

Chapter 5.15

Machine Learning in Personalized Anemia Treatment

Adam E. Gaweda
University of Louisville, USA

ABSTRACT

This chapter presents application of reinforcement learning to drug dosing personalization in treatment of chronic conditions. Reinforcement learning is a machine learning paradigm that mimics the trial-and-error skill acquisition typical for humans and animals. In treatment of chronic illnesses, finding the optimal dose amount for an individual is also a process that is usually based on trial-and-error. In this chapter, the author focuses on the challenge of personalized anemia treatment with recombinant human erythropoietin. The author demonstrates the application of a standard reinforcement learning method, called Q -learning, to guide the physician in selecting the optimal erythropoietin dose. The author further addresses the issue of random exploration in Q -learning

from the drug dosing perspective and proposes a “smart” exploration method. Finally, the author performs computer simulations to compare the outcomes from reinforcement learning-based anemia treatment to those achieved by a standard dosing protocol used at a dialysis unit.

INTRODUCTION

Pharmacological treatment of chronic illnesses often resembles a trial-and-error process. A physician usually initiates the treatment with a standard dose and observes the patient for a response and / or dangerous side effects. If the response is not sufficient, the dose is increased. On the other hand, if a side effect occurs, the dose is decreased. A standardized dosing strategy for a specific drug is usually established during Phase II of a clinical trial. However, such a standardized strategy is

DOI: 10.4018/978-1-60960-818-7.ch5.15

usually based on population dose-response characteristics and may not lead to optimal outcomes on individual basis. Optimal outcomes are more likely to be achieved if the dosing is adjusted over time based on the individual response. This is especially the case for chronic illnesses, such as anemia of End Stage Renal Disease (ESRD).

Anemia of ESRD is one of the common conditions in patients receiving hemodialysis (Eschbach and Adamson, 1985). It is caused by insufficient red blood cell production due to unavailability of a hormone called erythropoietin, produced primarily by the kidneys. Untreated, anemia leads to a number of conditions including cardiovascular disease (Harnett et al., 1995), as well as decreased quality of life (Wolcott et al., 1989) and increased mortality (Lowrie et al., 1994). External administration of recombinant human erythropoietin (rHuEPO) is a standard form of treatment for the anemia of ESRD. In the United States, the anemia treatment is governed by the National Kidney Foundation guidelines which recommend that the Hemoglobin (Hgb), a surrogate marker of red blood cell level, in patients receiving rHuEPO be maintained between 11 and 12 g/dL. Based on these recommendations, dialysis units develop their own sets of dose adjustment rules governing rHuEPO dosages in relation to Hgb levels. However, these rules are again often derived from population outcome measures and do not always provide the optimal outcome. For example, it is known that at a given point in time 2/3 of the ESRD patients in the U.S. are outside of the NKF recommended target range (Lacson et al., 2003).

To facilitate the personalization of drug administration, we have proposed the application of a Machine Learning paradigm called Reinforcement Learning to anemia treatment (Gaweda et al., 2005). Reinforcement Learning mimics the goal-oriented skill acquisition performed by humans (Dayan and Balleine, 2002). It represents an intelligent system, referred to as an “agent,” which discovers what actions to take through interaction

with its environment, by trying them and receiving rewards, in order to drive the environment to a “goal” state. Agent’s choices may affect the environment not only immediately, but also over the long term. The first reported attempt at using Reinforcement Learning for drug dosing has been described by Sinzinger and Moore (2005). They show that a Reinforcement Learning based control of patient sedation performs better than conventional closed-loop controllers, such as PID.

In (Gaweda et al., 2005), we applied a Reinforcement Learning method called *Q*-learning, to anemia treatment in an on-line fashion. We concluded that in its simplest form this method’s performance is comparable to that of a rHuEPO dosing protocol used at our dialysis unit. However, a pure trial-and-error based learning was found insufficient to warrant any further improvement in the primary outcome, defined as maintaining the Hgb within the 11-12 g/dL range. To alleviate this drawback, we focused our attention on adding an element of heuristic supervision to guide the *Q*-learning (Gaweda et al., 2007). Others (Martin-Guerrero et al., 2006) have also applied the concept Reinforcement Learning to anemia treatment in an off-line fashion, using *Q*-learning to elicit optimal dosing strategy for rHuEPO from historical data records.

One of the key issues of Reinforcement Learning is the trade-off between exploration and exploitation (Sutton and Barto, 1998). To maximize the reward, the agent should take actions that are known to result in large rewards (exploitation). However, to discover such actions, the agent must try actions that it has not selected before (exploration). Our efforts so far, as well as those by Martin-Guerrero et al. (2006) focused on a purely exploiting form of Reinforcement Learning. An obvious drawback of the exploiting learning is the likelihood of achieving suboptimal results.

In this chapter, building upon the work presented in (Gaweda et al., 2007), we describe a Reinforcement Learning approach to anemia treatment which combines the guided *Q*-learning

9 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage:

www.igi-global.com/chapter/machine-learning-personalized-anemia-treatment/56205

Related Content

Robust Stability Self-Tuning Fuzzy PID Digital Controller

Ginalber Luiz de Oliveira Serra and Edson B. M. Costa (2018). *Applied Computational Intelligence and Soft Computing in Engineering* (pp. 141-154).

www.irma-international.org/chapter/robust-stability-self-tuning-fuzzy-pid-digital-controller/189319

A Formal Knowledge Retrieval System for Cognitive Computers and Cognitive Robotics

Yingxu Wang and Yousheng Tian (2013). *International Journal of Software Science and Computational Intelligence* (pp. 37-57).

www.irma-international.org/article/a-formal-knowledge-retrieval-system-for-cognitive-computers-and-cognitive-robotics/101317

The Logical Rules of Commonsense Reasoning

Xenia Naidenova (2010). *Machine Learning Methods for Commonsense Reasoning Processes: Interactive Models* (pp. 88-101).

www.irma-international.org/chapter/logical-rules-commonsense-reasoning/38481

Supervision of Industrial Processes using Self Organizing Maps

Ignacio Díaz, Abel A. Cuadrado, Alberto B. Diez, Manuel Domínguez, Juan J. Fuertes and Miguel A. Prada (2012). *Intelligent Data Analysis for Real-Life Applications: Theory and Practice* (pp. 206-227).

www.irma-international.org/chapter/supervision-industrial-processes-using-self/67450

Inconsistency-Induced Learning for Perpetual Learners

Du Zhang and Meiliu Lu (2011). *International Journal of Software Science and Computational Intelligence* (pp. 33-51).

www.irma-international.org/article/inconsistency-induced-learning-perpetual-learners/64178