Neural Networks 33 (2012) 76–87

Contents lists available at SciVerse ScienceDirect

Neural Networks

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

Self-organizing high-order cognitive functions in artificial agents: Implications for possible prefrontal cortex mechanisms

Michail Maniadakis [a,b,*], Panos Trahanias [b], Jun Tani [a,c]

[a] Brain Science Institute, RIKEN, Wako-Shi, Saitama, Japan
[b] Institute of Computer Science, FORTH, Heraklion, Crete, Greece
[c] Electrical Engineering Department, KAIST, Yuseong-gu, Daejeon, Republic of Korea

ARTICLE INFO

Article history:
Received 5 January 2010
Received in revised form 23 October 2011
Accepted 6 April 2012

Keywords:
Neural networks
Executive control
Meta-cognition
Prefrontal cortex
Cognitive robot
Attractor
Self-organization

ABSTRACT

In our daily life, we often adapt plans and behaviors according to dynamically changing world circumstances, selecting activities that make us feel more confident about the future. In this adaptation, the prefrontal cortex (PFC) is believed to have an important role, applying executive control on other cognitive processes to achieve context switching and confidence monitoring; however, many questions remain open regarding the nature of neural processes supporting executive control. The current work explores possible mechanisms of this high-order cognitive function, transferring executing control in the domain of artificial cognitive systems. In particular, we study the self-organization of artificial neural networks accomplishing a robotic rule-switching task analogous to the Wisconsin Card Sorting Test. The obtained results show that behavioral rules may be encoded in neuro-dynamic attractors, with their geometric arrangements in phase space affecting the shaping of confidence. Analysis of the emergent dynamical structures suggests possible explanations of the interactions of high-level and low-level processes in the real brain.

1. Introduction

A well-known experiment investigating executive control functions and, more specifically rule switching is the Wisconsin Card Sorting Test (WCST), Berg (1948) and Milner (1963), where subjects are asked to discover and apply a card sorting rule based on reward and punishment feedback. At unpredictable times during the task, the rule is changed by the experimenter and must be re-discovered by the subjects. The ordinary WCST can be further enriched with the option of betting on behavioural outcomes (i.e., success or failure of sorting). The WCST-with-Betting (WCSTB) tests the capacity of subjects to monitor and implement confidence about the currently adopted rule, Koren et al. (2005) and Koren, Seidman, and Harvey (2006). This is a high-level cognitive task which requires coordinating a range of different processes, including the maintenance of working memory for the currently followed rule, the examination of conflicts between the adopted rule and the reward or punishment feedback, higher level executive control for rule adjustment, self-monitoring for confidence development, betting decisions and the generation of physical actions on the basis of the selected rule.

Existing modelling studies on WCST employ discrete and algorithmic computational processes, based on the common assumption that, although the posterior cortices can be characterized as fundamentally analog systems, the prefrontal cortex (PFC) has a more discrete, digital character, Dayan (2007) and O'Reilly (2006). Previous modelling studies, e.g. Dehaene and Changeux (1991) and Stemme, Deco, and Busch (2007a), employ local and discrete neural network representations where currently adopted rules are represented by separately activated local units. Rougier and O'Reilly (2002) proposed an on–off type gating operation that acts on working memory for storing the currently adopted rules. The essential idea is that the neural activation patterns representing the current rules in the working memory can be preserved by closing the gate until the rules are in conflict with the new rule selected by the experimenter. Dayan (2007) generalized this model to deal with various executive control functions, assumed to be present in PFC, employing a computational scheme of conditional rule matching and action execution. Other relevant models interpreting computationally human assumptions about rule switching work also on the basis of discrete states for rule representation based on either petrinets, Narayanan (2003), or pools of excitatory and inhibitory neurons, Stemme, Deco, and Busch (2007b), or Hopfield neural network with a separate hypothesis generation module, Kaplan, Sengor, Gurvit, Genc, and Guzelis (2006).

An alternative approach regards implementing cognitive capacities based on dynamic neural mechanisms. In this direction,

* Corresponding author at: Institute of Computer Science, FORTH, Heraklion, Crete, Greece.
E-mail address: mmaniada@ics.forth.gr (M. Maniadakis).

0893-6080/$ – see front matter © 2012 Elsevier Ltd. All rights reserved.
doi:10.1016/j.neunet.2012.04.005

a variety of computational models have interpreted computationally many of the well-known PFC functionalities. For example, dynamic working memory models are investigated in Botvinick and Plaut (2006), showing that recurrent neural networks can adequately accomplish serial recall tasks considering also the effects of background knowledge. Additionally, Machens, Romo, and Brody (2005) have investigated interval discrimination tasks showing that attractor dynamics can effectively combine memory maintenance and decision making, while Johnson, Spencer, Luck, and Schner (2009) have implemented an analogous model for visual working memory.

The discussion above highlights the main directions in the open debate regarding the discrete or dynamic nature of PFC processing (see also Brody, Romo, and Kepecs (2003)). The first type of models have been mainly inspired by experimental data showing active and relatively long lasting neural activity that may encode rules in working memory. However, an important drawback for the discrete models concerns how the static (usually bi-stable) representations can link executive control with the non-fully predefined and inherently continuous behaviors of the agent. The dynamic approach has gained significant support from experimental works showing that PFC processing is based on time-dependent activation patterns, Romo and Salinas (2003) and Singh and Eliasmith (2006), as well as dynamic interaction networks in the brain, Palva, Monto, Kulashekhar, and Palva (2010). Especially for rule switching, experimental electrophysiological data from monkeys trained to perform WCST showed that the DLPFC cells encode rules through dynamically changing neural activities, Mansouri, Matsumoto, and Tanaka (2006). The observed dynamical patterns may be ascribed to the cognitive processes taking place when accomplishing delayed response tasks such as external stimuli processing, decision making, response planning, motion execution monitoring etc. (see Jun et al. (2010) for multi-process coordination in PFC). The studies mentioned above indicate that PFC internal mechanisms are based on a dynamic rather than a stationary pattern of neural activity. In other words, active maintenance of neural activity does not necessarily mean static representations. The present work aims to examine whether the executive control functions involved in WCSTB can be implemented on the basis of dynamic processing and whether such a possibility provides the basis for a new understanding of high-level cognitive functions. More specifically, the dynamic modelling approach suggests the explanation for the confidence and preference people show for certain situations (i.e. we feel more confident when turning at a street intersection in our hometown than at a street intersection of another town), as well as how our minds organize rules into classes using some type of similarity criteria.

The computational exploration of alternative mechanisms can be based on evolutionary robotic experiments similar to Borrett, Jin, and Kwan (2005). This is because the real-time environmental interaction may provide more realistic and general explanations on executive control processes compared to the purely theoretically operating existing models. In particular, the current work employs a minimum constraint modelling approach to explore possible mechanisms of executive control functionality self-organized in simple neural network models achieving the WCSTB task. If the mechanisms for accomplishing the task consistently appear in statistically independent simulation runs, comparable principles may also operate in real brains, Ruppin (2002). In short, neural network models with recurrent connectivity are evolved to accomplish a robotic version of the WCSTB, using a standard genetic algorithm to search for optimal synaptic weights, Lipson (2005) and Nolfi and Floreano (2000). Following this approach, neural dynamics are free to self-organize in any appropriate way, revealing new and potentially more natural mechanisms for explaining high-level cognition, Baev (2007). We study the successfully evolved neural network models identifying their common internal characteristics, in order to provide suggestions of possible working principles in the brain.

In contrast to previous studies that focus on WCST exploring pure rule-switching, the option of betting that is additionally investigated in the present work provides a means for the deeper exploration of executive control functions. In particular, our experiments investigate the self-awareness capacity of the artificial agent that regards monitoring (i) the current behavioural context (i.e. the agent being in either a rule exploration or a rule following mode) and (ii) the confidence that the agent feels for each behaviour and how the latter affects its betting strategy.

The rest of the paper is organized as follows. In the next section we present the methodology followed in our work. In particular, we present the Continuous Time Recurrent Neural Network (CTRNN) model used in our study, how it is connected to the sensors and actuators of the simulated robotic agent, the computational counterpart of the WCSTB problem, and the evolutionary procedure used to explore configurations of CTRNN robot controllers. Experimental results addressing robot switching and betting on the basis of three alternative behavioural rules are presented in the following section. Then, a detailed discussion highlights the main finding of our computational experiments, formulating suggestions for the organization of biological executive control processes. Finally, conclusions and suggestions for further work are presented in the last section.

2. Experimental methodology

In order to investigate executive control dynamics, we have designed a robotic task that resembles the Wisconsin Card Sorting test, incorporating also a betting option for the artificial agent. The task investigates rule switching in a sample–response paradigm, similar to Joel, Weiner, and Feldon (1997). The agent has to learn three sample–response rules, selecting, applying and re-selecting each one of them, as indicated by reward and punishment signals provided by the experimenter. The three available rules, named Same Side (SS), Opposite Side (OS) and No Response (NR), are described briefly in Fig. 1. The robot starts always from the bottom of the T-maze environment, responding to the side of light presentation. According to the Same-Side (SS) rule, the agent must turn left if the light source appeared at its left side, and it must turn right if the light source appeared at its right side. According to the Opposite-Side (OS) rule, the robot has to turn to the opposite direction of the light side, i.e. right when light appears to the left, and left when light appears to the right. In the case of the No Response (NR) rule, the robot should ignore the side of light, staying close to the starting position.

The task explored in the current study is separated in phases, each one consisting of several sample–response trials. While performing the trials of a given phase, the robotic agent has to discover and follow the sample–response rule that is assigned to the phase. Correct responses are rewarded, while incorrect ones are punished (see Fig. 1). At the beginning of each trial, the agent bets on the success of the underlying response, having the opportunity to gain some profit.

Different phases correspond to different rules, which requires the agent to switch the adopted response strategy. Changes from one phase to another are performed by the experimenter in a random manner. This results in unpredictable rule changes that make the agent give spontaneous incorrect responses. Therefore, the agent has to develop mechanisms that consider rule changes, switch the adopted response strategy and efficiently control the amounts of betting, in accordance with the dynamically changing