Chapter 6

Summary Conclusions

6.1 Conclusions

This dissertation develops two novel algorithms, GP-BUCB and GP-AUCB, provides new theoretical guarantees about their performance and that of a class of similar algorithms, and applies a variant of GP-BUCB to successfully manage an SCI therapy problem in real animals. This sort of autonomous, online direction of real EES-enabled SCI therapy has not previously been successfully executed by an algorithm. These results represent a substantial step toward both autonomously adaptive neurostimulators for SCI and lower-cost EES-based therapies for SCI patients. These techniques may also be applicable to other multi-electrode neurostimulation problems, such as deep brain stimulation and retinal prostheses.

Section 1.4 lays out the major goal of this dissertation, the creation, implementation, and testing of an algorithm capable of directing epidural electrostimulation-based SCI therapy. Specifically, it was necessary to develop an algorithm which could:

- Incorporate and exploit structural assumptions about the problem;
- Learn the responses of the spinal cord and muscles in a query-efficient manner; and
- Perform effective therapy, as measured by on-line metrics, even if observations must be made in batches or with a delay.

This dissertation satisfies each one of these requirements directly.

In Chapter 3, this dissertation develops the simple and efficient GP-BUCB and GP-AUCB algorithms. It also develops theoretical, high-probability bounds on the regret (i.e., sub-optimality in performance) of a general class of algorithms which includes GP-BUCB and GP-AUCB. This result, Theorem 1, is the main theorem of the work. Further, this dissertation presents an improved bound, Theorem 4, which is asymptotically independent of batch size or delay length, if the algorithm is initialized with an easily constructed set of observations of a finite, batch-size-dependent size. These bounds also provide high-probability convergence guarantees.

In simulation studies in several problems, including both synthetic and real data, the GP-BUCB and GP-AUCB algorithms attained similar performance to other, state-of-the-art batch algorithms for Bayesian optimization, which do not have theoretical bounds. Further, GP-BUCB and GP-AUCB performed comparably to the sequential GP-UCB algorithm, indicating that they overcome the disadvantages inherent in batching their actions and observations.

In Chapter 4, a variant of GP-BUCB successfully controls experiments in four SCI rats, with the goal of choosing appropriate bipolar stimuli on a multi-electrode array so as to maximize the average reward obtained. The reward is measured as the amplitudes of evoked potentials elicited in the rat's left tibialis anterior muscle. These experiments are selected in batches of five. In all four of these animals, the algorithm finds high-performing stimuli (i.e., stimuli which resulted in large evoked potentials, of putative therapeutic benefit), indicating that the algorithm effectively and efficiently searches the space of stimuli. As the experiments progress, the algorithm gradually builds up a human-interpretable set of predictions about the responsiveness of the spinal cord. This posterior over the evoked potential amplitudes enables the selection of appropriate experiments, both for the purpose of exploring the responses of the spinal cord and exploiting those stimuli which respond strongly. The algorithm competes directly with an expert human researcher in three of these animals. In these trials, the algorithm's performance under two different metrics is respectively superior or comparable to that of the expert human, indicating that the algorithm manages the exploration and exploitation of the evoked potentials appropriately. This finding indicates that, with appropriate reward metrics, variants of GP-BUCB or GP-AUCB would be effective for autonomously controlling more complex SCI therapy problems. This capability potentially represents a substantial savings in cost and great improvement in availability of EES-based SCI therapy.

Finally, Chapter 5 sets out a roadmap to human experiments, lays substantial practical and theoretical foundations, and describes pilot experiments conducted with two patients.

6.2 Future Work

The most important extension of the work presented in this dissertation will be the execution of experiments using GP-BUCB, GP-AUCB, or their derivatives to control EES for SCI therapy in humans. Chapter 5 describes the steps taken, in progress, and to be taken in the future toward human experiments. Also described in Chapter 5 are a number of important algorithmic and theoretical challenges which are associated with the needs of these experiments, but which also would have wider applications. Another important component of the SCI therapy program and the evaluation of the GP-BUCB and GP-AUCB algorithms is continued experimentation in animals, discussed below.

The first and most vital extension to the current animal experiments is to perform a parylene array experiment which lasts longer than the 3-4 experimental sessions and approximately 10 batches

which were achieved in the first two parylene array animals. Laying aside the practical difficulties associated with the ongoing development of this technology, an experiment of substantial duration in a parylene array should be pursued and would allow improved evaluation of several components of the current animal experiments; most importantly, how the algorithm copes with a comparatively large decision set.

Another important extension of the animal experiments would be multi-muscle simultaneous optimization, i.e., experiments in which the reward function measures characteristics of several muscles. In this setting, the suboptimality of individual muscles must be balanced so as to obtain the highest possible aggregate reward. In animal 5, several experimental days (P34, 36, and 37) were devoted to two-muscle testing. Some promising results were obtained, but, due a coding error on P34 and to the slow failure of the array (which began on P36), consistent muscle responses were not obtained from day to day. Repetition of these experiments in a future animal, or perhaps on the archived data obtained from the single-muscle experiments, would be of great interest.

In future experiments using the wired array preparation, it would also be quite interesting to examine multipolar configurations, i.e., those for which there are a variable number cathodes and anodes, but at least one of each. Disallowing monopolar configurations or the configuration in which all electrodes are off, and restricting all electrodes to have the same master voltage, this would give 328 viable configurations on a 7 electrode array, substantially more than the 42 bipolar configurations. This experiment would have the advantages of both a simple and durable experimental preparation and a relatively large space of electrode configurations over which to search. This stimulus set would also mimic devices, like the Medtronic RestoreAdvanced neurostimulator used by Harkema et al. (2011), which have a master voltage control and can have arbitrarily many active electrodes, an important step toward using those devices.

Another significant advance in the animal experiments would be to make the entire system fully autonomous; all experiments described in this thesis were performed using a human researcher to control the stimulator and the recording system, while the algorithm performed an executive or directing role. This architecture has a number of advantages, among them that the human experimenter provides a fail-safe with respect to data acquisition (e.g., if an element of the data processing fails, the observations of the human experimenter can often be used to reconstruct the missing information) as well as with respect to animal safety (if an unexpected condition arises, the human experimenter can terminate stimulation, or, if a stimulus known to be painful is requested, the human experimenter can refuse to perform the requested experiment). This latter failsafe must somehow be maintained in an automatic system. Creating an integrated system in which the algorithm is controlling the data acquisition in (nearly) real time would allow a substantial acceleration of the testing process, and would constitute a very substantial step toward the goal of autonomy, a crucial requirement for a home-use or implantable device. Fortunately, the parylene arrays are controlled

through a software interface, to which the GP-BUCB algorithm could be connected for both data acquisition and control. A safety system would have to be created, allowing the user to veto stimuli and terminate the delivery of stimuli which were distressing to the animal.

Clearly, future animal experiments should switch away from evoked potentials to a more complex motor response, in particular, standing or stepping. Unfortunately, these activities are quite complex, and quantifying their performance is both difficult and imprecise. Many attempts at doing so have been made (e.g. Basso et al., 1995; Antri et al., 2002), but none of them are entirely satisfactory for these purposes. For gait in particular, these techniques generally require extensive manual annotation of video recordings, which would necessitate a long feedback loop. If a rough, but sufficient analysis could be automated (e.g., pattern recognition on the raw motion capture trajectories, or on the EMG activity), however, these activities could be used in real time. This, however, requires a substantial effort in terms of building an effective automated grading system, possibly on the basis of a classifier, and for which careful feature selection will be necessary.