Chapter 8

For Further Study

To deepen your understanding of the wide variety of neural network methods and applications, see [238], and for neural bio-foundations, see [268, 269, 558, 130]. More details on fuzzy control theory and applications are given in [412]. For more details on planning for autonomous robots, see [278] and a tutorial on model predictive control is given in [192]. The approach to the study of attentional systems was based on viewing the problem as a resource allocation problem and there is a large amount of literature on this topic. Here, the development was based on [418].

Neural Networks and Motor Control: We must emphasize that there are many topics in the area of neural networks that are not covered here. For instance, we do not discuss associative memories and recurrent networks, or Boltzmann machines. We refer the reader to [238, 262, 356] for treatment of these topics, or for references to sources where these topics can be studied. A recent theoretical treatment of modeling and analysis methods for theoretical neuroscience is in [130]. There, the authors cover a variety of topics including encoding and decoding information, neuron and neural network models, and adaptation and learning. By studying [130], the reader can gain a better understanding of "firing rate models," and "tuning curves," of neurons and hence, how accurate our neural network models are.

There has been extensive work on neurophysiological studies of motor control (e.g., on the hierarchy of the controls) in [206, 268, 269, 558], and the interested reader is recommended to see [274] and the current literature where close connections to control system methodology are found.

A method that has been popular in the control of robots has been the cerebellar model articulation controller (CMAC), which was first introduced in [9], and later applied in a different form in [362, 286]. Other neural network applications to robotic systems are contained in [279, 207].

Fuzzy and Expert Control: A recent, and relatively detailed, overview of the literature and methods of fuzzy control is given in [412]. You may also want to consider [155, 531, 262, 134]. Fuzzy logic is covered in [350, 280, 559, 440] and other places (here, we do not emphasize fuzzy logic formalisms for their own sake, but only use ideas from fuzzy logic needed for control systems, and in Section 11.6, for clustering methods in pattern recognition). For a detailed overview of how to perform stability analysis (absolute stability and via Lyapunov's first and second methods), how to study limit cycles via describing function analysis, and how to analyze tracking error for fuzzy control systems, see [263]. For an analysis of complexity of fuzzy systems and problems of the curse of dimensionality, see [235].

Expert systems are covered in many different books, so it is best, perhaps, to start with a general book on artificial intelligence, such as [444, 387, 136, 97, 82, 166]. The reader interested in studying stability analysis of expert control systems should consult [410, 338].

Planning Systems: The section on psychology and cognitive neuroscience of planning was based on [188, 206], and the work in [408] that was developed by using conventional control-theoretic concepts together with how the field of artificial intelligence views planning. Physiological foundations, particularly formation of cognitive maps for planning, are discussed in [239]. Planning systems are also discussed in [444, 387, 136] (a nice discussion on hierarchical and adaptive planning is in [136]) and in [137, 189]. Methods of anticipating the future (prediction) in microorganisms are discussed in [161]. These involve "circadian rhythms" and biological "clocks" [541]. Clearly, the use of regularly appearing events can be used to predict and hence react (early) quickly to stimuli, and these ideas are all related to the predictive nature of planning.

Planning in robotic applications has been studied extensively, and some papers you may want to study are [76, 77, 78, 44, 550], or the books [11, 258, 520]. For more information on path planning for robots, see [278, 90] where the "potential field method" is described (basically the method we introduce in this chapter where we use optimization over functions to guide a robot through a maze). In fact, some other functions that can be useful to build "obstacle functions" are given in [278]. A bibliography for heuristic search which has been used in planning is provided in [493] and an introductory treatment is given in [413]. The reader should be aware that there is a very large literature on combinatorial optimization methods, including many good books [400, 218] that provide methods to select a plan (e.g., by searching trees or graphs). One investigation on neural substrates for planning is given in [453].

A survey of model predictive control (which is also studied in Design Problem 6.2), what you can think of as the existing body of knowledge on planning methods in conventional control, is given in [192]. References on scheduling theory, which has many relevant concepts and techniques to planning, are provided below.

Attentional Systems: Discussions in cognitive neuroscience, clinical neuropsychology, and computational studies of attentional systems are contained in [404, 206, 522, 401]. The section in the chapter was developed using primarily [404, 206]. For other mathematical models of attentional systems that (unlike the one here) have been experimentally validated, at least to some extent, see [85, 482]. There is also a large literature on attention deficit disorder (a common disorder in children) that may provide insights into the operation and modeling of the human attentional system.

There are many books that treat the topic of scheduling and the topics treated there are relevant to both planning systems and attentional systems. Two books to consider are [419, 210]. Building on basic ideas in sequencing and scheduling [40, 109, 186], the approach to the development of attentional strategies here depends critically on the work in [418, 290] by using the time-based policies that were first studied in [81] and a discrete-event system [95, 244] theoretic framework for stability analysis that is described in detail in [411, 409].

Recently, some discussion on the use of attentional systems for control appeared in [11]. Earlier, it was shown in [296] how to augment adaptive controllers with attentional mechanisms.

The attentional systems approach of this part has been extended to the case where there are multiple agents cooperatively paying attention to multiple predators/prey in [212]. Also, the approach in [211] shows how such ideas can be used to allocate the focusing of multiple vehicle activities.

The section on multisensor integration is based on [339, 556, 158]. There are, however, many other papers on the theory and particularly the application of multisensor integration and management ideas. For more information on "world modeling," see [11].

Bayesian Belief Networks: One method that we did not introduce here, since to date it has found little use in control, is that of "Bayesian belief networks." This method has, however, found some use in a variety of engineering applications, such as diagnostic systems and decision-support applications. Moreover, it has potential for use in developing and implementing expert systems that reason under uncertainty and act as controllers. The reader interested in this method should consult [414, 96, 444] and the book [383] on learning Bayesian networks from data.

Part III

LEARNING

Part Contents

Sequence of Essential Concepts

- Learning theories from psychology and neuroscience form the foundations of biomimicry for incorporating learning into control and automation systems. The key underlying theories are classical conditioning, operant conditioning (reinforcement learning), and function approximation.

- Learning can be represented as an optimization process (e.g., gradient method) that involves sensing aspects of the environment and forming the associations or representations in memory that are best for trying to maximize performance and hence, survival chances. The type of association or representation sought depends on the situation. For instance, sometimes an organism learns to predict a stimulus from the occurrence of another stimulus (classical conditioning); other times, it learns the action that will result in a reward in a certain situation (operant conditioning, reinforcement learning).

- Neural networks and fuzzy systems can serve as tunable function approximators (interpolators) that hold associations or representations for making control decisions. We think of them as "approximators" for an unknown ideal mapping, one that we view as the target of our optimization process for incrementally learning the mapping. The neural or fuzzy systems can be trained (tuned) online to control a plant via reinforcement learning. For this, control decisions that lead to more reward (good closed-loop performance) are reinforced, and others are not. Under certain conditions, this iterative reinforcement leads to an appropriately shaped controller mapping after a long time period. If later the plant changes, then earlier "good actions" may not lead to rewards but other actions may lead to rewards. Then, the iterative reinforcement process reshapes the controller mapping in response to plant changes. Reinforcement learning control leads to what we call "adaptive control" for the plant.

- The key feature of using neural or fuzzy systems as tunable mappings for adaptive control is how to train them from data. There are a wide variety of training methods that you can use to learn functions from data. In linear least squares methods, we focus on tuning only a subset of the parameters of the approximator that enter linearly. Batch least squares methods focus on processing of data gathered offline, and recursive least squares methods focus on incrementally adjusting the approximator mapping as data are gathered in real time. The gradient methods that we discuss provide ways to tune all the parameters in an approximator structure, including the ones that enter in a nonlinear fashion, either in a batch or online mode.

- There are several fundamental issues involved in training approximators. The information in the training data is best in a specific form, but for control applications, it is often beyond the direct control of the training algorithm. The choice of the approximator structure, its complexity, and which parameters to tune has a significant impact on the quality of approximation and hence learning. The training method, initialization of the approximator parameters, and the parameters used to specify it (e.g., step size, termination method) can significantly affect learning performance. Moreover, there are fundamental issues to pay attention to in the training and testing process, including "generalization" (the ability of the trained approximator to respond similarly for similar inputs, or to discriminate between inputs that you would like it to), "local learning" (the ability of the method to learn the shape of the function in one region, and not to disturb what it has learned there when it learns in some other region), approximator "complexity" (while increases in approximator complexity generally give you an ability to approximate more complex functions, if the approximator is too complex, it can lead to poor generalization), and "overfitting" and "overtraining" (so that the approximator tries to match noise in the data, or tries to match the data gathered so closely that it does a poor job at generalization).

- The view of "learning as optimization" can be exploited to show how online optimization methods can be used to tune approximator structures to achieve adaptive control. The focus in such approaches is how to use data gathered online to shape functions (i.e., how to perform online function approximation). In the "indirect adaptive control" approach, the focus is on tuning approximators to match the nonlinear plant dynamics, and then using the approximations to specify the control inputs (using a "certainty equivalence approach"). In the "direct adaptive control" approach, the focus is on directly (i.e., without an approximation to the plant dynamics) tuning an approximator so that it approximates a controller that will achieve adaptive control. Optimization methods arising from learning (and foraging or evolutionary theory as studied later in this book) provide ways to adjust parameters for either the indirect or direct approaches. Since gradient-based adjustments reflect at least some biological learning/adaptation processes, online optimization approaches can also be viewed as biologically motivated. However, here we will depart somewhat from this focus to concentrate on what conventional optimization and approximation theory teaches us about the functionality and operation of adaptation mechanisms.

- Stability characterizes, for instance, how well a controller can achieve tracking of a desired reference input. Stable adaptive methods focus on how to construct, for example, online approximation-based controllers that will achieve stable and robust operation. While at the foundation of such approaches is Lyapunov stability theory, the methods essentially seek to minimize an instantaneous energy-based characterization of tracking performance. Hence, there is a close relationship to the online optimization methods. In stable adaptive control the focus is, however, on conditions under which the online optimization method will result in stable closed-loop control.