Spinal Cord Injury Therapy through Active Learning

Thesis by

Thomas Desautels

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

California Institute of Technology

Pasadena, California

2014

(Defended July 1, 2013)

To Caroline. We made it!

Acknowledgments

There are, of course, many people who deserve to be acknowledged for helping me in my studies and in the completion of this dissertation. Foremost among those who have helped me academically is my advisor, Dr. Joel Burdick, who has been insightful, patient, and wise, providing mentorship in many different ways. I am profoundly indebted to him. Now, if I could only get him to read my emails. I also owe a particular debt to Dr. Andreas Krause, who has offered his considerable expertise on diverse matters in the field of machine learning and who worked closely with me to crack many of the theoretical problems addressed here during a research stay at ETH Zürich.

The experimental work presented here would also not have been possible without the contributions of Dr. Reggie Edgerton of UCLA; he has provided excellent advice and a superb collaborative environment. Among his research group, I would particularly like to recognize Jaehoon Choe and Parag Gad, who carried out the animal experiments described here. Their contributions, both conceptual and in blood, sweat, and wiped-up rat droppings, brought my work to life.

From Caltech, I would also like to thank Yanan Sui, who helped nurse the code through some of the rat experiments, and the rest of the members of the Burdick group, especially Jeff Edlund and Zhao Liu, for many productive discussions. I would also like to extend my thanks to Maria Koeper, the Burdick group's administrative assistant. Without Maria, we would never get anything done. As I move on to new pursuits, I will miss her, her invaluable aid, and her kindness.

Also from Caltech, Dr. Y. C. Tai's input has always been very valuable. His group fabricated the parylene microelectrode arrays; Monty Nandra spent many long hours making these devices. I am also very thankful for the advice and support of Dr. James Beck.

A number of ideas which inspired the work in Chapter 3 are from the work of Dr. Daniel Golovin, with whom Dr. Krause has had a number of collaborations, and to whom I owe a debt of thanks. The theoretical work presented in that chapter was partially supported by SNSF grant 200021_137971, NSF IIS-0953413, and DARPA MSEE FA8650-11-1-7156. My personal support has been derived from NIH/NINDS grant R01NS062009, NIH/NIBIB grant R01 EB007615, CDRF contract ESH1-2012(JB), and the ThinkSwiss Research Scholarship.

Even with all of the aid I received academically, I never would have stayed sane enough to make it through this process without the support of my friends and family. I would particularly like to acknowledge the many contributions of my parents, who set me on the path to learning, drove their sons to the library over and over again, and always tolerated my desire to disassemble everything in sight, even when it placed the house in peril. They have continued to provide advice, love, and support, even now that I am off in the wider world and consequently present less danger to their home and sanity.

No list of those who have supported me would be complete without my wife, Caroline. She has kept me from going off the deep end many a time, and without her ever so subtle urging, I would never have been able to make myself attack this seemingly insurmountable task. Every day, she makes me challenge myself and keep growing. My dear, I love you more than I can say.

Abstract

Therapy employing epidural electrostimulation holds great potential for improving therapy for patients with spinal cord injury (SCI) (Harkema et al., 2011). Further promising results from combined therapies using electrostimulation have also been recently obtained (e.g., van den Brand et al., 2012). The devices being developed to deliver the stimulation are highly flexible, capable of delivering any individual stimulus among a combinatorially large set of stimuli (Gad et al., 2013). While this extreme flexibility is very useful for ensuring that the device can deliver an appropriate stimulus, the challenge of choosing good stimuli is quite substantial, even for expert human experimenters. To develop a fully implantable, autonomous device which can provide useful therapy, it is necessary to design an algorithmic method for choosing the stimulus parameters. Such a method can be used in a clinical setting, by caregivers who are not experts in the neurostimulator's use, and to allow the system to adapt autonomously between visits to the clinic. To create such an algorithm, this dissertation pursues the general class of active learning algorithms that includes Gaussian Process Upper Confidence Bound (GP-UCB, Srinivas et al., 2010), developing the Gaussian Process Batch Upper Confidence Bound (GP-BUCB, Desautels et al., 2012) and Gaussian Process Adaptive Upper Confidence Bound (GP-AUCB) algorithms. This dissertation develops new theoretical bounds for the performance of these and similar algorithms, empirically assesses these algorithms against a number of competitors in simulation, and applies a variant of the GP-BUCB algorithm in closed-loop to control SCI therapy via epidural electrostimulation in four live rats. The algorithm was tasked with maximizing the amplitude of evoked potentials in the rats' left tibialis anterior muscle. These experiments show that the algorithm is capable of directing these experiments sensibly, finding effective stimuli in all four animals. Further, in direct competition with an expert human experimenter, the algorithm produced superior performance in terms of average reward and comparable or superior performance in terms of maximum reward. These results indicate that variants of GP-BUCB may be suitable for autonomously directing SCI therapy.

Contents

List of Figures

List of Tables