Available online at www.sciencedirect.com

ScienceDirect

Artificial Intelligence 171 (2007) 417–422

Artificial Intelligence

www.elsevier.com/locate/artint

Multi-agent learning for engineers ☆

Shie Mannor [a], Jeff S. Shamma [b,*]

[a] Department of Electrical & Computer Engineering, McGill University, Canada
[b] Department of Mechanical and Electrical Engineering, University of California, Los Angeles, USA

Received 28 April 2006; received in revised form 16 October 2006; accepted 18 December 2006

Available online 26 January 2007

Abstract

As suggested by the title of Shoham, Powers, and Grenager's position paper [Y. Shoham, R. Powers, T. Grenager, If multi-agent learning is the answer, what is the question? Artificial Intelligence 171 (7) (2007) 365–377, this issue], the ultimate lens through which the multi-agent learning framework should be assessed is "what is the question?". In this paper, we address this question by presenting challenges motivated by engineering applications and discussing the potential appeal of multi-agent learning to meet these challenges. Moreover, we highlight various differences in the underlying assumptions and issues of concern that generally distinguish engineering applications from models that are typically considered in the economic game theory literature.
© 2007 Elsevier B.V. All rights reserved.

Keywords: Multi-agent systems; Cooperative control; Distributed control; Learning in games; Nash equilibrium

1. Introduction

In this paper we address the question "if multi-agent learning is the answer, what is the question?" posed in [34] by looking at the engineering agenda. As opposed to the descriptive agenda that tries to explain micro or macro economic phenomena using simple learning rules, or the predictive agenda that tries to forecast what could happen, the engineering agenda concerns designing systems that would satisfy certain pre-specified performance criteria. The purpose of multi-agent learning from an engineering perspective is therefore to assist in the design of a complex system that includes multiple agents.

From an engineering point of view, as we argue below, one of the main benefits of multi-agent learning is its potential applicability as a design methodology for distributed control, which is a branch of control theory that deals with design and analysis of multiple controllers that operate together to satisfy certain design requirements.

In this paper, we motivate the use of multi-agent learning in these domains in Section 2 as a means to simplify the design process of a distributed control system while reducing the complexity of each controller, taking uncertainty into account, and allowing for distributed interactions between the controllers. Issues that distinguish the engineering view of multi-agent learning from the descriptive and predictive views are discussed in Section 3. We finally provide an outlook to the engineering agenda and the challenges that it poses to multi-agent learning in Section 4.

☆ Research supported by Natural Sciences and Engineering Research Council of Canada, NSF grant #ECS-0501394, AFOSR/MURI grant #F49620-01-1-0361, and ARO grant #W911NF-04-1-0316.

* Corresponding author.
E-mail addresses: shie@ece.mcgill.ca (S. Mannor), shamma@ucla.edu (J.S. Shamma).

0004-3702/$ – see front matter © 2007 Elsevier B.V. All rights reserved.
doi:10.1016/j.artint.2007.01.003

2. Why multi-agent learning?

The problem of designing optimal (or even just reasonable) distributed control systems is notoriously difficult. Examples of complex engineering problems include scheduling in manufacturing systems [2,16,24,25], routing in data networks [3,4,26,30,31], and command and control of networked forces in adversarial environments [1,9,29]. These applications entail a collection of dispersed interacting components that seek to optimize a global collective objective through local decision making. The task is complicated by limited communication capabilities, local and dynamic information, faulty components, and an uncertain, if not hostile, environment. In general, it is not feasible to pass all information to a command center that could process this information and disseminate instructions. Furthermore, even if this were possible, the complexity of the overall system would make the problem of constructing a centralized optimal policy intractable.

It is both the complexity and distributed nature of these problems that motivate the use of game theoretic methods. A multi-agent viewpoint lets one look at an overall systems as a collection of simpler interacting components. Note that this means *choosing* to impose a multi-agent structure as a *design approach*. The result is that the decision making process for any single component is dictated by an optimization problem that is greatly simplified as compared to the centralized problem, but coupled to the decisions of other interconnected components. Accordingly, a Nash equilibrium reflects an optimality condition from the perspective of each individual component, but need not reflect optimal operation from as a collective. Nonetheless, the possibility of the system to self-organize into a suboptimal Nash equilibrium is less daunting than the prospect of constructing a centralized optimal policy.

3. Descriptive, predictive, or engineered?

It is safe to say that much of the research in multi-agent learning has its roots in the economic game theory literature. Accordingly, it is also safe to say that this literature did not intend to offer a design methodology for engineered systems. This does not mean that material, e.g., as in the many excellent monographs on learning in games [15,37, 38] or evolutionary games [32,35], cannot be a source of methods for engineered systems. Rather, it underscores the importance of recognizing "what is the problem?" when considering this material for a multi-agent approach for designing engineering systems.

That stated, this section presents selected issues that can be distinguishing characteristics of multi-agent learning in engineered designs. In particular, we will see that various notions that play a descriptive or predictive role in the economics literature become design considerations when considering a multi-agent learning approach.

3.1. Defining the game

We have suggested that multi-agent learning may be an effective approach to the problems described earlier. And yet, we have yet to define a specific game in which to apply multi-agent learning methods. Recall that an important motivation for taking a multi-agent perspective was alleviating complexity. This means that the basic elements of players, strategy spaces, and player utilities are all design considerations when imposing a multi-agent framework.

The operations of agents in reactive environments that change with time is an important reality a design methodology must face. The type of game that is chosen to capture such dynamics can be either a one-shot game, a repeated game, or a stochastic game. Choosing a one-shot game implies that we ignore the dynamics of the problem. A repeated game offers a richer structure by reflecting the consequences of a strategy over multiple time steps. This suggests using regret minimizing techniques such as [7]. Stochastic games are a more natural model to model dynamics. Learning in this context is quite complex, as shown by [20], in the context of getting to a Nash equilibrium and by [28] in the context of regret minimization.

Finally, once players and strategy spaces have been specified, another design consideration is defining agent utility functions. As with most design specifications, there is considerable latitude in specifying suitable utility functions. Even in the ideal case of an agreeable centralized objective, there are important considerations in "distributing" this objective among the different players. Ref. [36] contains a general discussion, which is applied to vehicle target assignment in [6].

Download English Version:

https://daneshyari.com/en/article/377476

Download Persian Version:

https://daneshyari.com/article/377476

Daneshyari.com