Foreword

This monograph addresses reinforcement learning with neural networks. Reinforcement learning is a machine learning technique that is used for solving sequential decision making problems. In essence, reinforcement learning is based on an iterative process of trial-and-error, very similar to how humans learn how to perform a task. A typical example of such a problem is game playing where, over many games, players learn moves that contribute to winning the game. A challenging caveat here is that during the game the utility of each individual move is unknown, and it is only at the end of the game that a winner is declared and the utility of all moves can be determined.

One way to implement reinforcement learning is by using neural networks, which can be trained with the temporal difference algorithm, a variation of the popular backpropagation algorithm. This approach implicitly solves the Bellman equation, allows for very large state spaces, and requires little *a priori* knowledge about the environment, which can plague other approaches such as models based on hidden Markov Models. However, there is one serious drawback to using neural networks for reinforcement learning: learning may not always convergence and a solution may not be found. For that reason, its extension to real-world problems is often not successful, and thus there are relatively few noteworthy applications.

The purpose of this monograph is to gain a better understanding of reinforcement learning in a variety of problems. This work contributes fundamental knowledge, basic science if you will, of what affects reinforcement learning and what contributes to a successful implementation. This work takes an empirical approach to understanding of the behavior and interactions between the two main components of a neural network based implementation of reinforcement learning: the learning algorithm and the functional representation of the learned knowledge.

This monograph serves as a textbook introduction to reinforcement learning with neural networks. A key original contribution of this work is the utilization of contemporary design of experiments methods, including a novel sequential experimentation procedure that finds convergent learning algorithm parameter subregions, and stochastic kriging for response surface metamodeling. Three single-agent problems are considered in this work: the mountain car problem, the truck backer-upper problem, and the tandem truck backer-upper problem. In the appendix, a similar approach to a game playing problem is illustrated as well.

The knowledge gained from this work provides insight as to what enables and what has an effect on successful reinforcement learning implementations so that this learning method can be applied routinely to more challenging problems.

Department of Industrial
and Systems Engineering
Rensselaer Polytechnic Institute
Clifton Park, New York—August 2014

Mark J. Embrechts, Professor

http://www.springer.com/978-3-319-12196-3

Design of Experiments for Reinforcement Learning
Gatti, C.
2015, XIII, 191 p. 46 illus., 25 illus. in color., Hardcover
ISBN: 978-3-319-12196-3