BioSystems 103 (2011) 45–56

Contents lists available at ScienceDirect

BioSystems

journal homepage: www.elsevier.com/locate/biosystems

Evolving experience-dependent robust behaviour in embodied agents

Jose A. Fernandez-Leon*

Centre for Computational Neuroscience and Robotics (CCNR) – University of Sussex, Brighton BN1 9QG, United Kingdom

ARTICLE INFO

Article history:
Received 14 August 2010
Received in revised form
19 September 2010
Accepted 22 September 2010

Keywords:
Behavioural robustness
Animat
Evolutionary robotics
Genetic algorithm
Walking behaviour
Transient dynamics

ABSTRACT

In this work, based on behavioural and dynamical evidence, a study of simulated agents with the capacity to change feedback from their bodies to accomplish a one-legged walking task is proposed to understand the emergence of coupled dynamics for robust behaviour. Agents evolve with evolutionary-defined biases that modify incoming body signals (sensory offsets). Analyses on whether these agents show further dependence to their environmental coupled dynamics than others with no feedback control is described in this article. The ability to sustain behaviours is tested during lifetime experiments with mutational and sensory perturbations after evolution. Using dynamical systems analysis, this work identifies conditions for the emergence of dynamical mechanisms that remain functional despite sensory perturbations. Results indicate that evolved agents with evolvable sensory offset depends not only on where in neural space the state of the neural system operates, but also on the transients to which the inner-system was being driven by sensory signals from its interactions with the environment, controller, and agent body. Experimental evidence here leads discussions on a dynamical systems perspective on behavioural robustness that goes beyond attractors of controller phase space.

1. Introduction

Recently, Macinnes and Di Paolo (2006) discussed the role that 'the indirect experience' of sensing the environment has to the production of agent's behaviours. They analyse the process of stimuli recognition in simulated agents and the 'meaning' that agents impose to achieve dynamical engagement. By meaning, they refer to the selections that agents make from their own stimuli and therefore find their own reference in how to process external information. Despite the importance of such research, we have very little idea about how an agent's own experience shapes sensory signals at the neuronal level and the effect of this shaping on robust behaviour. The majority of work in this area neither explains how these mechanisms emerge from sensorimotor interactions, nor analyses whether it promotes behaviour production and *robust* traits in different environmental conditions. Despite the lack of a formal definition, robustness usually refers to the continuation of function in the presence of perturbations (Kitano, 2004).

In von Uexküll (1957)'s terms, the selection of sensory stimuli can be seen as a process which can 'bring forth their own Umwelt', or relevance in the surrounding world of agents. This suggests that agent behavioural mechanisms can be thought as cognitively distributed between internal control, body, and environment. The description of 'functional circles' proposed by von Uexküll (1957) indicates that a 'cue' (or functional trigger) is distributed along the entire functional circle of which the organism is a part. Functional circles are "abstract structures that tie together a subjective experience or perception (termed a *perceptual cue*) and the effect that the perceptual cue has on the behaviour of the organism (called a *effector cue*)" (von Uexküll, 1957, from Macinnes and Di Paolo, 2006). It is meaningless to claim that a perceptual cue resides in a particular location in the agent's milieu. The ability to walk and the feedback that the nervous system receives during walking, for example, is not localized at neural level but is fully distributed throughout the agent and its dynamics, where part of the control task is 'outsourced' to the physical dynamics of the agent (Pfeifer et al., 2007). However, what sort of control-strategy emerges if an agent's own interactions with the environment shape its sensory capacity and its dynamics under sensorimotor perturbations?

An answer to this question has conceptual and practical interest for understanding robustness in neuroscience, e.g. in the consideration of controlled walking behaviour in humans under different scenarios (e.g. in different terrain conditions). Discussions regard an agent here for clarity refers as a dynamic system perturbed by, and hence responding to, a number of environmental cues and externally induced perturbations. Analyses in this article examine if the tuning of sensory offsets improves the agent's behavioural robustness in the presence of sensorimotor perturbations. The developed minimal model methodology described in Section 2 may illuminate how to answer the above question due to its relatively assumption-free paradigm. Studying the mechanisms that emerge can inform

* Tel.: +44 01273 872948; fax: +44 01273 678535.
E-mail addresses: jafphd@gmail.com, jf76@sussex.ac.uk.

0303-2647/$ – see front matter © 2010 Elsevier Ireland Ltd. All rights reserved.
doi:10.1016/j.biosystems.2010.09.010

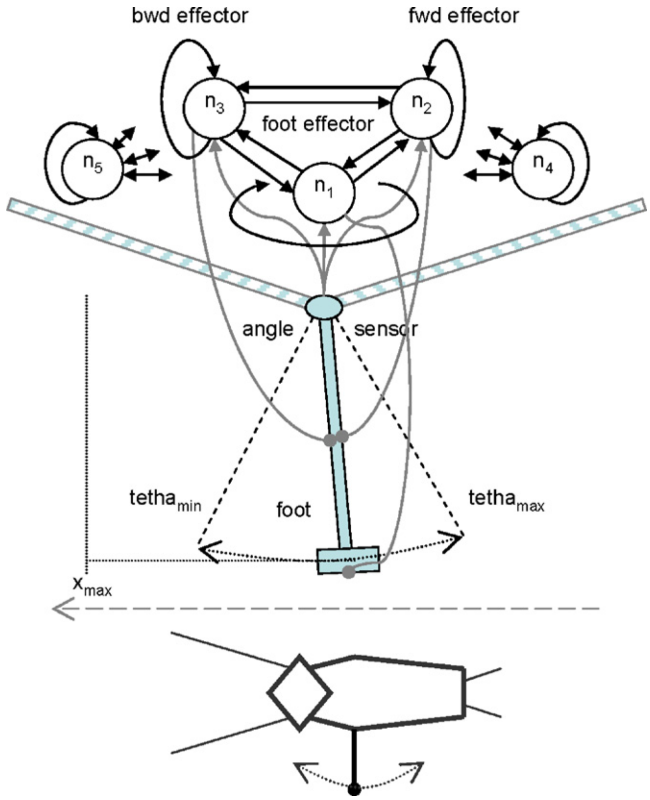

Fig. 1. (Top) Schematic representation of the agent's leg configuration for one-leg walking behaviour. Neurons are fully connected including self-connections. Three effectors controls the forward and backward force applied to the leg and the foot for walking. Effectors receive sensory stimuli of the leg angle during the ongoing task. (Bottom) The leg model of a simulated insect where the leg can swing about their single join with the body (figure based on Beer (1995a)).

our understanding of what to look for in natural systems and how to build better artificial ones which regulate the conditions of their exchange with the environment (Di Paolo and Iizuka, 2008). In Sections 2 and 3, the methods and experiments are introduced. Section 4 examines the consequences of results and discusses ideas that remain open.

2. Methods

The experimental part of this article first investigates how a one-legged agent model in Evolutionary Robotics (ER) (Nolfi and Floreano, 2000) can produce robust walking behaviour, controlled by a single neural network (neurocontroller). Artificial evolution is used to synthesize an embedded recurrent neural network enabling active regulation of agent-environment dynamical exchange. Experiments then inspect whether different populations of evolved agents with and without sensory offsets can exhibit a higher (or lower) qualitative dependence to environmental dynamics, but maintaining in all cases a quantitative high performance (lifetime fitness). Evidencing further dependency of neurocontrollers to body and environment dynamics can support discussions on behavioural mechanisms that distribute between nervous system, body and environment. The followed methodology is based on Beer's model software and use Mathematica's tools for dynamical analyses (Beer, 2010).

2.1. Agent and Structure of the Environment

The one-legged insect-like agent (Fig. 1) is used as experimental model in this work, which is a simplified model of the simulated hexapod agent described in (Gallagher et al., 1996). This agent model is based on Beer and Gallagher's (1992) (see also Beer et al.'s (1999), Izquierdo and Buhrmann's (2008), and Beer's (in press) works). In particular, Beer (1995a) has studied three variants of his legged-agent model, differing in whether sensory feedback is continuously available, only sporadically or completely absent for the agent.

Beer (1995a) observed three different types of controllers that produced the expected walking behaviour: the reflexive pattern generators (RPGs), central pattern generators (CPGs), and mixed pattern generators (MPGs). RPGs are reminiscent of Sherrington's proposed 'chained reflex' locomotion circuits that depend on the presence of external periodic timing signals (Sherrington, 1898) (Sherrington, 1898, from Gallagher, 2001). CPGs are defined as neural networks that can endogenously (i.e. without rhythmic sensory or central input) produce rhythmic patterned outputs (Marder and Calabrese, 1996). Finally, MPGs represent a combination of previous generators – i.e. like RPGs, MPGs can use sensory feedback when it is available to improve their operation, but like CPGs they can function in its absence if necessary (Beer, 2009). In this work, the focus is on RPGs requiring continuous sensory feedback from the leg's joint angle (Fig. 1). However, discussions qualitatively compare the trajectory of agent leg movements to the idealized trajectory obtained from a perfect CPG without sensory input.

Experiments here are based on the leg model in Fig. 1, which has two degrees of freedom, one for rotation and another for extension. The leg can swing through 45 degrees from vertical forward (*fwd*) or backward (*bwd*). The leg passively stretches between the joint and the foot as the body translates. The agent's leg has a foot that can be either up or down. The agent's body is considered stable as long as its foot is not too far back to enable the ongoing forward motion. In order to compute the force applied to the body, the model allows a supporting leg that has passed outside of the mechanical limits to apply force in a direction that moves it back toward that mechanical limit. This is not possible in a direction that would move it further away where these mechanical limits become one-way constraints for a supporting leg. When the leg's stability is lost, the agent falls and its forward velocity is immediately set to zero.

The leg's set-up is as follows: the leg's length is 15 units long; the maximum leg force, velocity, maximum torque and angular velocity are 0.05, 6.0, 0.5 and 1.0, in that order; forward and backward angle limits are $\theta_{min} = -\pi/6$ (or $-\theta = -0.5236$ radians) and $\theta_{max} = \pi/6$ (or $\theta = 0.5236$ radians), respectively. The leg is only able to generate force over a limited angular range of motion of $[-\theta, \theta]$ (see Fig. 1). In other words, when a stretched stancing leg lifts its foot, the leg immediately snaps back to the swing angular limits of $[-\theta, \theta]$. When a stancing leg reaches these limits, forward motion comes to an abrupt stop, which according to Beer's descriptions it produces a loss of postural stability. During the stance phase, the leg stretches between the body joint and the stationary foot as the body moves with a horizontal distance between the joint and the foot. A stancing leg exceeding the angular range of motion can still provide support, but only within vertical limits of $[x_{min}, x_{max}]$. Torque is controlled by two motor neurons (forward or backward neuron effectors onwards). When the foot is up (*swing phase*), torque produced by effectors serves to swing the leg along an arc relative to the body (Beer, in press). For this movement applies a limit constraint with a maximum angular acceleration of $\alpha_{max} = 1/40$. The binary state of the foot (*FT*) is up when the difference between effectors is lower or equal than 0.5, and down when such a difference is higher than 0.5.

The agent is given 220 units of time to walk and after this period, the algorithm measures the total walked distance during the trial (fitness measure). An agent performing a perfect walking behaviour will walk 305.7101 units of distance over these units of time, which represents more than 12 full walking steps. A successful agent must maximize the final walked distance. The overall performance of agents is averaged over all trials producing a value in range [0, 306].

2.2. The Implemented Network Topologies

The agent's leg is controlled by a fully connected five-neuron controller, where three of these neurons are effectors creating the force applied to the agent's body that generates translational motion. One effector (*n1*) governs the state of the foot, and the other two generate (*n2*) clockwise and (*n3*) counter clockwise torques to the leg's single joint producing forward and leftward movements. The remaining two units are interneurons with no-specified role in the agent's leg behaviour. Only effector neurons receive a weighted sensory input from the leg's angle sensor that measures the leg's angular position in radians. The angle sensor is proportional to the angular deviation of the leg from the perpendicular axis to the long one of the body. The neurocontroller supplies signals specifying what torques should be applied at each joint. These signals are summed, and depending on the state of the leg's foot will either move the body (foot down) or rotate the leg about its joint (foot up).

The embedded controller defines one or three mutable sensory attributes or offsets to the agent's genotype (further explained in the next section). These offsets are either all the same or different for every sensor–interneuron connection (Fig. 2), and are added to every sensory signal. The use of offsets means that we can no longer observe a signal of zero to neurons if we disrupt the angle sensor the controller. The agent consequently is able to co-evolve to some extent the capacity to sense the environment alongside the rest of its internal dynamics. By 'functional' in this context, the article means the capacity of agents to produce rhythmic stepping for the expected walking behaviour.

2.3. Agent's Controller Definition

A continuous-time recurrent neural network (CTRNN) (Beer, 1995b) controls the behaviour of the leg and finally the movement of the agent. The following equations define the implemented neuron-like units:

$$\tau_i \frac{dy_i}{dt} = -y_i + \sum_j^n w_{ji} z_j + I_i \quad (1)$$

$$z_j = \sigma(g_i(y_j + \theta_j)) \quad (2)$$