Using Landscape Theory to Measure Learning Difficulty for Adaptive Agents

Christopher H. Brooks[1] and Edmund H. Durfee[2]

[1] Computer Science Department
University of San Francisco
2130 Fulton St.
San Francisco, CA 94118
cbrooks@usfca.edu
[2] EECS Department
University of Michigan
1101 Beal Ave. Ann Arbor, MI 48109-2110
durfee@umich.edu

Abstract. In many real-world settings, particularly economic settings, an adaptive agent is interested in maximizing its cumulative reward. This may require a choice between different problems to learn, where the agent must trade optimal reward against learning difficulty. A landscape is one way of representing a learning problem, where highly rugged landscapes represent difficult problems. However, ruggedness is not directly measurable. Instead, a proxy is needed.

We compare the usefulness of three different metrics for estimating ruggedness on learning problems in an information economy domain. We empirically evaluate the ability of each metric to predict ruggedness and use these metrics to explain past results showing that problems that yield equal reward when completely learned yield different profits to an adaptive learning agent.

1 Introduction

In many problems, such as learning in an economic context, an adaptive agent that is attempting to learn how to act in a complex environment is interested in maximizing its cumulative payoff; that is, optimizing its performance over time. In such a case, the agent must make a tradeoff between the long-term value of information gained through learning and the short-term cost incurred in gathering information about the world. This tradeoff is typically referred to in the machine learning literature as the *exploration-exploitation tradeoff* [9]. If an agent can estimate the amount of learning needed to produce an improvement in performance, it can then decide whether to learn or, more generally, what it should learn. However, making this estimate requires that an adaptive agent know something about the relative difficulty of the problems it can choose to learn.

In this paper, we demonstrate how metrics from landscape theory can be applied to a particular agent learning problem, namely that of an agent learning

E. Alonso et al. (Eds.): Adaptive Agents and MAS, LNAI 2636, pp. 291–305, 2003.

the prices of information goods. A landscape is a way of representing the relative quality of solutions that lie near each other within some topology. We begin by describing our past work on the problem and appeal to a pictorial description to explain these results. Following this, we provide some background on landscapes and metrics for assessing their ruggedness, or difficulty. We then empirically evaluate two metrics, distribution of optima and autocorrelation, and show how these metrics can explain our previous results. We conclude by summarizing and discussing opportunities for future work.

2 Summarizing Price Schedule Learning Performance

In our previous work [1], we studied the problem of an adaptive agent selling information goods to an unknown consumer population. This agent acted as a monopolist and was interested in maximizing its cumulative profit. We assumed that the learning algorithm (amoeba [8], a direct search method) was a fixed feature of the agent. The adaptive agent's decision problem involved selecting a particular price schedule to learn, where this schedule served as an approximate model of consumer preferences. These schedules are summarized in Table 1

Table 1. This table presents the parameters of six pricing schedules, ordering in terms of increasing complexity. More complex schedules allow a producer to capture a greater fraction of potential consumer surplus by fitting demand more precisely, but require longer to learn, since they have more parameters.

Pricing Schedule	**Parameters**	**Description**
Pure Bundling	b	Consumers pay a fixed price b for access to all N articles.
Linear Pricing	p	Consumers pay a fixed price p for each article purchased.
Two-part Tariff	f, p	Consumers pay a subscription fee f, along with a fixed price p for each article.
Mixed Bundling	b, p	Consumers have a choice between a per-article price p and a bundle price. b
Block Pricing	p_1, p_2, m	Consumers pay a price p_1 for the first m articles ($m < N$), and a price p_2 for remaining articles.
Nonlinear Pricing	$p_1, ..., p_N$	Consumers pay a different price p_i for each article i.

We found that simple schedules were learned more easily, but yielded lower profit per period once learned. More complex schedules took longer to learn, but yielded higher profits per period after learning. We ran experiments comparing the performance of six different pricing schedules (a sample is shown in Figure 1) and found that moderately complex two-parameter schedules tended to perform