Available online at www.sciencedirect.com

ScienceDirect

Cognitive Systems Research 52 (2018) 429–449

Cognitive Systems RESEARCH

www.elsevier.com/locate/cogsys

Evolution of learning strategies in changing environments

John A. Bullinaria

School of Computer Science, University of Birmingham, Birmingham B15 2TT, UK

Received 24 January 2018; received in revised form 13 June 2018; accepted 17 July 2018
Available online 24 July 2018

Abstract

Learning is an important aspect of cognition that is crucial for the success of many species, and has been a factor involved in the evolution of distinct patterns of life history that depend on the environments in question. The extent to which different degrees of social and individual learning emerge follows from various species-dependent factors, such as the fidelity of information transmission between individuals, and that has previously been modelled in agent-based simulations with meme-based representations of learned knowledge and behaviours. A limitation of that previous work is that it was based on fixed environments, and it is known that different learning strategies will emerge depending on the variability of the environment. This paper will address that limitation by extending the existing modelling framework to allow the simulation of life history evolution and the emergence of appropriate learning strategies in changing environments.

Keywords: Social learning; Individual learning; Memes; Evolution; Life history; Changing environments

1. Introduction

There is considerable evidence of environmental variability, such as climate change, driving the evolution of versatile behaviours (e.g., Potts, 1996, 1998, 2013; Grove, 2011). In particular, the prevailing environment changes are known to affect the learning strategies that evolve (e.g., Rogers, 1988; Acerbi & Parisi, 2006; McElreath & Strimling, 2008; Ehn & Laland, 2012). There are several potential approaches for studying this issue, including empirical studies of real populations, detailed grounded models of real populations, abstract mathematical models, and simplified agent-based simulations. Then within the realm of agent-based simulations are distinct approaches for representing the learned behaviours, such as in terms of neural network weights or as simplified sets of abstract memes. All these approaches have their own advantages and disadvantages, and comparing their results is important for exploring their limitations and arriving at reliable conclusions. An agent-based simulation framework based on abstract memes has previously been formulated (Bullinaria, 2017) and shown to provide a powerful approach for modeling the interaction of direct individual and social learning in general, and in the specific context of life history evolution, but that has so far only been developed and explored for static environments. The aims of this paper are to extend that framework to accommodate changing environments and to investigate how the form and degree of variability affects what emerges. In particular, it will explore the interaction between social and individual learning, the optimal learning strategies that emerge, and how they can affect the resulting patterns of life history evolution.

The modeling framework here is deliberately simplified to allow computationally-feasible simulations, with all the environmental and behavioural details totally abstracted

E-mail address: j.a.bullinaria@cs.bham.ac.uk

https://doi.org/10.1016/j.cogsys.2018.07.024
1389-0417/

out, limited meme-sets, and small population sizes. Nevertheless, the results found previously for static environments (Bullinaria, 2017) indicate that it is sufficient to capture enough relevant details to provide useful models. Moreover, the approach is purposely kept very general so it has the important advantage of being easily applicable to modeling a diverse range of species with only a few changes to parameter settings being required. To establish whether the same general framework is also sufficient for modelling species in changing environments, a series of simulations need to be carried out to establish suitable implementational details and parameter value ranges for the models. In parallel with that, simulation results need to be generated to test its correspondence with known biological populations, existing mathematical models and more-grounded simulations, and reasonable intuitions about what should occur. The remainder of this paper will present the results of doing that.

The interaction of learning and evolution has already been extensively studied in the past, and it is well known why at least some lifetime learning is generally required, even when the environment is stable and perfect genetic assimilation of previously learned behaviour is possible. For example, most species grow after birth and precise good adult behaviours will therefore often be inappropriate for new-born individuals, in the same way that learned adult neural connection strengths will generally be suboptimal for new-borns (Bullinaria, 2003). Similarly, optimal innate behaviours for a particular environment will be inappropriate for spatially or temporally changing environments (Crispo, 2007). What is not so clear is which forms of learning are most appropriate for particular forms of environmental changes. That is one of the key questions this paper needs to address, specifically in relation to aspects of behaviour for which genetic acquisition is either inappropriate or impossible. In principle, direct individual learning based on the current environment should work well under all conditions, but that can be costly in terms of time and effort, and may be unreliable or even deadly in some circumstances, so various forms of social learning are often more effective and have consequently evolved for some species when sufficient transmission fidelity is possible (e.g., Boyd & Richerson, 1985; Tomasello, 1999; Kameda & Nakanishi, 2002; Rendell et al., 2010).

The issue of different learning strategies emerging in environments which change significantly over individual lifetimes, or over generational timescales, has previously been explored in the simplified mathematical models of Feldman, Aoki, and Kumm (1996), Wakano and Aoki (2006) and Ehn and Laland (2012), the neural network controlled agent simulations of Acerbi and Parisi (2006), and in the individual-based stochastic models of Whitehead (2007) and Whitehead and Richerson (2009), and different time-weighted learning strategies are known to be superior in such situations (Rendell et al., 2010). One of the key advantages of agent-based simulation approaches is that such strategies can be encoded genetically and optimized by evolution along with everything else. Moreover, factors that vary with location, such as environmental conditions, are also easily simulated in this approach, and incorporating them into the models can allow the emergence of cultural differences between groups to be studied in more detail (Henrich & Boyd, 1998). Modelling all the relevant issues clearly requires a well-defined, flexible and tested simulation framework that can accommodate the various interacting factors and trade-offs as accurately as possible in line with particular realistic scenarios, and facilitate reliable and reproducible comparisons with minimal confounding factors. Formulating and testing such a modeling framework based on agents and memes was presented in an earlier paper (Bullinaria, 2017) and the same approach will be followed here to accommodate a range of forms of environmental variability.

In the next section the general meme- and agent-based modeling framework of Bullinaria (2017) will be outlined and extended to incorporate changing environments. The following section then presents a series of simulation results obtained using it that explore how environmental variability affects the trade-off between social and individual learning, the resulting probabilities of population collapse and strategies for avoiding it, the potential social learning strategies and how they may evolve, and some representative life history evolution issues. The paper ends with some discussion and conclusions.

2. Simulation framework

The underlying agent and meme based simulation framework to be explored here for changing environments has already been described and tested in some detail for static environments (Bullinaria, 2017), so it should be sufficient to just outline the resulting framework here. The general idea is to maintain a population of individual agents that are each specified by a set of relevant innate parameters (such as brain size, individual learning rate, social learning rate, etc.) and have each acquired some subset of the available "memes" that represent the behaviours they have adopted for surviving in their environment. The overall performance of each individual is the sum of the "performance contributions" associated with their set of learned memes. Of course, mental representations are much more complex than small discrete sets of memes that can be copied directly, and memes will not really contribute to performance or fitness in a simple additive manner, but this simplification is a useful starting point for simulations that simply need a convenient way to "keep score" of how individually and socially learned information and behaviours are affecting individual performances (Aunger, 2002).

2.1. Meme-based agent simulations

The crucial underlying assumption here is that all the relevant knowledge and behaviours can be sufficiently well