Available online at www.sciencedirect.com
SCIENCE DIRECT®

Computers & Operations Research 34 (2007) 70–90

computers & operations research

www.elsevier.com/locate/cor

Neural network and regression spline value function approximations for stochastic dynamic programming

Cristiano Cervellera[a], Aihong Wen[b], Victoria C.P. Chen[b,*,1]

[a] *Institute of Intelligent Systems for Automation-ISSIA-CNR, National Research Council of Italy, Via De Marini 6, 16149 Genova, Italy*
[b] *Department of Industrial & Manufacturing Systems Engineering, The University of Texas at Arlington, Campus Box 19017, Arlington, TX 76019-0017, USA*

Abstract

Dynamic programming is a multi-stage optimization method that is applicable to many problems in engineering. A statistical perspective of value function approximation in high-dimensional, continuous-state stochastic dynamic programming (SDP) was first presented using orthogonal array (OA) experimental designs and multivariate adaptive regression splines (MARS). Given the popularity of artificial neural networks (ANNs) for high-dimensional modeling in engineering, this paper presents an implementation of ANNs as an alternative to MARS. Comparisons consider the differences in methodological objectives, computational complexity, model accuracy, and numerical SDP solutions. Two applications are presented: a nine-dimensional inventory forecasting problem and an eight-dimensional water reservoir problem. Both OAs and OA-based Latin hypercube experimental designs are explored, and OA space-filling quality is considered.

Keywords: Design of experiments; Statistical modeling; Markov decision process; Orthogonal array; Latin hypercube; Inventory forecasting; Water reservoir management

---

* Corresponding author. Tel.: +1 817 272 2342; fax: +1 817 272 3406.

E-mail addresses: cervellera@ge.issia.cnr.it (C.Cervellera), axw0932@omega.uta.edu (A.Wen), vchen@uta.edu (V.C.P. Chen).

[1] Supported by NSF Grants #INT 0098009 and #DMI 0100123 and a Technology for Sustainable Environment (TSE) grant under the US Environmental Protection Agency's Science to Achieve Results (STAR) program (Contract #R-82820701-0).

0305-0548/$ - see front matter © 2005 Elsevier Ltd. All rights reserved.
doi:10.1016/j.cor.2005.02.043

1. Introduction

The objective of dynamic programming (DP) is to minimize a "cost" subject to certain constraints over several stages, where the relevant "cost" is defined for a specific problem. For an inventory problem, the "cost" is the actual cost involved in holding inventory, having to backorder items, etc., and the constraints involve capacities for holding and ordering items. An equivalent objective is to maximize a "benefit", such as the robustness of a wastewater treatment system against extreme conditions. The *state* variables track the state of the system as it moves through the stages, and a *decision* is made in each stage to achieve the objective.

Recursive properties of the DP formulation permit a solution via the fundamental recurrence equation [1]. In stochastic dynamic programming (SDP), uncertainty is modeled in the form of random realizations of stochastic system variables, and the estimated expected "cost" ("benefit") is minimized (maximized). In particular, Markov decision processes may be modeled by a SDP formulation through which the state variables resulting from each decision comprise a Markov process. DP has been applied to many problems in engineering, such as water resources, revenue management, pest control, and wastewater treatment. For general background on DP, see Bellman [1], Puterman [2] and Bertsekas [3]. For some applications, see Gal [4], Shoemaker [5], White [6,7], Foufoula-Georgiou and Kitanidis [8], Culver and Shoemaker [9], Chen et al. [10], and Tsai et al. [11]. In high dimensions, traditional DP solution methods can become computationally intractable and attempts have been made to reduce the computational burden [12–14].

For continuous-state SDP, the statistical perspective of Chen et al. [15] and Chen [16] enabled the first truly practical numerical solution approach for high dimensions. Their method utilized orthogonal array (OA, [17]) experimental designs and multivariate adaptive regression splines (MARS, [18]). Although MARS has been applied in several areas (e.g., [19–22]), artificial neural network (ANN) models are much more prevalent in all areas of engineering (e.g., [23–30]). Given their popularity, an ANN-based SDP solution method would be more accessible to engineers. In this paper, we present:

- The statistical modeling approach to solving continuous-state SDP.
- Implementation of ANN models in continuous-state SDP as an alternative to MARS.
- Implementation of OA-based Latin hypercube (OA–LH) designs [31] as an alternative to pure OAs.
- Consideration of the "space-filling" quality of generated experimental designs. The designs are intended to distribute discretization points so that they "fill" the state space; however, generated designs of the same type do not necessarily have equivalent space-filling quality.
- Comparisons between ANN and MARS value function approximations, including methodological perspectives, computational considerations, and numerical SDP solution quality on two applications.

The novelty of our study is to simultaneously consider how the SDP solution is affected by different approximation methods and designs of different types and different space-filling qualities.

In the statistical perspective of continuous-state SDP, the two basic requirements are an experimental design to discretize the continuous state space and a modeling method to approximate the value function An appropriate design should have good space-filling quality, and an appropriate modeling method should be flexible in high dimensions, provide a smooth approximation, and be reasonably efficient. A review of several experimental designs and modeling methods is provided by Chen et al. [32] in the context of computer experiments. Of the possible modeling choices given in this review paper, only MARS, neural networks, and kriging satisfy the first two criteria. Some preliminary exploration of the kriging model