment determines subsequent electrical stimulation of the living neuronal network in an MEA. Below: One hour of the animat's path (*curved lines*), as it moves about within its environment under neural control, with feedback. The white boxes represent various environmental obstacles. (Figure by Tom DeMarse and Steve Potter.)

2.2.2 Living neurons control a mobile robot




Figure 0.4 Living neurons control a mobile robot.

Neural firings in response to paired electrical stimulations at various inter-stimulus intervals (ISI) are plotted. In the experiments, the ISI was proportional to the distance between the neurally controlled approaching animat and its target object. It was considered positive if the target was located to the right of the animat and negative if left of the robot. The neural response determined the magnitude of subsequent animat movement; the direction of movement was determined from which quadrant the ISI fell into (see the arrows and movement key, bottom). Inset: the neurally controlled animat's trajectory (Koala robot, represented by the triangle). The target object (Khepera robot, represented by the square) was held stationary until the robot approached, and then it was moved continuously under computer control (down and to the right in the figure). (Figure by Alec Shkolnik.)

One of the simplest forms of ‘intelligent’ behavior is that of approach and avoidance. The goal of the second system was to create a neural interface between neuron and robot that would approach a target object but not collide with it, maintaining a desired distance from the target. If a given neural reaction is repeatable with low

variance, then the response may be used to control a robot to handle a specific task. Using one of these response properties, we created a system that could achieve the goal (Shkolnik 2003).

Networks stimulated with pairs of electrical stimuli applied at different electrodes reliably produce a nonlinear response, as a function of inter-stimulus interval (ISI). Figure 2.4 shows averaged firing rate over all 60 electrodes following two stimulations separated by a time interval. At short ISI's, the response of the network following stimulation was enhanced; at longer intervals, the response was depressed. Furthermore, the variance of the data for each ISI was relatively small, indicating the effect is robust and thus qualifies as a good candidate for an input/output mapping to perform computation.

By mapping the neural response to a given ISI as a transformation of distance to an object, we created a robot that reacts to environmental stimuli (in this case sensory information about distance from an object) by approaching and avoiding that target. To construct our "approach and follow" hybrot, sensory information (the location of a reference object with respect to the robot) was encoded in an ISI stimulation as follows: the closer the robot is to the object, the smaller the ISI. The response of the neurons to a stimulation pair, measured as an averaged firing rate across all electrodes for 100 ms after the second stimulus, was used to control the robot's movements: a larger neural response corresponded to a longer movement (either forward or backward) of the robot.

When the robot was far away from the reference object, the ISI of the stimulation pair was long, and the neural response was large, moving the robot towards the object (Fig. 2.4, right). As the robot moved closer to the object, the stimulation interval decreased until it reached 150 ms. At this point, the neural response was minimal, and no movement was commanded. In other words, the robot reached its desired location with respect to the reference object. If the robot was closer to the object, the neural reaction was larger (a very short ISI), this time driving the robot away from the object. We

divided the input ISI into 4 quadrants (Fig. 2.4, left). Each of the 4 quadrants corresponded to a directional movement: forward/right, forward/left, backward/right, and backward/left. Then, a positive ISI caused movement in a direction opposite that for a negative ISI. Given the neural response to an ISI stimulation, we decoded which quadrant the response belonged to with good accuracy (>95%).

We used the Koala and Khepera robots (manufactured by K-Team) to embody the cultured network, and to provide an environment with a moving object. The Koala robot was used as the neurally controlled robot, while the Khepera served as the reference object, moving at random under computer control. Under neural control, the Koala successfully approached the Khepera and maintained a distance from it, moving forward if the Khepera moved away, or backing up if the Khepera approached

In addition to demonstrating the computational capacity inherent in cultured neurons, this hybrot can be used to study learning in cultured neural networks. In this case, learning would be manifested through changes in the neural activity and changes at the behavioral level of the robot (Chapter 5 and Appendix A).

2.2.3 Living neurons control a drawing arm

Meart (Multi-Electrode Array art) was a hybrot born from collaboration with the SymbioticA Research Group (discussed in greater detail in Appendix A). As an overview, the 'brain' of dissociated rat neurons in culture was grown on an MEA in our lab in Atlanta while the geographically detached 'body' resided in Perth. The body consisted of pneumatically actuated robotic arms moving pens on a piece of paper (Fig. 2.5). A camera located above the workspace captured the progress of drawings created by the neurally-controlled movement of the arms. The visual data then instructed stimulation frequencies for the 60 electrodes on the MEA. The brain and body interacted through the internet (TCP/IP) in real time providing closed loop communication for a neurally

controlled 'semi-living artist'. We see this as a medium from which to address various scientific, philosophical, and artistic questions (see *Discussion* in Appendix A).




Figure 0.5 Meart–The Semi-Living Artist.

Left: Meart's arms used markers to draw on a piece of paper, under live neural control. In the background was a projection of the MEA and cultured net, Meart's 'brain'. Right: one drawing created by Meart in an exhibition. (Photos by Guy Ben-Ary and Phil Gamblen.)

Meart has brought neurobiology research to numerous artistic events: *Biennale of Electronic Arts Perth* (Perth, 2003), *Artbots: the Robot Talent Show* at the Eyebeam gallery (New York, 2003), *Cyber@rts* at the Mercado del Ensanche (Bilbao, 2004), *2004: Australia Culture Now* at the Australian Center for the Moving Image and the National Gallery of Victoria (Melbourne, 2004), *1st Moscow Biennial for Contemporary Art* at the M'ARS gallery (Moscow, 2005), *Artrage* at the Black Box (Perth, 2005), at the Eyedrum gallery (Atlanta, 2006), *Strange Attractors: Australian Art & Science Exhibition* at the Zendai gallery (Shanghai, 2006) and *In the Presence of the Body*, the first permanent art exhibit of the Art and Biology Department at the Rensselaer Polytechnic Institute. The robotic arm and video sensors were shipped to each gallery while the living neurons sent and received signals from Atlanta. An overview of how

Meart worked may best be described by the artistic conception behind the early presentations: portrait drawing. First, a blank piece of paper was placed beneath the arm's end-effector and a digital photograph was taken of an audience member. Then, communication between the arm and the neurons was begun. The neural stimulation via the MEA was determined by a comparison of the actual drawing, found using a video camera taking images of the drawing paper, to the target image of a person's photograph. Both the actual image and the target image were reduced to 60 pixels, corresponding to the MEA electrodes, and the gray scale intensity of each pixel was found. Similar to how an artist continually compares her work to her subject, the gray scale percentages for corresponding pixels on the two images were continuously compared, in this case subtracted to produce a matrix of error values. The 60 error values determined in real-time the stimulation frequency per electrode using a custom stimulation circuit built by Thomas DeMarse. Arm movement was determined by the recorded neural activity, using averaged firing rates of the induced and spontaneous activity per stimulation. Stimulation affected this neural activity, and so the communication formed a loop, with a loop time of approximately one second.

In the prior example, the sensory-motor mappings used a stable neural property to reliably control the robot. With Meart's early experiments as described above, the sensory-motor mappings were less well defined, in the hope of demonstrating a micro-scale version of the brain's creative processes. The behavioral response of the robot sheds light on the properties of the neural network and directs further encoding refinements (Chapter 5 and Appendix A). An example drawing is shown in Figure 2.5. The drawings changed throughout the life of cultures (and were different for different cultures) demonstrating neural plasticity. Later experiments proved more promising, whereby goal-directed behavior was achieved in simulated animats (Chapter 5 and Appendix A).

2.3 Discussion

2.3.1 Embodying cultures: theory

2.3.1.1 A blank slate

Since the cultured neurons were first separated and allowed to settle onto the MEA at random, they start from a 'blank slate'. Neural structure is lost and the function of neural activity is no longer obvious, yet neural network processing remains, evidenced by the complex activity patterns we have observed. For traditional in vitro neural models, function is cloudy since activity no longer relates to or causes behavioral states or actions. One cannot say 'this neuron is involved in color perception' or 'this neural structure helps to coordinate balance' as could be said for in vivo experiments. Artificially embodying and situating cultured neurons redefines their behavioral function concretely.

The structure of neuronal networks is likely to be important in neuronal processing, and changes in structure are likely to underlie learning and memory (Engert and Bonhoeffer 1999; Trachtenberg, Chen et al. 2002). Our cultured neurons form two-dimensional monolayers; functional importance may lie in the affordances given by the three-dimensional layered nature of the cortex. We and others in the Neurolab are pursuing the construction of 3D MEAs to support three-dimensional cultures, as part of an NIH Bioengineering Research Partnership (Choi, Powers et al. 2003), (Blum, Ross et al. 2003). However, even cultured cortical monolayers (without 3D structure nor sub-cortical regions) have demonstrated an ability to adapt following stimulation via potentiation and/or depression (Jimbo, Tateno et al. 1999), (Tateno and Jimbo 1999), (Marom and Shahaf 2002), (Eytan, Brenner et al. 2003b). We are exploring using these plasticity mechanisms as a means to shape the network during development, within the Neurally-Controlled Animat paradigm, so it is no longer a blank slate.

2.3.1.2 Associations.

The biological brain makes associations between different phenomena observed through sensation, whether between various external stimuli or between the actions of a body and their consequences, and then commands movement accordingly. Our methods have been developed to study these processes in real time with enough resolution to capture the dynamics of these interactions. These processes can be expressed using dynamical systems theory (DST), a mathematical framework to describe systems that change in time. For example, the formation of certain functional structures (ocular dominance columns) in the visual cortex has been described using Alan Turing's reaction-diffusion equations (Turing 1953). Kuniyoshi and his group explore DST to connect sensory-motor control to the cognitive level (Yamamoto and Kuniyoshi 2002). As applied to cognition (Port and van Gelder 1995), DST describes the mind with a set of complex, recursive filters. This opposes the classical cognitive concept of neural processing being analogous to a digital computer, containing distinct storage and processing of symbols (Fodor 1980), (Vera and Simon 1993). DST contends that multiple feedback loops and transmission delays, both of which are widespread in the brain, provide a time dimension to allow higher-level cognition to emerge without the need for symbolic processing (Edelman and Tononi 2000). DST is a framework compatible with embodied perspectives. The dynamical systems perspective has too often been neglected in neurobiology and cognitive sciences.

In contrast to an intact brain in an animal, cultures of neurons are isolated because they do not contain the afferent sensory inputs or efferent motor outputs a body would provide and therefore no longer have a world with which to reference their activity. Under these conditions, what associations can the network make, and what would those associations mean? Moreover, what symbols are operated on? Because of this, any associations that are made must consequently be self-referential or circular and neural activity may be misleading. The network as a set of complex, recursive filters has no

external signals to filter, possibly leading to the abnormal barrage activity described below. To address this major shortcoming of in vitro systems, our neural cultures are embodied with sensory feedback systems, motor systems, and situated in an environment, providing a new frame of reference. New findings about the dynamics of living neural networks might be used to design more biological, less artificial AI.

2.3.1.3 Intelligence and Meaning.

By embodying cultured neurons, the ‘meaning’ of neural activity emerges, since this activity affects the external world and subsequent stimulation. Now the network has a body behaving and producing experiences, allowing for the study of concepts such as intelligence. We will take a behavioral definition of intelligence as our start: Rodney Brooks described intelligence in terms of how successfully an agent interacts with its world to achieve goal directed behavior (Brooks 1991). William James stated, "Intelligent beings find a way to reach their goal, even if circuitous," (James 1890). Neurons have inherent local goals (to transmit signals, integrate synaptic input, optimize synaptic strengths, and much more) that provide the foundation to intelligently achieve meaningful behavioral goals. No doubt the basis for intelligence is inherent at birth, but an interaction with a sufficiently complex environment (learning) is needed to develop it.

In our cultured networks, the local goals of neural interaction are subject to detailed optical and electrical observation, while the execution of higher-level behavioral goals are observed through the activities of the robotic body. (Note that the behavioral goals are artificially constrained by the stimulation and recording transformations chosen.) We hope this combination will lead to a clearer definition and a better understanding of the neurological basis of intelligence, in addition to explanations of other psychological terms: learning, memory, creativity, etc. Neurobiology has given inspiration to AI since the advent of the perceptron and consequent artificial neural networks, which are based on the local properties (goals) of individual neurons. We wish

to continue this trend by finding the principles of network processing by multiple neurons that lead to higher-level goals.

2.3.1.4 Network-wide Bursting.

The activity of cultured neurons tends towards the formation of dish-wide global bursts (barrages) (Kamioka, Maeda et al. 1996): sweeps of fast, multiple neural firings throughout the network lasting between hundreds of milliseconds to seconds in duration. These barrages have been observed often in cultured neurons (Nakanishi and Kukita 1998) but also in cortical slices (Corner and Ramakers 1991) and in computer models (Latham, Richmond et al. 2000a). Barrages of activity are reported in the cortex in vivo during early development, during epileptic seizures, while asleep, and when under anesthesia. These in vivo examples of barrages occur over finite periods of time. In contrast, barrages in vitro are continuous over the life of the culture. We consider the possibility that at some stage, dish-wide barrages of spiking activity are abnormal, a consequence of 'sensory deprivation' (Wagenaar, Madhavan et al. 2005), or a sign of arrested development (Tabak and Latham 2003).

For both a model system (Latham, Richmond et al. 2000a) and for cultured mouse spinal neurons (Latham, Richmond et al. 2000b), if more than 30% of the neurons are endogenously active, the neurons fire at a low steady rate of 1 to 5 Hz per neuron, while a reduction in the fraction of endogenously active cells leads to barrage activity. Endogenous activity is functionally similar to activity induced by afferent input, suggesting embodiment would lead to low steady firing rates. The hypothesis is then that the barrage activity may be due to the lack of an external environment with which to interact. We are developing animat mappings in which continuous sensory input quiets barrages, bringing the networks to a less 'sensory-deprived' state that allows more complex, localized activity patterns.

2.3.1.5 The World and the Brain.

Environmental deprivation leads to abnormal brain structure and function, and environmental exposure shapes neural development. Similarly, patterned stimulation supplied to cultured neurons may lead to more robust network structure and functioning than with trivial or no stimulation. The most dramatic examples of the importance of embodiment come from studies during development, when the brain is most malleable. Cognitive tests were performed on institutionalized children in Romania, children typically deprived of proper environmental and social interaction early in life (Rutter 1998), (O'Connor, Rutter et al. 2000). Compared to peers, the children showed severe developmental impairment that improved, however, after transplantation to a stable family. Those adopted prior to 6 months of age achieved nearly complete cognitive catch-up to similarly aged children, while those adopted after 6 months of age had significant but incomplete catch-up. Likewise, laboratory rats raised in environments with mazes and varied visual stimuli had 30% greater cortical synaptic density than those raised in minimalist environments, and performed better in various cognitive experiments (Black, Isaacs et al. 1990), (Diamond 1990). Synaptic morphology in adults (Weiler, Hawrylak et al. 1995) and adult neurogenesis is dependent on external cues (Gross 2000) demonstrating that environmental interaction is important throughout life.

A disembodied neural culture, whose activity never influences future stimulation, will not develop meaningful associations to an input. In the brain, if a sensation is not useful in influencing future behavior (no association is made between the two) the percept of the sensation fades. The environment triggers an enormous number of sensory signals, and the brain develops to filter out the excess while perceiving the behaviorally relevant. All one-month-old infants can distinguish between the English L and R sounds. Five months later, Japanese infants lose the ability while American infants maintain it, because the distinction is not needed to understand the Japanese language (Kuhl,

Andruski et al. 1997). Japanese adults consequently have great difficulty distinguishing these sounds, but perception of the distinction can be learned through targeted instruction. These studies further demonstrate how brain (re)wiring depends on environmental context and occurs throughout life: the brain focuses on perceiving the portions of the environment relevant to produce a meaningful interaction.

2.3.1.6 The Body and the Brain.

The choice of how to instantiate an animat or hybrot is important to processing in cultured neural networks. For example, the body, with its various sensory apparatus and motor output, is what detects and interacts with the environment. In addition to how different environments cause differences in the brain, differences in the body will have analogous effects on the brain. Changes in the frequency or type of sensory input via practice or surgical manipulation of the body causes gross shifts in the functional organization of corresponding cortical areas (the somatotopic maps) (Buonomano and Merzenich 1998). Amputation causes a sudden change to a body, and amputees later report having at times a sensation or impression that the limb is still attached. The impression lasts for days or weeks in most cases (years or decades in other cases) and then gradually fades from consciousness (Ramachandran and Hirstein 1998). These false 'phantom limb' sensations arise because the brain has wired itself for a given body that has now changed. This discrepancy further suggests the body and its interaction with the environment influence brain wiring and cognitive function. Neurally-Controlled Animats allow an unlimited variety of bodies to be studied; their structure and operating parameters can be easily varied to test effects on brain-body interactions.

2.3.2 Summary: integrating brain, body, and environment

The above paragraphs were worded as if the entities brain, body, and environment are independent. Finding physical boundaries between the three is easy, but since the

brain is so enmeshed in the states of the body (influencing mood, attention, and more), which in turn are so enmeshed in the body's interaction with its environment, finding functional boundaries between the three is difficult, if possible at all (Varela, Thompson et al. 1993; Clark 1997; Pfeifer and Scheier 1999; Pfeifer and Bongard 2007). Damasio contends that the mind depends on the complex interplay of the brain and the body, and consequently emotions and rationality cannot be segregated (Damasio 1994).

We have integrated our hybrot's brain (cultured network), body (robot or simulated animat), and environment (simulation, lab, or gallery) into a functional whole, even while the parts are sometimes 12,000 miles apart. Our experiments with these Neurally-Controlled Animats so far are rudimentary: we are still setting up the microscopic imaging systems to allow us to make correlations between changes in behavior and changes in neuron or network structure; we have only recently developed training algorithms that reliably result in learning (Chapter 5 and Appendix A). But in the process of creating this new research paradigm of embodied, situated cultured networks, we have already sparked a philosophical debate about the epistemological status of such semi-living systems (Manson 2004a), and have raised a number of issues about the validity of traditional (disembodied) in vitro neural research. We hope that others will make use of the tools we have developed such as our MeaBench software (Wagenaar, DeMarse et al. 2005), sealed-dish culture system (Potter and DeMarse 2001a), and multi-site stimulation tools (Wagenaar and Potter 2002a), to pursue a wide variety of questions about how neural systems function. We expect that these inquiries will lead to fundamentally different, more capable, and less artificial forms of AI.

CHAPTER 3

TETANUS INDUCED PLASTICITY AND

REGION SPECIFIC CHANGES IN NETWORK ACTIVITY²

Electrically interfaced cortical networks cultured *in vitro* can be used as a model for studying the network mechanisms of learning and memory. Lasting changes in functional connectivity have been difficult to detect with extracellular multi-electrode arrays using standard firing rate statistics. We compared the ability of various statistics to quantify functional plasticity at the network level. We compared three established statistical methods to one of our own design, called center of activity trajectory (CAT). CAT, which depicts dynamics of the location-weighted average of spatiotemporal patterns of action potentials across the physical space of the neuronal circuitry, was the most sensitive statistic for detecting tetanus-induced plasticity in both simulated and living networks. By reducing the dimensionality of multi-unit data while still including spatial information, CAT allows efficient real-time computation of spatiotemporal activity patterns. Thus, CAT will be useful for studies *in vivo* or *in vitro* in which the locations of recording sites on multi-electrode probes are important. In a separate preliminary study using similar tetanus protocols, Principal Components Analysis of the spatiotemporal patterns of action potentials also demonstrated network plasticity. In a previous publication (Wagenaar 2006a) on the same data set, standard firing rate statistics were not capable of exposing plasticity.

High frequency tetanic stimulation has been traditionally used to induce plasticity of neurons in MEAs. Therefore, to determine if plasticity could occur in our preparations, we applied tetanic stimulation. We also quantified the ability of various statistics to detect the induced plasticity. A version of the best statistic, the center of neural activity (CA), was subsequently used in the embodiment experiments in Chapter 5 and Appendix A.

² Contains sections from:

Chao ZC, Bakkum DB, Potter SM, “Region-specific network plasticity in simulated and living cortical networks: comparison of the center of activity trajectory (CAT) with other statistics” J. Neural Eng. 4 (2007) 294–308

3.1 Introduction

Modification of connectivity between cortical neurons plays an important role in the processes of learning (Ahissar, Vaadia et al. 1992; Buonomano 1998) and memory (Merzenich and Sameshima 1993). Connectivity at the synaptic level has been studied by administering stimuli while simultaneously recording neural activity, and then quantifying plasticity by analyzing the stimulus–response relationships. Culturing on multi-electrode arrays (MEAs) was introduced to help understand connectivity and plasticity in networks of neurons (Gross 1979; Pine 1980). This allows long-term (months), non-invasive observation of the electrical activity of multiple neurons simultaneously (Potter and DeMarse 2001b) in a system with less experimental complexity and greater control than preparations *in vivo*. External factors such as sensory inputs, attention and behavioral drives are absent, while many aspects of complex spatiotemporal spike patterns observed in animals remain (Gross and Kowalski 1999).

Many activity statistics have been used to quantify stimulus–response relationships from simultaneous recordings of multiple neurons (Brown, Kass et al. 2004). Most analyze the dependences between spike trains, such as the maximum likelihood method (Chornoboy, Schramm et al. 1988; Okatan, Wilson et al. 2005), product–moment correlation coefficient (Kudrimoti, Barnes et al. 1999), functional holography (Baruchi and Ben-Jacob 2004), etc. However, only a few were applied for measuring network plasticity. The most common of these was the firing rate (FR), which showed plastic modifications of network response induced by tetanic stimulation in cortical cultures (Reich, Victor et al. 1997; Jimbo, Robinson et al. 1998; Maeda, Kuroda et al. 1998; Jimbo, Tatenos et al. 1999) and dopamine regulated plasticity in anesthetized rats (Rosenkranz and Grace 1999). Firing rate histogram (FRH) uses firing rates integrated over successive sequential latency epochs to add detailed temporal information, and was applied to demonstrate adaptable image processing and pattern

recognition through training by tetanic stimulation in MEA cultures (Ruaro, Bonifazi et al. 2005). Mutual information (MI) characterized the statistical dependence between neuron pairs, exposing the strength of coupling between neurons and the functional connectivity among cortical areas (David, Cosmelli et al. 2004). Cross-correlation histograms (CCH) from pairs of neurons showed functional plasticity in the auditory cortex of behaving monkeys (Ahissar, Abeles et al. 1998), and the more advanced shift-predictor corrected cross-correlogram (SCCC) was used to quantify receptive field plasticity in the rat auditory cortex (Bao, Chang et al. 2003). Joint peristimulus time histogram (JPSTH) characterized the causality of firing between neuron pairs, and successfully demonstrated long-term facilitation of neural activity involved in respiratory control (Morris, Baekey et al. 2003). Robust neuronal computation and encoding is believed to involve the distribution of information over populations of neurons and synapses in a combination of spatial and temporal domains. Observing only pairs of neurons (MI, CCH, SCCC and JPSTH), neglecting temporal information (FR) and neglecting spatial information (all) limit the ability of these to measure the complex plasticity of the brain.

We applied tetanic stimulation to our cultures as a bench mark to test whether or not an MEA could induce and observe plasticity in dissociated cortical neurons. Plasticity of neural firing rate after electrical probe stimulation using in vitro MEAs has been reported (Jimbo, Robinson et al. 1998; Maeda, Kuroda et al. 1998; Jimbo, Tateno et al. 1999; Marom and Shahaf 2002; Ruaro, Bonifazi et al. 2005), but our lab has questioned the results as either the measured plasticity across the inducing stimulation did not exceed the inherent network drift (the measured plasticity outside of the inducing stimulation) or methodological flaws existed (Wagenaar, Pine et al. 2006b). Typically, firing rates were used to observe plasticity. However, rate-based plasticity metrics are not reliable: the presence of network drift (Madhavan, Chao et al. 2005), homeostatic mechanisms that re-adjust firing rates in response to plasticity (Turrigiano 1999), and inconsistent synaptic

transmission (Kandel et al., 2000) all lower reliability. By incorporating the timing of neural activity to better characterize the dynamics of the population activity (see also Chapter 4), we found tetanus induced plasticity could be better observed than with firing rate statistics alone. In a separate preliminary study, Principal Components Analysis (PCA) of the temporal patterns of action potentials demonstrated network plasticity on a data set where standard firing rate statistics were previously reported not capable of exposing the plasticity (Wagenaar 2006a; Wagenaar, Pine et al. 2006b).

We then further incorporated information of the physical locations of the recording sites and devised a statistic called the center of activity trajectory (CAT). CAT detected more pronounced changes in the network following tetanus than the FR, FRH, and SCCC. (In simulation, CAT showed the ability to detect smaller changes in the distribution of network synaptic weights than did FR, FRH, MI, SCCC or JPSTH (Chao, Bakkum et al. 2007).) Center of activity (CA) is analogous to the center of mass, where ‘mass’ at an electrode location is determined by the recorded firing rate, and CAT is the sequence of CAs over successive time intervals. CAT is advantageous because its values have a direct neural interpretation (see next paragraph). PCA reorients multidimensional data into dimensions ordered by the greatest variance. This is used to reduce noise while maintaining signals (plasticity here), but a consequence of its abstracted data representation is the possible amplification of stimulation artifacts (giving false plasticity).

By applying a shuffling method to the CAT analysis to erase the spatial information of recording locations in its calculation, we found that changes in activity patterns recorded from neighboring electrodes were not independent and contributed to the better performance of CAT to detect plasticity. The network plasticity was *region specific*: despite the apparent random connectivity of neurons, plasticity was not symmetrically distributed, and the location of neurons played a role in stimulus-induced plasticity.

No ground truth about network plasticity in living networks exists, because neuronal connectivity cannot be measured for more than a few pairs of neurons simultaneously. Therefore, the amount of plasticity detected by each statistic was cross-validated in a simulated network, in which the weights of *all* synapses were observable. Work with the simulated network was done by Zenas Chao (Chao, Bakkum et al. 2007) and therefore will be mentioned only briefly. To summarize, the CAT demonstrated the ability to detect the smallest changes in the distribution of network synaptic weights compared to alternative statistics.

3.2 Methods

Cell culturing and electrophysiology was done as reported in Appendix G. Culture ages during experiments ranged from 1 to 3 months.

3.2.1 Stimulation protocols

A neural response after a stimulation is due to a combination of the electrically evoked and ongoing neural activity. By measuring responses to repeated ‘probe’ stimulation, changes are due solely to changes in network state. Each experiment consisted of a 2h period of random probe stimulation (RPS) followed by a 15 min tetanic stimulation followed by another 2 h period of RPS (Protocol III.4, Table 7.3 in (Wagenaar 2006a)). In six experiments, the RPS periods consisted of six electrodes stimulated in a random order at an aggregate frequency of 0.5 Hz (in one experiment, the RPS periods consisted of only four probe electrodes). Two of these electrodes were used for the tetanic stimulation: 150 trains of 20 paired pulse stimuli with 10 ms intervals between paired pulses, 50 ms intervals between pairs and 6 s intervals between the start of each train. Prior to an experiment, every electrode was stimulated in a random order 20 times, and electrodes with six (or four) highest responses (the total number of spikes

counted within 100 ms latency after stimuli over recording electrodes) were selected as probe electrodes. The tetanus electrodes were randomly chosen from these.

3.2.2 Neural activity statistics

We used evoked responses within 100 ms after the stimuli of RPS for statistics calculations. We measured CAT from the evoked responses and compared it to the three most commonly used statistics: FR, FRH and SCCC. MI was not measured, due to its poor performance in detecting network plasticity in simulations (Chao, Bakkum et al. 2007). JPSTH was not measured because of its high dimensionality and computation time.

FR: The firing rate statistic summed the number of action potentials over 100 ms per each electrode e (dimension of $1 \times N_e$).

FRH: The firing rate histogram statistic summed firing rates per electrode over 5 ms moving time windows (stepped by $500 \mu s$ up to 100 ms; dimension of $20 \times N_e$).

CAT: The center of activity statistic calculated the center of activity of FRs in a 5 ms moving time windows (stepped by $500 \mu s$ up to 100 ms; dimension of 20×2) (see Fig. 3.1).

PCA: Principal component analysis was done on the FRH. As is standard, the first few principal components, representing the dimensions of largest variance, were used to calculate plasticity. We chose to use the first 2.




Figure 0.1 Whole-input-output (WIO) vectors for analyzing performances of different statistics.

WIO vectors calculated from each statistic were used to represent the network input-output function. As an example, the WIO vector of CAT calculated from probe responses to one RPS (a) delivered during an experiment is demonstrated. (b) CA was calculated for evoked responses to the stimulation electrode P_j ($j = 1$ to 6). Each frame indicates the firing rate over a 5 ms moving time window (with a 500 μ s time step) on an 8 by 8 grid of electrodes averaged over multiple stimuli at P_j (RPS might have multiple stimuli delivered at P_j , see (a)). The 2D trajectory of CAs from frame 1 to frame N (from 0 to 100 ms after the stimuli), CAT, can be represented by a 1D vector by joining CAT_x and CAT_y . This vector represents CAT of responses to stimuli P_j at the network state S_i (states at different times during an experiment). (c) CATs for responses to 60 different stimulation electrodes (P_1 to P_{60}) were joined together to form the WIO vector. This WIO vector represents the input-output function, in terms of CAT, of the network. Each statistic has a corresponding WIO vector to describe its input-output function. Sensitivity to changes in network state are shown by significantly different WIO vectors.

3.2.3 Quantifying plasticity

C/D: (See also a general description in Appendix G.) For each statistic, we calculated one WIO vector (see Fig 3.1) every 240 s (a ‘block’) for the experiments with six probe stimulation electrodes, and every 160 s for the experiments with four probe stimulation electrodes. Thus, there were 19.9 ± 4.2 (mean and standard deviation) stimuli delivered at each electrode for each WIO vector. Three periods were used for statistics:

Pre1, Pre2 and Post1 (see Fig. 3.2). Each period had a duration of 52.5 min, and the intervals between Pre1 and Pre2 and between Pre2 and Post1 were 15 min. The 15 min interval between Pre2 and Post1 was the tetanization. For each statistic, the mean distance of the WIO vectors in Pre1 to the centroid of the WIO vectors in Pre2 measured a quantity termed *Change*, C . A quantity termed *Drift*, D , was the mean distance of the WIO vectors in Pre2 to their own centroid. The ratio of change-to-drift, C/D , depicts the amount of neural plasticity between Pre1 and Pre2. No plasticity would give a ratio equal to 1. A similar measurement between Pre2 and Post1 was used to quantify the neural plasticity across the tetanus. The performance of each statistic to detect the tetanus-induced change was quantified by comparing the two C/D s ($n = 6$ experiments, Wilcoxon signed rank test).

3.3 Results

3.3.1 Preliminary principal components analysis of plasticity

PCA of FRH demonstrated tetanus-induced plasticity in a set of experiments done by Wagenaar (Wagenaar 2006b), in which FR statistics did not demonstrate plasticity (Fig. 3.2). Considering he excluded early phase responses (Fig. 4.2), additional data analysis was done on Wagenaar's data in collaboration with Zenas Chao to search for plasticity. Projecting the main principal components back onto the temporally binned data showed that the early phase, occurring within 20 ms were the main contributors to plasticity.

The difference between Wagenaar's four tetanus protocols are not pertinent to this chapter. The relevant point is that including temporal information now showed plasticity. The affect of burst quieting on tetanus-induced plasticity was pursued by labmate Radhika Madhavan as is further described in her thesis (Madhavan 2007). The above data supports her main conclusion that burst quieting was correlated to an increase in

plasticity. Here, quieting of bursts was achieved by distributing a 50 Hz stimulation over 20 to 40 electrodes, except during tetanization (Wagenaar 2006a).




Figure 0.2 Simultaneous tetanization of two electrodes produced a lasting change in network response properties.

PCA was applied to the FRH statistic. **(a)** The fluctuation of the first 2 principle components within period Pre1, see top of (b), was compared to the change from Pre1 to Pre2 to quantify the drift before the tetanus. The fluctuation within Pre2 was compared to the change from Pre2 to Post1 to quantify the change across the tetanus. The ratios of change versus fluctuation before and across the tetanus are shown in cyan and blue respectively for four different experimental protocols (N=4 cultures tested, 6 probe electrodes per culture, for each protocol). Significant differences between change and fluctuation are marked as (* $P < 10^{-4}$), Wilcoxon rank sum test). The change due to the tetanus was larger than the change from drift in all cases. No significant change from drift or tetanus occurred in the presence of elevated magnesium. Weaker tetani lasted 5 min while tetani used throughout the rest of the chapter lasted 15 min. See text for a discussion of the protocols. **(b)**, In an example experiment from the burst quieting protocol, the tetanus caused a new PC cluster to form. Data is from table 7.3 in (Wagenaar 2006a).

The same experiments conducted with the addition of +1 mM Mg^{++} above baseline did not show tetanus-induced plasticity nor network drift outside of the tetanus (Fig. 3.2). An increased level of Mg^{++} will occupy more NMDA-Rs, decreasing their effect in the transmission of synaptic signals. NMDA-R is considered important for neural plasticity, and this protocol served as a control that tetanus-induced changes were biological. Interestingly, the network plasticity that Jimbo observed was in the presence of elevated Mg^{++} , counter to Figure 3.2A and conclusions by Wagenaar (Wagenaar 2006a).

The use of PCA and further comparing it to FRH and FR statistics was not pursued further. The primary reason was that the decision of which PC to include is not well defined because the PCA's matrix transformation obscures the relationship between PCs and their underlying data: PCA is often used to analyze multi-dimensional data by reorganizing the component axes based on largest variance. As a consequence, the early principal components are the most significant as they hold the most variance, and later principal components can be ignored (O'Connell 1974). We felt uncomfortable arbitrarily dropping data, especially considering biological properties originally considered to be noise can turn out to be important in neural information processing (i.e. synaptic noise (Mainen and Sejnowski 1995) and variation in action potential propagation discussed in Chapter 4).

3.3.2 Incorporating spatial and temporal information: the center of activity trajectory (CAT).

We then focused our attention on a population statistic that did not drop data and additionally added spatial information, the center of activity trajectory (CAT). CAT was measured from the evoked responses to RPS in six experiments on living cultured cortical networks and compared to FR, FRH and SCCC. The change across the tetanus was significantly greater than the drift before the tetanus for CAT ($p < 1 \times 10^{-4}$,

Wilcoxon signed rank test), FRH ($p < 0.01$) and SCCC ($p < 0.01$), but not for FR ($p = 0.013$). *C/D* ratios (Methods and Appendix G) were used to quantify the change before the tetanus and the change across the tetanus. If the statistic could detect tetanus-induced plasticity, the *C/D* ratio across the tetanus should be significantly greater than the *C/D* ratio before the tetanus. The statistics of *C/D* from six experiments are shown in Figure 3.3.

In living networks, the change in synaptic weights throughout the network cannot be determined. Therefore, experiments were repeated in a simulated network (Chao, Bakkum et al. 2007) where synaptic weights could be directly measured. Results showed the CAT best exposed tetanus-induced plasticity (Figure 3.3C). Despite the reduced dimension over FRH, the addition of spatial information increased its sensitivity to tetanus induced plasticity. The location of the electrodes is relevant. The relationship between sensitivities of various statistics remained in the living network, and CAT was again the most sensitive statistic.

In order to get some idea of the degree of localization of function in cultured cortical networks, the performance of CAT statistic with electrode locations shuffled (CAT-ELS) was calculated. In CAT-ELSS, the information about the physical locations of the recording electrodes was removed. In both simulations (Chao, Bakkum et al. 2007) and experiments in living cultures, the electrode locations were shuffled ten times, and ten different corresponding CAT-ELSS were generated. The performance of these CAT-ELSS was evaluated and compared to the original CAT.

CAT, unlike the other statistics, incorporates the physical locations of the recording electrodes. This is the primary difference between methods, and we attribute CAT's superior performance in both living and simulated networks to this feature. For simulated networks (Chao, Bakkum et al. 2007), the comparison of the performance between CAT-ELS and original CAT is shown in Figure 3.4C. The detectable synaptic




Figure 0.3 Comparison of the performance of different statistics across a tetanization in living MEA cultures.

(a). The C/D of the statistics (Plasticity across the tetanus normalized by ongoing plasticity = Change / Drift) from 6 experiments showed that the change across the tetanus was significantly greater than the drift before the tetanus for CAT (**, $p < 1e-4$, Wilcoxon signed rank test), FRH (*, $p < 0.01$) and SCCC (*, $p < 0.01$), but not for FR ($p = 0.013$). The p-values indicate that CAT was more capable of detecting the change over the drift than FRH, SCCC and FR. (b). Different patterns of CATs were observed before and after tetanization from responses to probe. CATs from every block were overlaid (black trajectories), and the average CATs were shown by series of circles (from blue to red). (c). CAT showed the highest sensitivity to changes in synaptic state among 6 metrics in a simulated network (Chao, Bakkum et al. 2007). The percent synaptic change was a function of tetanus duration (S1 to S10) and could be observed in the model. Therefore, a known amount of plasticity was plotted and the minimum amount of plasticity significantly detectable by each statistic was found at the point the P-values reach a threshold of 0.05 (red arrows). Mean and standard deviation plotted. $N = 50$ model networks.




Figure 0.4 Comparison of CAT and CAT-ELS (Electrode Location Shuffled).

(a). The statistics of C/D for CAT-ELS in living networks ($n=60$, 6 experiments, 10 shuffles for each experiment). The change across the tetanus was no longer significantly different than the drift before the tetanus ($p=0.19$, Wilcoxon signed rank test), unlike CAT (**, $p<1e-4$). Thus, the shuffling of electrode locations greatly reduced the sensitivity of CAT. (b). The electrode locations shuffling “collapsed” the patterns of CAT-ELSs before and after tetanization. The difference between before and after tetanization trajectories (compared to Fig. 3.3) was reduced in CAT-ELS. (c). CAT-ELS from the simulated model with the same representation as Figure 3.3C (Chao, Bakkum et al. 2007). Ten sensitivity curves corresponding to different random shuffled electrode locations (CAT-ELS) and the mean of the ten curves (Mean CAT-ELS) are shown. The sensitivity curve of FRH is also shown for comparison. The decrease in sensitivity emphasized the importance of physical electrode locations in the sensitivity of CAT to synaptic change.

change was worse than CAT The decrease in performance indicated that electrode locations significantly affect the performance of CAT.

For living networks, the comparison between CAT and CAT-ELS is shown in Figure 3.4. The electrode location shuffling ‘collapsed’ the patterns of CAT-ELSs before and after tetanization (compared to Figure 3.3). The difference between pre-tetanization and post-tetanization clusters found in CAT was also reduced in CAT-ELS. The statistics of C/D for CAT-ELS ($n=60$, six experiments, ten shuffles for each experiment) are

shown in Figure 3.4. The change across the tetanus was significantly greater than the drift before the tetanus for CAT ($P < 1 \times 10^{-4}$, Wilcoxon signed rank test), but not for CAT-ELS ($P = 0.19$).

3.4 Discussion

3.4.1 Statistics of functional plasticity in extracellular multi-electrode recordings

While comparisons of firing rates show plasticity in intracellular recordings, more detailed statistics incorporating spatiotemporal population activity patterns are needed to reveal plasticity in extracellular multi-electrode recordings. Electrode spacing on the order of hundreds of microns means that any induced or observed plasticity will span pathways of multiple neurons instead of neighboring monosynaptic neurons (Jimbo, Tateno et al. 1999). Intracellularly, synaptic strength is directly observable by stimulating a pre-synaptic neuron while recording from an adjacent post-synaptic neuron. Extracellularly, synaptic noise across a chain of neurons and convergent pathways will obscure firing rate measures of stimulus-induced plasticity.

Alternatively, by incorporating the timing and spatial flow of activity, spatiotemporal patterns have been found both *in vivo* and *in vitro*. Spike sequences, imposed upon the network by behavioral manipulations, recur spontaneously during subsequent sleep episodes (Nadasdy, Hirase et al. 1999; Nadasdy 2000; Lee and Wilson 2002). Calcium imaging of cortical slices reveals reactivation of sequences of neurons, ‘cortical songs’, with distinct spatiotemporal structures over tens of seconds (Ikegaya, Aaron et al. 2004). Robust recurrent spike patterns were also found in a detailed cortical simulation (Ikegaya, Aaron et al. 2004) and in living slices (Fellous, Tiesinga et al. 2004). CAT provides a new and simple statistic to detect spatiotemporal patterns in networks and extends the previous studies by quantifiably demonstrating its ability to discern plasticity.

3.4.2 Region-specific plasticity

Although FRH included detailed temporal information about the activity dynamics at all electrodes, it was less capable of capturing network plasticity than CAT, which has the same temporal resolution as the FRH but ‘condenses’ the spatial dimension by linear combination. We hypothesize that this was due to the inclusion of spatial information of the electrode locations. The performance and the sensitivity of CAT with electrode locations shuffled were significantly worse than non-shuffled CAT, both in simulation and in living networks (the change across the tetanus was significantly greater than the drift before the tetanus for CAT, but not for CAT-ELS) (see Figure 3.4). This indicates that activity varied systematically with the electrode location, and also suggested that the observed network plasticity was *region specific*: the plasticity was not symmetrically distributed throughout the network. This further suggests that despite the apparent random connectivity of cultured neurons, neuron location played a role in tetanus-induced plasticity. Region specificity was not limited to plasticity induced by tetanization. In simulation, we also altered the weights of randomly selected synapses in reference networks to different degrees to generate different new network states. CAT still showed the highest sensitivity to changes, and furthermore, the sensitivity of CAT-ELS was still significantly lower (data not shown). Despite the synaptic plasticity not being region specific, the spatiotemporal flow of neural activity was region dependent, effectively making the plasticity of neural activity region specific. This result supports the notion of synfire chains or braids of neural activity (Ikegaya, Aaron et al. 2004; Izhikevich 2006), where information is transmitted in a pipeline of neighboring pathways as opposed to a single string of connections.

A common misconception regarding dissociated cultures is that they are random, homogeneous and lack structure, and thus cannot support stable changes to synaptic weights associated with memory formation. While plated from a random cell suspension,

microscopic observation reveals that a heterogeneous arrangement develops over time (Gross and Kowalski 1999). Although very different than structures found *in vivo*, the ability of neurons and glia to interact in complex ways remains, and a network having a diverse array of activity arises spontaneously (Wagenaar, Pine et al. 2006a). Altering sensory input of thalamic relays to cortical areas has demonstrated that the cortex develops structure according to the type of the sensory input (Sur, Garraghty et al. 1988), which suggests an important relationship between neural structure and function. CAT demonstrates that structure is also relevant to neural function in a cultured network, and that tetanic stimulation alters network function. Future experiments will incorporate closed-loop sensory-motor feedback and optical imaging to investigate the network mechanisms of our cultures to functionally and structurally adapt to environmental interaction (Potter, Wagenaar et al. 2006).

3.4.3 Plasticity versus spontaneous bursting

Without external stimulation, the most prominent feature of spontaneous activity found in MEA cultures and in simulated networks is synchronous bursting (Wong, Meister et al. 1993; Kamioka, Maeda et al. 1996; Gross and Kowalski 1999; Van Pelt, Wolters et al. 2004; Wagenaar, Pine et al. 2006a), and bursts were found to have effects on tetanus-induced synaptic plasticity in cortical neurons (Maeda, Kuroda et al. 1998). In simulation, the network synaptic state after tetanization was found to change gradually due to the presence of spontaneous bursts, which makes quantifying tetanus-induced plasticity difficult (Chao, Bakkum et al. 2005). In the six experiments we performed on living MEA cultures, 8.57 ± 3.33 spontaneous bursts per minute and 16.06 ± 4.55 stimulus-evoked bursts per minute were observed. Even with the presence of the spontaneous bursts, the tetanus-induced plasticity was still detected by using CAT. Since the level of bursting can be finely controlled in MEA cultures with multisite stimulation (Wagenaar, Madhavan et al. 2005), CAT proved useful for investigating how the degree

of bursting affects a network's ability to produce and/or maintain plasticity (Madhavan 2007).

3.4.4 Conclusion

Incorporating spatial and temporal information improved the performance of neural activity statistics in networks of cortical neurons grown over MEAs. CAT's superior performance, sensitivity and low computational load make it an attractive method for real-time applications. CAT can also be applied to *in vivo* multi-electrode or optical recording studies for neural activity aligned to behavioral or sensory cues. As techniques for observing distributed activity become faster and more fine-grained, studying the details of the spatial flow of activity through neuronal networks will reveal more and more about processes of learning and memory.

CHAPTER 4

PLASTICITY INDUCED BY PATTERNED STIMULATION:

CHANGES IN ACTION POTENTIAL PROPAGATION³

The precise temporal control of neuronal action potentials is essential for regulating many brain functions. From the viewpoint of a neuron, the specific timings of afferent input from the action potentials of its synaptic partners determines whether or not and when that neuron will fire its own action potential. Tuning such input would provide a powerful mechanism to adjust neuron behavior and in turn, that of the brain. However, axonal plasticity of action potential timing is counter to conventional notions of stable propagation and to the dominant theories of plasticity focusing on synaptic efficacies. Here we show the occurrence of activity-dependent plasticity of action potential propagation. We used a multi-electrode array to induce, detect, and track changes in propagation in multiple neurons while they adapted to different patterned stimuli in controlled neocortical networks in vitro. Changes occurred in action potential delays (up to 4ms or 40% after minutes and 13ms or 74% after hours) and amplitudes (up to 87%). The changes did not occur when the same stimulation was repeated while blocking synaptic activity. Even though the plasticity depended on synaptically evoked action potentials, its expression was non-synaptic: action potential propagation. We conclude that propagation plasticity is a cellular mechanism used to tune information processing in neuronal networks and potentially learning and memory in the brain.

Chapter 3 demonstrated that plasticity could occur and also provided methods to map neural activity into animat movement. However, the use of tetani for sensory feedback would allow only crude encoding of sensations. Therefore, we tested whether plasticity could arise for stimulation at the opposite extreme: low frequency patterned stimulation. Plasticity did indeed occur, suggesting the network could respond to most any type of electrical stimuli. In addition, we observed that plasticity occurred in the propagation of action potentials, a potentially impactful finding described further below. Similar versions of patterned stimuli were used as plasticity inducing stimuli in the embodiment experiments in Chapter 5 and Appendix A.

³ Manuscript submitted as:

Bakkum DJ, Chao ZC, Potter SM, “Long-term activity-dependent plasticity of action potential propagation in cortical networks.”

4.1 Introduction

The specific arrival times of afferent synaptic excitatory and inhibitory potentials determine how they summate as they converge towards a neuron's soma and whether or not and when to fire an action potential. Tuning such input timing would provide a powerful mechanism to adjust the output of a neuron, and potentially, could be a mechanism for learning and memory in the brain. We hypothesized that axonal plasticity of action potential propagation could vary how information is processed in the brain by regulating the timing and magnitude of synaptic input impinging on a neuron. This is fundamentally different than the dominant theories of neural plasticity that focus on the efficacy of synapses.

The specific arrival times of afferent synaptic excitatory and inhibitory potentials determine how they summate as they converge towards a neuron's soma and whether or not and when to fire an action potential. Tuning the timing of such input would provide a powerful mechanism to adjust the output of a neuron, and potentially, could be a mechanism for learning and memory in the brain. We hypothesized that axonal plasticity of action potential propagation could vary how information is processed in the brain by regulating the timing and amplitude of synaptic input impinging on a neuron. This is fundamentally different than the dominant theories of neural plasticity that focus on the efficacy of synaptic transmission.

Axons in the mammalian cortex have traditionally been regarded as stable transmission cables. However, this view is more likely due to a lack of, rather than support from, experimental evidence (Debanne 2004; Clark and Hausser 2006) because their small diameter ($< 1 \mu\text{m}$) makes direct recordings at multiple sites difficult. The action potential is often viewed as a binary signal, but recent experiments have found a novel analog mechanism: the ability to encode a neuron's background synaptic activity in the amplitude of its action potential in cortical slices from ferrets (Shu, Hasenstaub et al.

2006) and hippocampal slices from rats (Alle and Geiger 2006). Moreover, the temporal control of action potential propagation could encode a vast amount of information. For example, introducing non-uniform, although fixed, conduction delays in a model network produced a potentially unlimited number of “polychronous” groups of neurons capable of recognizing and classifying complex spatiotemporal stimuli (Izhikevich 2006), a theoretical canvas for memories; spike-timing dependent plasticity (STDP) organized the groups by potentiating afferent signals whose timings were close enough to integrate and produce an action potential in a post-synaptic neuron. Alternatively, modulating the propagation speeds themselves could also compose such groups (Eurich, Pawelzik et al. 1999). STDP causes long-term synaptic potentiation or depression based on the temporal order of pre and post-synaptic activation (Bi and Poo 1998), and the magnitude of plasticity is greatest at smaller activation intervals, with a sharp discontinuity from potentiation to depression as the post-synaptic neuron went from lagging to leading. This feature of the STDP rule makes it a sensitive detector of relative spike order and a means to create and maintain causal pathways or synfire chains (Abeles 1991). Finely tuning action potential propagation speed could provide a means to continuously mold functional circuits in the brain throughout life.

The precise control of action potential propagation has been shown to be important in both the central and peripheral nervous systems after development or many days of experience (Waxman and Bennett 1972; Swadlow 1985), but modulation of propagation has not been reported on the seconds to hours time scales relevant to learning and memory consolidation. For instance, the propagation of action potentials in the olivocerebellar pathway of rats can vary by 40% such that isochronic conduction occurred between neurons independent of axon length, a finding conserved even between animals (Sugihara, Lang et al. 1993). In addition, a decrease in the size of a spinal stretch reflex (H-reflex) after 40 days of operant conditioning was accompanied by an 8% (rat) or 6% (monkey) decrease in motoneuron conduction velocity (Carp, Chen et al. 2001).

The reported propagation velocities were consistent with myelinated fibers. Propagation speeds have been proposed to be tuned with millisecond precision through axon diameter (Rushton 1951; Sugihara, Lang et al. 1993), myelin thickness (Rushton 1951; Fields 2005), the location of nodes of Ranvier (Waxman 1997), and the kinetics of voltage gated sodium channels (Halter, Carp et al. 1995).

Research on neural intrinsic excitability has demonstrated that, in addition to the well studied synaptic plasticity, activity-dependent plasticity can be expressed outside of synapses on fast time scales (Daoudal and Debanne 2003; Zhang and Linden 2003; Debanne 2004; Xu, Kang et al. 2005). However, the proposed molecular mechanisms involved have not been investigated with respect to plasticity of action potential propagation. In embryonic rat hippocampal cultures, injecting depolarizing currents into a post-synaptic neuron such that it fired in synchrony with pre-synaptic activity for 80 seconds caused long-term plasticity of the excitability of the pre-synaptic neuron, requiring post-synaptic NMDA-R and pre-synaptic PKC activation and affecting the gating kinetics of voltage gated sodium channels (Ganguly, Kiss et al. 2000); while in part synaptic in origin, the plasticity was expressed outside of synapses. These channels actively propagate action potentials, and changing their gating kinetics could change action potential shape and propagation velocity within the central nervous system within minutes.

The above work demonstrates that action potential propagation can encode information (Milton and Mackey 2000; Izhikevich 2006), is regulated (Waxman and Bennett 1972; Swadlow 1985; Sugihara, Lang et al. 1993), and potential mechanisms for changes in AP propagation exist (Ganguly, Kiss et al. 2000). However, the potential mechanisms have not been related to propagation plasticity, and propagation has not been shown to be regulated on the minutes to hours time scale relevant to much learning and memory. Here, we show experimental evidence of the rapid induction of activity-dependent long-term plasticity of action potential propagation. Using networks of rat

neocortical neurons and glia grown over multi-electrode arrays (MEA), we observed the timing of direct electrically-evoked action potentials (dAPs) in multiple neurons changed up to 4 milliseconds (40%) after minutes of recording and up to 13 milliseconds (74%) after hours, adapting to different patterns of low frequency stimulation. The estimated propagation velocities suggest propagation occurred in unmyelinated axons (Waxman and Bennett 1972; Swadlow 1985; Manor, Koch et al. 1991; Debanne 2004; Patolsky, Timko et al. 2006). The plasticity was activity-dependent since no change occurred when synaptically-evoked action potentials (sAPs) were blocked using antagonists of the fast synaptic transmitter receptors NMDA-R, AMPA-R, and GABA-R. As with the intrinsic excitability experiments, the plasticity was expressed outside of the synapses. Action potential amplitude similarly adapted to the various patterned stimuli (up to 87% change). We conclude that propagation plasticity is a cellular mechanism used to tune temporal coding schemes and information processing in neural networks, and potentially learning and memory in the brain.

4.2 Results

This section begins with data describing the ability of our preparation to characterize dAP propagation (Figs. 4.1, 4.2, and 4.3). This is followed by experimental evidence of changes in the latencies and amplitudes of the dAPs in response to patterned stimulation (Figs. 4.4 and 4.5). Control experiments indicated that the plasticity, while requiring sAPs, was expressed independently of synaptic activity since the induced changes persisted when sAPs were blocked. The section ends with a theoretical consideration of our results intended to inspire others to appreciate the potential roles of action potential delays and timing in neural computation.

4.2.1 Propagation of direct electrically-evoked action potentials (dAPs) was robustly detected using planar multi-electrode arrays (MEAs).

By using cortical neurons cultured on extracellular multi-electrode arrays (MEAs) (Fig. 4.1A and Appendix G), action potentials can be sampled at multiple sites up to millimeters apart for extended durations. Stimulation by one electrode evokes neural responses that can be recorded in a subset of the rest of the electrodes (Fig. 4.1B). Of these, dAPs have been observed up to 25 ms later and can be distinguished from subsequent sAPs based on their high reliability of occurrence (>80%), low jitter (160 μ s), and consistency of waveform (Lipski 1981; Marom and Shahaf 2002; Wagenaar, Pine et al. 2004) (Fig 4.2). They are pre-synaptic as they persist when synaptic activity is blocked using fast neurotransmitter receptor antagonists, and they are propagating action potentials as they are eliminated with TTX (Wagenaar, Pine et al. 2004). In the following experiments, we quantified changes in dAP propagation by measuring their amplitudes and latencies (Fig. 4.2C) after the downswing of a biphasic stimulus pulse (Fig. 4.3), termed a “probe”.

Figure 4.1 shows that MEAs can be used to investigate the dAPs of many neurons in a cultured network. Our MEAs contained 59 functional electrodes (Fig. 4.1A), and many could evoke and/or record neural activity depending on the relative location of the neurons (Fig. 4.1B). On average, an electrode that evoked activity yielded 6.2 ± 5.7 dAPs detected elsewhere. Figure 1B contains a set of raster plots of dAP latencies (x axes), arranged topographically by recording electrode and color coded by stimulation electrode, in response to 24 minutes (y axes) of “whole-dish probing” of a 4 week old culture. Whole-dish probing consisted of stimulating each electrode 240 times in random order at an overall rate of 10Hz, with sAPs blocked. It was used here to quickly sample all accessible dAPs in each network and later to quantify induced plasticity.




Figure 0.1 Multi-electrode arrays (MEA) robustly detected dAPs.

(A) Neurons at 6 days in vitro grown over an MEA with 30 μm diameter electrodes spaced 0.2 mm apart. The large reference electrode is outside the field of view. Color represents the location of the stimulation electrode that evoked the dAPs plotted in B. (B) Recording electrodes (arranged topographically) detected dAPs (dots) in a 4-week-old culture in abundance. (C) Pie chart comparing the proportion of the electrodes showing and/or evoking dAPs in 5 three to four-week-old cultures from 3 dissociations (295 electrodes; data also in Fig. 4.5). Their incomplete overlap suggests that the location of an extracellularly recorded action potential can differ from the location of an extracellularly evoked action potential within a given neuron. The data from B are presented in the adjacent plot as an example. (D) Histograms of dAP latencies, distances from the stimuli site, and estimated velocities.

Extracellular stimulation and recording of neural activity may act on different sites of a neuron (McIntyre and Grill 2000). Interestingly, in data from 5 cultures (Figure 4.1C), 34% of the stimulating electrodes that could evoke activity (21% + 40%) did not record activity (21%), and 13% of the electrodes that recorded activity (6% + 40%) did not evoke activity (6%) when stimulated. Although specific numbers will depend on the density and location of the neurons (see Methods), this nevertheless suggests that action potentials tended to be recorded from different parts of a neuron than where they were

induced by electrical stimulation. For example, an action potential in a soma may produce a larger signal on an extracellular electrode than one in an axon because more ions would be needed to depolarize the larger surface area. Therefore, the recording range of an electrode would be larger for cell bodies than axons. Conversely, an axon would require less current than a soma to become depolarized, and the stimulation range of an electrode would be larger for axons than cell bodies. This hypothesis was supported implicitly by the greater number of electrodes that evoked activity but did not show activity compared to the number that showed activity but did not evoke activity: a cell body could be located near just one electrode, while an axon could branch and pass over or near multiple electrodes.




Figure 0.2 Neural activity recorded on one extracellular electrode in response to stimulation at another consists of an early directly-evoked action potential (dAP) phase and a later synaptically-evoked action potential (sAP) phase.

(A) The raster plot (1 dot per action potential) shows the first 100 ms of neural responses to 3 hours of periodic 1/4 Hz probe (red P) stimulation. (B) The peri-stimulus time histogram and (C) overlaid extracellular voltage traces across all trials emphasize the consistency of the early phase with respect to the later phase. The sharp peaks in the histogram arise from the trains of two dAPs. See also (Marom and Shahaf 2002).




Figure 0.3 DAPs were time-locked to biphasic voltage stimuli.

Latencies of dAPs from 13 stimulation-electrode / recording-electrode pairs (thin lines) were measured from the beginning of voltage stimuli with varying pulse widths (inset); sAPs were blocked. The thick line is a linear regression on all data points.

Figure 4.1D further characterizes the observed dAPs. A few were detected up to 25 ms after being evoked, but the majority were earlier. They could not be detected sooner than ~2 ms nor on the stimulating electrode due to the presence of electrical stimulation artifact (Wagenaar and Potter 2002). DAPs could have been evoked in the middle of axons that passed near an electrode, and thus the actual delay from the neuron to its post-synaptic targets would be longer than what was measured. Moreover, increased electrode spacing may lead to longer detectable latencies. Due to the geometry of the MEA, the majority of distances between stimulating and recording electrodes were closer to the minimum distance of 0.2 mm. Therefore, the plotted histogram of distances was normalized by the distribution of all possible inter-electrode distances. The estimated

velocities suggested the detected dAPs propagated through unmyelinated neurites (Waxman and Bennett 1972; Swadlow 1985; Manor, Koch et al. 1991; Debanne 2004; Patolsky, Timko et al. 2006).

By recording extracellular voltage spikes of dAPs, an MEA is able to measure propagation delays with high precision (Fig. 4.3 and Methods) and can reveal whether or not plasticity occurred after an experimental manipulation. Negative extracellular current most effectively depolarizes a neuron, which for a voltage step occurs during the pulse downswing ($I=C*dV/dt$). Thus, unlike during a current step, dAP timing can be precisely time-locked to the stimulus downswing, as confirmed by plotting dAP latency for different voltage step widths (Fig. 4.3) (Wagenaar, Pine et al. 2004). In Figure 4.3, the width of rectangular stimuli were varied while magnitude was kept constant at ± 500 mV; 30 stimuli were delivered per width in random order. A linear regression of the averaged latencies (dots) from each of the 13 pairs was used to align the data: the regressions' y-intercepts at the 200 mV phase width were set as 0 ms latency change. Another linear regression on all aligned data points produced a representative slope of 1.10 ms/ms, indicating most dAPs were time-locked to the voltage downswing. A slope $m = 0$ would indicate the dAPs were triggered at the beginning of stimuli; $m = 1$, at the downswing; $m = 2$, at the end; and $0 < m < 1$ or $1 < m < 2$, in the middle of a phase. See also (Wagenaar, Pine et al. 2004).

4.2.2 Action potential propagation depended on neural activity and variation of stimulation pattern.

By varying a simple low frequency stimulation pattern every 40 minutes, we induced changes in the time elapsed for dAPs to propagate from a probe electrode to a recording electrode and also in their extracellularly recorded amplitude (Fig. 4.4). Each pattern consisted of alternatively stimulating 2 electrodes at 2 second intervals. (1/4 Hz stimulation per electrode and 1/2 Hz overall stimulation; Fig. 4.4A). The second




Figure 0.4 Action potential propagation depended on ongoing neural activity and stimulation pattern.

(A) Experiment protocol. 1/4 Hz probe stimuli (red) produced dAPs whose latencies and amplitudes were investigated for plasticity. A context electrode (gray) was stimulated 2 sec prior to each probe stimulus, giving an overall 1/2 Hz stimulation frequency, and its location was shifted every 40 min to produce different patterns of stimulation (numbers and shaded bars). Right: electrode locations for data in B. (B) Example raster plots of a given dAP recorded on one electrode in response to probe stimulation of another electrode in culture media (left, Unblocked; sAPs are plotted with smaller markers) and when blocking sAPs (right, Blocked). Ongoing neural activity modified latency (x axis) and amplitude (color). Varying stimulation pattern significantly altered dAP latency (C) and amplitude (D) within 5 minutes (Mean $\pm$ s.e.m).

electrode, termed probe, was fixed and used throughout the duration of an experiment to consistently sample dAPs. The location of the first “context” electrode was moved spatially every 40 minutes to make each new pattern. The dAPs evoked by context stimuli were not analyzed. Interestingly, we found that the dAPs evoked by the probe stimuli changed via gradual shifts and jumps in latency (up to 4 ms or 40%) and amplitude (up to 20 mV or 80%) but not when the stimulation was repeated in the presence of antagonists of NMDA-R, AMPA-R, and GABA-R (Fig. 4.4B). These blocked all spontaneous activity (sAPs) except for an occasional self-firing neuron

(Latham, Richmond et al. 2000). Therefore, even though dAPs were detected pre-synaptically, changes in their latency and amplitude required successful synaptic transmission of glutamate and/or GABA. Electrical artifact and chemical interactions at the electrode interface did not contribute since changes were minimal for identical stimulation while blocking sAPs (Fig. 4.4).

The propagation of dAPs adapted to new stimulation patterns within minutes. Changes in dAP latencies were significantly greater just after a change in the patterned stimulation than at the end of a 40 minute block of a given patterned stimulation (Fig. 4.4C; Mean $\pm$ s.e.m. Unblocked: $**P < 1e-6$, $n = 130$ dAP trains. Blocked: $P = 0.24$, $n = 115$ dAP trains; Wilcoxon signed rank test for paired samples. 6 cultures from 4 dissociations). To a lesser significance, dAP amplitudes were similarly altered (Fig. 4.4D; Unblocked: $*P = 0.003$. Blockers: $P = 0.59$). The change “across” adjacent patterned stimulation blocks was calculated as the absolute value of the difference in mean latency between the 5 minute periods just prior to and just after shifting the location of the context electrode. Each 5 min period consisted of 75 probe stimuli. The change “within” a period was calculated between the 10 to 5 min and the 5 to 0 min periods prior to a shift. The patterned stimulation protocol could be considered a simplified analog for memory processes found in the brain. For example, repetitive cortical activation by the hippocampus during sleep consolidates memories (Wilson and McNaughton 1994) and repetitive body movements lead to cortical plasticity (Classen, Liepert et al. 1998).

4.2.3 Plasticity of action potential propagation had a non-synaptic expression.

Although induction of propagation plasticity depended on the occurrence of sAPs, long-term plasticity of action potential latency and amplitude was expressed outside of synapses (Fig. 4.5). Accompanying plasticity expressed at synapses is expected to have occurred also. The changes in dAP latencies and amplitudes in Figure 4.4 could reflect a

variation in propagation or instead a variation in the recent background synaptic activity. The technique of using extracellular electrodes to investigate axonal properties has a long history, and delays in antidromic propagation have been shown to depend on the somatic membrane potential (Lipski 1981). In particular, the impedance mismatch due to the change in volume from the axon into the soma causes a delay in propagation proportional to the somatic membrane potential, which varies with synaptic input (Lipski 1981). However, such changes in delay were elastic, recovering within 100 ms after altering the membrane potential.




Figure 0.5 The cause of plasticity was neural activity, and the site of plasticity was not synaptic.

(A) Experiment protocol. SAPs were blocked to eliminate the influence of ongoing synaptic activity, and 3 identical periods of whole-dish probing (shape and color coded) were applied before and after the 5 hours and 20 minutes of patterned stimulation (Fig. 4.4). (B) Example extracellular voltage traces for two separate dAPs during each whole-dish probing period (240 traces averaged). Changes that accrued during the patterned stimulation persisted (blue square to black triangle): they were not reflections of ongoing synaptic activity. Changes were minimal during patterned stimulation in the presence of blockers (red circle to blue square): the accrued changes were not artifacts from the electrical stimulation or replacing media. (C) Statistics for all observations (Mean $\pm$ s.e.m.). (D) Changes in latency were not monotonically correlated to changes in amplitude. The outlying data points, using an arbitrary cut-off at 10% of the distribution, were plotted with darker dots.

To experimentally eliminate variation in membrane potential due to background synaptic activity, additional whole-dish probing periods were conducted, in the presence of fast synaptic receptor blockers, before and after the patterned stimulation protocols (Fig. 4.5A), and long-term plasticity of action potential propagation delay and amplitude persisted (Fig. 4.5B). The whole-dish probing periods consisted of stimulating every electrode 240 times in random order at an overall frequency of 10 Hz (24 minutes) and recording the dAPs. If the changes in propagation induced by the patterned stimulation were expressed independently of synaptic activity, dAP latencies and amplitudes would vary between the periods of whole-dish probing before and after the 5 hour and 20 minute sequence of patterned stimulation blocks. Then as a control, dAP latencies and amplitudes should not vary between the periods of whole-dish probing before and after patterned stimulation conducted with blockers. Consequently, action potential propagation did change (in one case, latency decreased by 13 ms or 74%) significantly more after the patterned stimulation without blockers than with blockers (Fig. 4.5C; Mean $\pm$ s.e.m.; $**P < 1e-6$ for change in latency, $*P = 0.003$ for change in amplitude; Wilcoxon signed rank test for paired samples. $n = 904$ dAP trains. 5 cultures from 3 dissociations). Propagation plasticity was expressed outside of the synapses since it could be detected in the presence of synaptic blockers. Not all dAPs changed latency and amplitude (Fig. 4.5D), suggesting the plasticity induced by the patterned stimuli discriminated among different pathways of propagating neural activity (Jimbo, Tatenos et al. 1999). As opposed to Figs. 4C and 4D, the propagation plasticity averaged in Fig. 4.5C and plotted in Fig. 4.5D had accrued over the duration of the entire patterned stimulation protocol: 5 hours and 20 minutes instead of 5 minutes.

Increasing the overall stimulation frequency from 1/2 Hz, as with patterned stimulation, to 10 Hz during the whole-dish probing was done to decrease experiment durations. Using a different stimulation frequency was not of concern because the whole-dish probing was always done in the presence of blockers, where propagation plasticity

{{did not occur}}, and dAPs evoked by whole-dish probing sequences were compared only to those evoked by other whole-dish probing sequences. Moreover, the per electrode stimulation frequencies were comparable: whole-dish probing stimulated a given electrode every 5.9 seconds on average versus every 4 seconds during patterned stimulation.

4.2.4 Theoretical computational capacity.

A given neuron might have the ability to differentially modulate the timing and/or amplitude of action potentials impinging on multiple post-synaptic cells, greatly increasing the available computational capacity. Both increases and decreases in dAP latency (up to 200%) and amplitude (up to 600%) occurred, but interestingly, a monotonic correlation between the two was not found (Fig. 4.5D; $P = 0.22$, $r = 0.04$ Spearman's rank correlation coefficient). Homogeneous plasticity of cell properties along the length of a neurite, for example of voltage gated Na^+ channels, would be expected to cause a monotonic change between action potential latency and amplitude. Therefore, the plasticity occurred either via different mechanisms or in a locally discriminate manner throughout the neuron, perhaps by geometrical variations in axonal varicosities which could cause conduction delays up to 100s of ms at each synaptic bouton (Goldstein and Rall 1974; Manor, Koch et al. 1991).

As an example of the potential computational capacity for such a locally controllable propagation, a 1.2 mm axonal branch could achieve 1 billion temporal configurations for its synaptic output. Estimating the average conduction velocity to be 0.25 mm/ms (Fig. 4.1D, histogram peak multiplied by a safety factor of 2) and the discrete time resolution to be the average dAP jitter during patterned stimulation experiments without blockers, 160 ms, gives a discrete spatial resolution of 40 mm (distance = velocity * time). This spatial resolution would not be decreased by considering the magnitude of the observed latency changes nor by synapse density:

axonal synaptic boutons were spaced 5 to 10 mm apart on average in cat cortical neurons (Anderson, Binzegger et al. 2002). A billion or approximately 230 configurations give 1.2 mm in length (1.2 mm equals a 40 mm spatial resolution multiplied by 30 axonal sections). A neuron would have even more capacity considering (1) temporal encoding is analog, (2) axonal arbors have 2 to 3 orders of magnitude greater length (Anderson, Binzegger et al. 2002; Debanne 2004) and (3) extensive branching (Sik, Penttonen et al. 1995; Debanne 2004), (4) the capacity of separate branches are multiplicative, and (5) encoding via action potential amplitude was not considered. However, the actual temporal and spatial resolutions are not known, neither are the rules of plasticity induction, and computational capacity is only an upper limit to memory capacity. In particular, what governs an increase versus a decrease in latency and amplitude? An understanding of these rules is needed to apply propagation plasticity in computational models or in further exploring its role in cognition.

4.3 Discussion

We have identified plasticity mechanisms that depend on synaptic transmission for induction, but are nevertheless expressed without it. As with changes in synaptic strengths, the changes we observed in action potential propagation are also likely to influence computation, learning, and memory in neural systems. For example, changes in action potential delay alter the type and number of attractor states in recurrent delayed neural loops (Foss and Milton 2000; Ma and Wu 2007) and neural networks (Izhikevich 2006). Elucidating the cellular mechanisms of propagation plasticity is left to future work, but possibilities include non-uniform changes in ion channel properties (Ganguly, Kiss et al. 2000), in the geometry of varicosities and branch points (Goldstein and Rall 1974) or axonal arbors, in the proximity of glia (Ishibashi, Dakin et al. 2006), and in lipid membrane composition (Bedlack, Wei et al. 1994). The changes we observed in cortical neurons could be generalized to occur throughout the brain, although the role of glial

sheaths may dominate in faster-conducting myelinated axons (Ishibashi, Dakin et al. 2006). Gap junctions were not considered because, in a similar preparation, both electrical and dye coupling experiments did not reveal coupled neurons (Nakanishi and Kukita 1998). Patterning axon growth over a series of electrodes (Suzuki and Yasuda 2007) or nanowire transistor recording devices (Patolsky, Timko et al. 2006) and/or optical imaging (Kawaguchi and Fukunishi 1998) could expose the discrimination, resolution, and possible morphological correlates of propagation plasticity.

Past research has set the stage to discover the rules governing temporal coding in the brain. The neural orchestra is comprised of not only synapses but many instruments, in part tuned by propagation plasticity. By using an MEA to robustly detect and track changes in the propagation of electrically evoked action potentials, we found that variation of a low frequency patterned stimulation modulated action potential propagation delays and amplitudes. Even though the induction of plasticity depended on synaptically evoked action potentials, its expression was non-synaptic: action potential propagation. Propagation varied for different stimulation patterns and became more stable for unvarying patterns, attributes necessary for playing a role in encoding memories. Latencies and amplitudes increased and decreased in an un-correlated manner, potentially allowing a neuron to have variable synaptic transmission among multiple post-synaptic neurons. In summary, the results suggest that propagation plasticity could serve as a cellular mechanism to tune temporal coding schemes and information processing in neural networks. Plasticity mechanisms that regulate the timing and amplitude of synaptic input impinging on a neuron challenge the dogma that memories are stored solely as changes in synaptic strengths.

4.4 Materials and Methods

Cell culturing and electrophysiology was done as reported in Appendix G. Cultures matured for 3 weeks prior to experimentation.

4.4.1 Pharmacology.

To block synaptically-evoked action potentials (sAP), fast synaptic receptor antagonists were applied at concentrations of 50 μ M bicuculline methiodide (BMI), 100 μ M 2-amino-5-phosphonovaleric acid (APV), and 10 μ M 6-cyano-7-nitroquinoxaline-2,3-dione (CNQX) (from Sigma), dissolved in culture medium and stored at -80°C. These are antagonists of GABA-R, NMDA-R, and AMPA-R, respectively. Fresh 35°C culture medium was used at the start of experiments and when changing medium between patterned stimulation and whole-dish probing experiments. When replacing medium from with to without the antagonists, cultures were washed 4 times by applying and discarding 1 mL volumes of fresh medium. A medium change lasted a couple minutes and cultures were allowed to equilibrate for an additional 30 minutes prior to beginning stimulation.

4.4.2 Experiment parameters.

For the patterned stimulation experiments, a 40 minute duration was chosen for each stimulation pattern to allow enough time to stabilize plasticity induced after the patterns were changed. A slow 1/2 Hz overall stimulation rate (Fig. 4.4A) was chosen to avoid network fatigue or refractory periods (Darbon, Scicluna et al. 2002) from compromising the evoked responses. The stimulation electrode evoking the most dAPs was chosen as the probe electrode. The probe was paired with electrodes evoking varying degrees of neural activity to create patterned stimulation with diverse network activity responses.

4.4.3 Automated detection of direct electrically-evoked action potentials (dAPs).

For each recording electrode, detected electrically evoked spikes were sorted from peaks in a firing rate histogram (Fig. 4.6), and latencies were tracked in consecutive moving time windows throughout the duration of an experiment. The histograms were constructed from 10 min windows (140 probe stimuli), stepped by 1 min. A given

histogram had a bin width of 0.04 ms, corresponding to the sampling frequency of 25 kHz, and was smoothed in latency with a Gaussian kernel size of 31. All histogram peaks and valleys were found, up to 25 ms in latency. Directly-evoked action potentials (dAP) produced tall sharp peaks while synaptically-evoked action potentials (sAP) produced broader shallower peaks. Thus, a peak was considered to be caused by dAPs if the peak height was greater than 2 times the highest valley plus 0.5, which was an empirically determined threshold. The analysis was done separately for positive and negative height spikes. All assigned dAPs were verified manually in raster plots and by waveform. To track changes in the latency of a dAP, the peaks in consecutive histograms were compared.

- (1) If an assigned peak overlapped a peak in the previous histogram within a tolerance, then the peaks were considered to arise from the same dAP. The tolerance was the width of the Gaussian at the peak's half height plus 440 μ s (11 samples) on either side. The tolerance allowed tracking a dAP that changed latency.
- (2) If a peak did not overlap any prior peaks, then a new dAP was assigned.
- (3) On rare occasions, if a peak overlapped 2 prior peaks, then the new peak was matched to the closest prior peak.

Occasionally, dAPs would disappear and reappear. Therefore, a peak was kept in memory until overlapped by a new peak. However, only stable dAPs were considered in the paper.




Figure 0.6 DAPs were automatically detected and tracked from peaks and valleys in consecutive smoothed post-stimulus time histograms.

(A) The first 25 ms of neural activity in response to 1/4 Hz probe stimulation were searched for the occurrence of dAPs in 10 min windows (shaded). (B) Expanded view of the shaded 10 min window in A. (C) A firing rate histogram of the neural activity in B was first constructed. (D) Then, the histogram was smoothed in latency with a Gaussian kernel, and all peaks (red circles) and valleys (green pentagons) were found (only 2 peaks and their corresponding valleys are plotted for clarity). A peak was considered to contain dAPs if it exceeded an empirical threshold.

CHAPTER 5

SHAPING EMBODIED NEURAL NETWORKS FOR ADAPTIVE GOAL-DIRECTED BEHAVIOR⁴

Therapeutic techniques used to treat neurological movement disorders attempt to re-link neuronal activity to a goal movement via the use of directed mental attention or perceptible feedback signals of behavior, and a lot of practice. While helpful, the benefits of these techniques are limited and not guaranteed among different individuals or different disorders. Successful cases of physical therapy have been accompanied by plasticity of the motor cortex. In turn, a range of studies have shown plasticity arising from electrical stimulation of neuronal tissue. Thus theoretically, targeted stimulation of the brain may improve therapeutic results by directly treating malfunctioning neuronal circuits and their related pathways. Many steps are required to reach this ambition, primarily that of quantifying the ability, limitations, and protocol for electrical stimulation to induce functional or adaptive changes in neuronal activity. Here, we developed an adaptive training algorithm, whereby a controlled *in vitro* neocortical network learned to modulate its dynamics and achieve pre-determined activity states within tens of minutes through the application of patterned training stimuli using a multi-electrode array. *A priori* knowledge of connectivity was not necessary. Instead effective training sequences were continuously discovered and refined based on real-time feedback of performance. The short-term dynamics in response to training became engraved in the network, requiring fewer training stimuli later in time to achieve the same results. After 2 hours of training, plasticity was significantly greater than baseline for 80 min ($p < 0.05$). Interestingly, a given sequence of stimuli did not induce plasticity, let alone desired activity, when replayed to the network and no longer contingent on feedback. Our results encourage an *in vivo* investigation of how targeted electrical stimulation of the brain, contingent on the activity of the body or even of the brain itself, could treat neurological disorders.

Chapters 3 and 4 demonstrated that plasticity could be induced by electrical stimulation. In this chapter, we explored whether we could not only induce plasticity, but customize it to achieve goal-directed behavior in an animat. The center of activity (CA) from Chapter 3 and versions of the patterned stimulation from Chapter 4 (patterned training stimuli; PTS) were used to map neural activity into animat movement and behavioral feedback into neural stimulation, respectively. Discussion of our results are framed within potential applications to neuro-rehabilitation.

⁴ Manuscript to be submitted by: Bakkum DJ & Chao ZC, Potter SM (the first 2 authors contributed equally)

Related work includes:

Bakkum DJ & Chao ZC, Potter SM, "Adaptive goal-directed behavior in embodied cultured networks: living neuronal networks and a simulated model". IEEE EMBS Conf. On Neural Engineering, Hawaii, 2007. * first 2 authors contributed equally

5.1 Introduction

A life's experiences spur the brain to continuously rewire itself to best achieve behavioral goals. However, errors can occur when injury or a pathological condition causes aberrant neural activity, and often a disconnection arises between the activity of the brain and that of the body. Treating movement disorders using physical therapy has been shown to modify neural activity, and a range of studies have shown electrically stimulated neuronal tissue exhibits neuronal plasticity (see below). Thus theoretically, electrically-induced plasticity could allow the brain to be rewired to achieve a more desired behavioral state. Here in a preliminary experiment, we investigated how a neocortical network could learn to modulate its dynamics and achieve user-defined activity states through feedback training with electrical stimuli. Besides elucidating potential therapeutic roles for artificial stimulation of the brain, these experiments give insight into how the processes underlying learning and memory are expressed in and induced by network activity.

Electrical stimulation has been extensively used to artificially induce neuronal plasticity and study learning and memory. For example, cellular plasticity has been observed in a variety of functions, including in synaptic efficacy (Bliss and Lomo 1973; Bi and Poo 1998a; Bi and Poo 1998b), intrinsic neuronal excitability (Daoudal and Debanne 2003; Zhang and Linden 2003), neuronal (Uesaka, Hayano et al. 2007) and glial (Fields 2005; Ishibashi, Dakin et al. 2006) morphology, action potential propagation (Chapter 4), and neurogenesis (Kempermann 2002). A subsequent progression in the field is determining how cellular plasticity scales and integrates to influence neuronal network dynamics. In primate motor cortex, a neuron was repetitively stimulated 5 ms after the occurrence of an action potential on a different neuron using an electronics implant (Jackson, Mavoori et al. 2006); after halting the stimulation, subsequent activity of the recorded neuron caused an increase in the firing rates in the vicinity of the stimulated neuron. This "pathway-specific" plasticity (Jimbo, Tateno et al. 1999) and also

a “region-specific” variation in the flow of neuronal activity (Chao, Bakkum et al. 2007) have been observed in our cultured networks. Thus electrical stimulation can sculpt the flow of neuronal information through a variety of mechanisms, and holds promise for either retraining or bypassing malfunctioning neuronal circuits.

Therapeutic techniques used to treat neurological disorders create links between neuronal activity and behavior via the use of directed attention or perceptible feedback signals and a lot of practice. Physical therapy, including treadmill training and robotic-assisted or neuromuscular (functional) electrical stimulation (NMES / FES) -induced limb movement, has been used to combine natural motion with proprioceptive feedback (Dobkin 2004; Sheffler and Chae 2007). The feedback activates motor circuits, improving the control of paretic limbs. Further incorporating visual and aural cues, such as targets for foot steps or beats to maintain walking gaits, and also mental imagery of movement have been used to improve training by priming motor circuits in a top-down (cortical origin) manner (Morris 2000). Biofeedback therapy uses visual, aural, or tactile displays to improve performance by making underlying physiological or cognitive processes perceptible (Huang, Wolf et al. 2006). These examples all re-link neuronal activity to a movement, and moreover, physical therapies have led to cortical plasticity: improved hand movements after constraint-induced movement therapy was accompanied by an increased representational area in the motor cortex, observed by transcranial magnetic stimulation (Liepert, Bauder et al. 2000) and fMRI (Johansen-Berg, Dawes et al. 2002). However, while helpful, (1) benefits of these techniques are not guaranteed among different individuals or different disorders, (2) optimal therapeutic protocols have not been established, and (3) the relationship with neuronal plasticity is unclear (Dobkin 2004) (Huang, Wolf et al. 2006). Adaptive electrical stimulation may improve therapeutic results by directly treating the abnormal neuronal circuits and related pathways themselves.

Additionally, but without re-linking neural activity and behavior, electrical

stimulation inside the brain has successfully managed pathological symptoms. Deep brain stimulation (DBS) has been used to treat severe cases of essential tremor, dystonia, Parkinson's disease, Tourette syndrome, clinical depression, and epilepsy (Perlmutter and Mink 2006). However, its functional mechanisms are debated. It's not clear that plasticity plays a key role in DBS therapy: effects depend on continual stimulation, which causes serious side effects including attentional and learning impairments (Jahanshahi, Ardouin et al. 2000). Other methods to reduce epileptic seizures have been used, including repetitive vagus nerve stimulation (Schachter and Saper 1998) or electrical pulses at or near seizure foci applied prior to a predicted seizure onset (Martinerie, Adam et al. 1998), but their consequences on neuronal plasticity are not known. Therefore, designing an adaptive algorithm to select appropriate training stimuli, contingent on neuronal output, could optimize the effect of a treatment while also avoiding extraneous side effects from excessive stimuli.

We hypothesize that electrical training stimuli contingent on neuronal or motor output could provide a therapy by taking advantage of existing neural plasticity mechanisms to re-link the body with the brain. Many steps are required to reach this ambition, including quantifying the ability and limitations of electrical stimulation to induce functional or adaptive changes in neural activity. Here, we developed an adaptive training algorithm which reshaped the activity of a neocortical network into different desired motor outputs within tens of minutes through the application of patterned training stimulation (PTS) using an extracellular multi-electrode array (MEA). A priori knowledge of connectivity was not necessary. Instead, effective sequences of PTSs were continuously discovered and refined based on real-time performance. The short-term dynamics in response to PTSs became engraved in the network, requiring fewer PTS applications later in time to achieve the same results. Reducing the amount of stimulation would be expected to reduce the incidence of side effects if applied in patients, and with enough training, allow removal of stimulation hardware. Interestingly, a given training

sequence did not induce plasticity, let alone desired motor output, when it was replayed to the network and no longer contingent on neural activity. Results from our controlled in vitro model encourage an in vivo investigation of how targeted electrical stimulation of the brain, contingent on the activity of the body, could treat aberrant neural activity.

5.2 Methods:

Cell culturing and electrophysiology was done as reported in Appendix G. Cultures matured for 3 weeks prior to experimentation.

5.2.1 Closed-loop training algorithm

To train a cortical network to achieve a desired motor output, a feedback loop from neural activity to electrical stimulation needs to be created (Fig. 5.1). First, sequences of neuronal action potentials were transformed into movements. Understanding how such sequences encode movement and information in general is a subject of much scientific inquiry. Population coding is a candidate motor mapping found to occur in the motor cortex (Georgopoulos 1994), premotor cortex (Caminiti, Johnson et al. 1990), hippocampus (Wilson and McNaughton 1993), and other cortical areas: the firing rates of a group of broadly tuned neurons taken together provide an accurately




Figure 0.1 Schematic of the closed-loop feedback and adaptive training.

See Methods for details. **A.** **1** A single probe electrode was repetitively stimulated every 6 sec. After each stimulus, 100 ms of evoked responses was recorded to form the 2D center of activity (CA) vector. **2** The CA was transformed ($\hat{T}$) into incremental movement $\{dX, dY\}$. **3** If movement was within $\pm 30^\circ$ of the user-defined desired direction, a shuffled background stimulation (SBS) was delivered. Otherwise, a set of patterned training stimulation (PTS) was delivered. **4** Context-control probing sequences (CPS) were delivered after SBS or PTS and before each probe to reduce variability in the probe response from ongoing activity. **B.** For unsuccessful movement, a PTS was selected from a pool of 100 possibilities. The probability of each PTS (P_t) being chosen later (P_{t+1}) increased (blue) or decreased (red) depending on the success of the motor output. See Equations 5.2 and 5.3.




Figure 0.2 Stimulation sequence.

The various stimuli during a loop through the closed-loop protocol described in Fig. 5.1 is given, scaled in time. Approximately 14 PTS or SBS are applied within 100s of ms of movement.

tuned representation (e.g., to a preferred direction of arm movement). We used a related population coding to instruct motor output, termed the center of activity (CA). For training, plasticity was induced by repetitive stimulation of a set of electrodes, termed patterned training stimulation (PTS). If the correct movement occurred, a shuffled background stimulation (SBS) was used instead, which balanced overall stimulation rates. The sequence of stimulation events is given in Figure 5.2.

5.2.1.1 Motor mapping.

Action potentials evoked by repetitive stimulation (1/6 Hz) of a single electrode, termed probe, were recorded for 100 ms following each pulse on the remaining electrodes of an MEA. The same probe electrode was used throughout an experiment. The activity of cultured networks is characterized by global bursts lasting 100s of milliseconds interspersed by periods of quiescence (Wagenaar, Pine et al. 2006a), and a probe was needed to reliably and consistently sample network dynamics. A population code mapped the activity into incremental motor output $\{dX, dY\}$:

$$\begin{bmatrix} dX \\ dY \end{bmatrix} = \hat{T} * \vec{CA} = \hat{T} * \frac{\sum_e (N_e \cdot \vec{W}_e)}{\sum_e N_e} \quad (5.1)$$

Analogous to center of mass, the center of activity ($\vec{CA}$) is the vector summation of the number of action potentials recorded on each electrode e (N_e) weighted by the spatial location of the electrode ($\vec{W}_e$) (Fig. 5.1A ❶). The inclusion of spatial information was found more reliable for quantifying network functional plasticity than using firing rates alone (Chao, Bakkum et al. 2007). The CA was normalized by a fixed transformation matrix, $\hat{T}$, to remove any directional bias arising from different distributions of neurons in different MEAs. (Fig. 5.1A ❷).

5.2.1.2 Training algorithm.

For unsuccessful movement, plasticity of the probe response is desired. Paired stimulation of monosynaptically connected neurons evokes spike-timing dependent synaptic plasticity (STDP) dependent on the stimulation interval (Bi and Poo 1998a). Repetitive PTS using the extracellular electrodes of an MEA has the potential to induce STDP throughout any shared activation pathways. When movement was within $\pm 30^\circ$ of the desired direction, SBS was used instead of PTS (Fig. 5.1A**3**). Neurons at different electrodes can be connected through multiple neurons and pathways; some PTSs may give a desired neuronal plasticity while others may give the opposite or none. Therefore we compiled a pool of 100 possible PTSs, each with a different spatial stimulation sequence (see details below). Initially, each PTS had an equal probability of being chosen. If the current PTS (PTS_k) improved the performance, then the probability of PTS_k ($P_k(t)$) being chosen later ($P_k(t+1)$) increased, and the probability of other PTSs ($P_i(t+1)$) being chosen decreased (Figure 5.1B):

$$\begin{cases} P_k(t+1) = \frac{1.5 \cdot P_k(t)}{1 + 0.5 \cdot P_k(t)}, & k = \text{current } PTS \\ P_i(t+1) = \frac{P_i(t)}{1 + 0.5 \cdot P_k(t)}, & \text{for } i \neq k \end{cases} \quad (5.2)$$

Otherwise, if PTS_k worsened the performance, then $P_k(t+1)$ decreased from $P_k(t)$ and $P_i(t+1)$ for other PTSs increased:

$$\begin{cases} P_k(t+1) = \frac{0.5 \cdot P_k(t)}{1 - 0.5 \cdot P_k(t)}, & k = \text{current } PTS \\ P_i(t+1) = \frac{P_i(t)}{1 - 0.5 \cdot P_k(t)}, & \text{for } i \neq k \end{cases} \quad (5.3)$$

In this manner, neuronal plasticity could be directed. A maximum of 50% was set so that a high probability for one PTS would not saturate the pool of possible choices. A minimum of 0.2% was set to ensure each PTS remained available in the future. This allowed the flexibility to adapt to ongoing changes in neuronal network dynamics.

5.2.2 Experiment design

5.2.2.1 Determining PTS, SBS, CPS, and probe:

Before every experiment, biphasic voltage stimuli (Wagenaar, Pine et al. 2004) of ± 300 mV, 400 μ s per phase were delivered randomly at the 59 usable electrodes on our MEAs for 30 mins, one at a time, with random inter-stimulus-intervals between 200~400 ms. An electrode that evokes synaptic activity at multiple electrodes is capable of revealing changes in network-level synaptic connectivity, and one was selected as the probe electrode. Six electrodes that evoked only short responses (< 20 ms latency) were selected to create a context-control probing sequence (CPS). They were stimulated in order, prior to the probe stimuli, with inter-stimulus intervals between 200~400 ms in order to provide a consistent “context”. Once the spatiotemporal structure of the CPS and probe was determined, it was fixed throughout the experiment and used to sample the network state every 6 seconds to collect movement data.

The remaining electrodes that evoked activity were used as candidates for constructing PTS and SBS. The probe and the context electrodes were not used in PTS nor SBS. Each PTS consisted of 6 stimuli randomly selected from the candidate electrodes (repeated electrodes allowed). The inter-stimulus-interval between two consecutive stimuli was fixed at 10 ms. When training was required, a PTS was repetitively delivered until the beginning of the next CPS, with an inter-PTS-interval between 200-400 ms. When training was not required, SBS consisted of repetitively

delivering one randomly selected PTS with inter-SBS-intervals again between 200-400 ms, but with randomized electrode order each time delivered. The randomized order in SBS maintained comparable overall rates and distribution of stimuli as in PTS, while removing the repeating spatiotemporal pattern of neural activation. The probe-evoked response's overall firing rates, and in turn their distribution of CAs, depended on the amount and spatial distribution of stimulation (data not shown). Therefore, any changes in motor output would stem from changes in network connectivity. Stimulation voltages of ± 300 mV were lower than those in our previous investigations of plasticity magnitude (Chao, Bakkum et al. 2007) in order to better localize evoked neurons and PTS-induced plasticity. This and only a few seconds of training stimulation at a time were intended to induce plasticity incrementally with enough resolution to reach desired network states and minimize overshoot.

5.2.2.2 SBS-only stimulation:

At the beginning of an experiment and prior to each closed-loop training, CPS and probe stimuli were delivered with SBS interspersed for 6 hours. This allowed the network to habituate to the presence of electrical stimulation.

5.2.2.3 Closed-loop experiments with different desired movement directions:

One closed-loop experiment consisted of four 2-hr training periods, each with a different desired direction, and 2-hr SBS-only periods in between. After a training period, the duration of plasticity was measured during the 2-hr SBS-only period. The transformation matrix for each training period, $\hat{T}$, was calculated during the last 30 min in the preceding SBS-only period. Six closed-loop experiments were performed on 5 different cultures from 3 dissociations such that 23 training periods were analyzed (6*4 minus 1 where a technical error caused a loss of data). Two experiments were performed on one culture, with 13 days in between, with different CPS, and PTS set.

5.2.2.4 Open-loop stimulation experiments:

To test if the improvement in performance was an artifact of the electronics or electrode chemistry arising from a particular stimulation sequence, the stimulation sequence recorded from each closed-loop experiment was replayed to the same network about a day later. Since the same transformation might not successfully offset the CAs if activity changed, the transformations were recalculated as before. However, any similarity in the directional plasticity between the closed-loop and replayed open-loop stimuli would be comparable in the analysis in Figures 5.3 and 5.5 because the scaling transformations between closed-loop and open-loop trials were found not significantly different (P-value= 0.34 Wilcoxon sign rank test, N= 23 trials * 2 directions; transformation offsets do not affect movement direction).

5.3 Results and Discussion

5.3.1 Training shifted neural activity towards the desired activity

We designed a closed-loop algorithm to train cultured networks to learn user-defined motor outputs: moving $\pm 30^\circ$ within a pre-defined direction. The closed-loop training algorithm was tested in 5 cortical networks grown over MEAs. Action potentials evoked by repeating probe stimuli, at a single fixed electrode, commanded motor output, and electrical stimuli were fed back as training signals. The probability of selecting a PTS was updated based on how its application influenced the network's short-term activity dynamics during the following motor output (see Method). Four goals were applied sequentially to a network, that is, the desired movement direction was changed $\pm 90^\circ$ or 180° 3 times. Overall success was judged by the ability of the network to crystallize successful short-term changes into long-term plasticity and also by its ability to adapt to new desired motor outputs. We found that with training, motor output was able to head toward the predefined desired directions (one representative experiment is shown in

Figure 5.3A). The learning curves show that a greater proportion of movements were in the desired direction as training progressed (Figure 5.3C). Since a correct movement meant applying SBS instead of PTS, fewer training stimuli were needed in time, suggesting the network was learning the appropriate input/output function to allow successful behavior.

Learning curves for all experiments are shown in Figure 5.4. A random movement would give a 16.67% chance (horizontal line) of movement within $\pm 30^\circ$ of the desired direction ($60^\circ/360^\circ$). In 5 out of 23 trials (21.7%), the learning curves were below the 16.67% chance in the last 10 minutes of training (black arrows in Figure 5.4). This suggests that a more optimal training algorithm may exist. For example, using a different set of possible PTSs might improve success rates by inducing a different plasticity. The average normalized learning curve of closed loop experiments showed the success rate increased by a factor of 2.88 ± 0.08 (Mean $\pm$ s.e.m.) times after 2 hours of training for each desired direction ($n=23$ trials, from 6 experiments) (Figure 5.5).

The stimulation sequence delivered during a closed-loop experiment was recorded and replayed to the same network the next day. This ruled out artifactual changes in network responses due to non-biological causes, such as electrochemistry or electronic noise, and ensured that neuronal plasticity was responsible for the observed learning. With open-loop training, motor output was unable to move toward the desired directions (the motor outputs and the learning curves of the corresponding open-loop experiment of the closed-loop experiment shown in Figures 5.3A and 5.3C are shown in Figure 5.3B and 5.3D, respectively). The average learning curve of open-loop experiments was significantly lower than in closed-loop experiments ($p\text{-value}=9.9\text{e-}54$, $n=23$ trials, Wilcoxon signed rank test) (Figure 5.5). Changes in movement direction were distributed




Figure 0.3 Neural response to closed-loop and open-loop training.

A. Closed-loop training: Movement trajectory (left column) and the change in the probability distribution of movement directions (right column) demonstrated the motor output adapted to the desired direction (black arcs). Desired directions of -135° , -45° , 135° , and 45° were applied in random order for 2-hr periods (light to dark blue in time) interspersed by 2-hr SBS-only periods (see Section 5.2). Successful motor output was considered to be movement within $\pm 30^\circ$ of the desired direction. The smaller trajectory circle is a zoom-in of the beginning of the experiment and the 30-min SBS-only period (gray) used to calculate the transformation, $\hat{T}$. The probability distribution of movement directions during 10 minutes at the start of experiments was subtracted from that during the final 10 minutes, thus allowing negative values (red). **B. Open-loop training:** The closed-loop stimulation sequence was recorded and replayed to the same network. Movement trajectories (scaled to match the corresponding closed-loop experiment) changed but not necessarily towards the desired direction. The distribution of movement directions also changed but in a more distributed manner. Learning curves of **(C)** closed-loop and **(D)** open-loop examples shown in A and B: A learning curve was defined as the probability of movement in the desired direction within a 10-min moving time window (time step= 1 min). The probability of successful motor output increased in time when training was contingent on motor output. A random movement would give a 16.67% chance (horizontal line) of movement within $\pm 30^\circ$ of the desired direction ($60^\circ/360^\circ$).




Figure 0.4 Learning curves for all closed-loop experiments

For each experiment, the learning curves and the time courses of different PTS probabilities are shown. See more about the PTS probability in Section 5.3.3 and Figure 5.8B. In 5 out of 23 trials (21.7%), the learning curves were below the 16.67% chance in the last 10 minutes of training (black arrows).




Figure 0.5 Average normalized learning curves of all closed-loop and open-loop experiments.

To compare trends among different experiments, each learning curve was normalized by dividing by the probability of successful motor output when training began. The average normalized learning curve in the last 10 minutes of closed loop experiments was 2.88 ± 0.08 (SEM) times higher than at the start, which was significantly higher than 1.24 ± 0.03 for open-loop experiments ($p\text{-value} = 9.9\text{e-}54$, $n = 23$ trials from 6 experiments, Wilcoxon signed rank test). An exponential curve fit gives a time constant of 10.6 minutes and a learning curve asymptote of 3.13 (SSE = 2.95, R-square = 0.7814).

across a wider range of angles than with closed-loop training (compare Figures 5.3B and 5.3A). Therefore, we concluded that the successful learning reflected biological plasticity in the neuronal networks, and required closed-loop training in which stimuli were contingent on behavior.

5.3.2 Changes in motor output arose from neuronal plasticity, not an elastic dependency on stimulation history

The improved performance could be due to plastic changes in the neuronal network or, alternatively, due to an elastic dependency on the recent stimulation history. An elastic change in the neurons' responsiveness to stimuli was observed in dissociated cortical cultures, where the sensitivity of neurons selectively adapted to stimulation with different frequencies, and this change in the sensitivity faded away within several minutes after stimulation was removed (Eytan, Brenner et al. 2003a). In order to verify this, we followed the motor output after switching closed-loop training back to the SBS-

only stimulation, and quantified whether the learned movement was maintained, and if so, for how long.

By sampling the distribution of movement angles every 10 minutes from 1 hour before to 2 hours after the closed-loop experiment, we found that: (1) the movement angle gradually converged to the desired directions during closed-loop training, and (2) the learned directions were maintained after training during SBS-only stimulation. Results from the experiment used in Figure 5.3 are shown in Figure 5.6A. The distribution of the 10-min SBS-only period immediately before (*Pre*) closed-loop training was significantly different than that immediately after (*Post*) closed-loop training (histograms on the right side of Figure 5.6A). This training-induced plasticity led to desired motor outputs in 18 out of 23 trials (78.3%) (see Figure 5.4). This demonstrated that closed-loop training successfully directed network plasticity, that directed plasticity had occurred (Figure 5.6C). Moreover, the distributions were not significantly different between the last 10 minutes of the training and the next 10 minutes of the SBS-only control (Figure 5.6D). This indicated that motor outputs were preserved in the SBS-only period after training was turned off, and further demonstrated that the improved performance was due to network plasticity, not non-biological causes (such as stimulation artifacts of PTS).

For replayed open-loop stimuli, the distribution of movements was less focused (Figure 5.6B) and no significant changes in motor output occurred for most of the open-loop experiments (Figure 5.6C). Desired movement directions were found in 4 out of 23 open-loop trials (17.4%) during the last 10 minutes (see -45° trial in Figure 5.6B), which was close to the 16.67% chance but significantly lower than the 78.3% success rate for closed-loop experiments. Interestingly, the PTSs in open-loop training by themselves could not cause noticeable plasticity, but when contingent on neural activity, the set of PTSs were able to incrementally shift network dynamics until a significant functional change was detectable.




Figure 6.6 Long-term plasticity of movement direction.

A. Movement directions became concentrated within $\pm 30^\circ$ of the desired direction (red numbers and horizontal lines) during closed-loop training (CL) and persisted into the SBS-only periods. Data is from the same representative experiment as in Figures 5.3. The distribution of movement angles was sampled every 10 minutes from 1 hour before to 2 hours after the closed-loop experiment (gray scale). The distributions of the occurrence of different movement angles during 10-min SBS-only periods immediately before (*Pre*) and immediately after (*Post*) closed-loop training are shown in the histograms (right). **B.** Changes in movement direction were not observed in the corresponding open-loop experiment (OL). **C.** The distribution of movement angles in *Pre* periods was significantly different than that in *Post* periods for closed-loop training, suggesting directional plasticity occurred. This did not occur for replayed open-loop stimulation suggesting plasticity was not a stimulation artifact. P-values of the difference in movement angle distributions for 23 desired directions (two-sample Kolmogorov–Smirnov test, two-tailed) are represented in box plots showing the first (lower) quartile, the median, and the third (upper) quartile. Outliers are indicated as black dots, and the largest and smallest non-outlier observations are indicated as tic marks (whiskers). The median p-value for closed-loop experiments was below a significance level of 0.05 (0.016 median). **D.** The distribution of movement angles between the last 10 minutes in closed-loop (or open-loop trials; CL, OL) and that during the next 10 minutes of SBS-only (*Post*) were not significantly different, demonstrating the directed plasticity was not an artifact of PTS history.




Figure 0.7 Plasticity induced by closed-loop training lasted for 80 minutes.

Plasticity induced by closed-loop training was significantly greater than intrinsic plasticity before training, for 80 minutes on average for all experiments, but not for replayed open-loop stimulation. **A.** The change-to-drift ratio (C/D, see the definition in Section 3.2.3) between CAs in different 10-min periods after training (*Post*, see Figure 5.6) moving with 1-min time steps and CAs in a 10-min reference period immediately before training (*Pre*) was calculated. The mean and SEM of C/D across closed-loop and open-loop periods ($n = 23$ trials) was compared to C/D across SBS-only periods during the first 6 hours of SBS-only stimulation (Before training, $n = 12$ closed-loop and open-loop experiments). **B.** Closed-loop training-induced plasticity was significantly greater than intrinsic plasticity measured before training for 80 minutes. Time course of p-values were calculated by comparing C/D of closed-loop and open-loop experiments to C/D before training (shown in A). Plasticity in closed-loop experiments was significantly greater than intrinsic plasticity for 80 minutes (Wilcoxon rank sum test, $\alpha = 0.01$).

Although the training algorithm increased the probability of occurrence of a PTS based on the success of short-term “elastic” responses after the next probe stimuli, the elastic responses became consolidated as long-term plastic changes in time. This is demonstrated by the stability of the distribution of movement directions into the SBS-only period following training (Figure 5.6), by the learning curves (Figure 5.5), and by

progressively fewer PTSs being needed to maintain desired movement (Figure 5.8C). Initial changes in the probe responses were indeed short-term elastic responses because they were not maintained initially, requiring reapplication of PTSs.

In order to investigate the duration of training-induced plasticity, the change-to-drift ratio (C/D, see Equations G.1 and G.2) between CAs in different 10-min periods after training (*Post*, see Figure 5.6) moving with 1-min time steps and CAs in a 10-min reference period immediately before training (*Pre*) was calculated. The mean and SEM of C/D across closed-loop and open-loop periods (n= 23 trials) was compared to C/D across SBS-only periods during the first 6 hours of SBS-only stimulation (Before training, n= 12 closed-loop and open-loop experiments) (Figure 5.7A). Closed-loop training-induced plasticity was significantly greater than intrinsic plasticity (drift) measured before training for 80 minutes (Wilcoxon rank sum test, $\alpha = 0.01$) (Figure 5.7B). Replayed open-loop stimulation did not induce significant plasticity. The decrease in C/D about 70 minutes after closed-loop training indicates that CAs returned back to the distribution before training, and so did movement directions. The return could be due to an active “re-habitation” to the SBS-only stimulation. SBS, while less structured than PTS, still has a spatiotemporal structure and the ability to induce plasticity.

5.3.3 Training required different PBS at different times

A neural network is continuously plastic, being modified by both stimulus-evoked and spontaneous activity. The same PTS may therefore have different effects at different points in time, and successful adaptation to a desired motor output would require application of PTSs in a certain sequence. This is what we observed (Figure 5.8B), in agreement with our preliminary results from Appendix A (Figure A.6). Additionally, logic dictates that a finite amount of plasticity is needed to achieve the desired motor output. A given PTS could induce appropriate plasticity initially, but continued application of the PTS could be maladaptive.




Figure 0.8 Training required different PTS at different times.

A. A normalized version of the learning curve shown in Figure 5.3C. **B.** PTS probability in time for different desired directions, color coded to match A. Various series of PTS were needed to induce appropriate neural plasticity and successful motor output. For clarity, the PTS with lower probabilities were not plotted. Electrode locations and order (right) for the PTS indicated by red arrows are shown in 8 by 8 grids of electrode locations. **C.** Across the 23 trials, fewer PTS were needed in time to maintain successful motor output. The frequency of occurrence of a PTS was measured by a 10-min moving time bin with 1-min time step and normalized by its maximum value (set to 1). The initial rise occurred while appropriate PTSs were searched. PTS occurrence monotonically decreased over the last hour (Spearman correlation of means (dots), one-tailed, $\rho = -0.89$, $p\text{-value} = 1.2e-31$, $n = 231$ successful PTSs in 23 trials). A successful PTS was defined as one that improved performance at least one time. A cubic polynomial (red) was fit to the data to better visualize the trend.

We also found that fewer PTSs were needed across a training period to maintain successful motor output (Figure 5.8C). The trend of the PTS-delivering frequency was measured by counting successful PTSs in a 10-min moving time bin with 1-min time step and normalized by its maximum value. We defined a successful PTS as one that improved performance at least one time, and found that the occurrence of successful PTSs monotonically decreased over the last hour (Spearman correlation of means (dots), one-tailed, $\rho = -0.89$, $p\text{-value} = 1.2e-31$, $n = 231$ successful PTSs in 23 trials). This suggests that the training stimuli which were successful in the first hour were less often or no longer required during the last hour to maintain a high rate of correct motor outputs (see average closed-loop learning curve in Figure 5.5).

5.4 Conclusion

Following a footpath over a mountain range is quicker and more energy efficient than digging straight through middle. We hypothesize that directing plasticity using training stimuli contingent on the motor output is more efficient than blindly forcing plasticity, for example, via a large tetanic stimulation or open-loop DBS. Moreover, since neuronal activity is continuously plastic, each electrical stimulus and ongoing spontaneous activity alter the landscape, and routes cannot be plotted in advance. Learning is an ongoing and continuous process. Our training algorithm allowed the probabilities of the PTS pool to change, and “solutions” to achieve desired motor outputs were explored in real-time.

The learning curves increased (Figure 5.5), but success did not approach 100%, and some trials showed no learning (Figure 5.4). A more optimal training algorithm may exist, although learning may have continued if training was not stopped after 2 hours. Using a larger set of possible PTSs could improve success rates by inducing a greater range of plasticity; the tradeoff is potentially longer training to find an appropriate sequence of PTSs. Different spatiotemporally structured PTSs, e.g. using more

stimulation electrodes, could also produce different performance. Furthermore, the 0.5 maximum criterion on the probability of selecting a PTS allows randomly-applied inappropriate PTS to produce setbacks. In addition, the 0.002 minimum may keep unhelpful PTSs in the loop. Different optimization rules, such as updating the PTS pool with a genetic algorithm, or introducing a PID controller to govern the duration of training could improve performance. Future work includes determining the abilities and limitations of electrical stimuli to induce neuronal plasticity, optimizing training parameters, and applying closed-loop algorithms to achieve multiple simultaneous desired motor outputs.

For neurological disorders, targeted electrical stimulation of the brain contingent on the activity of the body or even of the brain itself could direct neuronal plasticity to bypass or accommodate aberrant neural activity. Initial candidate pathologies include those with (1) a focal neural source or related pathway at which to insert an MEA for electrical training and (2) a measurable physical manifestation from which to gather feedback on performance. As an example, to treat movement disorders, such as after stroke or with Parkinson's disease, electrical modification of the basal ganglia could be guided by physical measurements of changes in tremor, ataxia, rigidity, dystonia, etc. using electromyography. Directly measuring motor output negates the need for context and probe stimuli and allows training to be a continuous process. Ongoing afferent input from different brain areas would be expected to negate the need for SBS. Re-linking the body and the brain with electrical training stimuli would give existing brain mechanisms the potential to overcome neurological disorders.

CHAPTER 6

FUTURE STUDIES

My ultimate goal was to develop animats that could learn something about the environment and/or body given to them. Upon entering the lab, the technology to culture neurons for long durations on MEAs, robustly record neural activity, and stimulate an MEA's electrodes was recently achieved. However, the crucial ability to induce and detect neural plasticity was missing: methods were needed to determine appropriate sensory-motor mappings and training algorithms in order to produce any kind of adaptive behavior. Identifying what types of stimulation could induce plasticity and what kinds of activity statistics could identify plasticity was the focus of the open-loop experiments presented in Chapters 3 and 4. The results from these experiments were next used to develop animats that achieved adaptive goal-directed behavior in Chapter 5 and Appendix A. Now, network plasticity could be not just induced, but customized, suggesting that animats may yet be able to be conceived to behave in interesting and intelligent ways.

This chapter concludes with future directions, written as a guide to students who wish to build upon the work begun here. These include using light as an ideal stimulation interface with high spatiotemporal resolution in order to create better sensation-to-stimulation mappings for animats; finding new horizons for embodiment experiments; and, following the tangent in Chapter 4, investigating the molecular mechanisms and functional rules of changes in action potential propagation.

6.1 Programmable array microscopy (PAM) to evoke activity in channelrhodopsin-2 (ChR2) expressing neurons⁵




Figure 6.0.1 Schematic of programmable array microscopy to evoke channelrhodopsin-2 expressing neurons.

Cortical cells are important for learning and memory, and have been studied extensively at the cell and molecular level. On the other hand, little is known about how the activity of individual cells combine to produce behavior, learning, and memory. Our embodied cultures address this issue by allowing detailed observation of neuron population dynamics while the networks are expressing behaviors. However, our experiments, and those in neuroscience in general, are limited by the technology available to investigate brain cells. In particular, conventional electrodes can record and stimulate neurons only at predetermined locations and cannot access neurons apart from the electrodes. Subtleties in neuron information processing cannot be observed. An ideal

⁵ Submitted to the Japanese Society for the Promotion of Science (JSPS) postdoctoral fellowship for international researchers.

interface would access such subtleties by recording and stimulating each neuron individually.

The use of light could provide an ideal stimulator (Fig. 6.1). Additional motivation comes from our previous experiment in Chapter 5, where increasing the spatial resolution of stimuli (via lower voltage magnitudes) produced better learning results. This allowed electrodes to evoke finer tuned groups of neurons, and thus activity could be incrementally modified until the user-defined goal was reached. Further improving stimulation resolution is expected to improve learning performance (like treating a patient with acupuncture instead of a hammer) and allow more complex goals to be reached when integrated with MEA recording.

Table 1 compares conventional interfaces with the proposed interface and a light-addressable electrode interface developed by Takahashi's lab (Suzurikawa, Nakao et al. 2007) - red indicates non-ideal properties. As the table shows, the proposed interface is the most attractive: an ideal stimulator and the optimal recorder. (An ideal recording interface does not yet exist, but high spatial resolution was found not as important as temporal resolution in analyzing neuron network activity; interestingly, lower spatial resolution of recordings may better improve analysis (Chao, Bakkum et al. 2007).) A detailed discussion follows. Takahashi's lab combined a promising light-addressable electrode to stimulate neurons with special imaging dyes to record (Suzurikawa, Nakao et al. 2007). However, this technique sacrificed temporal resolution and experiment duration. Maintaining high temporal resolution of recordings is crucial because lower resolution measurements (firing rates recorded over 100s of milliseconds) were found to not adequately depict network plasticity (Chao, Bakkum et al. 2007). Moreover the calcium sensitive dyes can interfere with neuron learning

Table 6.1

Neuron interface comparison		Conventional	Takahashi Lab	Proposed
		MEA to stimulate and record	Light addressed electrodes + calcium imaging to record	PAM and ChR2 optical stimulation + MEA to record
Stimulation resolution	temporal	< 1 ms	< 1 ms	< 1 ms
	spatial	60 electrodes	~ 20 μm	~ 20 μm
Recording resolution	temporal	25 kHz	~ 10 Hz	25 kHz
	spatial	60 electrodes	~ 20 μm	60 electrodes
Stimulation artifact		yes	no	No
Max experiment duration		months	2 hours	months

The proposed method consists of optically evoking neurons expressing channelrhodopsin-2 (ChR2), a light activated ion channel found in algae⁶. ChR2 can be inserted into the DNA of mammalian neurons using lentiviral vectors. High intensity blue light activates the channel and evokes neuron activity. Programmable array microscopes (PAM) use an array of mirrors (up to 2 million; Texas Instrument's DLP originally developed for televisions) to pass a light beam onto addressable locations. PAM could evoke ChR2 expressing neurons in (1) any instantaneous pattern with (2) a spatial resolution of a neuron cell body and (3) temporal resolution of an action potential, (4) continuously, and (5) without electrical artifacts blanking recordings. In conventional MEAs, naturalistic stimulation cannot be applied while recording since stimulation artifacts would completely blank most recordings; using light to stimulate thus also increases the fidelity of electrode recordings. ChR2 expression persists for months allowing long-term experiments without photobleaching or phototoxicity. Both PAM and ChR2 technologies are relatively new and have not been used together.

Both the proposed PAM hardware and light-evoked mammalian brain cells expressing ChR2 have been used independently in other labs (Hanley, Verveer et al. 1999; Boyden, Zhang et al. 2005). Unknown is how networks of neurons will respond to

⁶ channelrhodopsin.org

patterned optical stimulation. Neuron activity causing plasticity is accepted as fact, and therefore optical stimulation will very likely affect neuron dynamics. If not, then understanding the difference between optical and electrical stimulation would be useful for physiologists. The biological brain was designed to control a biological body, and behaving within an environment is considered a necessity for a neuron network to function at its fullest (Pfeifer and Bongard 2007). Our embodiment experiments appear promising, allowing detailed analysis of cell and network dynamics and already demonstrating simple behavior and learning. The proposed interface will be a powerful tool to advance these experiments and neuroscience in general by providing better experimental access to neuron networks and the brain.

6.2 Future embodiments

Chapter 2 is entitled *Removing some ‘A’ from AI*, but adding back some of the ‘I’ from *AI* may likewise benefit neuroscience and our embodiment experiments. As discussed in the *Embodiment and Intelligence* section of Appendix A and elsewhere (Varela, Thompson et al. 1993; Pfeifer and Bongard 2007)], the fields AI and cognitive science began by assuming the brain manipulated concepts in order to produce cognition, a stance referred to as Good Old Fashioned AI (GOFAI). However, now becoming more accepted by scientists is the hypothesis that intelligence is not disembodied, but intimately entwined with the mechanics of the body and an interaction with the environment.

GOFAI has parallels in present-day neuroscience where many studies focus on “representation”, in particular the correlation of brain areas to mental processes as is done with many fMRI studies. However, correlation does not imply causation. In what Terry Sejnowski described as the “knee-jerk method of neuroscience” at a plenary talk for the EMBS Neural Engineering Conference 2007 in Hawaii, sensory information gets relayed through the thalamus to a primary sensory cortex, then to secondary sensory cortex to

frontal cortex to premotor cortex to motor cortex and finally culminating as muscle movement. This is true but vastly neglects the most prominent feature of the brain: about 1% of the synapses in cortical tissue make up this feedforward relay system, while the rest comes from feedback signals from a diverse array of brain areas. “Representation” is but an abstract concept (as with GOF AI) and not very descriptive of the underlying brain mechanics, which uses the the brain’s ongoing activity as a context to continually filtering relevant information, at vast range of spatial and temporal levels.

Following the recent trends in AI may lead neuroscience into more productive territories. The roboticist Rodney Brooks pioneered the use of robotics in AI and the idea to test algorithms with such real-world embodiments. He gives an interesting focal point for experiments: study the process of “abstraction” (filtering) instead of the process of “search”. Search in his sense is the matching of sensory input into a discrete set of outputs, while abstraction is the filtering of relevant sensory input. The later is a process he considered far more relevant/important for artificial or biological autonomous agents to successfully behave, where the real world is filled with possible sensory information that is often incomplete and noisy.

While this dissertation provided the important steps that neural plasticity can occur in our preparations (Chapters 3 and 4) and can be directed with the feedback electrical stimulation contingent on behavior (Chapter 5 and Appendix A), the capacity of the neuronal networks to achieve multiple and complex behaviors is not known. Furthermore, the goal activity in Chapter 5 would be considered solving a “search” problem by Brooks. Investigating the ability of future animats to instead “filter” sensory information may lead to better insight into how neural networks and brains use environmental interaction in their computations.

As an example experiment to discriminate patterns, a closed-loop algorithm could be created to give a continuous naturalistic stimuli, possibly using the PAM and ChR2 from the previous section. Relevant and extraneous sensory information could be

constructed by adjusting the frequency of occurrence of one (then more later) stimulation pattern contingent on neuron activity recorded at a set of electrodes. The network's ability to filter information could be judged by the network's ability to identify the contingent pattern, possibly by comparing the frequency of occurrence of the contingent pattern versus that of other non-contingent control patterns. Further exploring the AI literature may lead to other productive approaches.

6.3 Explaining dAPs

Chapter 4 provided experimental evidence of plasticity of action potential propagation. Further elucidating the cellular mechanisms involved would help make a more convincing argument of its role in neural computation. Possibilities include non-uniform changes in ion channel properties (Ganguly, Kiss et al. 2000), in the geometry of varicosities and branch points (Goldstein and Rall 1974) or axonal arbors, in the proximity of glia (Ishibashi, Dakin et al. 2006), and in lipid membrane composition (Bedlack, Wei et al. 1994). Moreover, besides the gradual shifts in latency described in Chapter 4, directly-evoked action potentials (dAPs) could appear and disappear (Fig. 6.2) and also jump to new latencies over multiple milliseconds (Fig. 3.4). Do these phenomenon arise from different mechanisms? A better characterization would make propagation plasticity better accepted and also provide novel targets for currently intractable pathologies. For example with schizophrenia, the coordination between different brain areas goes haywire, which could arise from a disorder in the temporal control of propagating action potentials. For neuroscientists, propagation plasticity offers a new perspective on neuron computation, learning, and memory in the brain. Theoreticians and investigators of artificial intelligence may consider new inquiries into the computational ability of the brain and into the engineering of artificial control systems.

Determining what causes a propagation latency to increase versus decrease would help computational neuroscientists and researchers of artificial intelligence to create new algorithms and explore the computational capacity of the brain in fundamentally new ways (Milton and Mackey 2000; Izhikevich 2006). Patterning axon growth over a series of electrodes (Suzuki and Yasuda 2007) or nanowire transistor recording devices (Patolsky, Timko et al. 2006) and/or optical imaging (Kawaguchi and Fukunishi 1998) could expose the discrimination, resolution, and possible morphological correlates of propagation plasticity. Further patterning and selectively stimulating afferent input could expose an input/output function to changes in propagation. Repeating experiments with various antagonists of molecules involved in long-term plasticity could expose the underlying cellular mechanisms.




Figure 6.0.2 Raster plot of evoked responses detected on one electrode in response to repeated stimulation on another electrode.

The red bars indicate periods where the directly-evoked action potentials (dAP) at 5 ms latency mostly disappear.

APPENDIX A

MEMOIRS OF A CYBORG ARTIST⁷

Here, we and others describe an unusual neurorobotic project, a merging of art and science called MEART, the Semi-living Artist. We built a pneumatically-actuated robotic arm to create drawings, as controlled by a living network of neurons from rat cortex grown on a multi-electrode array. Such embodied cultured networks formed a real-time closed-loop system which could now behave and receive electrical stimulation as feedback on its behavior. We used MEART and simulated embodiments, or animats, to study the network mechanisms that produce adaptive, goal-directed behavior. This approach to neural interfacing will help instruct the design of other hybrid neural-robotic systems we call hybrot. The interfacing technologies and algorithms developed have potential applications in responsive deep brain stimulation systems and for motor prosthetics using sensory components. In a broader context, MEART educates the public about neuroscience, neural interfaces, and robotics. It has paved the way for critical discussions on the future of bio-art and of biotechnology.

6.1 Introduction

“The most beautiful thing we can experience is the mysterious.

It is the source of all true art and all science.” – Albert Einstein, 1931

The mind’s emergence from the complex interactions of a brain is one of the greatest mysteries. Its contemplation has founded sciences and produced countless cultural artifacts in literature, film, and the arts. Here we present our own artifact, being a mixture of biology and technology, of metaphor and physicality, and of art and science, to question anew our notions and to seek a few answers.

⁷ Manuscript to be submitted as:

Bakkum DJ, Gamblen P, Ben-Ary G, Chao ZC, Potter SM, “Memoirs of a cyborg artist”, Invited submission for inaugural issue of *Frontiers in Neurorobotics*, 2007

See also:

Bakkum DJ, Chao ZC, Gamblen P, Ben-Ary G, Shkolnik A, DeMarse TD, Potter SM, “Embodying cultured neurons with a robotic drawing machine”. *IEEE EMBS Conf. Lyon*, 2007



Figure A.1 MEART's body.

Two arms cooperated to grip a set of colored pens and move them across a sheet of paper, according to neural activity in a culture dish that was up to 12,000 miles away. A CCD camera aimed at the drawing provided sensory feedback to the neuronal network. (Picture by Steve Potter.)

We built a robotic drawing machine and designed algorithms for it to converse with a network of cortical neurons grown in culture over a multi-electrode array (MEA), a cyborg named Meart (Multi-Electrode Array aRT) (Fig. A.1). The apparatus consisted of living neurons (Fig. A.2) in a laboratory connected by internet to a pen-wielding metal and plastic arms behaving in gallery exhibitions around the world. Neuronal action potentials recorded by an MEA determined movement, and video images of drawings determined the subsequent feedback of electrical stimuli delivered to the neurons. Artists and scientists collaborated to construct Meart, a concept stemming from prior artistic expressions of and scientific inquiries into hybrid bio-robotic technology, or “hybrots”. Our common ground was to explore the essence of creativity and intelligence. With an MEA, the underlying neural network mechanisms, including the manifestations of learning and memory, could then be quantified,. In neuroscience, a large gap exists between in vivo behavioral studies of learning and memory, and in vitro studies of cellular plasticity. With Meart, behavior and learning could be observed in concert with

the detailed and long-term electrophysiology available in vitro (Potter and DeMarse 2001b). We sought whether Meart could learn something about the environment given to it, and if an act of creativity could emerge.





Figure A.2 MEART's Brain.

(**Above**) A multi-electrode array culture dish. (**Below**) A culture of ~50,000 neurons and glia from embryonic rat cortex, growing in a multi-electrode array and forming a dense network 1-2 mm across. Fifty-nine 30 μm electrodes spaced at 200 μm intervals connect a few hundred of the network's neurons to the outside world, by allowing their activity to be extracellularly recorded or evoked by electrical stimulation.

Here we present, along with artistic and scientific dialogs, progress on engineering Meart's hardware, software, and wetware. In experiments, we applied patterned training stimuli (PTS) contingent on behavioral performance in order to achieve the goal-directed behavior of drawing geometrical shapes. The transformation from visual sensation into the delivery of a PTS was fixed, and while neural plasticity occurred, successful behavior did not. Prior knowledge of network connectivity was needed to determine the appropriate transformation. However, we modified the training algorithm using a living network connected to a simulated robot (an animat). Instead of a fixed transformation, behavioral performance was used to continuously discover and refine effective sequences of PTSs, and in a preliminary experiment, an animat repeatedly learned to draw in different desired directions. Encoding more detailed sensation and motor output could lead to increasingly complex and interesting behaviors. What questions would be posed if Meart was eventually deemed to be intelligent?

A.2 Methods: making the semi-living artist

Meart was comprised of living neurons, recording and stimulating electronics, robotic drawing arms, electronic control circuits for a pneumatic actuation system, a CCD camera to feedback images of drawings, and software communicating between the neurons and robot over the internet (Fig. A.3). The simulated animat was made of living neurons, recording and stimulating hardware, and a simple virtual embodiment on a computer. It was used to develop protocols in the intervals between Meart exhibitions. Three major topics needed to be addressed to embody the cultured networks:

- A. The care and feeding of the biological brain;
- B. The hardware (or software) implementation of the body; and
- C. The sensory transformation, motor transformation and training algorithms.

A.2.1 Preparing and caring for Meart's brain

Cell culturing and electrophysiology was done as reported in Appendix G.

Cultures matured for 3 weeks prior to experimentation.




Figure A.3 Schematic of the bio-robotic software algorithms and hardware, i.e. Meart's components.

◆ **Commanding movement.** The *Center of Activity* (CA) of neuronal action potentials was calculated from 100 ms of responses after a probe stimulation (8×8 box representing the MEA; increasing firing rate is black to white). Animat movement was instructed from a transformation (T) of the CA into a *Population vector*. The X, Y movement command was sent over the internet (yellow arrows) to the robotic arms every 4 sec. ◆ **Movement.** The robotic drawing machine consisted of 2 perpendicular arms actuated by braided pneumatic artificial muscles, allowing independent retraction (R) or extension (E) of the left (E_L/R_L) and right (E_R/R_R) arms within approximately a 30 cm by 30 cm workspace. Similarly, smaller muscles pressed the pens to the paper when at the target location (T), or optionally to trace movement trajectories (M). The supply line from an air compressor was split between 3 pressure regulators (green circles, 1 for each arm and 1 for the pens). 24 volt AC reticulation valves (light blue rectangles) controlled muscle air pressure. Joint encoders (purple arrows; 10k potentiometers) tracked arm location, and a BASIC Stamp microcontroller (BS2SX-IC) modulated the relay valves to provide accurate movement. ◆ **Sensory feedback.** A CCD camera located above the workspace captured an image of accumulating markings every 5 min. The images were pixelated into 8 bit grayscale values (isomorphic to the electrodes on the MEA) and sent back over the internet to command feedback stimulation of the neurons. ◆ **Training.** Animat behavior was compared to the goal behavior to determine the application of training stimulation. Feedback stimuli could change neuronal activity, in turn varying subsequent animat movement and sensory feedback, thus forming a closed-loop system. TCP/IP sockets were used to communicate between the drawing machine and the neurons, which were often located on separate continents.




Figure A.4 Life-support system for MEART's brain.

The microscope used for observing neural cultures in long-term experiments was wrapped in insulation and outfitted with systems for control of temperature and carbon dioxide levels to maintain normal cell culturing conditions.




Figure A.5 The body of Meart at the Moscow Biennale and drawings.

(Top) Metal and plastic arms rest on an 3x3 meter table. Plastic tubes fed pressurized air to pneumatic muscles. A CCD camera hung from the ceiling (not shown) captured images of the drawing. (Bottom right) Development of Meart. **New York** (July 2003). Video feedback was used for the first time to close the loop, but a “scribble” mode in effect randomized movement and pen placement. **Bilbao** (April 2004) Removing scribble demonstrated the arm moved between 4 points only, via 8 movement directions corresponding to the possible combinations of muscle activation. Pen placement remained random. **Melbourne** (June 2004) Joint encoders were added to read in arm positions and command movement in a feedforward manner: muscles were flexed for a duration proportional to the distance to reach the commanded location. Interior positions could be reached as in New York, however accuracy was low. Pen placement remained random. **Moscow** accuracy test (January 2005) A Basic STAMP microcontroller implemented feedback control of arm positions to achieve accurate movement. Outside pens were commanded down when at the target location. The middle pen was commanded down during arm movement.

A.2.2 Meart's body

Artistic design: The Meart data presented here were collected during the First Moscow Biennale of Contemporary Art at an exhibition entitled “art_digital_2004: I Click, Therefore I Am”, where Meart’s goal of filling a square at the center of the drawing workspace was inspired by the Russian artist Kazimir Malevich’s “Black Square” painting. From the art_digital_2004 program, “The action of Meart observing and drawing the Black Square explores the fundamentals of visual creativity and the way we communicate with the world through images, symbols and their underlying meanings.” This goal behavior was a simplification of the mappings used during our previous exhibitions, to improve experimental controllability. In a previous Meart exhibit, we added an element of interactivity by having Meart draw photographed faces of gallery attendees, entitled the “Portrait Series”. As with images from the drawing, the faces were pixilated and Meart’s goal was to shade the drawing to match the grayscale pixel values. To give viewers a better understanding of Meart’s brain, and the lab in which it was studied, live images from the laboratory, a close-up of the MEA, and a data display of neural activity were projected onto the exhibit walls. This, along with computers displaying the movement and feedback data streams, made the distributed nature of Meart more apparent.

Meart’s body was designed to closely resemble organic forms in function and aesthetics. Shapes were based on bones (influenced by the photographer Andreas Feininger (Feininger and Schlatter 2003)), and sanded Perspex offered an elegant look that referred to a skeletal structure. Similarly, the pneumatic muscles paralleled biological muscles. The design had no covering, never attempting to hide or deny the underlying technology. Analogously, the complex biology of the rat was reduced to a few thousand neurons and glia, grown in vitro. Meart was thus a symbol of the reductionist nature of science and of the stripping down to expose the physical substrates of the creative process

(a parallel to Malevich's stripping of conscious thought in an attempt to expose pure creativity in his work).

Below, Emma McRae (McRae 2004) paints a verbal picture of MEART:

"1. Introduction to a Cybernetic Entity

The soft popping sounds of air releasing, of the breaths taken between movements as the muscles contract and release on the mechanical structures at work on the table in the centre of the room, reach me first as I walk down the dark corridor in the Australian Centre for the Moving Image. I can see the plastic and metal arms and the tubes connected to two rows of valves – regular black garden hose valves – highlighted by a spotlight, that seem to create the movement of the arms. These arms (the creators call these structures arms, presumably because they hold pens and draw as human arms involved in drawing do) are busy drawing lines in apparently random directions with three different coloured pens on a large sheet of paper on the table. Behind the arms is a computer screen showing a photo of a man's face, a pixellated black and white image, a scrolling text box, and some graphs. The only other thing on the table is a camera which looks down over the arms at the picture they're drawing. A large screen on the wall behind the table shows a graph, a representation that looks like a glacial landscape and is constantly changing form, its peaks and troughs rising and falling in random motion, depicting varied intensities coloured in blue, yellow, white, and red. There are two smaller screens in the opposite corner of the room that intermittently display an image of a science laboratory, a close up of a petri-dish, a screen of 64 ECG-like blue tracking graphs, and a microscope view of cells."

Movement - The drawing machine consisted of two perpendicular, rigid, jointed arms (aluminum & acrylic Perspex) hinged at their ends to a 3 m x 3 m table actuating the *X* and *Y* positions of a group of pens over a sheet of paper (Figs. A.1 and A.5).

Similar to biceps and triceps, McKibben braided pneumatic artificial muscles could contract individually, allowing independent flexion or extension of each arm within approximately a 30 cm by 30 cm workspace. Similarly, activation of smaller muscles pressed pens to the paper; a dark pen marked target locations, while an optional lighter colored pen traced the movement trajectories. The supply line from an air compressor was split between 3 pressure regulators, 1 for each arm and 1 for the pens, to isolate pressure fluctuations. Air pressure and thus arm and pen movement was controlled by opening and closing 24 volt AC pneumatic valves. Pneumatic muscles, while offering a

high power to weight ratio, produce nonlinear motion difficult to predict. Therefore, arm location was tracked using joint encoders (10k potentiometers), and a BASIC Stamp microcontroller (BS2SX-IC) modulated valve opening to provide accurate movement as commanded by the living network (Fig. A.6).




Figure A.6 Accuracy test of the robotic drawing machine.

Movements between 7 locations were commanded 200 times in random order. A dark pen marked the target locations, while an offset lighter colored pen traced the movement trajectory. 3 cm x 3 cm resolution targets could be reached within 4 sec and a 1 cm x 1 cm target around 10 sec (not shown). A photograph of Malevich's "Black Square" painting can be seen projected on the gallery wall.

Sensory feedback - A digital camera located above the movement workspace captured images of the drawing in progress. Fluctuations in light from shadows and clouds could strongly influence the image quality. Therefore, ambient and natural light sources were eliminated except for bright spotlights on the drawing itself. Image inhomogeneity was corrected by subtracting from the captured images an image of the sheet of paper when blank, prior to a drawing. The accumulation of markings was recorded every 5 min by retracting the arms out of view and capturing an image.

Internet communication - TCP/IP sockets were used to send motor commands to the drawing machine and return images of the progression of a drawing for feedback. To reduce internet bandwidth, 8 bit grayscale values of an 8 x 8 grid of pixels (isomorphic to the electrodes on the MEA) were sent back over the internet to command feedback stimulation of the neurons.

A.2.3 Software development and Experimental design

Motor transformation - For an animat to behave, sequences of neuronal action potentials need to be transformed into body movements, but understanding how such sequences encode information is a subject of much scientific inquiry. Population vector coding is a candidate motor mapping found to occur in the motor cortex (Georgopoulos 1994), premotor cortex (Caminiti, Johnson et al. 1990), hippocampus (Wilson and McNaughton 1993), and other cortical areas: the vector sum of firing rates of a group of broadly tuned neurons taken together provide an accurately tuned representation (e.g., to a preferred direction of arm movement). We found that a related population calculation of the Center of neural Activity (*CA*, analogous to the *center of mass*) could reliably quantify neuronal network plasticity on a MEA by including spatial information, whereas measuring firing rates alone could not (Chao, Bakkum et al. 2007). Therefore, animat movement was calculated from the *CA* of 100 ms of responses after each probe stimulus:

$$\text{Meart:} \quad \begin{bmatrix} X \\ Y \end{bmatrix} = \hat{T} * \overrightarrow{CA} = \hat{T} * \frac{\sum_e (N_e \cdot \vec{W}_e)}{\sum_e N_e} \quad (\text{A.1})$$

$$\text{Simulated animat:} \quad \begin{bmatrix} dX \\ dY \end{bmatrix} = \hat{T} * \overrightarrow{CA} = \hat{T} * \frac{\sum_e (N_e \cdot \vec{W}_e)}{\sum_e N_e} \quad (\text{A.2})$$

The *CA* is the vector summation of action potentials at each electrode e (N_e) weighted by the spatial location of the electrode (W_e). The transformation, $\hat{T}$, is a normalization matrix found prior to the closed-loop experiment to offset and scale the *CAs* (in electrode space) such that animat movement could produce a uniform distribution and the ability to place pen marks throughout the workspace (Meart) or move in any direction (simulated animat). Achieving a goal for either Meart or the animat required shifting the distribution of normalized *CAs*. Therefore, plasticity results were comparable. The responses to 1 Hz stimulation on a probe electrode were averaged

between consecutive movements (every 4 sec or 1/4 Hz) and used to command Meart pen location, while the responses to 1/4 Hz stimulation on a probe electrode were used to command the simulated animat. A single repeating probe electrode was used throughout an experiment.

Movement could be commanded by absolute location (Meart) or in relative increments (simulated animat). For each case, the activity was normalized to equally distribute the distribution of CAs prior to experiments. For absolute location, this set the possible pen locations to be distributed throughout the whole workspace. For incremental movement, this set the possible movement directions to be distributed throughout 360°. Absolute pen location was used with Meart to avoid movement exceeding the workspace, which would introduce discontinuities in behavior. Incremental movement (eqn. 2) was later used for the simulated animat as workspace size was not physically limited, and we were more interested in direction of movement than position.

Training and sensory feedback – Successful behavior was determined from comparisons between consecutive feedback images. If a larger proportion of markings occurred inside the target geometrical area than outside, behavior was considered successful. Otherwise, a change in the probe response was desired. For training, plasticity was induced by repetitive stimulation of paired electrodes, termed Patterned Training Stimulation (PTS). A PTS was constructed by pairing the probe electrode with another active electrode (one that evokes network responses) 20 ms later, repetitively stimulated for 3 sec with an inter-pair-interval of 100 ms.

For the simulated animat, the training algorithm was modified in 2 ways. A pool of candidate PTSs was formed by pairing the probe electrode with other electrodes ($N_E = 58$) and inter-pulse intervals $\{-80, -40, -10, 10, 40, 80 \text{ ms}\}$ ($N_{PTS} = 58 \times 6$). The probabilities of choosing a given PTS were initially uniform and increased or decreased based on whether subsequent animat performance was successful or not. This allowed an iterative search for an appropriate training "solution" to direct neuronal plasticity.

Second, plasticity can arise from both the PTS stimuli and ongoing spontaneous activity occurring between probes. In a model network, a random stimulation stabilized neural synaptic weights (Chao et al., 2005). Therefore, when animat behavior was successful (no PTS application), a random background stimulation was used between probes such that the plasticity accumulated from a series of PTSs was maintained. The goal of the simulated animat was now to learn to move within $\pm 30^\circ$ of a goal angle.

A.3 Results:

Meart was first exhibited in August 2002 at the Biennale of Electronic Arts Perth (BEAP). However, the precursor to Meart, Fish & Chips, was shown in 2001 at Ars Electronica in Austria. For this ground-breaking bio-art exhibit, SymbioticA Research Group created Meart's drawing arm and used it as the embodiment of a semi-living artist. This was called Fish & Chips because an acute goldfish brain slice was maintained and electrically interfaced on a silicon chip, and used as the controlling "brain" of the arm. From the collaboration between SymbioticA in Perth and the Potter lab in Atlanta, Meart was born: the first robot controlled by a network of neurons in a culture dish, with a 2-way interface via a multi-electrode array. To the existing drawing arm, we added a sensory system, where images from a CCD camera were translated into electrical stimuli for the cultured network. It was also the first neurally-controlled robot whose brain lay far away from its body, with the internet in some ways serving as a very long nerve connecting brain to body. It was the first physical embodiment for a cultured network that remained continuously connected for extended periods of several days, creating numerous drawings during exhibitions.

Early exhibitions were devoted to debugging the communication software and robot mechanics (Fig. A.5), and the most recent exhibitions allowed experimentation. We noticed early that continuous sensory input over the course of days tended to reduce the number of spontaneously occurring network-wide bursts. This led to a hypothesis that

other types of bursting, such as epileptic seizures, might be treated by continuous multi-electrode stimulation. We quantified the short term “quieting” effects of distributed multi-site stimulation on cortical cultures (Wagenaar et al., 2005b), and we are now pursuing the longer-term, or homeostatic effects of continuous stimulation that comes as a consequence of embodiment.

For the data presented here, Meart’s behavioral goal was to draw a solid 12 cm x 12 cm square within the center of its 30 cm x 30 cm workspace. The simulated animat was used to test training algorithms between Meart exhibitions in order to improve behavioral performance. The simulated animat’s behavioral goal was to incrementally move within $\pm 30^\circ$ of a desired angle. (Note that this differed from Meart’s goal behavior of producing pen markings, commanded by absolute location.) For both Meart and the simulated animat, the relationship between changes in neuronal activity and the decision whether or not to apply feedback training stimulation were identical, and thus results about plasticity and learning were comparable.

Meart behavior (unsuccessful), during an exhibition at the M’ARS gallery as part of the 1st Moscow Biennial for Contemporary Art, and animat behavior (successful) are presented in Figure A.7. Electrical stimulation can be an artificial inducer of neuronal plasticity, changing a network’s input-output function. Bi and Poo found that for mono-synaptically connected cultured neurons firing within a few tens of milliseconds of each other, directional spike timing dependent synaptic plasticity occurred (Bi and Poo 1998a). Repetitive stimulation of pairs of electrodes in a PTS could therefore cause plasticity in shared pathways of neural activation. Fetz and co-workers (Jackson, Mavoori et al. 2006) provided supporting evidence of pathway plasticity in vivo: they repetitively stimulated a neuron in the primate motor cortex 5 milliseconds after the occurrence of an action potential on a different poly-synaptically connected neuron using a chronically implanted neural interface. After halting the stimulation, subsequent activity of the recorded neuron caused an increase in the firing rates in the vicinity of the stimulated

neuron. In this manner, we hypothesized the PTS would lead to potentiation of the probe response in the vicinity of the second paired electrode, modifying the CA and population vector such that arm movements would approach the target area.




Figure A.7 Plastic changes in Meart and animat behavior.

Unsuccessful and successful training of goal-directed animat behavior. Meart. Training with predetermined patterned training stimuli (PTS) caused a shift in the probability distribution of commanded movement directions in 2 experiments (circles, bottom row), but in an uncontrolled manner. The presence of a movement bias (see Methods) caused marks to accumulate on a side of the drawing's workspace (CCD camera image of the drawing and pixelized feedback), but successful PTS training should shift the markings back towards the center (red arrow middle row; black arc bottom row). The probability distribution of movement directions during 10 min at the start of 2 hr experiments was subtracted from that during the final 10 min, thus allowing negative values (red). Simulated animat. Iteratively updating the probability of selecting a given PTS for training allowed an animat to learn to move in multiple directions (circles; see Methods). Desired angles of 0° , 90° , and -45° (black arcs) were applied in consecutive 2 hr periods. Successful behavior was considered to be movement within the desired angle $\pm 30^\circ$. Notice the changes in probability distribution of movement direction were now more likely to be in the appropriate direction and more focused than for Meart.

For Meart, the transformation from visual sensation into the delivery of a PTS was fixed. For example, if previous movements occurred below the target area, the probe was paired with a predetermined electrode at the top of the MEA. While successful behavior did not occur (Fig. A.7), neural plasticity did (Figs. A.7 and A.8), suggesting training stimuli had the potential to modify behavior. Normalized plasticity was defined as the difference in distribution of movement-controlling output (the CAs) in a given 10 minute period (CA_{Post}) to those of the first 10 minutes (CA_{Pre}). It is equivalent to the C/D calculation in Equations G.1 and G.2. A value of 1 indicates no change.

We concluded that since neurons at different electrodes are connected through multiple intermediate neurons and pathways, the effect of a given PTS can not be predicted. By using feedback of behavioral performance to select and refine effective sequences of PTSs, instead of using Meart's fixed PTSs, the simulated animat could now achieve its goal-directed behavior (Fig. A.7). To demonstrate that the successful behavior was a consequence of the biological changes in the neural network and not an artifact of the algorithms, the desired movement angle was switched between 3 angles every 2 hours. Even though movement was commanded by absolute location for Meart and incremental movement for the animat, training was intended to produce the same effect on neural activity: shift the distribution of CAs (and in turn movement angles) towards a desired goal direction.

The adaptive training algorithm allowed a search for "solutions" to achieve goal directed behavior (Figs. A.7 and A.8). Some PTSs may give desired neuronal plasticity while others may give the opposite or none. Furthermore, a neural network is continuously plastic, and the same PTS may have different effects at different times. The training algorithm commanded the application of a sequence of PTSs to produce the appropriate neural plasticity for successful adaptation. The learning curve in Fig. A.9 shows the percentage of successful movements in time; progressively fewer PTSs were

needed to maintain the desired behavior, suggesting that the animat was learning the appropriate behavior.




Figure A.8 Neuronal plasticity.

A. An experiment with Meart (data is the same as Fig. A.7, left) run for 2 hours and compared to 1 hour probe-only periods before and after. “Normalized change” is a comparison of the movement outputs (the CAs) in any 10 min period to those of the first 10 min. At time = 0, the same periods were compared, giving no change (a value of 1), and the 10-minute window for subsequent values was stepped by 1 minute. The drop below 1 in the control periods meant the variability in CAs decreased, possibly indicating a habituation to the stimulation. The addition of training stimuli caused plasticity, but not behavioral success (Fig. A.7). **B.** The experiments with the animat (data is the same as Fig. A.7) run for 2 hours. The adaptive training algorithm caused plasticity. For 90°, change hovered around 1 because this was the direction of bias, a 60°/360° chance, (see Methods).




Figure A.9 Training series and learning curve for the simulated animat.

Animat learning curve and training history in living culture. (Data from the -45° desired angle trial in Fig. A.7.) **A.** A greater portion of animat movement was toward the desired direction after 30 min. An animat moving randomly would give a 16.67% chance that the movement was within the $\pm 30^\circ$ range of the desired angle ($60^\circ/360^\circ$). **B.** Training was designed to select the PTS that induced appropriate neural plasticity as determined by subsequent animat behavior. The improved performance at 30 min corresponded to an increase in the occurrence of PTS205, whose paired pulse pattern is shown below; its electrode location is shown in the 8x8 grid (blue dot; the probe electrode is a black X). A different PTS pattern increased the RBS occurrence at 80 min (red).

A.4 Discussion

To view Meart is to witness a collage of contradictions. It offers us the actual biological substance of the thinking brain yet out of its biological context and system of developmental ordering. What is visible to us as Meart in the space of public display is a visualization of and/or window into ongoing experiments occurring thousands of miles away in a laboratory. The outcomes are neither pre-defined, nor are their meanings fully understood. Indeed, any of the aforementioned skeptical questions place us as viewers firmly in the midst of vigorous scientific debates—a fact underscored by the “real-time” nature of the Meart performance.

Like a work of science fiction, Meart stimulates broad inquiry into our own lived contexts. However, unlike sci-fi, it is not simply a representational text, but also an operational one. It cannot be dismissed as a mere illustrative flight of fancy, but must be interrogated as a concrete example. Meart is an “operational fiction”—a cyborg of representation and reality, art and science, and of course flesh and transistor.

-- Paul Vanouse, Excerpt from the Strange Attractors exhibition catalogue, Zendai gallery, Shanghai, 2006

Gallery visitors were first captivated by the aesthetics of the kinetic sculpture. Meart’s organic movement and the “breathing” sound of the pneumatic relay valves intermittently popping and hissing, not quite structured and not quite random, gave an intriguing sense of calm, maybe similar to watching trees sway in a gentle breeze. This hinted at the presence of an underlying natural process. A subtle curiosity to figure out what was happening turned into apprehension of the uniqueness of this semi-living artist, and then intense questions about the nature of the mind, the body, life, and about the artistic and scientific messages.

A.4.1 Art vs Science

In our society, art and science usually occupy distinct disciplines. Humans are very adept at forming categories, and this is useful in making sense of the world, but convention is tailored by culture’s current mood. The wide influence of 15th century

artist and scientist Leonardo Da Vinci gives reason for pause and reminds us of the many connections between the artistic and scientific. After working on Meart, we have come to appreciate that both developing a work of art and making a scientific discovery require a curiosity and a passion to find new ideas, an ability to recognize a void in human understanding, and the creativity to form a solution. Does this comprise the ‘mysterious’ in Einstein’s quote?

Of course, tensions exist. The scientist needs to add precision and controllability to the project, then objectively document the results, constraints an artist may consider extraneous. In turn, the artist needs to conceptualize the project’s importance and perfect its aesthetics, details a scientist may consider superficial. However, art and science also share the same goal: to expose new perspectives or forgotten truths about the world – to expand wisdom. Their presentation differs, but viewing an object of study from multiple angles broadens perspectives to new, possibly fertile ground. Exposure to the other’s discourse can lead to a clash of cultures, but also a mirror to critically reassess one’s own perspective. If nothing else, Meart certainly got artists thinking more about science, and scientists thinking more about art. Since 2002, “Meart, the semi-living artist” has exhibited at galleries in Shanghai, Moscow, Atlanta, Melbourne, Bilbao, New York, and Perth, often as part of larger exhibitions that focused on the use of new technology in art. The galleries became laboratories, as exhibitions were nearly the only time when experimentation was possible, and the scientific method became performance art. Meart has been presented at scientific conferences on artificial intelligence, neuroscience, and bioengineering in Switzerland⁸, Germany⁹, Italy¹⁰, and France¹¹, in addition to numerous other lectures to scientists and college students.

⁸ 50th Anniversary Summit of Artificial Intelligence, Monte Verita, Switzerland, 9-14 July 2006

⁹ Embodied Artificial Intelligence, International Seminar, Dagstuhl Castle, Germany, July 7-11, 2003

¹⁰ European School of Neuro-IT and Neuroengineering, Genova, Italy, June 13-17, 2006

A.4.2 Intelligence and embodiment

The desire to breathe life into sculpted clay, or today into silicon microchips, has been around for thousands of years (www.a-i.com ; (Kac 1997)). This desire in part formed the scientific fields of artificial intelligence, cognitive science, and robotics. Their inquiries into the nature of intelligence began in the middle of the last century without a concern for its substrate: intelligent thought was considered the manipulation of abstract concepts, synonymous to how digital computations could be run on any manner of Turing machine. Digital computers (computers used to be humans who computed) have accomplished impressive feats, solving equations and defeating chess champions by relaying bits of information through discrete logic gates within nanoseconds. However, intelligence has not yet been attributed to computers or the robots they have been used to control. Tasks trivial to humans have proven difficult for computers, such as adaptation, multitasking, pattern recognition, fault tolerance, etc. This is likely due to significant differences in computational implementation, with brains using massively parallel processing, feedback loops on many scales, and relay switches (neurons) that learn and change function (Potter 2007). Digital computers were predicted to revolutionize calculators and the control of traffic lights. They did that, but obviously have embedded themselves in almost every aspect of our modern lives and technology. A better understanding of biological intelligence would have its own exceedingly unimaginable impact.

Now becoming more accepted by scientists is the hypothesis that intelligence is not disembodied, but intimately entwined with the mechanics of the body and an interaction with the environment (Varela, Thompson et al. 1993; Clark 1997; Pfeifer and

¹¹ 29th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, Lyon, France, August 23-26, 2007

Bongard 2007). The act of walking combines roles for neural signaling, proprioceptive feedback, the spring tension of muscles, the friction of shoes contacting pavement, and gravity to assist leg swing: both our brains and bodies were designed to take advantage of the physics in the world. With Meart and also biological movement, the presence of friction improved precision and stability by damping overshoot. Meart's muscles and other nonlinear components were not considered negatives, and our experiments tested the neuronal network's ability to learn the dynamics of its body to achieve goal-directed behaviors.

So Meart is embodied and *situated* in the real world. Does Meart manipulate abstract concepts of the external world in its small brain of a few thousand neurons? We doubt it, agreeing with the anti-representationalist stance Neil Manson and his interpretation of our work, whether the cultured network is embodied in a simulated neurally-controlled animat or an actual robot (Manson 2004b):

"Anti-representationalist theorists propose an alternative model: an embodied agent conception of cognition (Varela, et al., 1991; Franklin, 1995; Clark, 1997). On this conception the creature is viewed as part of the causal flux of its environment. Its success in satisfying its needs depends upon its competence in shaping its trajectory through the environment. Successful action requires creatures to use the information present in their environment (i.e., the causal regularities that actually obtain in their environment). This does not require the formation of an internal representation of the environment, it simply requires the creature to stand in the right kind of causal relations to its environment. Cognition on this view is an embodied, situated affair."

"The NCA (neurally-controlled animat) experiment has its background in this model of cognition. Earlier, I talked of the cognitive aspirations of the Potter Group. This can be read in two ways. If we assume the traditional model of cognition, the NCA methodology will only be of use for cognitive neuroscience if the cluster of neural cells gives rise to internal representations of the virtual environment. If we reject this model and situate the NCA methodology in its proper home—artificial life and embodied-agent AI—the cognitive aspirations look quite different. Some of the explananda of cognitive neuroscience (e.g., the brain's role in learning, adaptive behaviour, and linking perception and action) are amenable to embodied-agent modelling, and this is exactly what the Potter Group seem to be doing with the NCA experiment. On this second interpretation it need not be assumed that the neural cells subserve internal representations of the objects in the artificial world,"

A natural extension is the use of external tools to scaffold intelligence (Clark 1997). People have learned to extend memories with notebooks, voices with cell phones, vision with microscopes, and more. The intelligent task of multiplying large numbers often combines cognition with a pencil and paper. The distinction between technology and biology in ‘being human’ is becoming more tenuous (Clark 2003). Many humans now live symbiotically with heart pacemakers and cochlear neural interfaces, and extend life spans with medicine. Meart continues this conversation and further questions the body-space of living agents by including the internet as an ‘extended nervous system’: its biological brain and artificial body were often located on different continents.

On the other hand, behavior is constrained by the limitations of the brain and the body. With Meart, movement was confined to a 2 dimensional plane and constrained by the machine's speed and accuracy. The choice of how to map neuronal activity into motion and sensory feedback into electrical stimuli constrains which neuronal plasticity mechanisms could be observed behaviorally. This can be an advantage if investigating an individual mechanism or a disadvantage by limiting the available neuronal computational capacity.

A.4.3 Nature of art and being an artist

Meart has many of the characteristics of a “real” artist. It lives, it dies, it leaves behind a body of work for others to contemplate, but can rat neurons and a mechanical body be labeled an artist? Maybe Meart is disqualified by being manmade. However, fillings for cavities in teeth and artificial hips make people part manmade, but no less human. Meart would have to be disqualified in some other sense. Does it possess sufficient creativity and intelligence to produce a work of art? Maybe not, but if so, would this suggest art is not solely a human endeavor; have we made an artist? If it possesses intention, maybe we have infringed on its intellectual property rights when

drawings were purchased by a gallery as discussed by Hughes (Hughes 2007) (Fig. A.10). Will the training algorithm enslave biology in order to steal from it, or are such goals natural: does the body enslave the brain in order to live, by demanding it learn how to find and eat food?




Figure A.10 Does Meart create valuable art? Meart drawing and notes from an early accuracy test. This and four other drawings were purchased by a museum in Spain (MEIAC, Museo Ibero Americano de Arte Contemporânea) for their permanent collection.

Of course, Meart is a primitive construction, and much scientific inquiry remains to be done. But the continued merging of biology and technology give substance to such questions. The answers given for the potential offspring of the Meart project may be more controversial. For now, the tangible debate centers on what is the creative output: the drawings, the machine (if so then why not the brain?), a performance piece, conceptual art, not art because it was infected by science, or the system as a whole.

A.4.4 Fear and the future

"Within thirty years, we will have the technological means to create superhuman intelligence. Shortly thereafter, the human era will be ended."

Vernor Vinge -- 1993 essay "The Coming Technological Singularity"

After addressing viewer's typical first questions during exhibitions: 'Is it alive?', 'Is it thinking?', 'Is it creating art?' ('Partly.', 'That is the scientific question.', 'What do you think?'), a next question is often "Will this turn into Terminator II?", a robotic harbinger of the apocalypse in a doomsday movie. One of the goals of Meart is to provide a public forum for education and dialogue to address the "fear of the unknown" and to critically examine the paths to be paved by biotechnology: living with the semi-living, becoming the semi-living as we incorporate more technology into our bodies and lives. Further understanding biological intelligence is expected to improve artificial intelligence (Bakkum, Shkolnik et al. 2004b; Potter 2007), but Meart remains rudimentary, and as mentioned above, digital computers and robotics lag behind the capabilities of biological agents.

The ethics of any technology lie not in the technology itself, but in how it is used. For example, nuclear energy can both level cities and create a nuclear winter or power cities and create life. Plague bearing rats themselves have now become indispensable tools to advance medical technology. An understanding of how networks of neurons process information and how they can be best interfaced to achieve goal-directed behavior could influence future neural prosthetics for sensory deprived or paralyzed patients. Currently, prosthetics are being developed to restore hearing, vision, motion and even anatomical parts of the brain itself (Berger and Glanzman 2005). Will giving a bit back to those who have lost outweigh potential negatives and sacrificed animals?

More immediate are concerns about the continued melding of biology and technology and the role humans will have in creating life and death, especially if 'semi-living' agents ever learn human qualities: intentionality, memory, irony, interpretation,

creativity, etc. Moreover, the use of biology as an artistic pallet shifts art from imitation of nature to one that subsumes nature in its expression: partly alive artwork. Meart required constant care and attention. During the 2002 BEAP - Biofeel exhibition at PICA in Perth, Meart stopped moving when the neuronal culture died from insufficient environmental control, since improved (see Methods). The gallery went silent with the sudden realization that Meart had been somehow alive. The implications of such technology to manipulate life had been presented through the irony of a death, and highlighted the need for compassion and a greater understanding of life.

A.4.5 Conclusion

Meart and other hybrot provide a platform to continue philosophical and begin experimental inquiry into the fundamental makeup of intelligence and existence. Our next step is to apply the algorithms developed with the simulated animat to control the drawings of Meart, creating a real-world biologically-based agent exhibiting goal-directed behavior.

APPENDIX B

HOMEOSTASIS OF GLOBAL FIRING RATE

B.1 Introduction

Neurons modify their activity through synaptic and intrinsic plasticity mechanisms induced by efferent input and temper these changes through homeostatic mechanisms. Consequently, most cortical neurons tend to fire within a frequency range, 1 to 5 Hz *in vivo* (Abeles 1991), as maintained by homeostatic mechanisms (Turrigiano and Nelson 2000). This implies that a network of neurons will tend to operate within a corresponding set of network activity states. Indeed, we found that while comparisons of firing rates have shown plasticity in intracellular recordings, firing rate statistics alone could not depict plasticity across networks of neurons (Chao, Bakkum et al. 2007). Synaptic noise across a chain of neurons (Kandel, Schwartz et al. 2000), convergent and divergent pathways (Abeles 1991), and homeostatic mechanisms that re-adjust firing rates in response to plasticity (Turrigiano and Nelson 2000; Spitzer, Borodinsky et al. 2005) all obscure firing rate measures of plasticity detected by extracellular MEAs. Therefore, we investigated the relationship between homeostasis of firing rate in our cultures and the activity dependent plasticity of action potential propagation (Chapter 4). We found that network firing rate was perturbed by application of a patterned stimulation, and as plasticity accumulated, network firing rate re-approached a “homeostatic set point”.

B.2 Methods

From the patterned stimulation protocol in Chapter 4, a simple low frequency stimulation pattern was varied every 40 minutes. Each pattern consisted of alternatively

stimulating 2 electrodes at 2 second intervals. The second electrode, termed probe, was fixed and used throughout, while the location of a preceding context electrode was moved spatially every 40 minutes to make each new pattern. An experiment consisted of 8 consecutive sets of patterned stimulation. 40 minutes was chosen to allow enough time to stabilize plasticity induced after the patterns were changed. A slow 1/2 Hz overall stimulation rate was chosen to avoid network fatigue or refractory periods (Darbon, Scicluna et al. 2002) from compromising the evoked responses. The stimulation electrode evoking the most dAPs was chosen as the probe electrode. The probe was paired with electrodes evoking varying degrees of neural activity to create patterned stimulation with diverse network activity responses. See Methods in Chapter 4 for how to detect dAPs and Appendix G for cell culture methods and electrophysiology.

B.3 Results

The original motivation for the pattern stimulation protocol in Chapter 4 was to test if the cortical networks exhibit homeostasis of firing rate and then characterize the mechanism and rate of adaptation (Fig. B.1). Changing the pattern or location of stimulation changes which neurons and which order those neurons produce action potentials. Some synapses may become more active and others less. Some neurons may fire more and others less. These induced changes in the spatiotemporal flow of network activity may place the network state outside of its ‘comfort area’ or area of ‘homeostatic stability’, depending on the effectiveness and amount of stimulation. Given homeostasis mechanisms, synaptic and/or intrinsic plasticity are expected to occur in order to move the network state back into an area of homeostatic stability.

The best estimate for the homeostatic network firing rate is the end of the experiment where the network would have had the most time to habituate to 5 hours of 1/2 Hz stimulation, and after the last half of the final period to minimize the effect of the last change in patterned stimulation. Since 6 experiments were conducted, 6 firing rate set

points were found, and the firing rates within each of the 8 pattern stimulation periods were indeed found to significantly move closer to the estimated set rate. Figure B.1a suggests the network firing rate homeostatically adapted to changes in stimulation pattern.

Many cellular mechanisms could contribute, including the modulation of action potential latency. Despite the large variance in Figure B.1b, changes in latencies after a change in stimulation pattern relative to the end of a stimulation pattern were significant (Fig. 4.4c). This supports changes in dAP latency being an adaptive plasticity mechanism that may have a role in the homeostasis of network firing rate.




Figure B.11 Network firing rate (FR) and dAP latency adapted to stimulation pattern.

(a). The FR is plotted in time for each 40 minute stimulation period. FR is plotted as a percent of an estimated FR set point, the FR measured at the end of the experiments (see text). The FR difference was closer to zero at the end of a period than the beginning indicating FRs approached the set point ($p = 0.02$, Wilcoxon signed rank test. $N = 48$ periods; 6 cultures from 4 platings). 3σ was fit with an exponential curve ($\tau = 15$ min; $R^2 = 0.88$). (b) The accumulation of plasticity in dAP latency slows in time. Similar to the calculation of normalized change in Chapter 5 (also the C/D calculation in Appendix G), the average difference from each latency for 5 minutes (150 probe stimuli; time within period denoted on the x-axis) to the mean latency in the 5 minutes prior to the period, was calculated. The mean and s.e.m. are plotted and an exponential curve is fit through the mean values ($\tau = 33$ min; $R^2 = 0.92$; $N = 130$ PTS trains; 6 cultures from 4 dissociations).

APPENDIX C

CLOSED-LOOP MOTH EXPERIMENT

My first closed-loop experiment consisted of a simple embodiment with a single association between neural activity and the frequency of feedback stimulation on one electrode. This was preliminary work for the embodiment experiments described in Chapters 2 and 5 and Appendix A in order to test the closed-loop interface and software. The experimental design was inspired by a moth's tendency to fly towards a light bulb: the closer it gets the stronger the stimulation. Thus, the inter-pulse-interval, $T(x)$, of stimulation on one electrode was adjusted based on the animat's location in an artificial environment, x :

$$T(x) = \frac{1}{\text{frequency}} = b - kx \quad (\text{C.1})$$

where k and b are constants equal to 4.9 and 5 seconds. The constants were chosen such that the maximum stimulation frequency was 10Hz for values of x constrained within ± 1 . The x location determined stimulation, but the artificial environment was a 2D plane $\{x, y\}$. (This was done for controls, where later experiments would switch to the y location to command $T(y)$) Eight electrodes with high activity were chosen to contribute to the population vector, representing eight cardinal directions (Fig. C.1). Locations x and y were calculated from the firing rates of these electrodes:

$$\begin{aligned} x &\leftarrow x + \sum_{i=1}^8 R_i \cdot Q_{x,i} \\ y &\leftarrow y + \sum_{i=1}^8 R_i \cdot Q_{y,i} \\ R_i &= \frac{\text{FiringRate}_i}{R_{\text{norm},i}} \end{aligned} \quad (\text{C.2})$$

where i is an electrode index, R is a normalized electrode firing rate, and Q is a unit vector with x and y components. The normalization constants, R_{norm} , were determined in advance from averaging firing rates of each electrode during 5 minutes of 0.2 Hz stimulation (corresponded to $x = 0$). The normalization of the electrode corresponding to positive x movement was slightly decreased to bias the increments of x to be in the positive direction at the start. Changes in the trajectory of movement indicated a change in the ratio of firing rates (Fig. C.1). The first occurrence of such a change was around 27 min and was accompanied by a period of high activity sustained for 40 to 60 seconds. The elevated activity appeared to be different than a typical network burst in that it had a longer duration and lower firing rate.




Figure C.1 Closed-loop moth experiment.

(left) The trajectory of movement. **(right)** The firing rates of the electrodes contributing to the movement. A period of high activity sustained for 40 to 60 seconds was observed (peaks in the right figure) upon reaching ~ 2 Hz stimulation, after which the network activity patterns changed such that the movement became relatively stationary with respect to the direction of adjusting frequency (the upward shift in the left figure).

Unfortunately, the data for this and a large number of similar experiments was corrupted and so cannot be analyzed further. However, the results suggest differences could exist between open-loop and closed-loop protocols in how stimulation can encode

plasticity. For example, neural networks are characterized by feedback loops, and adding an external (environmental) loop may better allow the electrical stimuli to tap into neural population dynamics. I hypothesized that such effects would manifest themselves as differences in the time course, magnitude, or stability of plasticity (see also Fig.5.5). I expect similar simple closed-loop experiments could tease apart a few details.

APPENDIX D

USING THE NEUROMANCER STIMULATION BOARD FOR ARTIFACT SUPPRESSION

D.1 Introduction

The Neuromancer stimulation and recording board was developed by Edgar Brown, Jim Ross, and Richard Blum in Steve DeWeerth's laboratory in order to actively suppress electrical artifacts arising from electrical stimulation (Brown, Ross et al. 2007; R.A. Blum, J.D. Ross et al. 2007). Artifacts dominate recorded voltages traces for a period of time after a stimulus, masking the detection of action potentials and thus limiting the maximum stimulation frequency. To use the Neuromancer stimulation board in experiments, its performance on real MEAs plated with neurons needed to be understood. Artifact suppression here *is different* than results using test MEAs (resistors connected to ground) because electrodes have varying properties dependent on wear, and the MCS amplifier introduces 3rd order dynamics. In addition, the communication speed of the board was also tested to determine if it could be used in closed-loop embodiment experiments.

To summarize, the Neuromancer suppressed artifacts on the stimulation electrode but offered no advantages for suppressing artifact on non-stimulation electrodes, where a few ms blanking was unavoidable. Moreover, the artifact suppression parameters needed to be hand adjusted for each electrode, which can be very tedious. An advantage was multiple electrodes could be stimulated at the same time, although the board currently supports half (32) of the MEA electrodes. Edgar Brown made a new microchip to allow stimulation encoding on an analog channel, but this was not yet tested. The board could

handle closed-loop speeds of 20 Hz now and 100 Hz if code is adapted to use pthreads in rtlinux. The main points are:

1. The Neuromancer stimulation board successfully suppressed artifact to (on average) within 1 ms of evoking neural activity.
2. The Neuromancer stimulation board successfully evoked neural activity at comparable rates to the RACS stimulation board.
3. The Neuromancer stimulation board could record neural activity on the stimulation electrode after filtering through Salpa on average within 3.12 ms after evoking activity. However, the discharge phases distorted the voltage waveform and introduced false but repeatable spikes.
4. The Neuromancer stimulation board did not reduce the time until other non-stimulated electrodes could be recorded, compared to stimulation with the RACS board.

D.2 Artifact suppression of the Neuromancer board compared to the RACS board, on MEAs plated with neurons.

Artifact suppression was quantified on the Neuromancer stimulation board and compared to the performance of the RACS stimulation board (designed by Daniel Wagenaar). 5 electrodes were stimulated in random order 24 times each at 2Hz, giving 1 minute total stimulation time. Three setups were compared:

1. RACS board with a biphasic voltage stimulation:
 - a. 400us, 500mV per phase (Potter lab standard setup).
2. Neuromancer board with the same biphasic voltage stimulation + a single discharge phase:
 - a. Discharge phase D1 of 400us length, $R_{ej} = -1.65V$, and $V_{ref} = 0V$.
 - b. 400us, 500mV per phase.
3. Neuromancer board with voltage pulse optimized (by hand):
 - a. 350us 400mV up phase; 250us -600mV down phase.
 - b. Discharge phase D1 of 2.5ms length, $R_{ej} = -1.65V$, $V_{ref} = 0V$.
 - c. Discharge phase D2 of 5 ms length, $R_{ej} = -0.05V$.

Voltage traces were recorded using Meabench via the MCCard data acquisition card.

Four signals were extracted:

1. Raw voltage trace.
2. Raw voltage trace filtered through Salpa.
3. Spikes detected from the raw voltage trace.
4. Spikes detected from the Salpa voltage trace.

See also the movies each showing a close up of the stimulation artifact, raw voltage responses, and salpa filtered voltage responses. The stimulation electrodes were in random order, but for clarity are aligned in the movies (t-square: Douglas_Bakkum / Progress Reports / Progress Report 092706 - artifact suppress Neuromancer vs RACS):

- RACSmovie.mov
- NMmovie.mov
- NMoptmovie.mov

In case 2, the voltage trace recorded on an oscilloscope would not return to zero without a D1 phase included. The RACS board voltage trace did return to zero. However, the Neuromancer board's artifact suppression outperformed the RACS board. (The RACS board did not have active artifact suppression built in to its circuit: increasing the time until switching off of the stimulation electrode did not boost artifact suppression of the RACS board, and instead made the artifact bigger.)


When using a test MEA, which connects each electrode contact pad to a common ground through resistors, the artifact produced by the Neuromancer board was consistent among channels and was able to be suppressed. In contrast, when using a standard MEA (8x8 grid with 30 um electrodes spaced 200 um center to center), the artifact produced by the Neuromancer board was inconsistent among channels, requiring the discharge parameters to be set for each stimulation electrode independently. Not all artifacts were able to be suppressed fully. Failure here is defined as the voltage trace recorded through the MCCard remained on the rail after the commanded biphasic voltage pulse ended. (In

Case 3, the optimal discharge parameters were chosen such that they could be used for each of 5 stimulation electrodes. These parameters were not optimal for all electrodes.)

Setting the discharge parameters for each channel by hand makes using the Neuromancer stimulation board time consuming and tedious for experiments using many stimulation electrodes. The DeWeerth group is working on an automated system to set the parameters.

Results

A raw voltage trace of an artifact for a stimulation electrode for the Neuromancer board with optimized discharge parameters:



Artifact lengths:

	NMopt				NM			
	Raw rail edge (ms from stim)		Salpa blank edge (ms from stim)		Raw rail edge (ms from stim)		Salpa blank edge (ms from edge)	
Probe HW	Probe	Others	Probe	Others	Probe	Others	Probe	Others
57	0.39	0.89	2.02	2.51	48.16	0.93	48.42	2.22
51	0.40	0.68	4.66	2.04	43.07	0.83	43.37	1.94
6	0.38	0.54	3.25	1.99	42.27	0.75	42.55	1.92
1	0.56	0.81	1.83	2.33	38.53	0.89	38.80	2.12
0	1.85	1.02	3.83	2.70	52.88	1.00	53.14	2.48
Average	0.72	0.79	3.12	2.31	44.98	0.88	45.26	2.14

	RACS			
	Raw rail edge (ms from stim)		Salpa blank edge (ms from stim)	
Probe HW	Probe	Others	Probe	Others
55	54.14	0.48	54.39	1.31
51	49.35	0.47	49.61	1.42
24	65.35	0.50	68.73	1.31
1	41.69	0.48	41.97	1.31
0	58.70	0.48	58.92	1.31
Average	58.85	0.48	54.72	1.33

Raw rail edge: longest latency until the recorded voltage is within the recordable range. The raw rail edge indicates the speed of artifact suppression.


Salpa blank edge: salpa will return values of 0 when the raw trace is on a rail OR when the raw trace is too steep. The salpa blank edge indicates the earliest time we can record spikes.

Noise levels over 5 sec interval (uV std).

	NMopt	NM	RACS
South Plug (16 pins)	2.59	2.61	2.57
East Plug (14 pins)	2.57	2.57	2.55
Non-plugged	2.55	2.55	2.54

Global firing rates on recording electrodes, excluding the probe, within 200 ms latency (Hz) are comparable.

	NMopt	NM	RACS
Probe			
NM:57 RACS:54	127	125	113
51	102	103	86
NM:6 RACS:24	14	2	1
1	65	77	53
0	24	34	27



The Neuromancer stimulation board successfully suppressed artifact to (on average) within 1 ms of evoking neural activity, and recording neural activity on the stimulation channel after filtering through Salpa could be done on average within 3.12 ms after evoking activity.

However, even though the Neuromancer with optimized stimulation gave an electrode voltage trace that is off the rail within a few milliseconds, the 2 discharge phases distorted the voltage trace waveform and can produce false positive spikes. The effective recording delay would then be after the discharge phases for the stimulation electrode. Minimizing the discharge times by modifying the biphasic stimulation pulse or smoothing the transitions between different stimulation phases using the ramp (/R)

command may help eliminate false spikes. Nevertheless, this time is much less than for non-optimized cases.

The Neuromancer stimulation board successfully evoked neural activity (see movies). For non-stimulated recording electrodes, The Neuromancer board did not allow earlier recordings than the RACS board used in the Potter lab. The Neuromancer board and RACS board did not add noise to the system, and noise levels remained around 2.5 μV .

D.3 Update speed and timing accuracy of the Neuromancer stimulation board for closed-loop experiments.

The stimulation part of the Neuromancer stimulation and recording board was further tested on our setup for use in closed-loop embodiment experiments. (Recording continued to be done with MCard and Meabench.) 60Hz noise was eliminated using a lot of tin foil to wrap around the ribbon cables extending to the MEA and the box containing the board. As the current implementation of the board required updating a microcontroller via a serial cord for each change in stimulation (i.e. choosing a different electrode or voltage), the speed of update provided a limiting factor to maximum stimulation frequency. However, testing indicated speeds of 20 Hz were possible and up to 100 Hz if code is adapted to use pthreads in rtlinux: see details below. Also I wrote a C driver to control the board, as opposed to Matlab.

Testing the Neuromancer board for stimulation.

I received the board with a Matlab GUI as the control interface where buttons needed to be pushed to send consecutive stimuli. To test how fast the board could update stimuli, as would be needed for background stimulation for example, I modified the Matlab code to automatically send a series of stimulations. Later, a TCP/IP socket was added to communicate with a rtlinux computer and Daniel's stimulation board for accurate timing.

To test speed, code was developed to loop through:

> Start

- Update the PIC microcontroller (serial port) stimulation setting (i.e. stim channel).
- Send a trigger to stimulate (either via the serial port or an external pulse on Daniel's stim board).
- Send a wait period until next stimulation (inter-stim interval, ISI).
- Read data returned from PIC (serial port) signaling execution of trigger.

< Goto start

The resulting stimulation pulses were sent to the analog A1 channel of the MCCard, recorded with Meabench, and speed and accuracy were measured. For accurate timing, rtlinux was used (Daniel's racs module containing pthreads). The external triggers were commanded via a TCP/IP socket from the recording computer to an rtlinux computer controlling the RACS stimulator. RACS was commanded to send a digital output via a BNC cable to the external trigger of Neuromancer to trigger stimulation. In time, rtlinux code similar to the Racs.o driver can be developed to control Neuromancer – meaning the results given below should be representative of a standalone Neuromancer application.

Setup 1: ***Matlab*** interface communicating with the serial port for Neuromancer and a ***TCP/IP*** socket for Racs.

For commanded ISIs of 100 ms, the mean actual ISI was 102.8 ms with nearly all values within one sample rate. The sample rate was 25 kHz or every 0.04 ms; 14/100 varied by 1 sample rate. The maximum speed while maintaining a high accuracy was estimated to be 80 ms.

Setup 2: ***C-code*** interface communicating with the serial port for Neuromancer and a ***TCP/IP*** socket for Racs.

For commanded ISIs of 50 ms, the mean actual ISI was 52.4 ms with nearly all values within a sample rate (25 kHz or 0.04 ms; 9/100 varied by 1 sample rate). The maximum speed while maintaining a high accuracy was estimated to be 40 ms.

Setup 3: ***C-code*** interface communicating with the serial port for Neuromancer without Racs. Trigger was instead sent by the serial cable – faster but less accurate until optimized with rtlinux / pthreads.

For the usleep() function call to delay 10 ms, the mean actual ISI was 15.1 ms with 85% within 1 sample rate. The maximum spread was 1.28 ms. For usleep() set to 1 ms, the ISI was around 6-7 ms; this is the likely upper bound on stimulation frequency. I expect that implementing timing control, for example with rtlinux and pthreads, would improve the accuracy and stability of this setup.

A C-code driver is on the file server along with the movies.

D.4 Methods

See Appendix G for cell culturing and general electrophysiology methods. The neural cultures were E-18 rat cortices plated 4/13/6 (4 ½ mo old) on dish 7113 (Multichannel systems MEA). A total of 55k cells were plated: 5uL/k*6.5uL.

Salpa filtering and spike detection were done using built in functions in Meabench. Salpa is a software implementation of artifact suppression (smoothing). It creates a series of polynomials to fit the artifact curve then subtracts these from the original data to get a flat trace (Wagenaar, Pine, Potter 2004). A drawback is large spikes or a series of close spikes as in a burst can be fit by Salpa, distorting the waveform and inducing false positive and negative spikes.

Spike detector type 8 “LimAda-70%-98%-20ms-1s-5.0” was used. It is a bandpass filter that calculates noise levels in a local time slice such that it adapts in real time to sudden changes in electrode noise levels. It also ignores spike occurrence when calculating noise levels (see online Meabench manual). A drawback is the noise adaptation can be fooled and consider bursts a change in noise level, and thus miss spikes during bursts.

Since the raw voltage traces were recorded, they can be passed through alternative filters and spike detectors if desired.

APPENDIX E

HOW TO CODE CLOSED-LOOP EXPERIMENTS

MENTIONED FILES ARE ON THE FILE SERVER (or accompany Meabench)
www.t-square.gatech.edu
Potter Lab Resources / Potter / Douglas_Bakkum / CODE / Generic Closed-loop / .

I have used C/C++ to code my closed-loop experiments. The program *generic_closedloop.cpp* is provided as an example to create new closed-loop experiments. This program accesses Meabench's (Wagenaar, DeMarse et al. 2005) *spikedet* spike stream using *EasyClient.H* (a part of the Meabench download). Therefore, Meabench must be run, as done for normal recording, to create the spike stream. Neural spikes are read into *generic_closedloop.cpp* and will be used to determine motor output. For Meart, TCP/IP sockets were used to send the output to move the arms and to receive feedback. For simulated animats, the program (or a companion program run on the same computer) can hold its own virtual world. The feedback determined which subsequent stimulation was sent via another TCP/IP socket to the stimulation computer, completing the closed-loop. I used Daniel Wagenaar's stimulator which accepted text files coded in C as:

```
string COMMANDstring="";

sprintf( tmp,"%i  aux    %3i    0  \"
        \"      aux    %3i   400  \"
        \"      aux   127   800  \"
        \"      aux    %3i 1200  \"
        \"      aux     0 1600  \"
        \"      sw     e%i 1700  \"
        \"      dac %3.2f 1800  \"
        \"      dac %3.2f 2200  \"
        \"      dac     0 2600  \"
        \"      sw      0 2700\n\", \
IPI,-13*(CR(0)-1),-13*(CR(1)-1),EPOCH*3,CR(0)*10+CR(1),VOLT,-VOLT );

COMMANDstring= tmp;
```

where *IP* is the inter-pulse-interval in milliseconds, *CR* is a two value int array holding the Column and Row values of the electrode to stimulate, *EPOCH* is an int to encode user defined epochs, and *VOLT* is the stimulation voltage in Volts. The third column indicates the time of application of the line item command in microseconds. *aux* commands the auxiliary output on the stimulation board (to connect to BNC analog inputs of the MCard data acquisition card, used here for encoding the stimulation channel, time, and epoch. *dac* is the digital to analog voltage supply. *sw* switches the *dac* to the correct electrode.

The stimulator board runs on an rtlinux computer. I modified DW's *accstep* perl script to run a TCP/IP server, *stim_server.perl*, which needs to be started prior to attempting a connection.

Companion TCP/IP libraries were modified from open-source code found online, and need to be compiled before hand: *ClientSocket.cpp*, *ServerSocket.cpp*, *Socket.cpp*, *ClientSocket.H*, *ServerSocket.H*, *Socket.H*, *SocketException.H*. Note that other TCP/IP software can be used, such as built in code in Matlab.

To compile *generic_closed_loop.cpp*, include the *Makefile* in the current directory and type 'make' on the command line. Be sure to specify the correct version of Meabench in the *Makefile* such that the *EasyClient.H* can be found. Note, sometimes the Meabench library *libmea-1.0.##.so* needs to be copied into the directory for the spike stream to be accessed correctly, where '##' is the Meabench version number.

To run:

- 1- Run Meabench's *spikedet*
- 2- Run *stim_server.perl*
- 3- Run the closed-loop program adapted from *generic_closedloop.cpp*

If Meabench has been recently installed, an error may occur when running the program stating `libmea-1.0.##.so` cannot be found. In this case, the library needs to be added to a configuration file:

1. Login as the *root* (administrator) user.
2. In the `/etc/ld.so.conf` file, add the line `"/usr/local/lib"`. This is the location of the Meabench libraries.
3. Then on the command line, type `"/sbin/ldconfig -v"`.

A sample spike data file is provided for testing purposes. It consists of 2.5 hours of neural spikes from a tetanus experiment. The first 20 minutes has spontaneous activity, followed by an hour of $\frac{1}{2}$ Hz “probe” stimulation, then a large tetanus, then another hour of $\frac{1}{2}$ Hz probe stimulation. Use Meabench’s *replay* program to rerun the spike data file, and use Meabench’s *spikedet* program to create the spikestream accessed by *generic_closedloop.cpp* program. See the Meabench manual for further details.

The spike file is also located at:

www.t-square.gatech.edu

Potter Lab Resources / Potter / Douglas_Bakkum / CODE / 061605_D.spike

APPENDIX F

FORMATING AND ANALYZING LARGE SPIKE DATA SETS

MENTIONED FILES ARE ON THE FILE SERVER:

www.t-square.gatech.edu

Potter Lab Resources / Potter / Douglas_Bakkum / CODE / Format Spike File / .

Long experiments produce large spike files that can take a lot of time to load into Matlab for analysis, and an especially long time to format spike latency from a stimulation and to decode stimulation channel and epoch from the analog signal. Therefore, I made a C program to do this quickly, called *FormatSpikeData.cpp*, and a Matlab program to load the formatted data, called *loadspikeform.m*.

To compile *FormatSpikeData.cpp* in Linux, include the *Makefile* in a directory folder (along with all other files downloaded from the website) and type ‘make’ on the command line of a terminal window. Be sure to specify the correct version of Meabench that was installed on the computer in the *Makefile* such that the header files can be found during compilation. The companion library *Tools.H* contains the following functions:

- *int hw2cr(int hw)* to convert electrodes notation from hardware to column/row format.
- *int cr2hw(int cr)* to convert electrodes notation from column/row to hardware format.
- *int cleanlite(float *context)* to use cleanlite to remove ‘dirty’ spikes (likely false positives).
- *short context2probe(float *context)* to decode probe hardware electrode from analog channel.
- *short context2epoch(float *context)* to decode epoch from analog channel.

Tools.H also creates a data type to hold very large 3 dimensioned matrices called *Matrix3*.

To run, type in the terminal window:

```
./FormatSpikeData -o outputFileName -s in.spike
```

Optional flags are:

-a analogHW (default 60),

-p number (1 to find dAPs; default 0 to ignore dAPs; see Chapter 4)

The default value for *analogHW* is hardware channel 60, corresponding to analog channel 1. The default value for *number* is 0. Using the value of 1 will find and number dAPs (see Chapter 4; dAPs used to be called precisely-timed spikes, hence the *p* notation).

This program outputs a binary file that can be quickly loaded into Matlab with *loadspikeform.m*. To load into Matlab, in the Matlab editor type:

```
y = loadspikeform(outputFileName);
```

The following is now available:

y.T	spike time	(sec)
y.L	spike latency	(msec)
y.P_t	probe time	(sec)
y.C	spike hardware channel	(0 to 63)
y.H	height	(uVolts)
y.W	width	(msec)
y.Th	threshold	(uVolts)
y.E	epoch	(integer)
y.clid	clean spike ID	(1 or 0)
y.CLID	clean spike ID	(ID of clean spikes)
y.PTS	dAP number	(integer)
y.P_hw	probe hardware channel	(0 to 63)
y.P_num	probe number	(integer)

Each variable in the *y* structure is an array with values for each spike detected in Meabench (except *y.CLID*). Use the *find* command in Matlab to parse different sections of data. For example:

$$ID = find(y.T > 10 \& y.T < 30)$$

will assign the indices of spikes that occurred between 10 and 30 seconds in *ID*, and

$$ID = find(y.L > 0 \& y.L < 20)$$

will assign *ID* with the indices of spikes that occurred less than 20 milliseconds after a stimulation. *y.CLID* is equal to *find(y.clid == 1)*.

A sample spike data file is provided for testing purposes. It consists of 2.5 hours of neural spikes from a tetanus experiment. The first 20 minutes has spontaneous activity, followed by an hour of ½ Hz “probe” stimulation, then a large tetanus, then another hour of ½ Hz probe stimulation.

The spike file is also located at:

www.t-square.gatech.edu

Potter Lab Resources / Potter / Douglas_Bakkum / CODE / 061605_D.spike

APPENDIX G

GENERAL METHODS

G.1 Cell culturing

We have developed techniques to maintain neuronal cultures and conduct experiments for many months (Potter and DeMarse 2001b). Briefly, 50k cells from E18 rat cortices were dissociated using papain and trituration and plated over an approximately 3 mm diameter area on top of multi-electrode arrays (MEA; from Multi Channel Systems). A thin layer of polyethyleneimine followed by a 15 μ L drop of laminin were used for cell adhesion. The cultures were grown in 1 mL of DMEM containing 10% horse serum with glutamax, insulin, and sodium pyruvate additives. Experiments were conducted inside an incubator to control environmental conditions (35°C, 65% humidity, 9% O₂, 5%CO₂). The MEAs were sealed with a hydrophobic membrane (fluorinated ethylene–propylene) that is selectively permeable to O₂ and CO₂, and relatively impermeable to water vapor, bacteria, and fungus, allowing us to conduct long-term, non-invasive experiments.

G.2 Electrophysiology and data acquisition

Electrically evoked activity was induced using symmetric positive then negative voltage pulses of 400 μ s duration and 500 mV magnitude per phase (Wagenaar, Pine et al. 2004) using a custom built all-channel stimulation circuit board (Wagenaar and Potter 2004a). Data was collected through Multi Channel Systems' pre-amplifier and data acquisition card (MCCard), which had a 25 kHz sampling frequency and could accurately record microvolt signals. Data processing, visualization, artifact suppression, and spike detection were controlled using Meabench (Wagenaar, DeMarse et al. 2005) (www.its.caltech.edu/~pinelab/wagenaar/meabench.html). Neural action potentials were

detected if the absolute value of a voltage spike exceeded 5 standard deviations rms noise in amplitude. Algorithms to suppress the electrical artifacts from stimuli allowed us to detect neural action potentials ~2 ms after being evoked (Wagenaar and Potter 2002b).

G.3 Change to drift ratio, C/D

The change-to-drift ratio (C/D) is used to quantify the amount of plasticity that a neural activity statistic (such as the center of activity trajectory) detected between two periods of interest, for example between the periods before (P_i) and after (P_j) a plasticity-inducing stimulation. *Change*, C , is the average of the Euclidean distances of each activity statistic (x) in P_j to the mean activity statistic in P_i . *Drift*, D , is the average of the Euclidean distances of each activity statistic in P_i to its own mean activity statistic. The change-to-drift ratio, C/D , then quantifies an amount of neural plasticity. A ratio of 1 would indicate no plasticity.

$$\text{Change} = \frac{\sum_{j=1}^J (x_j - \bar{x}_i)}{J} \quad (\text{G.1})$$

$$\text{Drift} = \frac{\sum_{i=1}^I (x_i - \bar{x}_i)}{I} \quad (\text{G.2})$$

Note that the C/D nomenclature can lead to confusion. Change is often used as a synonym to plasticity, and drift can also refer to the inherent background plasticity ever present in neural activity. Finding a new name may be appropriate.

REFERENCES

- Abeles, M. (1991). Corticonics: Neural circuits of the cerebral cortex. Cambridge, Cambridge University Press.
- Ahissar, E., E. Vaadia, M. Ahissar, H. Bergman, A. Arieli and M. Abeles (1992). "Dependence of cortical plasticity on correlated activity of single neurons and on behavioral context." Science **257**(5075): 1412-1415.
- Ahissar, E., M. Abeles, M. Ahissar, S. Haidarliu and E. Vaadia (1998). "Hebbian-like functional plasticity in the auditory cortex of the behaving monkey." Neuropharmacology **37**(4-5): 633-655.
- Arkin, R. C. (1999). Behavior-based robotics. Cambridge, MIT Press.
- Bakkum, D. J., A. C. Shkolnik, G. Ben-Ary, P. Gamblen, T. B. Demarse and S. M. Potter (2004). Removing some 'a' from ai: Embodied cultured networks. Embodied artificial intelligence. F. Iida, L. Steels and R. Pfeifer, Springer-Verlag.
- Bao, S., E. F. Chang, J. D. Davis, K. T. Gobeske and M. M. Merzenich (2003). "Progressive degradation and subsequent refinement of acoustic representations in the adult auditory cortex." J Neurosci **23**(34): 10765-75.
- Baruchi, I. and E. Ben-Jacob (2004). "Functional holography of recorded neuronal networks activity." J Neuroinformatics **2**(3): 333-52.
- Bedlack, R. S., Jr., M. D. Wei, S. H. Fox, E. Gross and L. M. Loew (1994). "Distinct electric potentials in soma and neurite membranes." Neuron **13**(5): 1187-93.
- Beer, R. D. (2000). "Dynamical approaches to cognitive science." Trends in Cognitive Science **4**: 91-99.
- Berger, T. W. and D. L. Glanzman (2005). Toward replacement parts for the brain: Implantable biomimetic electronics as neural prostheses. Boston, The MIT Press.
- Bi, G. and M. Poo (1998). "Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type." J. Neurosci. **18**(10464-10472).
- Black, J. E., K. R. Isaacs, B. J. Anderson, A. A. Alcantara and W. T. Greenough (1990). "Learning causes synaptogenesis, whereas motor-activity causes angiogenesis, in cerebellar cortex of adult-rats." Proceedings of the National Academy of Sciences of the United States of America **87**(14): 5568-5572.
- Bliss, T. V. and T. Lomo (1973). "Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path." J Physiol **232**(2): 331-56.

- Blum, R. A., J. D. Ross, C. M. Simon, E. A. Brown, R. R. Harrison and S. P. DeWeerth (2003). A custom multielectrode array with integrated low-noise preamplifiers. Proceedings of the IEEE Engineering in Medicine and Biology Conference. M. Akay: 3396-3399.
- Boyden, E. S., F. Zhang, E. Bamberg, G. Nagel and K. Deisseroth (2005). "Millisecond-timescale, genetically targeted optical control of neural activity." Nat Neurosci **8**(9): 1263-8.
- Braitenberg, V. (1984). Vehicles, experiments in synthetic psychology. Cambridge, Mass., MIT Press.
- Brooks, R. A. (1991). "Intelligence without representation." Artificial Intelligence **47**: 139-159.
- Brooks, R. A. (1999). Cambrian intelligence: The early history of the new ai. Cambridge, MA, MIT Press.
- Brown, E. A., J. D. Ross, R. A. Blum, Y. Nam, B. C. Wheeler and S. P. DeWeerth (2007). Stimulus-artifact elimination in a multi-electrode array. IEEE Transactions on Biomedical Circuits and Systems.
- Brown, E. N., R. E. Kass and P. P. Mitra (2004). "Multiple neural spike train data analysis: State-of-the-art and future challenges." Nature Neuroscience **7**(5): 456-461.
- Buonomano, D. V. and M. M. Merzenich (1998). "Cortical plasticity: From synapses to maps." Annual Review of Neuroscience **21**: 149-186.
- Butera, R. J., C. G. Wilson, C. A. DelNegro and J. C. Smith (2001). "A methodology for achieving high-speed rates for artificial conductance injection in electrically excitable biological cells." Ieee Transactions on Biomedical Engineering **48**(12): 1460-1470.
- Caminiti, R., P. B. Johnson, Y. Burnod, C. Galli and S. Ferraina (1990). "Shift of preferred directions of premotor cortical cells with arm movements performed across the workspace." Exp Brain Res **83**(1): 228-32.
- Chao, Z. C., D. J. Bakkum, D. A. Wagenaar and S. M. Potter (2005). "Effects of random external background stimulation on network synaptic stability after tetanization: A modeling study." Neuroinformatics **3**(3): 263-80.
- Chao, Z. C., D. J. Bakkum and S. M. Potter (2007). "Region-specific network plasticity in simulated and living cortical networks: Comparison of the center of activity trajectory (cat) with other statistics." Journal of Neural Engineering **4**(3).

- Choi, Y., R. Powers, V. Vernekar, A. B. Frazier, M. C. LaPlaca and M. G. Allen (2003). High aspect ratio su-8 structures for 3-d culturing of neurons. ASME International Mechanical Engineering Congress and RD&D Expo, Washington, D. C.
- Chornoboy, E. S., L. P. Schramm and A. F. Karr (1988). "Maximum likelihood identification of neural point process systems." Biol Cybern **59**(4-5): 265-75.
- Clark, A. (1997). Being there: Putting brain, body, and the world together again. Cambridge, MIT Press.
- Clark, A. (2003). Natural-born cyborgs: Minds, technologies, and the future of human intelligence. New York, Oxford University Press, USA.
- Clarkson, T. G. (1999). Applications of neural networks in telecommunication. ERUDIT workshop on applications of computational intelligence techniques in telecommunication, London.
- Corner, M. A. and G. J. Ramakers (1991). "Spontaneous bioelectric activity as both dependent and independent variable in cortical maturation. Chronic tetrodotoxin versus picrotoxin effects on spike-train patterns in developing rat neocortex neurons during long-term culture." Ann N Y Acad Sci **627**: 349-53.
- Damasio, A. R. (1994). Descartes' error: Emotion, reason, and the human brain. New York, Gosset/Putnam Press.
- Daoudal, G. and D. Debanne (2003). "Long-term plasticity of intrinsic excitability: Learning rules and mechanisms." Learn Mem **10**(6): 456-65.
- Darbon, P., L. Scicluna, A. Tscherter and J. Streit (2002). "Mechanisms controlling bursting activity induced by disinhibition in spinal cord networks." Eur J Neurosci **15**(4): 671-83.
- David, O., D. Cosmelli and K. J. Friston (2004). "Evaluation of different measures of functional connectivity using a neural mass model." Neuroimage **21**(2): 659-73.
- DeMarse, T. B., D. A. Wagenaar, A. W. Blau and S. M. Potter (2001). "The neurally controlled animat: Biological brains acting with simulated bodies." Autonomous Robots **11**: 305-310.
- DeMarse, T. B., D. A. Wagenaar and S. M. Potter (2002). "The neurally-controlled artificial animal: A neural-computer interface between cultured neural networks and a robotic body." Society for Neuroscience Abstracts **28**: 347.1.
- Dennett, D. C. (1991). Consciousness explained. Boston, Little, Brown and Co.
- Diamond, M. (1990). Morphological cortical changes as a consequence of learning and experience. Neurobiology of higher cognitive function. A. B. Scheibel and A. F. Wechsler. New York, Guilford Press: 370.

- Dobkin, B. H. (2004). "Neurobiology of rehabilitation." Ann N Y Acad Sci **1038**: 148-70.
- Edelman, G. M. and G. Tononi (2000). A universe of consciousness: How matter becomes imagination. New York, Basic Books.
- Engert, F. and T. Bonhoeffer (1999). "Dendritic spine changes associated with hippocampal long-term synaptic plasticity." Nature **399**(6731): 66-70.
- Eytan, D., N. Brenner and S. Marom (2003). "Selective adaptation in networks of cortical neurons." J. Neurosci. **23** 9349-9356.
- Feininger, A. and N. E. Schlatter (2003). Structures of nature: Photographs by andreas feininger, University of Richmond Museums.
- Fellous, J.-M., P. H. E. Tiesinga, P. J. Thomas and T. J. Sejnowski (2004). "Discovering spike patterns in neuronal responses." The Journal of Neuroscience **24**(12): 2989-3001.
- Fields, R. D. (2005). "Myelination: An overlooked mechanism of synaptic plasticity?" Neuroscientist **11**(6): 528-31.
- Fodor, J. A. (1980). "Methodological solipsism considered as a research strategy in cognitive-psychology." Behavioral and Brain Sciences **3**(1): 63-73.
- Ganguly, K., L. Kiss and M. Poo (2000). "Enhancement of presynaptic neuronal excitability by correlated presynaptic and postsynaptic spiking." Nat Neurosci **3**(10): 1018-26.
- Geake, J. G. and P. W. Cooper (2003). "Implications of cognitive neuroscience for education." Westminster Studies in Education **26**(10): 7-20.
- Georgopoulos, A. (1994). Population activity in the control of movement. Selectionism and the brain. San Diego, Academic Press: 103-119.
- Goldstein, S. S. and W. Rall (1974). "Changes of action potential shape and velocity for changing core conductor geometry." Biophys J **14**(10): 731-57.
- Gross, C. G. (2000). "Neurogenesis in the adult brain: Death of a dogma." Nature Reviews Neuroscience **1**(1): 67-73.
- Gross, G. W. (1979). "Simultaneous single unit recording in vitro with a photoetched laser deinsulated gold multimicroelectrode surface." IEEE Transactions on Biomedical Engineering **26**: 273-279.
- Gross, G. W., A. N. Williams and J. H. Lucas (1982). "Recording of spontaneous activity with photoetched microelectrode surfaces from mouse spinal neurons in culture." J Neurosci Methods **5**: 13-22.

- Gross, G. W., B. K. Rhoades and J. K. Kowalski (1993). Dynamics of burst patterns generated by monolayer networks in culture. Neurobionics: An interdisciplinary approach to substitute impaired functions of the human nervous system. H. W. Bothe, M. Samii and R. Eckmiller. Amsterdam, North-Holland: 89-121.
- Gross, G. W. and J. M. Kowalski (1999). "Origins of activity patterns in self-organizing neuronal networks in vitro." Journal of Intelligent Material Systems and Structures **10**(7): 558-564.
- Hanley, Q. S., P. J. Verveer, M. J. Gemkow, D. Arndt-Jovin and T. M. Jovin (1999). "An optical sectioning programmable array microscope implemented with a digital micromirror device." J Microsc **196**(Pt 3): 317-31.
- Huang, H., S. L. Wolf and J. He (2006). "Recent developments in biofeedback for neuromotor rehabilitation." J Neuroengineering Rehabil **3**: 11.
- Hughes, R. (2007). The semi-living author: Post-human creative agency. Architecture and authorship. T. Anstey, K. Grillner and R. Hughes. London, Black Dog Publishing.
- Ikegaya, Y., G. Aaron, R. Cossart, D. Aronov, et al. (2004). "Synfire chains and cortical songs: Temporal modules of cortical activity." Science **304**(5670): 559-64.
- Ishibashi, T., K. A. Dakin, B. Stevens, P. R. Lee, et al. (2006). "Astrocytes promote myelination in response to electrical impulses." Neuron **49**(6): 823-32.
- Izhikevich, E. M. (2006). "Polychronization: Computation with spikes." Neural Comput **18**(2): 245-82.
- Jackson, A., J. Mavoori and E. E. Fetz (2006). "Long-term motor cortex plasticity induced by an electronic neural implant." Nature **444**(7115): 56-60.
- Jahanshahi, M., C. M. Ardouin, R. G. Brown, J. C. Rothwell, et al. (2000). "The impact of deep brain stimulation on executive function in parkinson's disease." Brain **123** (Pt 6): 1142-54.
- James, W. (1890). The principles of psychology. New York, H. Holt.
- Jimbo, Y., H. P. C. Robinson and A. Kawana (1998). "Strengthening of synchronized activity by tetanic stimulation in cortical cultures: Application of planar electrode arrays." IEEE Transactions On Biomedical Engineering **45**: 1297-1304.
- Jimbo, Y., T. Tateno and H. P. C. Robinson (1999). "Simultaneous induction of pathway-specific potentiation and depression in networks of cortical neurons." Biophysical Journal **76**: 670-678.

- Johansen-Berg, H., H. Dawes, C. Guy, S. M. Smith, D. T. Wade and P. M. Matthews (2002). "Correlation between motor improvements and altered fmri activity after rehabilitative therapy." Brain **125**(Pt 12): 2731-42.
- Kac, E. (1997). "Digital reflections: The dialogue of art and technology, special issue on electronic art." Art Journal **56**(3): 60-67.
- Kamioka, H., E. Maeda, Y. Jimbo, H. P. C. Robinson and A. Kawana (1996). "Spontaneous periodic synchronized bursting during formation of mature patterns of connections in cortical cultures." Neuroscience Letters **206**: 109-112.
- Kandel, E. R., J. H. Schwartz and T. M. Jessell (2000). Principles of neural science, McGraw-Hill.
- Kawaguchi, H. and K. Fukunishi (1998). "Dendrite classification in rat hippocampal neurons according to signal propagation properties. Observation by multichannel optical recording in cultured neuronal networks." Exp Brain Res **122**(4): 378-92.
- Kempermann, G. (2002). "Why new neurons? Possible functions for adult hippocampal neurogenesis." J Neurosci **22**(3): 635-8.
- Kudrimoti, H. S., C. A. Barnes and B. L. McNaughton (1999). "Reactivation of hippocampal cell assemblies: Effects of behavioral state, experience, and eeg dynamics." J Neurosci **19**(10): 4090-101.
- Kuhl, P. K., J. E. Andruski, I. A. Chistovich, L. A. Chistovich, et al. (1997). "Cross-language analysis of phonetic units in language addressed to infants." Science **277**(5326): 684-686.
- Latham, P. E., B. J. Richmond, P. G. Nelson and S. Nirenberg (2000). "Intrinsic dynamics in neuronal networks. I. Theory." Journal of Neurophysiology **83**(2): 808-827.
- Lee, A. and M. Wilson (2002). "Memory of sequential experience in the hippocampus during slow wave sleep." Neuron **36**(6): 1183-94.
- Lee, D. D. and H. S. Seung (1999). "Learning the parts of objects by non-negative matrix factorization." Nature **401**: 788-791.
- Liepert, J., H. Bauder, H. R. Wolfgang, W. H. Miltner, E. Taub and C. Weiller (2000). "Treatment-induced cortical reorganization after stroke in humans." Stroke **31**(6): 1210-6.
- Lipski, J. (1981). "Antidromic activation of neurones as an analytic tool in the study of the central nervous system." J Neurosci Methods **4**(1): 1-32.

- Madhavan, R., Z. Chao and S. Potter (2005). Spontaneous bursts are better indicators of tetanus-induced plasticity than responses to probe stimuli. The 2nd International IEEE EMBS, Arlington, VA, USA.
- Madhavan, R. (2007). Role of spontaneous bursts in functional plasticity and spatiotemporal dynamics of dissociated cortical cultures. Biomedical Engineering. Atlanta, Georgia Institute of Technology.
- Maeda, E., Y. Kuroda, H. P. C. Robinson and A. Kawana (1998). "Modification of parallel activity elicited by propagating bursts in developing networks of rat cortical neurones." European Journal of Neuroscience **10**: 488-496.
- Mainen, Z. F. and T. J. Sejnowski (1995). "Reliability of spike timing in neocortical neurons." Science **268**: 1503-1506.
- Manson, N. (2004). "Brains, vats, and neurally-controlled animats." Studies in the history and philosophy of biology and the biomedical sciences **35**(2): 249-268.
- Marom, S. and G. Shahaf (2002). "Development, learning and memory in large random networks of cortical neurons: Lessons beyond anatomy." Quarterly Reviews of Biophysics **35**: 63-87.
- Martinerie, J., C. Adam, M. Le Van Quyen, M. Baulac, et al. (1998). "Epileptic seizures can be anticipated by non-linear analysis." Nat Med **4**(10): 1173-6.
- McRae, E. (2004). A report on the practice of symbiotica research group in their creation of meart - the semi living artist. . Department of History and Philosophy of Science, University of Melbourne, <http://www.fishandchips.uwa.edu.au/project/publications.html>. Accessed November 2007.
- Merzenich, M. and K. Sameshima (1993). "Cortical plasticity and memory." Curr Opin Neurobiol. **3**(2): 187-96.
- Meyer, J. A. and S. W. Wilson (1991). From animals to animats: Proceedings of the first international conference on simulation of adaptive behavior. Cambridge, MIT Press.
- Milton, J. G. and M. C. Mackey (2000). "Neural ensemble coding and statistical periodicity: Speculations on the operation of the mind's eye." Journal of Physiology-Paris **94**: 489-503.
- Morris, C. G. (1973). Psychology: An introduction. New York, Appleton-Century-Crofts.
- Morris, K. F., D. M. Baekey, S. C. Nuding, T. E. Dick, R. Shannon and B. G. Lindsey (2003). "Invited review: Neural network plasticity in respiratory control." J Appl Physiol **94**(3): 1242-52.

- Morris, M. E. (2000). "Movement disorders in people with parkinson disease: A model for physical therapy." Phys Ther **80**(6): 578-97.
- Nadasdy, Z., H. Hirase, A. Czurko, J. Csicsvari and G. Buzsaki (1999). "Replay and time compression of recurring spike sequences in the hippocampus." Journal of Neuroscience **19**: 9497-9507.
- Nadasdy, Z. (2000). "Spike sequences and their consequences." Journal of Physiology-Paris **94**: 505-524.
- Nakanishi, K. and F. Kukita (1998). "Functional synapses in synchronized bursting of neocortical neurons in culture." Brain Research **795**(1-2): 137-146.
- O'Connel, M. J. (1974). "Search program for significant variables." Comp. Phys. Comm. **8**: 49.
- O'Connor, T. G., M. Rutter, C. Beckett, L. Keaveney and J. M. Kreppner (2000). "The effects of global severe privation on cognitive competence: Extension and longitudinal follow-up." Child Development **71**(2): 376-390.
- Okatan, M., M. A. Wilson and E. N. Brown (2005). "Analyzing functional connectivity using a network likelihood model of ensemble neural spiking activity." Neural Comput **17**(9): 1927-61.
- Patolsky, F., B. P. Timko, G. Yu, Y. Fang, et al. (2006). "Detection, stimulation, and inhibition of neuronal signals with high-density nanowire transistor arrays." Science **313**(5790): 1100-4.
- Perlmutter, J. S. and J. W. Mink (2006). "Deep brain stimulation." Annu Rev Neurosci **29**: 229-57.
- Pfeifer, R. and C. Scheier (1999). Understanding intelligence. Cambridge, Massachusetts, The MIT Press.
- Pfeifer, R. and J. Bongard (2007). How the body shapes the way we think: A new view of intelligence. Boston, MIT Press.
- Pine, J. (1980). "Recording action potentials from cultured neurons with extracellular microcircuit electrodes." Journal of Neuroscience Methods **2**: 19-31.
- Port, R. F. and T. van Gelder, Eds. (1995). Mind as motion: Explorations in the dynamics of cognition. Cambridge, MA, MIT Press.
- Potter, S. M. (1996). "Vital imaging: Two photons are better than one." Current Biology **6**: 1595-1598.

- Potter, S. M., S. E. Fraser and J. Pine (1997). "Animat in a petri dish: Cultured neural networks for studying neural computation." Proc. 4th Joint Symposium on Neural Computation, UCSD: 167-174.
- Potter, S. M. (2000). Two-photon microscopy for 4d imaging of living neurons. Imaging neurons: A laboratory manual. R. Yuste, F. Lanni and A. Konnerth. Cold Spring Harbor, CSHL Press: 20.1-20.16.
- Potter, S. M. (2001). Distributed processing in cultured neuronal networks. Progress in brain research: Advances in neural population coding. M. A. L. Nicolelis. Amsterdam, Elsevier. **130**: 49-62.
- Potter, S. M. and T. B. DeMarse (2001). "A new approach to neural cell culture for long-term studies." J. Neurosci. Methods **110**: 17-24.
- Potter, S. M., N. Lukina, K. J. Longmuir and Y. Wu (2001). "Multi-site two-photon imaging of neurons on multi-electrode arrays." SPIE Proceedings **4262**: 104-110.
- Potter, S. M., D. A. Wagenaar and T. B. DeMarse (2006). Closing the loop: Stimulation feedback systems for embodied mea cultures. Advances in network electrophysiology using multi-electrode arrays. M. Taketani and M. Baudry. New York Springer.
- Potter, S. M. (2007). What can ai get from neuroscience. 50th anniversary of artificial intelligence. M. Lungarella, F. Iida, J. Bongard and R. Pfeifer.
- R.A. Blum, J.D. Ross, E.A. Brown and S. P. DeWeerth (2007). An integrated system for simultaneous, multichannel stimulation and recording. IEEE Transactions on Circuits and Systems (TCAS).
- Ramachandran, V. S. and W. Hirstein (1998). "The perception of phantom limbs - the d.O. Hebb lecture." Brain **121**: 1603-1630.
- Reger, B. D., K. M. Fleming, V. Sanguineti, S. Alford and F. A. Mussa-Ivaldi (2000). Connecting brains to robots: The development of a hybrid system for the study of learning in neural tissues. Proc. of the VIIth Intl. Conf. on Artificial Life.
- Reich, D. S., J. D. Victor, B. W. Knight, T. Ozaki and E. Kaplan (1997). "Response variability and timing precision of neuronal spike trains in vivo." J Neurophysiol **77**(5): 2836-41.
- Rosenkranz, J. A. and A. A. Grace (1999). "Modulation of basolateral amygdala neuronal firing and afferent drive by dopamine receptor activation in vivo." J Neurosci **19**(24): 11027-39.
- Ruaro, M. E., P. Bonifazi and V. Torre (2005). "Toward the neurocomputer: Image processing and pattern recognition with neuronal cultures." IEEE Trans. Biomed. Eng. **52**: 371-383.

- Rutter, M. (1998). "Developmental catch-up, and deficit, following adoption after severe global early privation." Journal of Child Psychology and Psychiatry and Allied Disciplines **39**(4): 465-476.
- Schachter, S. C. and C. B. Saper (1998). "Vagus nerve stimulation." Epilepsia **39**(7): 677-86.
- Sharp, A. A., L. F. Abbott and e. al (1992). "The dynamic clamp: Computer-generated conductances in real neurons." Pre-print.
- Sheffler, L. R. and J. Chae (2007). "Neuromuscular electrical stimulation in neurorehabilitation." Muscle Nerve **35**(5): 562-90.
- Shkolnik, A. C. (2003). Neurally controlled simulated robot: Applying cultured neurons to handle and approach/avoidance task in real time, and a framework for studying learning in vitro. Dept. of Mathematics and Computer Science. S. M. Potter and J. Lu. Atlanta, Emory University.
- Simoni, M., G. Cymbaluyk, M. Sorensen, R. Calabrese and S. DeWeerth (2001). Development of hybrid systems: Interfacing a silicon neuron to a leech heart interneuron. Advances in neural information processing systems 13, nips 2000. T. K. Leen, T. G. Dietterich and V. Tresp. Boston, MIT Press: 173-179.
- Spitzer, N. C., L. N. Borodinsky and C. M. Root (2005). "Homeostatic activity-dependent paradigm for neurotransmitter specification." Cell Calcium **37**(5): 417-23.
- Sternberg, R. and W. Salter (1988). Conceptions of intelligence. Handbook of human intelligence. R. Sternberg. Cambridge, Cambridge University Press.
- Sur, M., P. Garraghty and A. Roe (1988). "Experimentally induced visual projections into auditory thalamus and cortex." Science **242**(4884): 1437-1441.
- Suzuki, I. and K. Yasuda (2007). "Detection of tetanus-induced effects in linearly lined-up micropatterned neuronal networks: Application of a multi-electrode array chip combined with agarose microstructures." Biochem Biophys Res Commun **356**(2): 470-5.
- Suzurikawa, J., M. Nakao, R. Kanzaki, Y. Jimbo and H. Takahashi (2007). Light-addressed stimulation and simultaneous calcium imaging for probing spatio-temporal activity of cultured neural network. IEEE EMBS Conference on Neural Eng (NER). Hawaii.
- Tabak, J. and P. E. Latham (2003). "Analysis of spontaneous bursting activity in random neural networks." Neuroreport **14**(11): 1445-1449.
- Tateno, T. and Y. Jimbo (1999). "Activity-dependent enhancement in the reliability of correlated spike timings in cultured cortical neurons." Biological Cybernetics **80**: 45-55.

- Trachtenberg, J. T., B. E. Chen, G. W. Knott, G. P. Feng, et al. (2002). "Long-term in vivo imaging of experience-dependent synaptic plasticity in adult cortex." Nature **420**(6917): 788-794.
- Turing, A. M. (1953). "The chemical basis of morphogenesis." Philosophical Transactions of the Royal Society of London B **237**: 37-72.
- Turrigiano, G. G. (1999). "Homeostatic plasticity in neuronal networks: The more things change, the more they stay the same." Trends Neurosci **22**(5): 221-7.
- Turrigiano, G. G. and S. B. Nelson (2000). "Hebb and homeostasis in neuronal plasticity." Current Opinion in Neurobiology **10**(3): 358-364.
- Uesaka, N., Y. Hayano, A. Yamada and N. Yamamoto (2007). "Interplay between laminar specificity and activity-dependent mechanisms of thalamocortical axon branching." J Neurosci **27**(19): 5215-23.
- Van Pelt, J., P. S. Wolters, M. A. Corner, W. L. C. Rutten and G. J. A. Ramakers (2004). "Longterm characterization of firing dynamics of spontaneous bursts in cultured neural networks." IEEE Trans. Biomed. Eng. **51**: 2051-2062.
- Varela, F. J., E. Thompson and E. Rosch (1993). The embodied mind: Cognitive science and human experience. Cambridge, Mass., MIT Press.
- Vera, A. H. and H. A. Simon (1993). "Situated action - a symbolic interpretation." Cognitive Science **17**(1): 7-48.
- Wagenaar, D. A. and S. M. Potter (2002). "Real-time multi-channel stimulus artifact suppression by local curve fitting." J. Neurosci. Methods **120**: 113-120.
- Wagenaar, D. A., J. Pine and S. M. Potter (2004). "Effective parameters for stimulation of dissociated cultures using multi-electrode arrays." J Neurosci Methods **138**(1-2): 27-37.
- Wagenaar, D. A. and S. M. Potter (2004). "A versatile all-channel stimulator for electrode arrays, with real-time control." J Neural Eng **1**(1): 39-45.
- Wagenaar, D. A., T. B. DeMarse and S. M. Potter (2005). Meabench: A toolset for multi-electrode data acquisition and on-line analysis. IEEE EMBS Conference on Neural Engineering. Arlington, VA.
- Wagenaar, D. A., R. Madhavan, J. Pine and S. M. Potter (2005). "Controlling bursting in cortical cultures with closed-loop multi-electrode stimulation." J Neurosci **25**(3): 680-8.
- Wagenaar, D. A. (2006). Development and control of epileptiform bursting in dissociated cortical cultures. Pasadena, California Institute of Technology.

- Wagenaar, D. A., J. Pine and S. M. Potter (2006a). "An extremely rich repertoire of bursting patterns during the development of cortical cultures." BMC Neurosci **7**: 11.
- Wagenaar, D. A., J. Pine and S. M. Potter (2006b). "Searching for plasticity in dissociated cortical cultures on multi-electrode arrays." J Negat Results Biomed **5**: 16.
- Watanabe, S., Y. Jimbo, H. Kamioka, Y. Kirino and A. Kawana (1996). "Development of low magnesium-induced spontaneous synchronized bursting and gabaergic modulation in cultured rat neocortical neurons." Neuroscience Letters **210**: 41-44.
- Weiler, I. J., N. Hawrylak and W. T. Greenough (1995). "Morphogenesis in memory formation - synaptic and cellular mechanisms." Behavioural Brain Research **66**(1-2): 1-6.
- Wilson, M. A. and B. L. McNaughton (1993). "Dynamics of the hippocampal ensemble code for space." Science **261**(5124): 1055-8.
- Wong, R. O. L., M. Meister and C. J. Shatz (1993). "Transient period of correlated bursting activity during development of the mammalian retina." Neuron **11**: 923-938.
- Yamamoto, T. and Y. Kuniyoshi (2002). Stability and controllability in a rising motion: A global dynamics approach. International Conference on Intelligent Robots and Systems (IROS), Lausanne, Switzerland.
- Zhang, W. and D. J. Linden (2003). "The other side of the engram: Experience-driven changes in neuronal intrinsic excitability." Nat Rev Neurosci **4**(11): 885-900.