RESEARCH STATEMENT

Alan J. Lockett

Complex autonomous systems such as self-driving cars, unmanned aircraft, and humanoid robots promise great benefits for society in the near future. However, before these technologies can be widely adopted, general-purpose control methodologies for systems with heterogeneous state spaces and many degrees of freedom need to be developed. I believe that theoretical inquiry plays an important role in filling this gap. Therefore, my goal is to develop a theoretical and mathematical understanding of feedback controllers for such systems that enables versatile automated systems in human environments.

Many methods exist for controlling autonomous systems at a high level, but many of them are inefficient in practical terms, not interactive, or not well understood theoretically. For example, calculus of variations approaches can yield excellent results in planning trajectories for high-degree-of-freedom humanoid robots. However, it takes several minutes to plan an action such as moving a book from a shelf to a table, and these plans cannot be adjusted if, for example, a child steps in the robot's way in the middle of the action. Techniques such as artificial neural networks can be used to develop controllers that respond to changing sensor information online, and promising experimental results on perceptual tasks have been achieved using deep neural network models with millions of free parameters. However, there is no general theory for how to estimate parameters in these models efficiently. Additionally, because no effective methods for setting parameters in large networks exist, neural networks for control tasks are simpler than those used for state filtering, limiting applications.

Control methods such as neural networks are driven by a parameterized model for which the correct parameters cannot be determined analytically. Consequently, the parameters are estimated by optimizing some performance measurement. Optimization thus lies at the core of the search for effective controllers. In my dissertation, I studied the theoretical structure of stochastic iterative optimization methods in Hausdorff topological spaces, finding that the set of such methods can be embedded as a closed, convex subset of a Banach space. I examined optimizer performance in terms of linear functionals on this space obtained by integrating over the stochastic trajectories of the optimizers. I conjectured and subsequently proved an analytic expression for the optimal trajectory through the parameter space. This expression states that the best choice for the next iterate can be found by optimizing the conditional expectation of optimizer performance given proposed iterates. In some special cases, it may be possible to convert this analytic expression into a feasible calculation, which is one subject of my ongoing work. In the final part of the dissertation, I used qualitative insights drawn from this theory in order to propose a new method called *evolutionary annealing* that actively chooses trajectories in parameter space with good expected performance given prior parameter evaluations. I developed an instance of this method called *neuroannealing* to train neural network controllers; this method exhibited stronger ability to learn complex controllers than some previous methods. Overall, my dissertation incorporated the belief that good theory should lead to good practice and that theoretical inquiry should be employed as a tool for developing better practical applications. A 400-page monograph built on the dissertation is under preparation, to be published by Springer-Verlag.

To generalize the results of the dissertation, I developed a theoretical characterization of iterative optimization as a two-player game between optimization methods and problems. This setting more closely reflects that concept of an automaton trying to find the best course of action in an unknown and uncertain environment that may include other intelligent agents. The setting also generalizes objective functions of various types, including randomly chosen static objectives as well as objectives that vary stochastically, dynamically and even adaptively during the optimization process. Within this context, I demonstrated that the sets of players and problems can be modeled as closed, convex sets within two

Banach spaces that are isometrically embedded inside of each other's topological duals. As an unexpected consequence, there is no advantage to using stochastic optimization methods over a well-chosen deterministic optimization method. One interesting feature of approach is that the state and action spaces may consist of continuous signals and controls. Such large spaces have not typically been studied in the machine learning community, but by applying techniques from functional analysis, this work transfers many common intuitions about optimization in small spaces over to optimization in larger spaces, with implications for optimizing parameters of control models.

In order to apply theoretical ideas in practice, I applied for and received a two-year postdoctoral grant from the US National Science Foundation to study the integration of sensory input and actuator control using deep neural network controllers for a 43-degree of freedom humanoid iCub robot. Under this grant, I am working to train the robot to play chess on a typical chessboard. This task is important because it requires a high-degree-of-freedom robot to manipulate small objects to achieve a goal; consider a robotic chef or maid as two practical analogues. The robot must use stereo camera vision to extract the state of the chessboard, including its position in space and the location of pieces on the board. In addition, the robot must plan movements to pick up and place the chess pieces. The problem is challenging due to the high-dimensionality of both the perception and control tasks. I have formed a small team with two Ph.D. students in the host lab in order to complete the project and recently co-authored a paper describing a low-level controller based on the natural gradient that transitions through the state space more effectively than a Jacobian-based controller. A similar method was previously used by the co-author to efficiently solve inverse kinematics problems. In my view, practical research of this form provides an important test-bed to both stimulate and direct the theoretical work.

For the next phase of my career, I plan to establish a lab to study the theory behind the control of autonomous systems, especially those with heterogeneous sensors and numerous actuators. This theory will be oriented towards enabling practical applications in collaboration with colleagues in engineering and other departments. Ideas could flow in both directions. On the one hand, the theory might suggest the controllability of a particular system, and I would seek partners to implement and test the theory experimentally. On the other hand, partners might have built a system that has the potential to perform well, but for which the controls are poorly understood. On some occasions, both the theory and the experiments might be completed solely within the lab using third-party hardware.

Graduate students in this lab would study related topics with a focus on either theory or application. Students on a theoretical track might develop a framework for analyzing abstract heterogeneous control systems, or they might extend my previous results on optimization theory to more general or specific settings. Students focusing on applications might develop a formal analysis of how to control a particular system with a view towards robustness, stability, and flexibility. This analysis would then be verified experimentally. Other possible thesis topics include: 1) stability analysis of deep neural networks in perceptual tasks; 2) efficient estimation of inverse kinematics in high-degree-of-freedom systems; 3) a theoretical analysis of smoothing mechanisms to set velocity and acceleration when implementing discrete movement plans in high-dimensional configuration spaces; and 4) control methodologies for deep neural networks with millions of parameter. The goal is to have at least one graduate student in theory and one in applications at all times after the first year. All work within the lab will be unified under the topic of control of autonomous systems with an emphasis on theory.

I believe that this research could yield substantial practical results for controlling general-purpose autonomous systems within the next decade. Such advancements would benefit many individuals and society as a whole, in conformity with the mission of the research university.