Toward Spinozist Robotics: Exploring the Minimal Dynamics of Behavioural Preference

Hiroyuki Iizuka and Ezequiel A. Di Paolo

CSRP 590

May 2007

ISSN 1350-3162

University of Sussex

Cognitive Science
Research Papers

Toward Spinozist Robotics: Exploring the Minimal Dynamics of Behavioural Preference

Hiroyuki Iizuka [1;2] and Ezequiel A. Di Paolo [1]

[1]Centre for Computational Neuroscience and Robotics,
Department of Informatics, University of Sussex
Brighton, BN1 9QH, UK

[2]Department of Media Architecture, Future University-Hakodate
116-2 Kamedanakano-cho, Hakodate,
Hokkaido, 041-8655, Japan

Contact author:
Dr. Hiroyuki Iizuka
e-mail: ezca@sacral.c.u-tokyo.ac.jp
phone: +44-1273-87-2948

To appear in *Adaptive Behavior*

Abstract

A preference is not located anywhere in the agent's cognitive architecture, but it is rather a constraining of behaviour which is in turn shaped by behaviour. Based on this idea, a minimal model of behavioural preference is proposed. A simulated mobile agent is modelled with a plastic neurocontroller, which holds two separate high dimensional homeostatic boxes in the space of neural dynamics. An evolutionary algorithm is used for creating a link between the boxes and the performance of two different phototactic behaviours. After evolution, the agent's performance exhibits some important aspects of behavioural preferences such as durability and transitions. This paper demonstrates 1) the logical consistency of the multi-causal view by producing a case study of its viability and providing insights into its dynamical basis and 2) how durability and transitions arise through the mutual constraining of internal and external dynamics in the flow of alternating high and low susceptibility to environmental variations. Implications for modelling autonomy are discussed.

keywords : behavioural preference, homeostatic adaptation, dynamical systems approach to cognition, evolutionary robotics

1 Introduction

How does an embodied agent develop a stable behavioural preference such as a habit of movement, a certain posture, or a predilection for spicy food? Is this development largely driven by a history of environmental contingencies or is it endogenously generated? Kurt Goldstein (1934) described preferred behaviour as the realization of a reduced subset of all the possible performances available to an organism (in motility, perception, posture, etc.) which are characterized by a feeling of comfort and correctness as a contrast to non-preferred behaviour which is often difficult and clumsy. Merleau-Ponty, following insights from Gestalt psychology, postulated that bodily habits are formed by resolving tensions along an intentional arc where the external situation solicits bodily responses and the meaning of theses solicitations depend on the body's history and dynamics. The overall tendency is thus towards an optimum or maximal grip on the situation (e.g., like finding just the right distance to appreciate a painting), (Merleau-Ponty, 1962, p. 153). In these views, the fact that a preferred behaviour is observed more often would be derivative and not central to its definition (preferred behaviour is often efficient but not necessarily optimal in any objective sense). Following this idea, we define a *preference* as the strength or commitment with which a behavioural choice is enacted, which is measurable in terms of its robustness to different kinds of perturbations (internal or external). Such a preference is typically sustained through time without necessarily being fully invariant, i.e., in time it may develop or it may be transformed into a different preference. In order to understand such durable states from a dynamical systems perspective, it is therefore convenient to study under what conditions these preferences may change since this will reveal more clearly what are the factors that contribute to their generation.

The word *preference* has many higher-level cognitive connotations that may not be captured by this minimal definition. On the one hand, we recognize that the dynamical systems picture portrayed in this paper does not do full justice to the richness of the concept (e.g., human preferences involve a complex interaction between habits, needs, cultural context and sometimes conflicting values). On the other hand, our purpose is precisely to explore the *minimal* dynamical properties that might be shared by many instances of preferred behaviours. We follow the directions of synthetic minimalism which has been defended as a useful route towards clarifying complex ideas in cognitive science (Beer, 1999; Harvey, Di Paolo, Wood, Quinn, & Tuci, 2005). Our objective is to achieve such a clarification of the term *preference* and related terms such as *disposition*, *tendency*, *commitment*, *conation*, etc.

using the language of dynamical systems.

Dynamical systems approaches to cognition have typically examined processes at the behavioural timescale such as discrimination, coordination, and learning, (e.g., Beer, 2003; Kelso, 1995) and they have also been deployed to describe changes at developmental timescales (Thelen & Smith, 1994). But of course, the strength of the approach lies in its potential to *unify* phenomena at a large range of timescales. In particular, behavioural preferences and their changes lie between the two scales just mentioned (the behavioural and developmental) and share properties with both of them. Goldstein has argued that we cannot really find the originating factors of a preferred behaviour purely in central or purely in peripheral processes, but that both the organism's internal dynamics and its whole situation participate in determining preferences (Goldstein, 1934). In this view, it becomes clear that a preference is never going to be captured if it is modelled as an internal variable (typically a module called "Motivation") as in traditional and many modern approaches, but that a dynamical model needs to encapsulate the mutual constraining between higher levels of function, such as performance, and lower processes, such as neural dynamics (Varela & Thompson, 2003). A preference is not "located" anywhere in the agent's cognitive architecture, but it is rather a constraining of behaviour which is in turn shaped by behaviour. A goal of this work is to explore and possibly operationalize how this might be achieved in concrete terms suitable for further hypothesis generation.

Preferences as described here are not typically addressed in minimally cognitive or robotic models. Most artificial agents are designed so as to have *no* preferences, or rather a single preference: that of adequate performance. Straying from the assigned task into a different behaviour chosen by the agent itself may be a sign of increased autonomy but not a frequent or explicit goal in current robotics. In this paper we present an exploratory model of behavioural preference with the objective of exploring the assertion of multi-causality. We consider that the minimal requirements to capture the phenomenon of preference is a situation with two mutually exclusive options of behavioural choice. An agent should be able to perform either of these options but the choice should not be random, but stable, and durable. The choice should not be invariant either but it should eventually switch (in order to study the factors that contribute to switching). There should be a correspondence between internal dynamical modes and stable behaviours. For this we use the methods of homeostatic adaptation to design not only the agent's performance but to put additional requirements on the corresponding internal dynamics. We then examine the dynamical factors that play a role in sustaining and changing preferences.

2 A Spinozist approach

In contrast to functional/computational approaches, a dynamical systems perspective on cognition makes it harder to conceptualize intentional terms (such as motivations, tendencies, goals, emotions, etc.) as functional states in the cognitive architecture of an agent (often implemented in computational modules; a practice some people refer to as "boxology"). However, there isn't as yet a clear and generally accepted alternative way to deal with intentional concepts from the dynamical camp. Attempts are not lacking. For instance, Kelso (1995) has suggested that it is possible to describe intentional behavioural changes in terms of transitions between dynamical attractors which correspond to different behaviours. The proposed process would be achieved through the successive stabilisation and destabilisation of attractors. In this view, intentions emerge as the order parameters of self-organised bodily and neural dynamics leading to behaviour when coupled with the environment. A similar idea has been proposed by Thelen and Smith (1994) who use the changes of attractor stability to explain the development of behaviours both at the developmental and behavioural timescales. In their view, the landscape of the stability changes depending on ontogenetic processes. Juarrero (1999) offers a related view where prior intentions are understood as the setting up of a dynamical landscape of attractors (e.g., through recursive self-organizing activity on the system's constraints) and proximate intentions imply the local selection of such dynamical alternatives. All these views share some problems, such as the problem of how the subject or agent to whom intentions might be attributed is itself dynamically constituted, i.e., who is the agent that re-shapes a dynamical landscape so that we may speak of intentions as belonging to it. Nevertheless, the central intuition of dynamical re-shaping is appealing and worthy of expansion and clarification.

An important point shared by both intentional behaviour for Kelso and Juarrero and developmental processes for Thelen and Smith is that they cannot be separated from the ongoing behavioural dynamics, i.e., the systems engagement with the environment. Given that these are processes that themselves may alter the dynamics of interaction with the environment, the emerging picture is one of mutual modulation, or circular causality (Clark, 1997; Thompson & Varela, 2001) where intentional and developmental processes re-parameterize behaviour and, in turn, interaction with the environment constrains and modulates intentions and development.

Dynamical views similar to those of Kelso and Juarrero have been explored in robotics. For instance, in the work of Ito et al. a humanoid robot shows transitions between two different ball handling behaviours using a special kind of neurocontroller in which specific nodes are trained to have an association with each behaviour (Ito, Noda, Hoshino, & Tani, 2006). Transitions are demonstrated to occur when the robot is interrupted by a person changing the position of the ball. Although the functioning of the system may be conceived in terms of a re-organization of the attractor landscape, transitions are demonstrated through human guidance, i.e., triggered externally. This, unfortunately, does not allow the evaluation of more spontaneous forms of switching in terms of the dynamical account of mutual constraining between interactive and neural levels that we want to explore in this paper.

Following this idea, in this paper we suggest that a way of approaching a study of preference is by embedding the circular causality between internal organization and interaction in an embodied agent through the application of homeostatically-driven neural plasticity (further justified below). Even though this idea is not disconnected from those of Kelso, Juarrero and Thelen and Smith, the view in our proposed approach is different from the attractor model in the following important sense.

Suppose that a dynamical system has two attractors, each of which corresponds to a different behaviour. If a transition between them occurs, it could only be caused when noise or perturbations to the dynamics are strong enough to detach the trajectory from the current attractor. As Kelso discusses, this would not constitute a proper intentional change because the behavioural transition depends only on a random event and would not result in durable behaviour as in the case of a preference. Conversely, the system cannot change attractors without noise, which means that in such a case it would produce a same behaviour permanently.

Therefore, there are problems in explaining preferences in behaviour as the switching between attractors in the proper dynamical systems sense. It should be noted that we are not denying noise and random events may play a role in triggering the transitions but insisting on the significance of transitions caused by internal factors, which can be seen as higher-level dynamics as in Kelso's and Thelen and Smith's explanations. A proper intention, motivation, or alteration to a preference should also be dependent on endogenous conditions, such as changes to the dynamical landscape itself as a result of a history of behavioural interactions with the environment in non-arbitrary ways. By explicitly implementing the circular causality between inner stability and external behaviour into our model through the use of homeostatic mechanisms (as only one possibility for achieving this), important aspects of preferences such as durability and transitions between behaviours can be captured.

However, the choice of homeostatic adaptation to model the dynamics of preference still requires a stronger justification.

A view on preferences, dispositions and tendencies that is amenable to a dynamical interpretation is Spinoza's doctrine of *conatus* or *striving* developed in his *Ethics* (IIIp4-7). This view serves as an inspiration for how these issues are addressed in our model. We read in part III proposition 6: "Each thing, as far as it can by its own power, strives to persevere in its being". There are several issues concerning this doctrine discussed by Spinoza scholars, especially the unargued proposition IIIp4 that states that "No thing can be destroyed except through an external cause" (Matson, 1977); one might think of a bomb as a counter-example. However, in the context of cognition a dynamical systems approach is highly compatible with views that seek to establish a continuity between life and mind, (Jonas, 1966; Maturana & Varela, 1980; Stewart, 1992; Wheeler, 1997; Bourgine & Stewart, 2004; Weber & Varela, 2002; Di Paolo, 2005). The *essence* or *being* of a living system is its self-producing organization (autopoiesis) and it so happens that IIIp4 is indeed true of living systems at least at a minimal level of organization, (Di Paolo, 2005)– even though this is not such an obvious statement for more complex systems beyond bare autopoiesis; think of autoimmune diseases. Hence, we may assume that the striving doctrine is applicable to minimal cognitive systems as well and so we can use this idea as a regulative principle informing our approach (complex cases such as culturally-embedded human cognition may present special kinds of problem such as the origin of conflicting values; we are not considering these possibilities here).

So the question of preference becomes the question of understanding an agent's changing conatus. For Spinoza, this would directly relate to the agent's current being. A hungry animal strives for food and a thirsty one for water. Conatus resolves in the interactional domain (actions and perceptions) a tension that originates in the internal constitutive domain (an internal need). To cash out such ideas in dynamical terms we must define in the dynamical organization of the agent a condition of *tension* and *satisfaction* where the first is defined as being in a state that leads to alterations in the system's organization (its essence or being in Spinozist terms) and the second as being in a state where the system's organization remains invariant. This closely follows Ashby's idea of an ultrastable system

(Ashby, 1960). Conatus is therefore, on a first approximation, the interactive and internal striving of an ultra-stable system to remain in those states that satisfy the condition of invariant organization. This need not be interpreted teleologically yet even though the terms tension and satisfaction have not been chosen innocently, (for a discussion regarding the conatus doctrine and teleology see (Bennett, 1984, 1992; Curley, 1990)). This view will be refined after understanding the results from our model.

Finally, in order to sharpen the contrast of the view we are proposing to previous ones, a relation to action selection should be mentioned. The action selection literature is concerned with how animals, robots or simulated agents can solve the ongoing problem of choosing what to do next in order to achieve their objectives (Humphrys, 1997; Bryson, 2003). Agents are given several options of action units in advance and they must decide which one should be taken next and how they should be ordered and combined. In other words, in action selection models the sensorimotor coupling at the lower level is fixed and discretised in order to set the action units and the aim is to decide how they should be chosen for the planning of the higher-level goals. Roughly, this means that planning as a higher-level description is separable from the ongoing actions as a lower-level description although this may depend on how the action units are defined. It might be rather precise to say that this kind of work tries to illuminate our planning intelligence in an ecological multi-objective context at the cost of simplifying (even removing) the effects of the sensorimotor dynamics on the system's changing organization. This can be seen even in work on embodied robotics which models intentional terms such as preferences, motivations and goals as certain variables or "compartmentalized" functions named emotions or motivations (Breazeal(Ferrell), 1998; Velásquez, 1997). Our view is different from the idea expressed in most action selection models in that planning is *not* separable from actions. Producing actions based on sensorimotor coupling organizes and preserves plans, and in turn plans regulate the sensorimotor flows. It should be clear that our motivations are also different although we would expect that future developments of our work may interact more closely with problems in action selection.

3 Homeostatic adaptation in neural controllers

A salient feature of the homeostatic adaptation model proposed in (Di Paolo, 2000) is that local plastic adaptive mechanisms work only when neural activations move out of a prescribed region; an idea inspired in Ashby's homeostat (Ashby, 1960). Such a mechanism has been implemented in a neurocontrolled simulated vehicle evolved with a fitness function rewarding phototaxis and the maintenance of neural activations within the homeostatic region. The use of intermittent plasticity in combination with this selective pressure allows for the evolution of a novel kind of coupling between internal and environmental dynamics. Once the neurocontroller gives rise to behavioural coordination within a given environmental situation that results in internal stability, synaptic weight changes no longer happen. If the situation changes, such as in an inversion of the visual field or some other perturbation, this causes a breakdown of coordination, which means that the neural activations cannot in general be maintained within the homeostatic region. As this happens, the local adaptive mechanism is activated until it finds a new structure (synaptic weight values) which can sustain the activations within the homeostatic region and (very likely, though not necessarily) re-form the behavioural coordination. As a result, the agent can adapt to perturbations it has never experienced before.

This approach aims at creating a high dimensional bounded set, or box, in phase space that corresponds to neural homeostasis often linked to suitable performance. The dynamics within the box is stable in that if trajectories go out of the bounds, the network's own configuration, by design, will change plastically until they come back again into the box. If trajectories remain within the box, the system's configuration no longer changes. This plasticity-dynamics relation makes it possible for a coordination to occur (under suitable circumstances) between a higher function, phototactic behaviour in this case, and the process that regulates the sensorimotor flows. If the behaviour cannot be achieved, homeostatic adaptation attempts to find new sensorimotor flows. This is therefore a concrete example of the circular constraining between levels mentioned above[1].

In the current context, we extend this property to create a model of

---

[1] The description of the homeostatic bounded sets as boxes could cause some confusion with the traditional perspective that we have criticised above describing it as "boxology". We would like to clarify that we use this term to refer to the practice of putting a cognitive function that belongs to the level of the whole agent into a computational module inside the agent's architecture. This is different from creating areas (boxes) in the space of the agent's internal dynamics that are both 1) generally stable and 2) linked to a particular behaviour. Nothing is specified about the nature of this link in functional terms. In fact, the behavioural function and the internal organization are linked bi-directionally, since behaviour will impact on the internal dynamics, provoking plastic changes, and the internal landscape will constraint the domain of possible behaviours. So, the functional nature of the dynamics within a box is not fixed which is a general feature of the boxology approach.

behavioural preference. Our idea is that if the system holds *two* separate high dimensional boxes in the space of neural dynamics which are associated with performing different behaviours, for simplicity's sake phototaxis towards different light sources A and B, a preference could be formed by the dynamical transitions that select which box the dynamics go into and stay in. This provides one of our requirements for talking about preference, that of *durability* (bottom-up construction of the stability). Once a behaviour is formed, due to the stability in a box, the system keeps doing the behaviour while ignoring other behavioural possibilities. This plays the role of a spontaneous top-down constraint that regulates the sensorimotor flows. However, some disturbances might eventually cause a breakdown of the stability and then another behaviour can be reconstructed through the homeostatic adaptive mechanisms. Since by design, the system has another region of high stability, the system will be likely to switch into it and then start enacting its other behavioural option. In this way, behaviour can switch due to the corresponding transitions between two boxes. We expect to see both spontaneous and externally-induced transitions from the viewpoint of the top-down and bottom-up construction or destruction of durable but impermanent dynamical modes. Here we find our second requirement, that of the possibility of *transformation*, or change in preference. Except for considerations of symmetry, the placement of these high-dimensional boxes is arbitrary in the present model (see discussion).

4 Model

Our proposed minimal model of behavioural preference extends the homeostatic adaptation model. The idea is implemented in a simulated mobile agent with a plastic neural controller containing two separated, high-dimensional homeostatic regions. The simulated robot is faced with two different kinds of light as mutually exclusive options of behavioural choice. It must visit one of them and this “choice” must correspond to the internal dynamics being contained in the corresponding homeostatic box. An evolutionary algorithm is used to design the neurocontroller. Agents are evaluated on each type of light separately, and on both types simultaneously, in which case one of them is presented as blinking. The agent must approach the non-blinking light, (see below). After evolution, this cue is removed and the agent's preference is investigated by presenting two constant, non-blinking lights of each type simultaneously.

Agent. An agent is modelled as a simulated wheeled robot with a circular body of radius 4 and two diametrically opposed motors. The motors can drive the agent backwards and forwards in a 2-D unlimited plane. Agents have a very small mass, so that the motor output directly determines the tangential velocity at the point of the body where the motor is located. The translational movement of the robot is calculated using the velocity of its center of mass, which is simply the sum of the two tangential velocities and the rotational movement is calculated by dividing the difference of the two motor outputs by the diameter (leading to a maximum translation speed of 2.0 units and rotational speed of 0.25 radians at each Euler time step). The agent has two pairs of sensors for two different light sources, A and B, mounted at angles of $\pi/3$ radians to the forward direction. The two lights do not interfere with each other and the model includes the shadows produced by the agent's body.

Plastic controller. A fully connected continuous-time recurrent neural network (CTRNN) (Beer, 1990) is used as the agent's controller. The time evolution of the states of neurons is expressed by:

$$\tau_i \dot{y_i} = -y_i + \sum_{j=1}^{N} w_{ji} z_j(y_j) + I_i, \quad z_i(x) = 1/(1 + e^{-x-b_i}), \tag{1}$$

where y_i represents the cell potential of neuron i, z_i is the firing rate, τ_i (range $[0.4, 4]$) is its time constant, b_i (range $[-3, 3]$) is a bias term, and w_{ji} (range $[-8, 8]$) is the strength of the connection from the neuron, j, to i. I_i represents the sensory input, which is given to only sensory neurons. The number of neurons, N, is set to 10 in this paper, 4 of which are assigned to sensory neurons, i.e., 2 neurons for each light sensor. The sensory input is calculated by multiplying the local light intensity by a gain parameter (range $[0.01, 10]$). The decay of the light intensity follows a sigmoid function in order to avoid the situation in which it becomes too strong when the agent is very close to the light. There is one effector neuron for controlling the activity of each motor. Similarly, the motor output is calculated from the firing rate of the effector neuron, which is mapped into a range $[-1, 1]$ and is then multiplied by a gain parameter (range $[0.01, 10]$). All free parameters are determined genetically.

A plastic mechanism allows for the lifetime modification of the connection weights between neurons. The homeostatic regions are described by a plasticity function of the firing rate of the post-synaptic neuron; this function determines the strength of change of all incoming weights. The plastic function is 0 in the homeostatic regions, which means no plasticity. To as-sign the different phototactic behaviours to different boxes in our extended model, two separated regions are arbitrarily set corresponding to firing rates of $[0.15, 0.4]$ and $[0.6, 0.85]$ as shown in Fig. 1 and for each neuron each region is arbitrarily assigned for phototaxis A or B at the start of the evolutionary run. To reduce bias, the upper (bottom) region of half of the internal neurons is assigned for phototaxis A (B), and the other region is for phototaxis B (A). These assignments remain the same for the evaluations during the evolutionary run. However, the scheme with two separate regions is not applied for input and output neurons because this would introduce biases in the input sensitivity as well as prescribe particular styles of movement. For input and output neurons, a function which has a single homeostatic region is applied (Fig. 1 (right)).

Weight change follows a Hebbian rule which also depends linearly on the firing rate of the pre-synaptic neuron and a learning rate parameter. Weights from neuron i to j are updated according to :

$$\Delta w_{ji} = \eta_{ji} z_i p(z_j) \tag{2}$$

where z_i and z_j are the firing rates of pre- and post-synaptic neurons, respectively, Δw_{ji} is the change per unit of time to w_{ji}, $p(x)$ is the plastic function (see Fig. 1), and η_{ji} is a rate of change (range $[0, 0.9]$), which is genetically set for each connection. For simplicity, we restrict this parameter to a positive range, so that the product of this value and the plastic function always works in the direction of returning the flow into the homeostatic region. For example, if the firing rate of neuron j, z_j, is between $[0, 0.15]$ or $[0.5, 0.6]$, the plastic function returns a positive value and the weight change is calculated by multiplying the firing rate of pre-synaptic neuron i and η_{ji}, which are positive. This means that the weight change works to increase the firing rate z_j. If the firing rate is between $[0.4, 0.5]$ or $[0.85, 1.0]$, the plastic function returns a negative value. In this case, the weight change works in the opposite direction to decrease the firing rate.

This implementation for the plasticity rule is not the only possible way to realize neural homeostasis. Other studies show different forms of neural plasticity (Di Paolo, 2000; Williams, 2004). However, we use the simplest implementation for out current purposes in that plasticity always works in the direction of stability within the boxes. The validity of stipulating the homeostatic regions in this way will be addressed later in the discussion.

The time evolution of agent's navigation and plastic neural network are computed using an Euler method with a time step of 0.1.

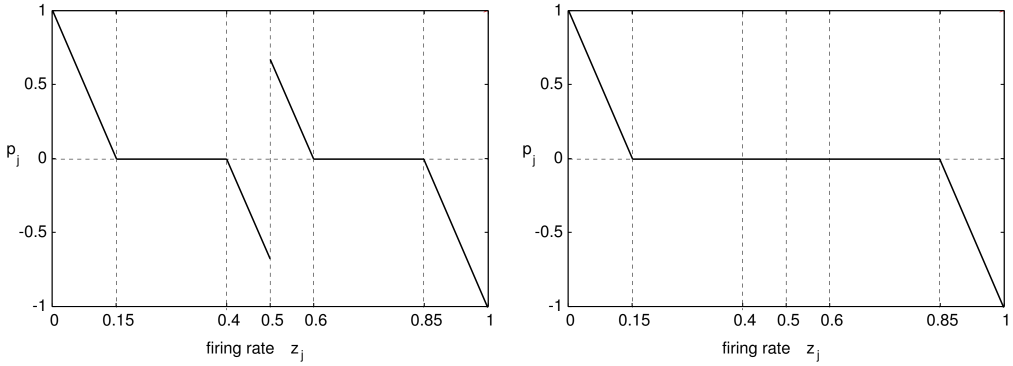

Figure 1: Plastic facilitation as a function of firing rate for internal neurons (left), and for input/output neurons (right).

4.1 Evolutionary setup

A population of 60 agents is evolved using a rank-based genetic algorithm with elitism. All network parameters, w_{ji}, τ_i, b_i,, η_{ji} and the gains are represented by a real-valued vector ([0,1]) which is decoded linearly to the range corresponding to the parameters (with the exception of gain values which are exponentially scaled). Crossover and vector mutation operators, which adds a small random vector to the real-valued genotype (Beer, 1996), are used. The best 6 (10%) agents of the population are kept without change. Half of the remaining slots are filled in by randomly mating agents from the previous generation according to rank, and the other half by mutated copies of agents from the previous generation also selected according to rank.

The agents are evaluated under 4 different situations: a single light A, a single light B, two-lights-A, two-lights-B. The task of a single light A (B) consists of the serial presentation of 8 distant light sources of type A (B) which the agent must approach in turn and remain close to. Only one source is presented at a time for a period, called a trial, of random duration drawn from the interval $[700, 1000]$ update steps. In contrast, under the task of two-lights-A and two-lights-B, two different light sources, A and B, are presented simultaneously, one is blinking, and the other is constant. The agent gains fitness by approaching the latter. Our aim in putting a blinking light as a dummy is to encourage the agent to get to the target while in the presence of distracting stimuli that it must ignore. Otherwise, the agent would not get a chance to experience the simultaneous presence of both sources of stimuli and face the problem of picking one as the target.

The blinking light flickers with a 15% probability at every timestep. As well as the single light task, the two-lights task consists of the serial presentation of 8 pairs of distant light sources, which are $\pi/2$ apart from each other from the agent's point of view. The length of the trial is chosen in the same way as in the single light task. After a trial, lights are extinguished and new ones appear at a random distance, $[100, 150]$.

Each individual agent is tested for 12 independent runs in total, i.e., 3 independent runs for each task. At the beginning of each run, the synaptic weights are reset to the initial values. Each run consists of 8 trials and only the last 3 of those are evaluated in order not to penalize slow plastic changes.

Fitness is calculated based on three terms. F_D corresponds to the proportion of reduction between the final (D_f) and initial (D_i) distance to a target, $1 - D_f/D_i$. F_p indicates the proportion of time that the agent spends within a distance less than 15 to a target during a trial. F_H represents the average score of neural homeostasis. For each timestep that a neuron fires homeostatically within the region corresponding to the target light, a counter is incremented by 1. If it is within the homeostatic region assigned for the non-target light, the counter is not incremented and for all other regions the counter is incremented by 0.5. These counters are then averaged for all neurons and over the whole trial. Selecting for high F_H will tend to create an association between each homeostatic region and the corresponding phototactic behaviours. For each trial, the total fitness $(F_D + F_p)F_H$ is calculated and then averaged over all 12 runs.

5 Results

The two phototactic behaviours can be easily evolved, however, it is more difficult to obtain agents that are able to associate the behaviours with the two homeostatic regions and to maintain the internal dynamics within the regions. The evolved agents can be highly sensitive and dependent on the history of the interactions. Since our purpose in this paper is to explore how the high level concept of preference can be described in terms of neural and sensorimotor dynamics and to demonstrate the logical consistency of a multi-causal view, we restrict our analysis to the in-depth study of a single successful agent. This will allow us to generate a more concrete hypothesis that further modelling (including statistical analysis across many runs) and specific empirical studies can investigate.

5.1 Basic phototactic behaviours

First, in order to check for long-term stability of the two phototactic behaviours and maintenance of the internal dynamics, the agent is tested for longer successions of lights (only 8 were evaluated during evolution). In the case of interacting with a single light A, or two-lights-A (constant light A and blinking light B), the agent shows a long-term stability (more than 100 lights) for phototaxis A and the maintenance of the internal dynamics. On the other hand, in the case of interacting with a single light B, or two-lights-B (blinking light A and constant light B), the stability is less than in the former case. One or two of internal neurons sometimes stay within the homeostatic region for light A even when approaching light B. Even so, this is not a major problem for our purposes because there is still a difference in stable states of the other neurons between two types of phototaxis. In terms of behavioural patterns, the agent typically goes straight to the target light.

Due to the weaker stability of phototaxis B, the following experiments are run for less than 100 successive lights, which was long enough to show the preference.

5.2 Transitions

When lights are presented simultaneously during evolution, a cue telling the agent which light to visit is given by the difference between the lights (constant or blinking). A choice under this condition can be regarded as being environmentally pre-determined before the agent starts interacting. The agent is now tested under a situation (unseen during evolution) where no cue is given, i.e., by making both lights constant.

Figure 2 shows, for 100 consecutive trials, the final distances to both light sources at the end of each trial (very short distance means the agent has approached the light at the end) and how long the internal dynamics stay within each homeostatic region on average over the whole trial. As can be seen, the agent always “chooses” one of the two lights, it never stays in between the two. This is not trivial because the two-light situation has not been experienced by the agent and it might have produced oscillations or deadlocks as a result of competition between sensory inputs. It is likely that the conditions for such behaviours are rare, however. It is also shown that the selection is neither random nor permanent but durable or quasi-stable (only 4 changes are counted in this run of 100 presentations).

Examples of behavioural patterns are shown in Fig. 3 when interacting with the first, the 10th, and the 30th pair of lights. At the first pair, the agent takes a side trip while ignoring both lights. This is normal and corresponds to an initial period where plastic rules are settling the initially random weight values. The 10th and 30th presentation are typical examples of approaching patterns to light *A* or *B* while ignoring the other. These are similar patterns to those observed in the task situations during the evolution (single lights, and 2 lights, 1 blinking).

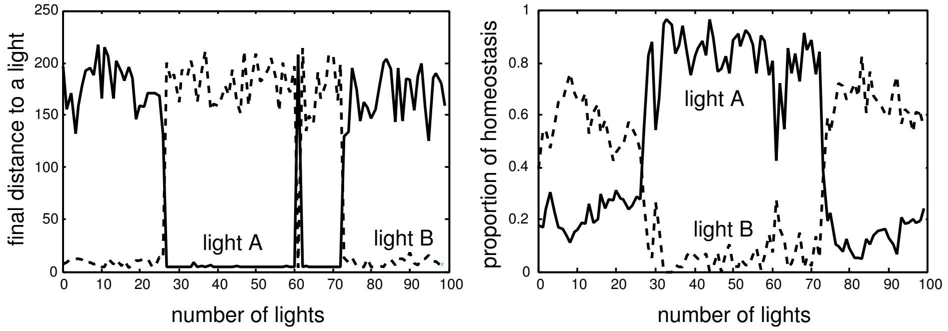

Figure 2: Left: Final distance to each light at the end of trials on serial presentations of 100 pairs of constant lights. Right: Proportion of neurons that have stayed within the homeostatic region for each light in correspondence with trials on the left.

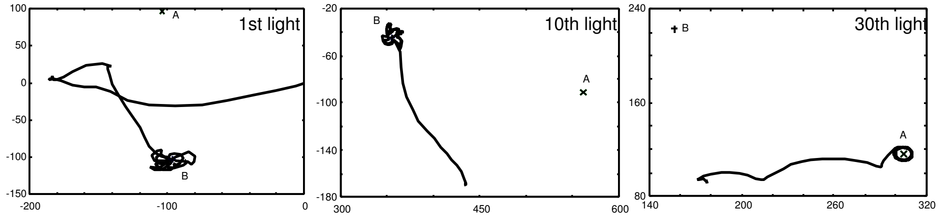

Figure 3: Spatial trajectories at the 1st, the 10th and the 30th presentation in Fig. 2.

Stopping plasticity. In order to confirm whether the plastic mechanism plays a role in the transition from one type of phototaxis to the other, plasticity is turned off at different points during the same 100 sequence of light presentations. The experiment is started with the same initial configuration of Fig. 2 and plasticity is turned off (weights remain at their current value) at the 20th or 50th presentation before the transitions take place. Figure 4 shows the distance to both light sources and the proportion of homeostasis.

As can be seen in both cases, the agent's behaviour does not switch in the period tested; it keeps approaching the light it preferred before stopping plasticity. While the behaviours are sustained, the internal dynamics is also maintained within the each region as much as before plasticity was stopped.

It should be noted that it is possible for a CTRNN to switch the internal dynamics into another region without changing the network weights (e.g., Phattanasri, Chiel & Beer, submitted). If so, it means that the agent's behaviour transition might be happening without explicit synaptic plasticity. However, it is clear that, at least in the examples explored, non-plastic transitions are not common. After stopping the plastic mechanism, neither the behaviours nor the internal dynamics seem to change. This shows that the homeostatic adaptation works as required and the homeostatic regions are associated with the phototactic behaviours as we expected even in novel situation with both sources constantly emitting light.

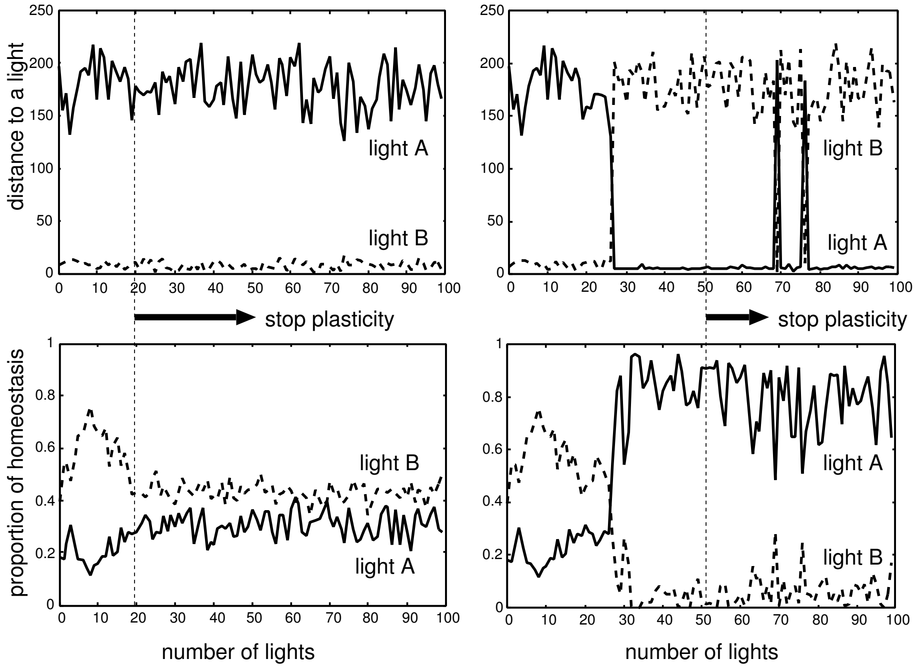

Figure 4: Final distances and proportion of homeostatic neurons when plasticity is stopped after the 20th or the 50th presentation. Other settings are same as Fig. 2.

5.3 What makes a preference change?

It was shown that plasticity drastically affects the possibility of behavioural transitions but this does not mean that the transitions are caused solely by factors internal to the agent. Generally, it is difficult to discuss causation in the context of complex embodied dynamical systems. However, in this section, we try to make a distinction between endogenous dynamics and externally-driven interactions in terms of the susceptibility to the environment in order to study the effect of different factors affecting the switch of preference.

5.3.1 Persistence of preference

We investigate the persistence of the preference for a light type. In this experiment, the positions of the two lights are swapped at some point during the approaching behaviour. If the agent has a consistent preference, the agent should seek the light it prefers and re-approach the target in its new position. This procedure resembles the one applied by Beer to study the agent's commitment to a behavioural outcome in his work on minimal categorical discrimination, (Beer, 2003).

The experiment is performed by using the same agent as before, which has interacted with the same environment of Fig. 2 until the 50th presentation at a point where it exhibits a preference for light A. In this trial, it takes approximately 750 timesteps for the agent to reach the target.

Starting from the same conditions, we investigate in detail the effect of a single position swap occurring at a timestep t during this trial. The approached light is recorded for all tested values of t. Figure 5 shows the final distance to the lights as a function of the timing where positions are swapped. It also shows the agent's trajectory when the light sources were swapped at $t = 150$ as an example.

This trial is taken from a period where the agent has a preference of light A, which means if nothing happens, the agent will approach light A. The result shows that swapping at the earlier timesteps allows the agent to persist in its preference and approach the target light. It can be considered that there is no big difference between the stimulus strength of both lights at earlier timesteps and the agent can still "pick" which light to go, following its preference. For position swaps occurring later, the change in environmental factors become stronger and the agent will approach A or B depending on its current orientation. If positions are swapped after 350 timesteps, the agent is close enough to receive a very strong stimulus from light B after swapping which seems to produce a transition and so the agent's behaviour is now to stay around light B. This single instance is enough to show that strong enough variations in environmental factors (e.g., orientation, strength of stimuli) can produce a change of preference in the agent's behaviour. Interestingly, there is an indeterminate period around $t = 340$ where phototaxis A and B are mixed. We have not investigated this period deeply yet but it could be expected that ambiguous environmental factors and/or internal dynamics produce a conflict.

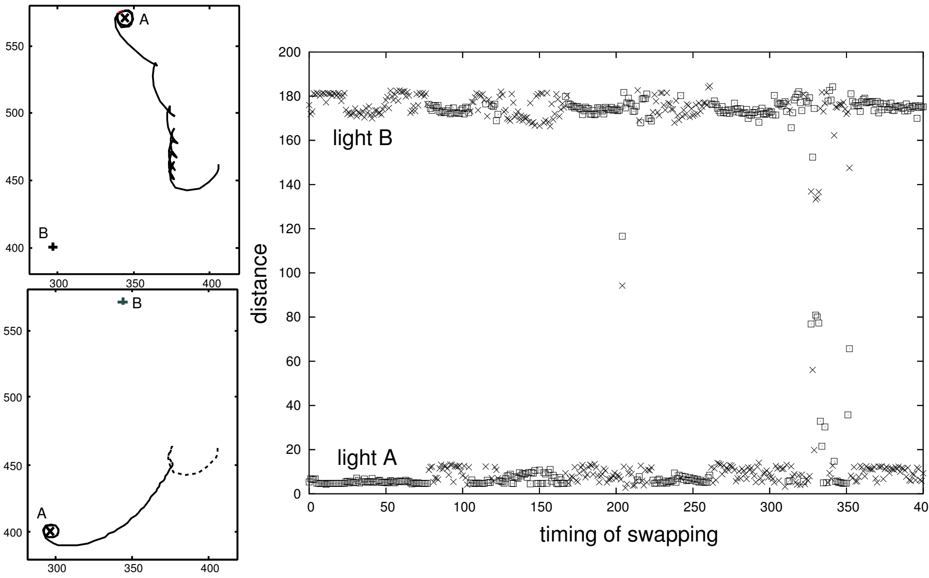

Figure 5: Right: Distance to the lights (squares for light A and crosses for B) as a function of the time when light positions are swapped. The agent has interacted until the 50th presentation shown in Fig. 2. The distances are recorded after 800 timesteps. Left: Trajectory when the light is swapped at 150 timesteps. The upper plot shows the original trajectory (lack of smootheness indicates sharp angular turns).

The swapping experiments show both the persistence of the preference and the influence of the environment. To study the persistence more clearly, another experiment is performed. The swapping experiments are reproduced as before but in addition one of the two lights vanishes. The following three cases are tested: (B→A, A→x), (A→B, B→x), and (A→x, B→B). The distances to the only remaining light after 800 timesteps are shown in Fig. 6.

The notation (B→A, A→x) indicates that, at time t, the position of light B is changed to that of light A and the original light A disappears. Therefore, the agent suddenly sees light B on the way to light A (and no other lights). As with the swapping experiments, if the agent is close enough to get strong stimulus from light B, the agent changes the preference to light B and remains close to it, which can be seen in the corresponding plot. However, at earlier values of the swapping time t, the agent *does not approach* light B even if the agent is on the way to the position where A was. Notice that the behaviour of approaching a single light has been explicitly selected during evolution. This is a situation where the agent should approach whatever single light is available. In spite of this, the agent does not approach the solitary light B. This is very strong evidence in support for an endogenous sustaining of the A preference.

In the case of (A→B, B→x), the agent keeps approaching light A while following the preference even if the light is placed in a different place. This is expected. In contrast, in the case of (A→x, B→B), light A disappears and light B remains in the same position. Both in the previous case and in this one only one light remains at the original B position. The only difference is that this light is A in the previous case and B in this one. Any other neural states, body direction, etc. are same. Therefore the agent receives the same intensity of light B in the latter case as of A in the former, where the agent could eventually locate light A. This implies that the agent must now be able to sense light B. However, the agent does not approach it for most values of t. As in the first case, this result shows a strong persistence of the preference even when the external situation should be expected to change it.

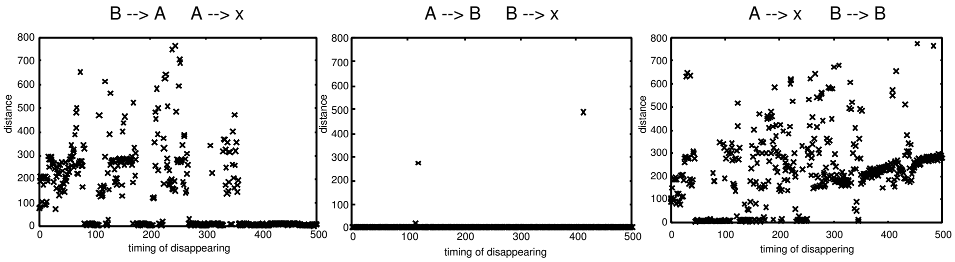
B --> A   A --> x

A --> B   B --> x

A --> x   B --> B

Figure 6: Distances to the remaining light as a function of the timing where one of two lights vanishes (and positions are swapped). The graphs are made in the same way of Fig. 5. The notation above the plots is explained in the text.

5.3.2 Effects of reducing external variability

Although the importance of endogenous factors in the formation and persistence of preference has been established, the internal dynamics are not independent of the history of environmental coupling. The agent creates the preferences *through interactions*. Random elements such as the position of the lights with respect to the orientation of the agent that can have an effect on the probability of switching to a different behaviour. To further show that the environmental factors do affect the preference, the agent is tested in an environment with fewer fluctuations where each new pair of lights always appears in a same position relative to the body, which is at distance of 130 on a bearing of $\pi/4$ to the right (left) of the agent's current heading for light A (light B). Under this condition, we counted how many transitions between preferences take place during 100 trials (presentation of a pair of light). Successive visits of more than 3 to the same light type was defined as evidence of a sustained preference. On an average over 100 independent runs, 0.25 transitions took place during 100 trials in this condition in contrast to 1.16 transitions in the normal condition. A reduction in environmental fluctuations produces the stabilization of the sensorimotor flow and then the preference is also stabilized. We should also notice that the number of transitions is reduced but not to zero, thus indicating that changes do occur due to the interaction between internal dynamics and environment even when external uncertainty (in light position) is removed.

A different test was carried out in an environment where the lights sources remain fixed on the same place. We test the agent in this configuration for a time corresponding to the 100 consecutive presentations of lights as in the above experiments. On average over 1000 independent runs, transitions take place 0.039 times during this period. Again we obtain a drastic reduction of transitions by removing external variability. In this particular case, the reduction is achieved by the agent remaining close to a very strong source of stimulation, which in some sense seems like the most difficult condition to switch away from. However, we notice that as before the number of transitions is not fully reduced to zero even in these extreme conditions.

5.3.3 Transitions between high and low susceptibility to perturbations

The emerging picture is one where *both* endogenous dynamics and environmental factors play a role in the sustaining and the changing of a preference. Does it still make sense to ask the question of whether the choice that an agent actually makes corresponds to a spontaneous or externally-driven "decision"? Put in these traditional terms, the answer is no. However, it is possible to capture in more detail the dynamical relation between the different factors in order to formulate a clearer distinction. This distinction is made operational by observing the agent's potential behaviours in different situations departing from a same initial state. If the agent "decides" to go to one of the lights by a preference that is endogenously sustained, its behaviour must be robust to variations in environmental factors. On the contrary, "decisions" that are highly-affected by environmental variation can be attributed to the role played by external factors. We label the two possibilities respectively as *strong* and *weak commitment* to a choice.

Based on this idea, we select the agent's states (neural and bodily) corresponding to different times in Fig. 2. For each selection of initial states, we record which light is approached by the agent as a function of different initial angular positions of the two lights (both placed at the same distance). This is the closest we can get in the present setting to a quantitative measure of preference. The results are shown by different shades of gray for the final destination in Fig. 7. In the case of (a), in which the agent originally has the preference of light B, the "decision" is stable against the various initial positions of the lights. The agent robustly approaches light B for practically all the angular positions tested. Therefore, the "decision" to approach B does not depend on environmental variation in this case. (As demonstrated in the swapping experiments, in some of these cases, the agent actively avoids the light that is presented directly in front of it and searches for the alternative light even if its position is such that no stimulus from that light is directly impinging on the sensors). The same is also true in the case of (c). Except for the small region where it selects light B, the agent approaches light A wherever else the lights are placed. By contrast, in cases (b) and (d) the agent is rather "uncommitted" because the approached target changes depending on the lights' position.

In order to see how this dependency of the "decision" on environmental variability changes during the history of interactions, the proportions of dark grey (light A) and light grey (light B) circles in the plots of Fig. 7 is calculated (Fig. 8) for different times corresponding to times in Fig. 2 (which shows the actual choices taken). It is shown that when the agent has the preference for light B since the proportion of B-circles at the beginning of this period is high, which means that the agent effectively ignores light A and keeps approaching light B. It could be said that the agent's behaviour is committed in the sense that it does not depend on the environmental factors as mentioned above. When the preference changes from light B to A and a while after that, the proportions stay around 0.5, which means that the agent does not have a strong commitment to which light should be selected as target. There is ample scope for environmental factors to alter the agent's behaviour. Then, the choice of light A changes towards a more stable (or committed) dynamics.

We are not implying with these results that during periods of weak environmental dependence, the endogenous dynamics are solely responsible for the agent's performance. In *all* cases, behaviour is the outcome of a tightly coupled sensorimotor loop. It is clear that the mode of environmental dependence, whether weak or strong, changes over time and that this is a property of the agent's own internal dynamics as well as the history of interaction. During the periods of high susceptibility to external variations, the agent is highly responsive to environmental variability resulting in less commitment towards a given target. By contrast, during periods of weak susceptibility, the consistent selection of a target is a consequence of low responsiveness to environmental variability.

The important point is that the autonomy of the agent's behaviour can be seen as the flow of alternating high and low susceptibility, which is an emergent property of the homeostatic mechanism in this case (but might be the result of other mechanisms in general). There is nothing apart from the flow of neural and sensorimotor dynamics that stands for a mode of commitment to a preference or other. No internal functional modules, no external instructions. Nevertheless, the existence of the different modes can be determined and measured. It should be made clear that this picture is quite in contrast with the idea that autonomy may be simply measured as how much of behaviour is determined internally vs. how much is externally-driven. Strong autonomy (in this context the capability of defining one's own goals) is orthogonal to this issue since simply all of behaviour is conditioned by both internal and external factors at all times. It is the mode of responsiveness to variations in such factors that can be described as committed or open, and it would be a property of strong autonomous systems that they can transit between these modes (maybe in less contingent ways as this agent, e.g., in terms of needs, longer-term goals, etc.).

6 Discussion

Though minimal, the dynamical model discussed in this paper exhibits some important aspects of behavioural preferences such as durability and transi-

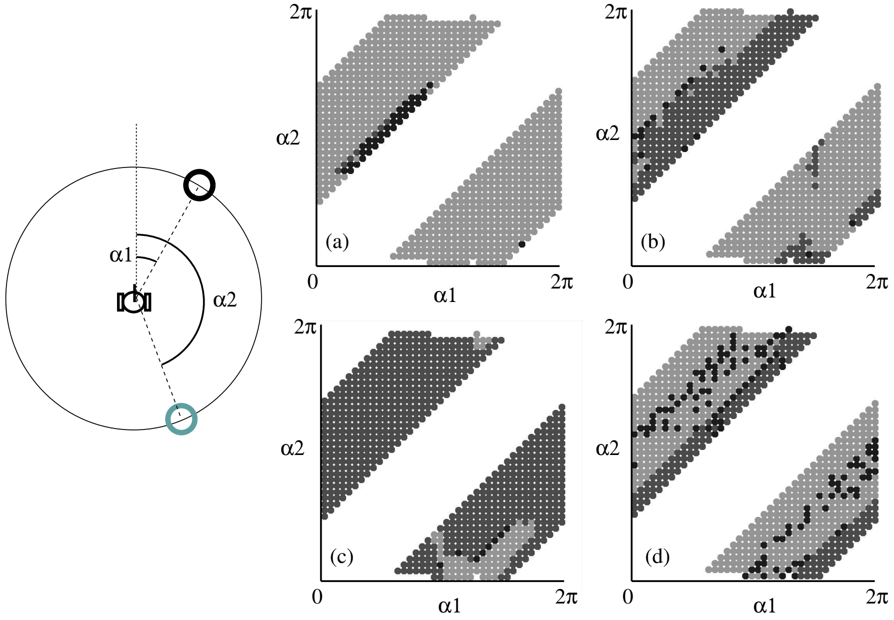

Figure 7: Light preference of the agent corresponding to the states of (a)20, (b)25, (c)50 or (d)95 in Fig. 2, against different light positions. Horizontal and vertical axes indicate the initial angles of lights A and B relative to the agent's orientation respectively. The positions of lights whose difference is less than $\frac{\pi}{2}$ are removed in order to better determine which light the agent is approaching. The dark grey circles show that the agent approaches light A. The light grey circles correspond to light B and black shows the agent does not approach either of lights.

tions through mutual constraining of internal and external dynamics. We hasten to say that the understanding of preferred behaviours is not exhausted by the dynamical properties explored in this paper. In particular, preference in natural cognitive systems implies a capability to appreciate the *value* of a choice. The current model does not capture this extra level of complexity, though a dynamical approach to value-generation may indeed be possible (Di Paolo, 2005) and it may be close to the constraining of dynamical modes shown in the present model. Much less justice is done to the problem by traditional approaches where preferences (or motivations, moods, etc.) are reified as internal variables and with which the view presented here should be contrasted.

The model shows the effectiveness of the dynamical approach in allow-

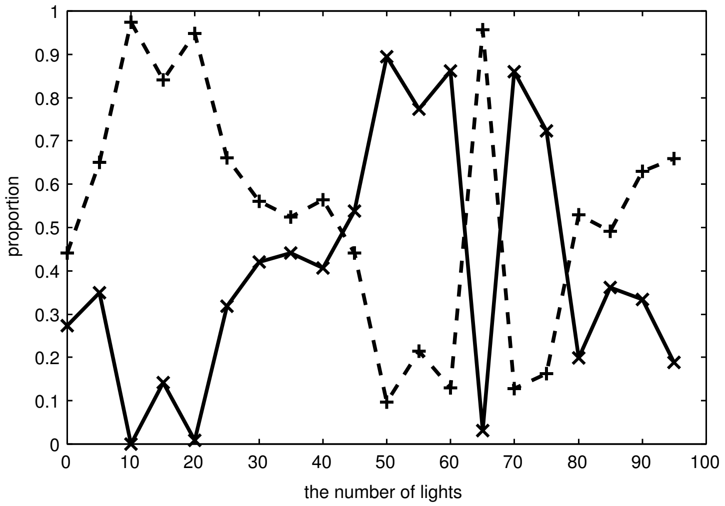

Figure 8: The proportion of how much dark grey (solid line: light A) and light grey (dotted line: light B) circles appear in the plots of Fig. 7 and others like them corresponding to different times. The values of the horizontal axis corresponds to the number of the lights the agent has interacted with in Fig. 2.

ing the posing of the right questions and suggesting solutions. It is possible to clearly formulate some constitutive properties of preference and the specification immediately suggests a road towards a minimal design. In particular, it seems that what is necessary is a process that can link more than one mode of internal stable dynamical flows with the corresponding interactional dynamics. This link can successfully produce the property of persistence without permanence in our model. It also makes concrete the description of the mutual constraining between two levels (neural dynamics, agent/environment interactions). Consequently, the model lends support to Goldstein's and Merleau-Ponty's multi-causal view of preferred behaviour by producing a case study of its viability and providing insights into its dynamical basis. We have shown the logical consistency of the view that persistence of preference and their transitions cannot be attributed either to internal or external factors on their own. And yet, there is a sense in which we can say that internal and interactive dynamics are implicated in generating stronger or weaker degrees of commitment.

The model allows us to make some observations.

It is important to notice that in several cases the agent keeps searching for a target light even when it is no longer detectable and avoids approaching the alternative target that is present. Here it is possible to draw a parallel between the persistence of preference and the similar phenomenon of object permanence observed in infant experiments by Piaget (1954). According to Piaget, the concept of an object as something that has an ongoing existence independent of the observer is not immediately given. The lack of such concept is his explanation of the famous A-not-B error in which 7-12 month old infants search for a toy, not in the location they have just seen it being hidden, but in the location where they searched for it in the past. However, there are alternative more parsimonious explanations. The perseverance observed in such cases may respond to simpler dynamical mechanisms (often described as motor memories) (Thelen, Schoener, Scheier, & Smith, 2001). A recent study by Wood and Di Paolo (Wood & Di Paolo, 2007) actually demonstrates that the *same* homeostatic mechanisms used in the current study (two homeostatic regions corresponding to two behavioural options) is sufficient to reproduce the pattern of errors observed in infants and their disappearance with age, lending support to the idea that motor memories are at the basis of perseverance and providing a plausible sensorimotor explanation for the origins of the higher cognitive capacity of object permanence.

The method shows another use for homeostatic adaptation in combination with evolutionary techniques in shaping both behavioural and internal requirements for the neural/body/interaction system. However, we may ask whether this is the simplest model of behavioural preference. It seems *a priori* a good idea to attempt the same experiments using non-plastic CTRNNs with homeostasis, or even without it. We predict that the latter case would be brittle, i.e., largely driven by the environmental configurations, and so show little or no persistence of preference. The other case, non-plastic CTRNNs with homeostasis is less easy to predict. Systematic comparisons along these lines are planned in order to establish the role played by structural plasticity.

On analyzing the role of environmental factors (position of lights) on the choice of target we find the agent presents different degrees of openness to environmental variability, or put differently, commitment to a target goal. These assertions can be made operational in terms of the stability of the dynamics. The results allow us to formulate the following dynamical hypothesis: *A switch in preference will take place not necessarily as specific internal variables acquire specific values, but rather as a result of changes in the stability landscape of the neural and sensorimotor dynamical flows between committed and open modes.* In some modes, the landscape will be highly stable to environmental variability. Alternative choices and opportunities for behavioural change will not affect the behavioural and neural flow. The agent may even be "blind" to stimulations corresponding to these alternative behaviours. Such are the committed modes. In other cases, even if the actual behaviour shows a stable trajectory towards a target, the sensitivity to environmental variability may be higher. These are the open or un-committed modes which may result, in the appropriate circumstances in a change of preference. Between the two modes lies a spectrum of intermediate possibilities.

The results justify the choice of the Spinozist inspiration for the design of our model. This view allows us to pose the problem of preferred behaviour in dynamical terms, and of the change of preference in terms of change of conatus. In turn, the operationalization of commitment proposed in this paper feeds back into the task of understanding conatus dynamically. In this way, a dynamical systems approach to preferences (and associated cognitive phenomena such as decision making) looks for global dynamical properties at the internal and interactive levels to determine whether a behaviour is preferred or not, chosen with strong or weak commitment. Of course, in many cases the specific determination of preference carried out in this paper (resetting the agent to a given state and altering its environment) may be hard or impossible to achieve. In such cases, alternative or derivative operationalizations will be required.

As a final note on autonomy, it is clear that in our model the achievement of committed or open modes of sensorimotor flows is done through the history of interaction by the agent itself. However, the fact remains that its autonomy is severely limited by the arbitrary imposition of the two internal homeostatic regions. We believe that in reality the condition of using a strict region corresponding to zero plasticity may be relaxed and that the dynamics may consist of moving gradients of plasticity and the spontaneous formation of highly stable regions where plastic change is small and in general pointing back into the same stable region. Designing an agent where such homeostatic regions are themselves the consequence of the agent's own activity will be a further step towards strongly autonomous behaviour. In a sense, such an agent will not only be switching spontaneously between a choice of externally-provided goals, it will be creating its own goals as a consequence of its history of interactions.

Acknowledgments: We are very grateful to two anonymous reviewers for their helpful comments. This research was partially supported by the Japanese Ministry of Education, Science, Sports and Culture, Grant-in-Aid for JSPS Fellows, 17-4443.

References

Ashby, W. R. (1960). *Design for a brain: The origin of adaptive behaviour (Second edition).* London, Chapman and Hall.

Beer, R. D. (1990). *Intelligence as adaptive behavior: An experiment in computational neuroscience.* San Diego: Academic Press.

Beer, R. D. (1996). Toward the evolution of dynamical neural networks for minimally cognitive behavior. In Maes, P., Mataric, M. J., Meyer, J.-A., Pollack, J. B., & Wilson, S. W. (Eds.), *From Animals to Animats 4: Proceedings of the 4th International Conference on Simulation of Adaptive Behavior*, pp. 421–429. Cambridge, MA: MIT Press.

Beer, R. D. (1999). Arches and stones in cognitive architecture. *Adaptive Behavior*, *11*(4), 299–305.

Beer, R. (2003). The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior*, *11*(4), 209–243.

Bennett, J. (1984). *A study of Spinoza's Ethics.* Cambridge University Press.

Bennett, J. (1992). Spinoza and teleology: A reply to Curley. In Curley, E., & Morean, P. F. (Eds.), *Spinoza, issues and directions*, pp. 53–57. Leiden:E.J. Brill.

Bourgine, P., & Stewart, J. (2004). Autopoiesis and cognition. *Artificial Life*, *10*, 327–346.

Breazeal(Ferrell), C. (1998). A motivational system for regulating human-robot interaction. In *Proceedings of AAAI-98*, pp. 54–61. Madison, WI.

Bryson, J. (2003). Action selection and individuation in agent based modelling. In Sajlach, D. L., & Macal, C. (Eds.), *The proceedings of Agent 2003: Challenges of Social Simulation*, pp. 317–330.

Clark, A. (1997). *Being there.* Cambridge, MA: MIT Press.

Curley, E. (1990). On Bennett's Spinoza: The issue of teleology. In Curley, E., & Morean, P. F. (Eds.), *Spinoza, issues and directions*, pp. 39–52. Leiden:E.J. Brill.

Di Paolo, E. A. (2000). Homeostatic adaptation to inversion in the visual field and other sensorimotor disruptions. In Meyer, J., Berthoz, A., Floreano, D., Roitblat, H., & Wilson, S. (Eds.), *From Animals to Animats VI: Proceedings of the 6th International Conference on Simulation of Adaptive Behavior*, pp. 440–449. Cambridge, MA: MIT Press.

Di Paolo, E. A. (2005). Autopoiesis, adaptivity, teleology, agency. *Phenomenology and the Cognitive Sciences*, *4*(4), 429–452.

Goldstein, K. (1995/1934). *The organism.* New York: Zone Books.

Harvey, I., Di Paolo, E., Wood, R., Quinn, M., & Tuci, E. A. (2005). Evolutionary Robotics: A new scientific tool for studying cognition. *Artificial Life*, *11*(1-2), 79–98.

Humphrys, M. (1997). *Action selection methods using reinforcement learning (PhD thesis).* University of Cambridge.

Ito, M., Noda, K., Hoshino, Y., & Tani, J. (2006). Dynamic and interactive generation of object handling behaviors by a small humanoid robot using a dynamic neural network model. *Neural Networks*, *19*, 323–337.

Jonas, H. (1966). *The phenomenon of life: Towards a philosophical biology.* Northwestern University Press.

Juarrero, A. (1999). *Dynamics in action: Intentional behavior as a complex system.* MIT Press.

Kelso, J. A. S. (1995). *Dynamic patterns: The self-organization of brain and behavior.* MIT Press.

Matson, W. (1977). Death and destruction in Spinoza's ethics. *Inquiry*, *20*, 403–417.

Maturana, H., & Varela, F. (1980). *Autopoiesis and cognition: The realization of the living.* Boston, Reidel.

Merleau-Ponty, M. (1962). *Phenomenology of perception.* (Colin Smith, Trans.). London: Routledge & Kegan paul.

Phattanasri, P., Chiel, H., & Beer, R. The dynamics of associative learning in evolved model circuits. *submitted.*

Piaget, J. (1954). *The construction of reality in the child.* New York: Basic Books.

Stewart, J. (1992). Life=cognition: The epistemological and ontological significance of artificial life. In Bourgine, P., & Varela, F. (Eds.), *Toward a Practice of Autonomous Systems: proceedings of the first European conference on Artificial life*, pp. 475–483. MIT Press.

Thelen, E., Schoener, G., Scheier, C., & Smith, L. (2001). The dynamics of embodiment: A dynamic field theory of infant perseverative reaching. *Behavioural and Brain Sciences*, *24*, 1–86.

Thelen, E., & Smith, L. B. (1994). *A dynamic systems approach to the development of cognition and action.* Cambridge, MA: MIT Press.

Thompson, E., & Varela, F. (2001). Radical embodiment: Neural dynamics and conscious experience. *Trends in Cognitive Science*, *5*(10), 418–425.

Varela, F., & Thompson, E. (2003). Neural synchrony and the unity of mind: A neurophenomenological perspective. In Cleeremans, A. (Ed.), *The Unity of Consciousness Binding, Integration, and Dissociation*, pp. 266–287. Oxford: Oxford University Press.

Velásquez, J. (1997). Modeling emotions and other motivations in synthetic agents. In *Proceedings of AAAI-97*, pp. 10–15.

Weber, A., & Varela, F. (2002). Life after Kant: Natural purposes and the autopoietic foundations of biological individuality. *Phenomenology and the Cognitive Sciences*, *1*, 97–125.

Wheeler, M. (1997). Cognition's coming home: The reunion of life and mind. In Husbands, P., & Harvey, I. (Eds.), *Proceedings of the Fourth European Conference on Artificial Life*, pp. 10–19. MIT Press, Cambridge, Mass.

Williams, H. (2004). Homeostatic plasticity in recurrent neural networks. In Schaal, S., Ijspeert, A., Billard, A., Vijayakumar, S., & Meyer, J. (Eds.), *From Animals to Animats 8: Proceedings of the 8th International Conference on the Simulation of Adaptive Behavior*, pp. 344–353. Cambridge MA: MIT Press.

Wood, R., & Di Paolo, E. (2007). New models for old questions: Evolutionary robotics and the 'A not B' error. In *Proceedings of the 9th European Conference on Artificial life (in press)*. Springer-Verlag.