

What evolutionary game theory tells us about multiagent learning

Karl Tuyls^{a,*}, Simon Parsons^b

^a *Institute for Knowledge and Agent Technology, Maastricht University, The Netherlands*

^b *Department of Computer and Information Science, Brooklyn College, City University of New York, 2900 Bedford Avenue, Brooklyn, 11210 NY, USA*

Received 1 May 2006; received in revised form 8 January 2007; accepted 9 January 2007

Available online 26 January 2007

Abstract

This paper discusses *If multi-agent learning is the answer, what is the question?* [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] from the perspective of evolutionary game theory. We briefly discuss the concepts of evolutionary game theory, and examine the main conclusions from [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] with respect to some of our previous work. Overall we find much to agree with, concluding, however, that the central concerns of multiagent learning are rather narrow compared with the broad variety of work identified in [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].

© 2007 Elsevier B.V. All rights reserved.

Keywords: Evolutionary game theory; Replicator dynamics; Multiagent learning

1. Introduction

In *If multi-agent learning is the answer, what is the question?* by Shoham, Powers and Grenager [20], the authors make a valiant effort to analyse the state of the field of multiagent learning, to summarise the results that have been achieved within the field, to discern the major research directions that have been followed, and to issue a call to arms. In short, Shoham et al. conclude that most work in multiagent learning can be placed into one of five “buckets” each of which is associated with a distinct research agenda (these descriptions are taken directly from the “caricatures” in Section 5 of [20]):

- (1) Computational: learning algorithms are a way to compute the properties of a game.
- (2) Descriptive: learning algorithms describe how natural agents learn in the context of other learners.
- (3) Normative: learning algorithms give a means to determine which sets of learning rules are in equilibrium with one another.
- (4) Prescriptive, cooperative: learning algorithms describe how agents should learn in order to achieve distributed control of dynamic systems.
- (5) Prescriptive, non-cooperative: learning algorithms describe how agents should act to obtain high rewards.

* Corresponding author.

E-mail addresses: k.tuyls@micc.unimaas.nl (K. Tuyls), parsons@sci.brooklyn.cuny.edu (S. Parsons).

In addition [20]:

Not all work in the field falls into one of these buckets. This means that either we need more buckets, or some work needs to be revisited or reconstructed so as to be well-grounded.

The authors also point out that research in multiagent learning is often unclear about which of these agendas it is pursuing—and that, in contrast one needs to be very clear about one’s aims—that the field cannot progress by defining “arbitrary learning strategies and analys(ing) whether the resulting dynamics converge in certain cases to a Nash equilibrium or some other solution concept”, and that the field needs to take evaluation more seriously, especially on some set of standard problems.

We basically agree with all of the points that we have quoted above, and below will expand on those that we feel our background leaves us best qualified to discuss. However, we do have one major point of disagreement with the views expressed in [20]. The point we disagree with is the idea that there is some overall *taxonomy* of research agendas into which all work can be slotted. This is not explicitly stated—all that Shoham et al. say is what we quoted above, that there are five distinct agendas to which some additional ones may need to be added—but the existence of an underlying taxonomy into which these additional categories can be slotted seems to be implied.

Our objection is neatly summarised by the following passage from McWhorter’s *The Power of Babel: A natural history of language* [15] in which the author tries to express how the original language, spoken by our common ancestors, became the many thousands of languages that descended from it. He starts by saying “I have implied that speech varieties have developed like a bush, starting from a single sprout and branching in all directions, each branch then developing subbranches, and so on . . .” before going on to explain that the inter-relationships between languages, the constant process of adoption of terms from one language into another, and the formation of dialects, creoles and intertwined languages means that [15, page 94]:

we might do just as well with another analogy, say stewing a good spring lamb stew without much juice (because the juice messes up the analogy). Clearly, one can distinguish the lamb from the peas from the potatoes from the carrots from the leeks from the rosemary leaves. Yet all of these ingredients, if it’s a good stew, are suffused with juice and flavor from the other items as well. Every now and then, you even encounter a piece of something that, covered with liquid and cooked out of its original shape and consistency, you have to work to figure out the original identity of . . . Overall, there is nothing in this stew that tastes or even looks like it would if you had just dumped it into a pot of boiling water by itself.

It seems to us that multiagent learning is such a stew and though it is very helpful to identify the various ingredients—especially if, to stretch the metaphor, some of them would be better taken out of the pot—to concentrate on the constituents misses some of the essence. It is the places in which the agendas that make up the stew meld into new things, things that cannot be put into a taxonomy because they are a mixture, that we often find the most interesting work.¹

That said, we should reiterate that we are largely in agreement with the agendas identified in [20], and in the next section amplify our agreement by examining all five agendas through the lens of evolutionary game theory (EGT), which is an area of multiagent learning in which we have been working, and one that is not much discussed in [20]. Following that exploration, we return to our point about the interplay between agendas, illustrating our discussion with some of our recent work.

2. The five research agendas from the perspective of evolutionary game theory

In this section we will use EGT to illustrate our reaction to the analysis of multiagent learning presented in [20]. To do this, we first consider what each of the five research agendas means in terms of EGT, taking them in the order that best fits our argument, rather than the order in which they are presented by Shoham et al.

¹ Though one has to be especially careful to be clear what one is doing at these junctures.

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات