Synchronisation effects on the behavioural performance and information dynamics of a simulated minimally cognitive robotic agent

Renan C. Moioli, Patricia A. Vargas, Phil Husbands

Abstract Oscillatory activity is ubiquitous in nervous systems, with solid evidence that synchronisation mechanisms underpin cognitive processes. Nevertheless, its informational content and relationship with behaviour are still to be fully understood. Additionally, cognitive systems cannot be properly appreciated without taking into account brain - body - environment interactions. In this paper, we developed a model based on the Kuramoto Model of coupled phase oscillators to explore the role of neural synchronisation in the performance of a simulated robotic agent in two different minimally cognitive tasks. We show that there is a statistically significant difference in performance and evolvability depending on the synchronisation regime of the network. In both tasks, a combination of information flow and dynamical analyses show that networks with a definite, but not too strong, propensity for synchronisation are more able to reconfigure, to organise themselves functionally and to adapt to different behavioural conditions. The results highlight the asymmetry of information flow and its behavioural correspondence. Importantly it also shows that neural synchronisation dynamics, when suitably flexible and reconfigurable, can generate minimally cognitive embodied behaviour.

Keywords Synchronisation · Evolutionary Robotics · Oscillatory networks · Transfer Entropy · Kuramoto Model

1 Introduction

Oscillatory neural activity is closely related to cognitive processes and behaviour (Engel et al., 2001; Buzsaki, 2006).

Renan C. Moioli and Phil Husbands are with the Centre for Computational Neuroscience and Robotics (CCNR), Department of Informatics, University of Sussex, Falmer, Brighton, BN1 9QH, United Kingdom, (email: {r.moioli, p.husbands}@sussex.ac.uk). Patricia A. Vargas is with the School of Mathematical and Computer Science, Heriot-Watt University, Edinburgh, Scotland, EH14 4AS, United Kingdom, (email: p.a.vargas@hw.ac.uk)

More specifically, it has been claimed that the flexible synchronisation of firing neurons is a fundamental brain mechanism (von der Malsburg, 1981; Singer, 1993, 1999; Womelsdorf et al., 2007). Synchronisation may mediate the interactions between neurons and be an active mechanism of large-scale integration of neuronal assemblies, impacting on motor control and cognitive performance (Hatsopoulos et al., 1998; Varela et al., 2001; Jackson et al., 2003). Moreover, a range of pathological states are related to abnormal neuronal synchronisation regimes (Glass, 2001; Brown, 2003; Arthuis et al., 2009).

Whereas increasingly sophisticated computational and imaging techniques help us to observe and record from different physiological aspects of the nervous system, a great part of the challenge lies in the comprehension of this avalanche of data and its relationship with behaviour (Chialvo, 2010). In addition, the brain presents spontaneous electrochemical activity that is constantly shaped by the body's constraints and time-varying environmental stimuli (Katz, 1999). The brain, therefore, not only processes information but also produces it, and cognitive phenomena are a product of brain-body-environment interactions (Stewart et al., 2010). Situated artificial agents, in this sense, provide an appropriate experimental scenario to study the principles of intelligent behaviour (Boden, 2006).

There is a rich literature exploring the relationship between neural synchronisation and complex motor control in embodied (robotic) rhythmical behaviours (Taga, 1994; Ijspeert et al., 2005; Pitti et al., 2009), where synchronisation appears as a more intuitive underlying mechanism; however, to date there has been very little research on the wider issues of neuronal synchronisation in the generation of embodied cognitive behaviours. Therefore, using concepts from Evolutionary Robotics (Harvey et al., 2005; Floreano et al., 2008; Floreano & Mattiussi, 2008; Floreano & Keller, 2010), we explore the role of synchronisation in the performance of a simulated robotic agent during the execution of

two different minimally cognitive tasks: the first, a categorical perception task (Beer, 2003; Izquierdo, 2008; Dale & Husbands, 2010), in which the robot has to discriminate between moving circles and squares; the second, an orientation task, where the robotic agent has to approach moving circles with both normal and inverted vision, adapting to both conditions. These tasks were chosen for being currently regarded as benchmarks in the evolutionary robotics and adaptive behaviour communities, with categorical perception underpinning cognitive systems (Harnad, 1987).

The work reported here is intended to shed some light on the role of neural synchronisation in simple embodied cognitive behaviours. More specifically it aims to 1) test whether different degrees of coupling in the agent's oscillatory neural network, which will encourage more or less synchrony in the network dynamics, have an effect on the performance of the agent and 2) determine if there are circumstances in which more (or less) synchrony is better suited to the generation of adaptive behaviour in the context of the tasks studied.

The neuronal model employed is based on the Kuramoto Model of coupled phase oscillators (Kuramoto, 1984), which has been extensively studied in the Statistical Physics literature, with recent applications in a biological context due to its relatively simple and abstract mathematical formulation yet complex activity that can be exploited to clarify fundamental mechanisms of neuro-oscillatory phenomena without making too many a priori assumptions (Ermentrout & Kleinfeld, 2001; Cumin & Unsworth, 2007; Kitzbichler et al., 2009; Breakspear et al., 2010). The model explicitly captures the phase dynamics of units that alone have spontaneous oscillatory activity and once connected can generate emergent rhythmic patterns. Its synchronisation regime can be adjusted by one parameter (Strogatz, 2000), suiting our study, whilst also avoiding any problems in obtaining phase information (an issue in other models which consider frequency and amplitude dynamics (Pikovsky et al., 2001)). Phase relationships contain a great deal of information on the temporal structure of neural signals, are associated with cognition and relate to memory formation and retrieval (Li & Hopfield, 1989; Izhikevich, 1999; Varela et al., 2001; Kunyosi & Monteiro, 2009). Furthermore, it has been shown that firing and bursting neurons can be modelled as oscillators (Murray, 1989). Hence the Kuramoto model is highly relevant, at a certain level of abstraction, to modelling neural mechanisms underlying adaptive and cognitive behaviours.

Analysis of results is centred on how the information dynamics between the nodes of the network, the agent's body and the environment vary depending on the current synchronisation status of the system and how this is reflected in the behaviour being displayed. Information Theory provides a framework for quantifying and emphasizing the non-linear relationships between variables of the system, hence its suitability in Biology and Robotics studies (Rieke et al., 1997; Rolls & Treves, 1998; Lungarella & Sporns, 2006). In this context, agent-environment systems pose extra challenges in devising and interpreting a sensible measurement of information flow, for they normally have noisy and limited data samples, asymmetrical relationships among elements of the system and temporal variance (i.e. sensory and motor patterns may vary over time). Transfer Entropy, in this scenario, is suggested as a suitable and robust information theoretic tool (Lungarella et al., 2007b,a), and has also been applied to investigate real neural assemblies and other neuroscience problems (Borst & Theunissen, 1999; Gourévitch & Eggermont, 2007; Buehlmann & Deco, 2010; McDonnell et al., 2011; Vicente et al., 2011); it will, thus, be used in our analysis.

The paper is organised as follows: the next section presents some theoretical background to the main concepts explored in this work: oscillators and synchronisation, focusing on the Kuramoto Model, evolutionary robotics (ER) and the two minimally cognitive tasks studied, and Information Theory in an agent-environment context, describing Transfer Entropy. We then present the methods adopted to develop the experiments and analysis, covering the ER framework, the details of the active categorical perception and orientation under normal and inverted vision tasks, a description of the genetic algorithm used to optimize the parameters of the system, and concluding with details of the time-series analysis using Transfer Entropy. Following this, we present the results, which show a difference in performance and behaviour depending on the synchronisation regime of the network, relating the observed sensorimotor strategies with the dynamics and the information flow of the system. The paper closes with a discussion on the results obtained and future work proposals.

2 Theoretical Background

2.1 Oscillators, Synchronisation and The Kuramoto Model

In a study pioneered by Winfree (Winfree, 1980), the dynamics of a population of interacting limit-cycle oscillators have been approximated by a population of interacting phase oscillators, leading to a mean-field approximation model extensively explored by Kuramoto (Kuramoto, 1984). In his approach, the phase of each oscillator is determined by its natural frequency (drawn from some distribution) modulated according to a function that represents its sensitivity to the phase in every other node (see Equation 1)

$$\frac{d\theta_i}{dt} = \omega_i + \sum_{j=1}^{N} \Gamma_{ij}(\theta_j - \theta_i), i = 1, \ldots, N. \tag{1}$$

where: θ_i is the phase of the ith oscillator, ω_i is the natural frequency of the ith oscillator, Γ_{ij} represents the interaction between nodes and N is the total number of oscillators.

If one considers the interaction function to be periodic, i.e. $\Gamma_{ij}(x + 2\pi) = \Gamma_{ij}(x)$, it can then be expanded into a Fourier series. Considering only the first term of this series, and specifying that the network is symmetrically coupled, the previous model reduces to Equation 2, which is known as the Kuramoto Model:

$$\frac{d\theta_i}{dt} = \omega_i + \frac{K}{N} \sum_{j=1}^{N} sin(\theta_j - \theta_i), i = 1, \ldots, N. \qquad (2)$$

Basically, if the frequency of all possible pairs of nodes i and j $(i, j = 1, 2 \ldots n)$ are equal, i.e. $d\theta_i - d\theta_j = 0$ or $\theta_i - \theta_j = constant$, the model is said to be globally synchronised. It is possible to calculate a synchronisation index, which gives a good idea of how synchronised the set of oscillators are (Kuramoto, 1984). Consider Equation 3, where r is the synchronisation index (1 meaning high synchronisation, 0 meaning incoherent oscillatory behaviour) and ψ is the mean phase of the system.

$$re^{i\psi} = \frac{1}{N} \sum_{j=1}^{N} e^{i\theta_j} \qquad (3)$$

One of the most interesting properties of the Kuramoto Model is that by varying the value of the coupling between units the network can be tuned in subcritical, critical or supercritical regions. In the subcritical region, the interactions between oscillators are very weak, leading to an almost total lack of any collective behaviour. In the supercritical region, the coupling is so strong that the whole system synchronises and behaves like a single giant oscillator. The critical region is characterized by a second-order phase transition region between the subcritical and the supercritical regions; it tends to exhibit complex, unstable dynamics. It has been claimed in recent years that systems operating in this critical region - and there is some evidence to suggest this includes biological neuronal networks - have better performance on a variety of tasks and process information in a highly efficient way (Beggs, 2008; Kitzbichler et al., 2009).

In the Kuramoto Model, the critical coupling value i.e. the value of the coupling between units that will induce the network to remain in this critical region is given by Equation 4:

$$K_c = \frac{2}{\pi g(0)} \qquad (4)$$

where $g(0)$ is the value of the distribution of natural frequencies $g(\omega)$ formed by each oscillator's natural frequency ω calculated at $\omega = 0$.

In this paper, the probability densities of the discrete frequency distributions are estimated using a smoothing kernel approach with automatic choice of window width, but other techniques may be used (Silverman, 1998).

The details of how the Kuramoto Model is incorporated into the framework developed here will be presented in Section 3. The next subsection introduces the information theoretic tool that is used to support the analysis of experimental results.

2.2 Information Theory and Transfer Entropy

Previous work has explored the relationship between information dynamics and behaviour in real and simulated systems, supporting our approach. Using a chaotic oscillator coupled to different robotic architectures, Pitti and collaborators (Pitti et al., 2009, 2010) explored the mechanisms underlying the control of motor synergies and showed a correspondence between synchronisation and robust behaviour. Lungarella & Sporns (2006) used different information theoretic tools to analyse the information flow in sensorimotor networks of various robotics systems, stressing the relationship between body, environment and information processing. In the same sense, Williams & Beer (2010) used an evolved agent in a relational categorization task to introduce a new information-theoretic approach to study the dynamics of information flow in embodied systems, complementing the more established dynamical systems analysis.

According to the definition, information is not an absolute value obtained from a measurement but rather a relative estimation of how much you can still improve in your current knowledge about a variable. Commonly, transmitter-receiver modelling involves random variables and the inherent uncertainty in trying to describe them is termed entropy.

Transfer Entropy (TE) (Schreiber, 2000) is based on Shannon's work and allows one to estimate the directional exchange of information between two given systems. The choice of TE in this work is based on a study conducted by Lungarella et al. (2007a), who compared the performance of different IT tools in bivariate time-series analysis, which will be the case here, and concluded that TE is in general more stable and robust than the other tools explored. The next paragraphs describe the technique.

Let I be a stationary higher-order Markov process (with memory) with transition probability p from one state to another i.e. $p(i_{t+1}|i_t, \ldots, i_{t-k+1}) = p(i_{t+1}|i_t, \ldots, i_{t-k})$ where i is a state from I at a given time t and k is the order of the process. Using Shannon Entropy, the optimal number of bits necessary to encode one more observation in the above time series is given by $h_I(k) = -\sum p(i_{t+1}, i_t^k) \log p(i_{t+1}|i_t^k)$. If, instead of the optimal probability function $p(i)$, we had a different function $q(i)$, using

the Kullback Entropy (Kullback, 1959) one can obtain the number of bits in excess when encoding one more observation, given by $K_I(k) = -\sum p(i_{t+1}, i_t^k) \log \frac{p(i_{t+1}|i_t^k)}{q(i_{t+1}|i_t^k)}$).

Consider now two time series, $X = x_t$ and $Y = y_t$, and assume they can be represented as a stationary higher-order Markov process. Transfer Entropy calculates the deviation from the generalised Markov property $p(y_{t+1}|y_t^n, x_t^m) = p(y_{t+1}|y_t^n)$ where $x_t^m \equiv (x_t, x_{t-1}, \ldots, x_{t-m+1})^T$, $y_t^n \equiv (y_t, y_{t-1}, \ldots, y_{t-n+1})^T$ and m and n are the orders of the higher-order Markov process (note that the above property holds only if there is no causal link between the time series). Based on the concepts described in the previous paragraph, Schreiber (Schreiber, 2000) defines Transfer Entropy as:

$$TE(X \rightarrow Y) =$$
$$\sum_{y_{t+1}} \sum_{x_t} \sum_{y_t} p(y_{t+1}, x_t^m, y_t^n) \log \frac{p(y_{t+1}|y_t^n, x_t^m)}{p(y_{t+1}|y_t^n)} \quad (5)$$

Therefore, from Equation 5 one can estimate the information about a future observation y_{t+1} given the available observations x_t^m and y_t^n that goes beyond the information of the future state y_{t+1} provided by y_t^n alone. It is thus a directional, non-symmetrical estimate of the influence of one time series on another.

The original formulation of Transfer Entropy suffers from finite sample effects when the available data is limited, and the results obtained may not be correctly estimated. To attenuate these limitations, Marschinski & Kantz (2002) introduced an improved estimator, "Effective Transfer Entropy" (ETE), which is calculated as the difference between the usual Transfer Entropy (Equation 5) and the Transfer Entropy calculated after shuffling the elements of the time series X, resulting in the following equation:

$$ETE(X \rightarrow Y) \equiv TE(X \rightarrow Y) - TE(X_{shuffled} \rightarrow Y)$$
$$(6)$$

The ETE formulation is the one used in this paper.

Section 3 will provide the implementation details of ETE in this work. The next subsection presents the key concepts of Evolutionary Robotics, which underpins the experiments conducted.

2.3 Evolutionary Robotics and Minimally Cognitive Tasks

Evolutionary Robotics (ER) is a relatively new field of interdisciplinary research grounded in concepts from Computer Science and Evolutionary Biology (Harvey et al., 2005; Floreano et al., 2008; Floreano & Mattiussi, 2008; Floreano

& Keller, 2010). Originally devised as an engineering approach to automatically generate efficient robot controllers in challenging scenarios, where traditional control techniques have limited performance, ER is now well regarded among biologists, cognitive scientists and neuroscientists, as it provides means to simulate and investigate brain-body- environment interactions that underlie the generation of behaviour in a relatively unconstrained way, thus penetrating areas that disembodied studies cannot reach (Noe, 2004; Pfeifer et al., 2007; Engel, 2010).

Consider a real or simulated robot, with sensors and actuators, situated in a scenario with a certain task to accomplish. Each solution candidate (individual) is represented by a genotype, which contains the basic information of the agent's body and/or its controller's parameters (e.g. the number of wheels the robot has and/or the values of the weights of an artificial neural network acting as its controller). According to some criteria, normally the previous performance of that individual in solving the task (fitness), parents are selected and undergo a process of mutation and recombination, generating new individuals which are then evaluated in the task. This process is repeated through the generations, eventually obtaining individuals with a higher performance in the given task.

In this sense, ER is a reasonable approach to studying embodied and situated behaviour generation, because it can be used as a powerful model synthesis technique (Beer, 2003; Husbands, 2009). Relatively simple, tractable models can be produced and studied in the context of what have been called Minimally Cognitive Tasks (Beer, 2003), which are tasks that are simple enough to allow detailed analysis and yet are complex enough to motivate some kind of cognitive interest.

In this paper, the agent is submitted to a minimally cognitive task related to active categorical perception. It has been chosen, among other tasks, for being considered fundamental to cognition (Harnad, 1987) and encapsulating basic aspects of cognitive systems.

The next section presents the implementation details of this task, the evolutionary robotics framework and the information theoretic tool adopted.

3 Methods

3.1 Framework for application in evolutionary robotics

The model studied here is inspired by the Kuramoto Model, adapted so that it could be applied to control a simulated robotic agent.

The framework, illustrated in Figure 1, is composed of 15 fully connected oscillators, with even-numbered nodes connected to the robot's sensors (1 sensor per node). The frequency of each node is the sum of its natural frequency
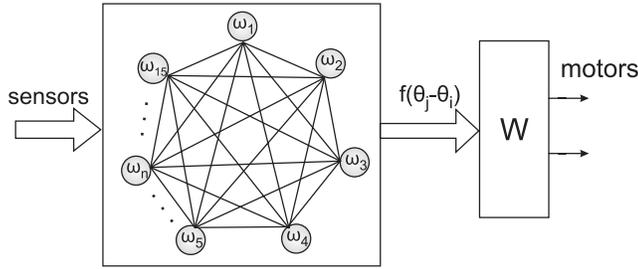
Fig. 1 Framework for application in evolutionary robotics. The oscillatory network is composed of 15 fully connected oscillators, with even nodes connected to the robot's sensors (1 sensor per node). A phase-sensitivity function $f(\phi)$ is applied to each of the phase differences $\phi_{i,i-1}, i = 2\ldots15$, then linearly combined by an output weight matrix, W, resulting in two signals that will command the left and right motors of the agent.
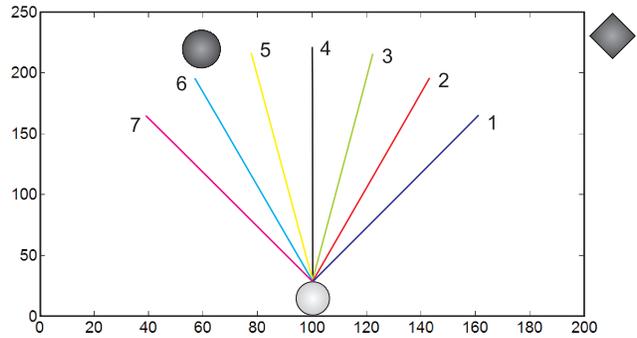

Fig. 2 Experiments 1 and 2 scenario. The agent (grey circle at the bottom) has to catch falling circles and avoid squares in Task 1 and catch falling circles with normal and inverted vision in Task 2. The robotic agent has 7 ray sensors, symmetrically displaced with relation to the central ray in intervals of $\pm pi/12$ radians, and two motors that can move it horizontally.

of oscillation, w_i, and the value of the sensory input $I_i(t)$ (calculated according to Equation 9 in the following section) related to that node (0 if there is no input), scaled by a factor z_i (Equation 7).

$$\frac{d\theta_i}{dt} = (\omega_i + z_i I_i(t)) + \frac{K}{N} \sum_{j=1}^{N} sin(\theta_j - \theta_i) \qquad (7)$$

The natural frequency w_i can be associated with the natural firing rate of a neuron or a group of neurons, and the sensory inputs mediated by z_i alters its oscillatory behaviour according to environmental interactions, thus improving the flexibility of the model to study neuronal synchronisation (Cumin & Unsworth, 2007) within a behavioural context.

At each iteration the phase differences from a node i to nodes $i-1, i = 2\ldots15$, are calculated following Equation 7 (modified from Equation 2 as described in the previous paragraph). The rationale for a network with 15 nodes relates to richer dynamical behaviour in the Kuramoto Model with this number of nodes (Maistrenko et al., 2005). A phase-sensitivity function $f(\phi)$ is then applied to each of the phase differences $\phi_{i,i-1}$ in order to reduce instabilities in the phase differences caused by the resetting of the phase of each oscillator when it exceeds 2π. In this work the sin function is used. The modified phase differences are then linearly combined by an output weight matrix, W, resulting in two signals that will command the left and right motors of the agent (Equation 8).

$$\boldsymbol{M} = \boldsymbol{W}' sin(\boldsymbol{\phi}) \qquad (8)$$

where $\boldsymbol{M} = [M_1, M_2]^T$ is the motor state space, with M_1 corresponding to the left motor command and M_2 to the right motor command.

In this way, the phase dynamics and the environmental input to the robotic agent will determine its behaviour. It is

important to stress that nodes that receive no input participate in the overall dynamics of the network, hence their natural activity can modulate the global activity of the network.

The experiments reported later explore the relationship between the network synchronisation state and the agent's behaviour. Different synchronisation regimes in the network are obtained by varying the coupling strength between units (K). During a given simulation, the value of the coupling is set by multiplying the value of K_c (calculated according to Equation 4) by a scaling factor to tune the nodes in a subcritical ($K < K_c$), supercritical ($K > K_c$) or critical zone ($K = K_c$). The next section describes the minimally cognitive task studied.

3.2 Active Categorical Perception

The active categorical perception task studied here consists of a circular robotic agent, able to move horizontally along the bottom of a 250×200 rectangular environment (Figure 2), which has to discriminate between circles and squares as they move from the top of the arena to the bottom (only one object on each trial) (Beer, 2003; Izquierdo, 2008). The robotic agent's body has 7 ray sensors, symmetrically displaced in relation to the central ray in intervals of $\pm pi/12$ radians, and two motors to move it left and right along a straight line. Sensory inputs $I_i(t)$ for each distance sensor i at a given time t are calculated following Equation 9:

$$I_i(t) = \begin{cases} 10(1 - \frac{l_i(t)}{200}), & \text{if } l_i(t) < 200 \\ 0, & \text{otherwise.} \end{cases} \qquad (9)$$

where l_i is the length of the ith ray between the robot's body and the object, $i = 1, \ldots, 7$.

The diagonal of the square, and the radii of both the robotic agent and the circle all measure 15 units. Each sensor

is attached to an even-numbered node of the network (Figure 1). The velocity adjustment V (Equation 10) is obtained subtracting the left from the right motor command value (a positive value drives the robot to the right, a negative value to the left).

$$V = s(M_2 - M_1) \tag{10}$$

where s is a motor output weight, M_1 is the left motor command and M_2 is the right motor command (Equation 8).

At the beginning of each trial, a circle or a square is dropped from the top of the environment at a random horizontal position within a maximum of 50 units from the robotic agent, and moves vertically with a velocity of 3 units per timestep. The robotic agent has to approach the circles and avoid the squares, adjusting its horizontal velocity accordingly (limited to 5 units per timestep).

A variation on this experiment consists of an orientation task. In the same environmental set-up, the robotic agent has to adjust its horizontal position and catch circles with normal and inverted vision (catching is considered to occur automatically when the object is within the agent's body radius). A task, comprised of many trials, is composed of an equal number of trials under normal and inverted vision. When submitted to visual inversion, sensory readings from an object on the right side of the agent are perceived by the agent's left set of sensors, and vice-versa. Therefore, to successfully execute this task the agent has to devise a context-dependent strategy to cope with the ambiguity of having equal stimuli in different scenario and vision configurations.

Taken together, the two tasks present challenging and important scenarios for a cognitive agent (Di Paolo, 2000; Izquierdo, 2008), justifying their choice as an experimental scenario: the first task requires the engagement of, and discrimination between, two different objects whereas in the second task the robotic agent has to develop a strategy to overcome the disruption caused by the inversion of the visual field.

The next section presents the details of the evolutionary process.

3.3 Genetic Algorithm

A geographically distributed genetic algorithm with local selection and replacement (Husbands et al., 1998) is used to determine the parameters of the system: the frequency of each of the 15 nodes, $w_i \in [0, 10]$, the 7 input weights $z_i \in [-5, 5]$ to odd nodes (0 otherwise), the matrix $W_{N-1,o}$ with $14 \times 2 = 28$ elements in the interval $[-1, 1]$, a motor output weight $s \in [0, 100]$ and the network update time $t_u \in [0, 50]$ (the number of time steps Equation 7 is updated before the phase-differences are used to calculate the motor

output), resulting in a genotype of length 52 for a network with $N = 15$.

The network's genotype consists of an array of integer variables lying in the range $[0, 999]$ (each variable occupies a gene locus), which are mapped to values determined by the range of their respective parameters. For all the experiments in this paper, the population size was 49, arranged in a 7×7 toroidal grid. A generation is defined as 100 breeding events and the evolutionary algorithm runs for a maximum of 100 generations. There are two mutation operators: the first operator is applied to 20% of the genes and produces a change at each locus by an amount within the $[-10, +10]$ range according to a normal distribution. The second mutation operator has a probability of 10% and is applied to 40% of the genotype, replacing a randomly chosen gene locus with a new value within the $[0, 999]$ range in an uniform distribution. There is no crossover.

In a breeding event, a mating pool is formed by choosing a random point in the grid together with its 8 neighbours. A single parent is then chosen through rank-based roulette selection, and the mutation operators are applied, producing a new individual, which is evaluated and placed back in the mating pool in a position determined by inverse rank-based roulette selection. For further details about the genetic algorithm, the reader should refer to Husbands et al. (1998).

In the active categorical perception task, fitness is evaluated from a set of 34 trials with randomly chosen objects (circles or squares), starting at an uniformly distributed horizontal offset in the interval of ± 50 units from the robotic agent. Fitness is defined as the robotic agent's ability to catch circles and avoid squares, and is calculated according to the following equation: $fitness = \sum_{i=1}^{N} i f_i / \sum_{i=1}^{N} i$, where f_i is the ith value in a descending ordered vector $F_{1,N}$, and is given by $1 - d_i$, in the case of a circle, or by d_i in the case of a square. d_i is the horizontal distance from the robotic agent to the object at the end of the ith trial (when the object reaches the bottom of the environment), limited to 50 and normalized between 0 and 1. Therefore, a robotic agent with good fitness maximizes its distance from squares and minimizes its distance from circles. Notice that the form of the fitness function creates pressure for good performance in *all* trials in a given fitness evaluation, instead of just averaging the performance across the trials, which could bias the mean fitness of a given robotic agent leading to poor generalisation behaviour.

In the orientation task under normal and inverted vision, fitness is evaluated from 17 trials with the normal vision scheme followed by 17 trials with inverted vision. The circles are dropped at an uniformly distributed horizontal offset in the interval of ± 50 units from the robotic agent. Fitness for each part of the evaluation is defined as in the previous paragraph, but considering just the circle catching scenario. Therefore, a robotic agent with good fitness minimizes its

distance from circles, in both normal and inverted vision situations.

The next section describes the methods for Transfer Entropy analysis, the information theoretic tool used to investigate information transfer dynamics between variables of the system.

3.4 Transfer Entropy

Following the notation and the theoretical background presented in Section 2.2, the analysis will focus on the information transfer between pairs of variables of the system. For all experiments, we adopt the orders of the higher-order Markov processes as $m = n = 1$ (Equation 5), but see the Discussion section for further comments. The conditional probabilities are calculated by rewriting them as joint probabilities which are then estimated using histograms.

To collect the data, the evolved agent under analysis was evaluated on a single trial, from a predetermined initial position (equal for every test), having its sensory information, its motor commands and the phase dynamics of each node recorded. Details of agent selection are presented in the next section. Each data point interval (obtained according to an Euler integration time-step of $15ms$) is then linearly interpolated (in $1.5ms$ time intervals) resulting in a coarse grained time-series, which facilitates a more robust ETE analysis. If the data were collected in shorter time intervals instead, the computational time needed for the evolutionary process to complete would be prohibitive. Moreover, given the structure of the data observed in the experiments, the changes in the qualitative aspects of the original time series after interpolation were minimal.

The time series for the seven sensors (one time series for each sensor), obtained as described above, are submitted to a dimension reduction using a principal component analysis (PCA) (Jolliffe, 2002). First, the mean of each sensor's data vector is subtracted from its correspondent data set to ensure that the dynamics of the first component describes the direction of maximum variance in the data. We then calculate the principal components and project the original seven time series data onto the first principal component, reducing the original data to a single time series that captures the most significant features of the multidimensional input space. The motor commands are also combined to generate a single time series by subtracting the value of the left wheel command from the right wheel command. The phase dynamics are recorded directly from the nodes. Finally, the time series used in the calculations of Effective Transfer Entropy are obtained from the first derivative of each of the above described time series, discretised in 6 equiprobable states, which improves the robustness of the statistics (Marschinski & Kantz, 2002; Lungarella et al., 2005).

Two ETE analyses are conducted across the experiments: one considering the information flow between the agent's sensors and motors, and another considering the information flow between the nodes of the network (represented by the phase dynamics). We are mainly interested in studying how ETE between these variables vary as the task progresses, given that the agent continually engages with the object during the discrimination and orientation processes. For this purpose, a sliding window technique is used (Staniek & Lehnertz, 2008; Szczepanski et al., 2011), with a window size of 200 data points. Therefore, at every time step of the task, the ETE is estimated (according to Equations 5 and 6) considering the time series contained in that window.

4 Results

In the first experiment, individuals are evolved to perform categorical perception under different coupling configurations, calculated by re-scaling the current value of K_c (see the $Methods$ section for details). Figure 3(a) presents the fitness statistics for the best evolved agents obtained in 100 evolutionary runs at different coupling configurations. Notice that the agent's overall performance is poor in the no coupling case ($0K_c$), smoothly increases as the coupling strengthens (peaking at $K = K_c$) and then sharply falls for strong coupling ($3K_c$ onwards).

These results indicate that tuning the oscillatory networks in a certain region of their synchronisation regime may have a direct impact on the performance and evolvability of the agent, but would this result hold had we had a different task? We therefore performed a second set of experiments by changing from categorical perception to the circle catching (orientation) task under normal and inverted vision, as described in Section 3. Figure 3(b) presents the fitness statistics for the best evolved agents obtained in 100 evolutionary runs at different coupling configurations. Notice that the performance increases with the coupling and reaches its peak with values of K around $2K_c$, whilst in the previous experiments we observed a decrease in fitness with coupling values larger than K_c.

It is clear from these two figures that the synchronisation regime the network is tuned to influences the behavioural performance of the agent. In order to clarifying why this happens and to explore the underlying mechanisms, we investigate the performance, system dynamics and information flow properties of three evolved agents with different couplings strengths, which encourage different degrees of synchronisation, within each of the above experiments (there is no parameter change after the evolutionary process is over). These were chosen to be $0K_c$, K_c and $5K_c$ for Analysis 1 (categorical perception) and $0K_c$, $2K_c$ and $10K_c$ for Analysis 2 (orientation), to represent weak, 'optimal' and strong coupling for those behaviours. All experiments are carried

(a)   (b)
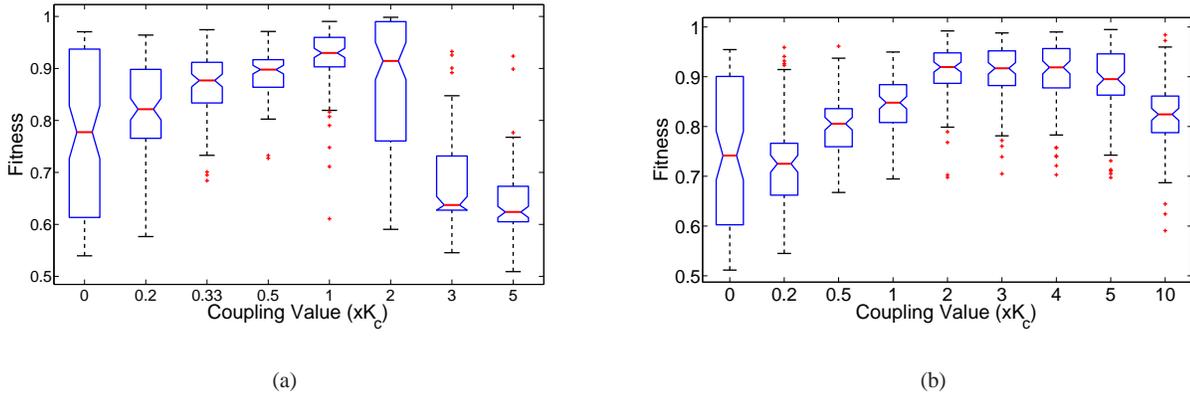
Fig. 3 Fitness statistics for the best evolved agent obtained in 100 evolutionary runs at different coupling configurations for a categorical perception task (a) and an orientation under normal and inverted vision task (b). The boxes represent the lower, median and upper quartile. The central mark is the median and whiskers (dashed lines emerging from the boxes) extend to the most extreme data points not considered outliers (outliers are shown as red crosses).

out with agents chosen to have similar t_u (network update time, see Section 3.3) and with final fitness values as close as possible to the median of the fitnesses obtained in the evolutionary process for the given coupling scenario.

4.1 Analysis 1: Categorical Perception Task

Figures 4(a), 4(b) and 4(c) present the generalisation performance (performance in 100 aleatory trials after evolution) of three agents evolved in the $0K_c$, K_c and $5K_c$ conditions, respectively. The plots illustrate the agent-object horizontal separation.

Notice that all agents present some level of discrimination between circles and squares, characterized by the different trajectories adopted by the agent as the task progresses; however, their accuracies differ which seems to be the main factor responsible for the variations in fitness scores. Additionally, both $0K_c$ and $5K_c$ agents present a high proportion of trajectories similar to the ones observed with intrinsic autonomous network dynamics (i.e. sensory inputs fixed at 0), whilst in the K_c agent this is only observed when the object is a square and has its initial positioning far from the agent - the general strategy elsewhere is to initially move to centralize the object and then, approximately half way through the task, orient with the circles and move away from the squares.

Figure 5 presents, for each of the three agents, the time course of various system variables from a particular scenario where the agent (object) starts at the horizontal coordinate 100 (60). The agent-object trajectory, the frequency of each node and the synchronisation index are shown. Recall that in our model the sensory inputs, proportional to the agent-object distance, impact on the natural frequency of some nodes. That, in turn, affects the phase dynamics of the net-

work, producing the different motor outputs that determine the agent's position and hence its performance on the task.

For the $0K_c$ agent, the frequency plot shows that the oscillatory behaviour of each node is only influenced by its natural frequency and the sensory input. For this reason, nodes without sensory connections don't change their frequency as the task progresses and the network does not present any kind of internal coupled dynamics. The synchronisation index remains low, below 0.4, and its amplitude variation can be explained by observing Equation 3: the index reflects a sum of vectors, each representing the phase of an oscillator; with different evolved natural phase speeds (frequencies), the total sum varies over time, unless these frequencies maintain a constant relation, reflecting synchronisation.

In the K_c coupling condition, the frequency of each node presents a rich dynamics due to the critical internal coupling and the external sensory stimulus. With the agent approaching the object, the sensory inputs cause a large variation in the nodes' frequencies and as a consequence different node clusters emerge (see the black arrow in the circle catching row of plots). Conversely, moving away from the squares diminishes the external stimulus, and most of the network's nodes synchronise. This can be further observed by the synchronisation index plot, which approaches the dynamics observed in the autonomous condition in the square avoidance scenario (see the black arrow in the corresponding plot) but diverges from it in the circle catching sequence. Yet, the dynamics of both indexes reflect a higher synchronisation level (around 0.7 throughout the task).

The last figure of the set shows the results for the $5K_c$ network. The agent's behaviour is practically the same for both circle and square scenarios: the agent moves away from the object towards one side of the environment at constant
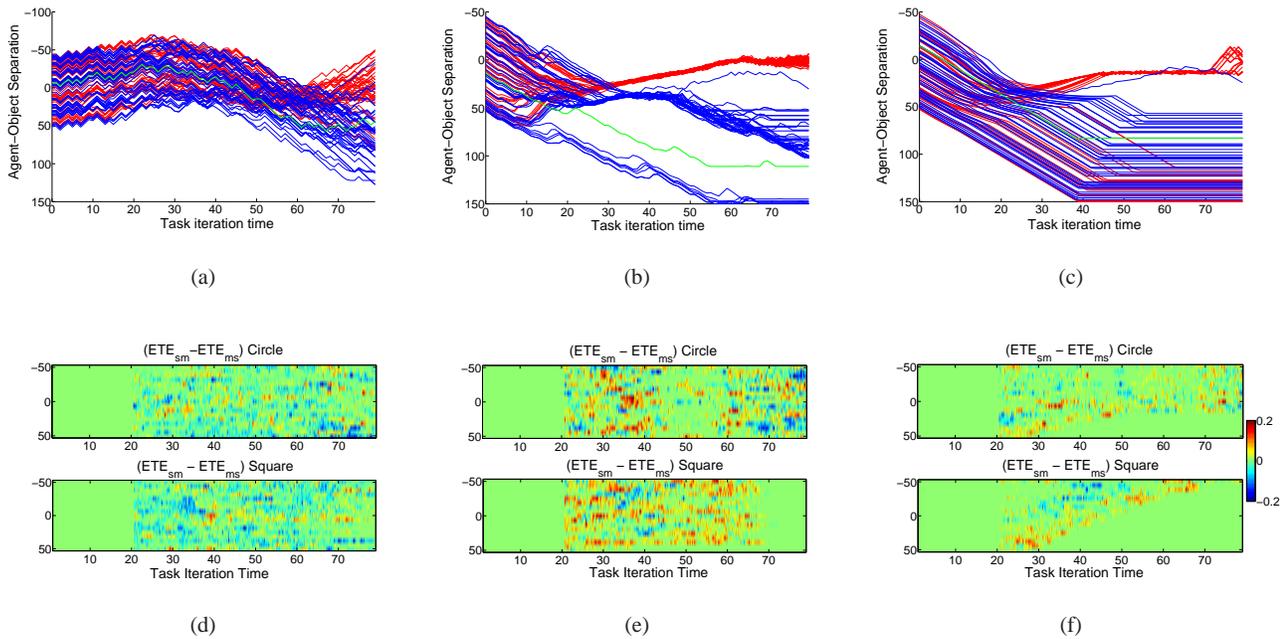
Fig. 4 Top plots: generalisation performance of the agent over 100 aleatory runs. The red colour is related to the circle catching behaviour and the blue colour to the square avoidance behaviour. Green lines reflect the behaviour under autonomous network dynamics. The plots illustrate the value of the horizontal separation of the agent and the object. Bottom plots: difference between Effective Transfer Entropy from sensors to motors, ETE_{sm}, and from motors to sensors, ETE_{ms}, for each of the 17 different initial horizontal displacements ($[-50, 50]$) that the agent is evaluated during evolution (Section 3). ETE is calculated according to Equation 6. Figures refer to (a)(d)$0K_c$, (b)(e)K_c and (c)(f) $5K_c$.

speed, failing to discriminate. Inspecting the frequency plot, we notice that the network rapidly entrains with the synchronisation index remaining high throughout the task (in contrast to the K_c coupling condition). The small variations in the frequency dynamics (marked by a black arrow on the plot) relate to the phase resetting phenomena, where some nodes, although synchronised, reset at slightly different times, perturbing the ongoing synchronisation status. With these plots in mind, the agent's behaviour can thus be explained by a specific property of the model (Figure 1), where the synchronised phase dynamics result in stable phase differences, which culminate in constant motor speeds that cause a linear movement of the agent. Therefore, looking back at Figures 4(b) and 4(c), one can speculate why the $5K_c$ network does not achieve a high fitness: the tendency to fully synchronise reduces the sensitivity of the system in discriminating objects when the stimulus is not strong enough to modulate the ongoing phase dynamics, whereas the K_c network, although synchronising at some stages, is still flexible enough to escape the entrained state and adapt its phase dynamics in order to generate the appropriate motor commands.

In the above paragraphs and in the analyses to follow, the behaviour being displayed by the agent is compared to the one observed in the absence of sensory readings (behaviour under autonomous network dynamics - henceforth often referred to as the autonomous trajectory for short). This is to

emphasize that the behaviour being displayed by the agent is not simply a response to its sensory readings or to the ongoing network dynamics alone, but a more intricate interplay between the two. The actions taken by the robot at the beginning of the task, when the sensory readings are small, are mainly a result of the intrinsic network dynamics, but they will modulate and affect the sensory readings that the robot will face in the future, which will also affect and shape the sequence of actions.

The previous results stress the relationship between the sensory inputs (environmental context) and the internal dynamics of the network. From these plots alone, however, it is not easy to determine to what extent the behaviour displayed by the robot is influenced by the network's internal dynamics or a response to the sensory readings. The Transfer Entropy analysis, as described in Section 2.2, may offer some insights into the temporal flow of information within the brain-body-environment system.

Consider Figures 4(d), 4(e) and 4(f). For each of the 17 different initial positions that the agent is evaluated on during evolution (Section 3), we calculate the flow of information from sensors to motors, ETE_{sm}, and from motors to sensors, ETE_{ms}. The first corresponds to the influence of the sensors in determining the behaviour of the agent whereas the latter corresponds to the changes in sensory inputs brought about by the agent's motion. Notice that we are
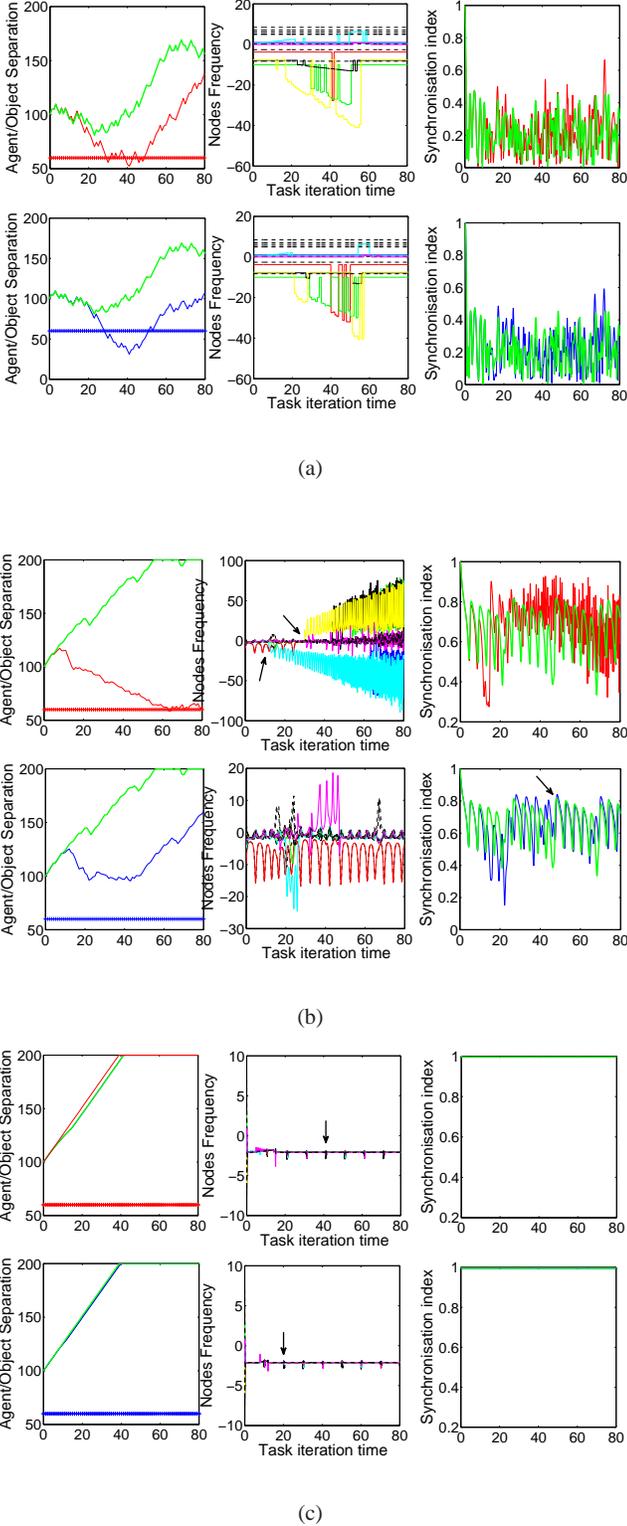
(a)

(b)

(c)

Fig. 5 Detailed behaviour of the agent's internal and external dynamics for a given starting position. The top three graphics of each figure refer to the circle catching behaviour and the bottom ones to the square avoidance behaviour. The leftmost column illustrates the horizontal coordinate of the agent and the object (straight horizontal lines), the middle one shows the frequency of each node of the network as the task progresses (colour lines for nodes with inputs, dashed lines otherwise) and the rightmost ones present the synchronisation index. Green lines reflect the autonomous behaviour dynamics (clamped sensors). (a) $0K_c$, (b) K_c and (c) $5K_c$

not primarily concerned with the magnitudes of the information flows but rather on their form of interaction (causal influence); therefore, we depict the difference between ETE_{sm} and ETE_{ms} (positive values indicate a higher information flow from sensors to motors than from motors to sensors, negative values indicate the opposite). The plots show the difference between the results obtained for the circle catching and the square avoidance scenarios, and, together with the respective Figures 4(a), 4(b) and 4(c), allow for a wider perspective on how the information flow develops given different scenario configurations.

Figure 4(d) presents the ETE for the $0K_c$ coupling condition. Notice the widespread (wide deviations from zero) flow of information in both circle and square scenarios, with measurements being recorded in both directions (motors to sensors and sensors to motors) at all times, which may relate to a constant interaction between the agent and the object - indeed, the lack of internal coupling between nodes would promote a highly reactive agent resulting in a constant mutual interaction with the environment to perform the categorization. The performance, however, is poor.

Figure 4(e) presents the corresponding analysis for the K_c agent. The first observation, contrasting with the uncoupled network, is the higher magnitude of the difference between the ETE from sensors to motors and the ETE from motors to sensors during diverse moments of the task execution for both objects. Following the argument in the previous paragraph, we would therefore expect that in the current coupling condition the agent is more able than with the uncoupled architecture to convey sensory information into motor responses, adjusting the agent's position in order to maximize its performance. In the circle catching sequence, there are mainly two intervals of high difference in ETE, coinciding with major adjustments in the agent's trajectory when compared with the autonomous trajectory. Between iterations 20 and 40, the flow from sensors to motors is more noticeable for most initial conditions, which may relate to the scanning behaviour and identification that the object is a circle and should be captured. At the end of the task, after iteration 60, however, the difference in the flow increases towards a greater ETE from motors to sensors for many initial conditions, while at a behavioural level we observe that the agent is centred on the circle and stays relatively still keeping the horizontal separation near 0. In the square scenario, the ETE from sensors to motors is higher than the ETE from motors to sensors through most of the task for most initial conditions, while we observe the agent moving away from the object. At the end of the task, when the object is out of sensory reach, the ETE_{sm} and consequently the ETE_{ms} tend to 0, which explain the lack of information flow after iteration 70.

Turning to the $5K_c$ plot (Figure 4(f)), the immediate observation is the sparse and reduced amount of ETE through-

out the task when compared to the K_c condition. Although some significant values occur, there is no consistency in the flow with many null values recorded. Observing the behavioural strategy, this is easily explained: the agent succeeds in discriminating between the objects in just a few situations, which have a corresponding reading in the ETE plot, whereas from most of the initial displacements the agent moves to a far side of the environment, where all sensory readings are absent. In this sense, the poor performance can be associated with a lack of sensitivity to sensory readings, explained by the high synchronisation level of the network despite the external perturbation to its nodes' frequencies. The external perturbations are not able to modulate the highly synchronised network dynamics.

The analysis of the flow of sensorimotor information provides insights into how the brain-body-environment system interacts over time and engages in determining the observed behaviour of the agent for different objects; insights which may not be evident solely from the system variables' time series. In order to add further depth, the next set of results explore the information flow between the nodes of the network, as described in Section 3.

Consider Figure 6(a), which presents in its first row the average over time of the information flow between the oscillatory nodes of the uncoupled network for the circle catching (first column) and square avoidance (second column) scenarios, respectively, according to Equation 11:

$$\bar{T}(n,n') = \frac{1}{T_w} \sum_{w=1}^{T_w} [ETE_w(n,n') - ETE_w^A(n,n')] \quad (11)$$

where $\bar{T}(n,n')$ is the average Effective Transfer Entropy ETE between the pair of nodes (n,n'), estimated according to Equation 6, w is the time window and T_w is total number of windows. The superscript A stands for the ETE calculated in the autonomous case.

As expected, the uncoupled network has almost no ETE between nodes during most of the task, except when sensory readings modulate the ongoing phase activity, which explains why the few deviations from zero that occur are with relation to nodes that have sensory inputs (recall that only even nodes receive sensory inputs, see Figure 3.1 for details). Observe that the flow between any two given nodes is not necessarily symmetric (e.g. there is a flow from node 14 to 12 but the opposite is not true), as explained in Section 2.2.

In the K_c condition (Figure 6(a), second row), however, the coupling elicits a much broader modulation of the nodes' phase activity by the sensory inputs, stressing the internal communication between oscillators. The network is fully connected but nodes do not interact in the same way: the ETE values between pairs of nodes vary according to the corresponding sensory stimulus, the previous phase activity

of the nodes and its natural frequency. Also, notice that the sensory inputs not only increase the information flow within the network but also decrease it, reflecting a modulation of the spontaneous network phase activity (the activity the network would have without sensory inputs). Although Transfer Entropy always results in a positive measurement, recall from Section 3 that to highlight the sensory influence the above results are obtained by subtracting the ETE calculated for the autonomous condition, hence the negative readings in the plots.

Compare the previous analysis with the $5K_c$ condition (Figure 6(a), third row). As the network is strongly coupled and acts as a single giant oscillator, weak sensory stimuli of one or more nodes cannot promote a significant perturbation or change in the ongoing dynamics; in other words, the sensory readings do not contribute to the discrimination of the object given that different incoming stimuli through other nodes would have little impact on the entrained network. Even though the system is able to detect the object and generate a motor response to it, the categorization performance is impaired.

This diversity of information flow between the different coupling configurations follows the variation in the synchronisation index between nodes of the network, portrayed in Figure 6(b). The small level of synchronisation in the $0K_c$ network has a corresponding small flow of information between nodes, and the large level of synchronisation in the $5K_c$ network is accompanied by a reduced flow of ETE between its nodes. The intermediate point, the K_c network, has both a synchronisation regime and diverse and highly asymmetric information dynamics.

In summary, although all configurations are fully connected, the K_c network is able to use a great variety of flexible, reconfigurable functional connections. Figure 6(c) illustrate this fact by redrawing the connection scheme according to the magnitude of the ETE between nodes - only nodes with mean information flow 50% above the global average are connected. Notice the much richer connectivity in the K_c condition than in the other two. The higher performance and evolvability achieved by this configuration can thus be attributed to its nodes' superior ability to communicate and organise themselves functionally, resulting in a more efficient information flow between the ongoing network activity, the agent's body and the environmental context.

Figure 7 shows the temporal evolution of the net Transfer Entropy from a given node to all remaining nodes of the network in 50 equally spaced time windows, defined as the preferred direction of information flow (Staniek & Lehnertz, 2008):

$$T(n,n',w) = \frac{1}{N-1} \sum_{n \neq n'} [ETE_{n,n'}(w) - ETE_{n',n}(w)]$$

$$(12)$$

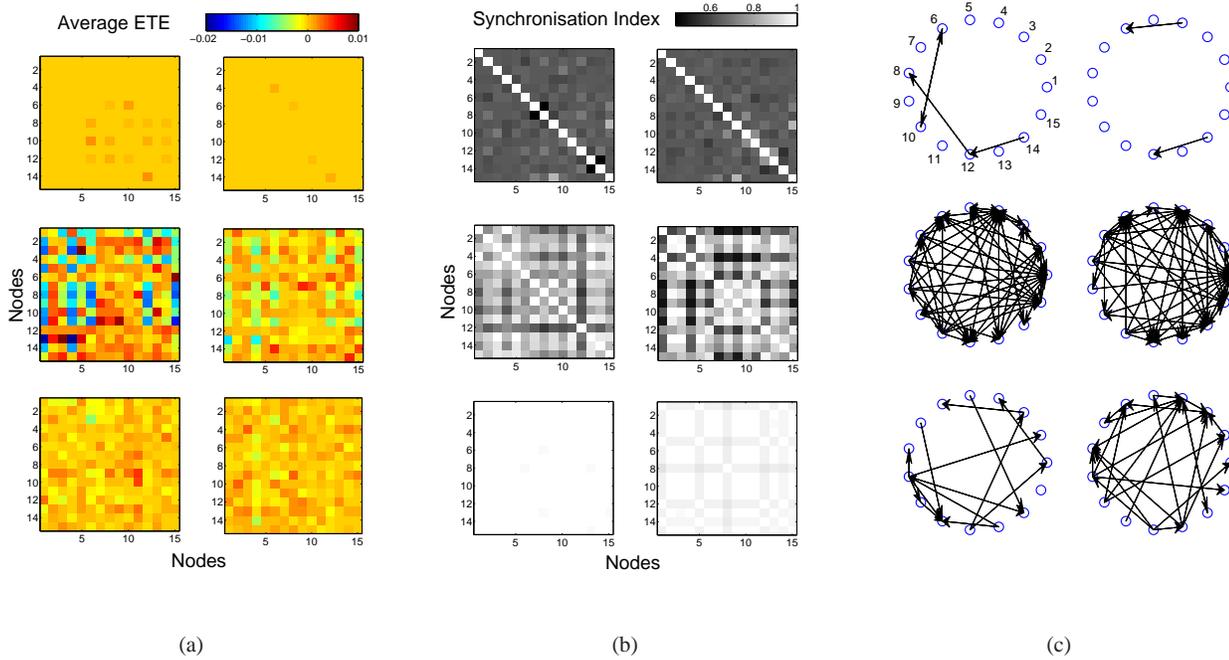(a)   (b)   (c)

Fig. 6 (a) Averaged information flow between the oscillatory nodes of the network (Equation 11). The autonomous flow is subtracted in all plots to highlight the task-specific activity. (b) Same as (a) but for the synchronisation index. (c) functional network - a directed arrow is drawn whenever the Transfer Entropy between two given nodes exceeds 50% of the maximum ETE value obtained in the whole task. The rows correspond to the $0K_c$, K_c and $5K_c$ agents, the left side of each column corresponds to circle catching and the right side to square avoidance.
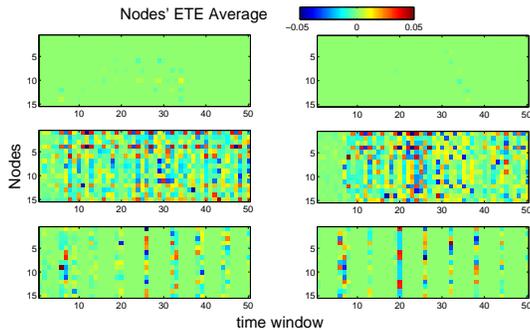

Fig. 7 Net Transfer Entropy from a given node to all remaining nodes of the network in 50 equally spaced time windows (Equation 12). The autonomous flow is subtracted in all plots to highlight the task-specific activity. Top row: $0K_c$ agent, middle: K_c, bottom: $5K_c$.

where ETE is the Effective Transfer Entropy between the pair of nodes (n, n') estimated according to Equation 6, w is the time window and N is the total number of nodes in the network.

The first configuration (top row), the $0K_c$ network, shows almost no information flow between nodes, whereas in the K_c network (middle row) there is a widespread activity, with driving nodes (nodes with higher information flow) alternating in time. Notice, however, the increased activity in the square avoidance scenario between time windows 20 and 35.

Observing Figure 4(b), this iteration interval corresponds exactly to major adjustments in the agent's trajectory when compared with the autonomous trajectory. This could help explain the more widespread activity observed in the circle catching scenario, as the agent constantly adjusts its trajectory in order to capture the approaching object. The $5K_c$ network (bottom row), as with the $0K_c$, presents a reduced flow in both scenarios, with few exceptions which correspond to the phase resetting behaviour of the model described earlier.

The experiments analysed in this section show that tuning the networks in a certain region of their synchronisation regime appears to have a direct impact on the performance and evolvability of the agent. The information flow analysis together with that of the system variables' dynamics indicate that extreme scenarios, e.g. no coupling or strong coupling, are not suitable for the categorical perception task. Importantly it also shows that neural synchronisation dynamics, when suitably flexible and reconfigurable, can generate minimally cognitive embodied behaviour.

In order to test the generality of these results, following the same methodology as above, the next set of experiments explores the performance of different coupling configurations on a different minimally cognitive task where the agent has to orient itself to an approaching object (circle) under
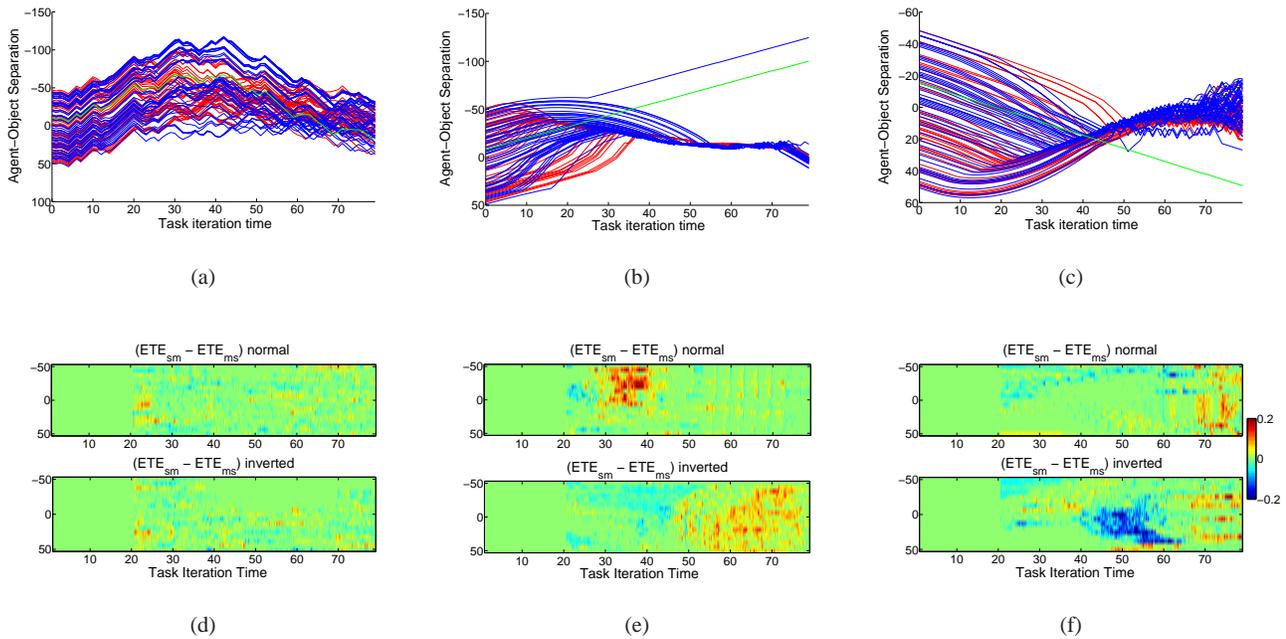
Fig. 8 Top plots: generalisation performance of the agent over 100 aleatory runs. The red colour is related to circle catching under normal vision and the blue colour to the inverted vision scenario. Green lines reflect the behaviour under autonomous network dynamics (clamped sensors). The plots illustrate the value of the horizontal separation of the agent and the object. Bottom plots: difference between Effective Transfer Entropy from sensors to motors, ETE_{sm}, and from motors to sensors, ETE_{ms}, for each of the 17 different initial horizontal displacements ($[-50, 50]$) that the agent is evaluated on during evolution (Section 3). ETE is calculated according to Equation 6. Figures refer to (a)(d)$0K_c$, (b)(e)$2K_c$ and (c)(f) $10K_c$.

normal and inverted vision (see Section 3 for experimental procedures).

4.2 Analysis 2: Orientation under Normal and Inverted Vision

Recall Figure 3(b), which presents the fitness statistics for the best evolved agents obtained in 100 evolutionary runs at different coupling configurations for task 2 (circle catching with normal and inverted vision). The performance increases with the coupling and peaks at $K = 2K_c$, presenting a noticeable fall in performance for couplings greater than $5K_c$. The following analysis will therefore consider three individuals: one drawn from the $0K_c$ scenario, one from the $2K_c$ and one from the $10K_c$ scenario, all with fitness values close to the fitness median obtained on each respective coupling scenario set of results.

Observe Figures 8(a), 8(b) and 8(c), which present the generalisation performance of the agents (performance in 100 aleatory trials after evolution). The agent controlled by the uncoupled network ($0K_c$) cannot succeed in catching circles with either normal or inverted vision and all trajectories seem to be a variation of the autonomous trajectory (illustrated by the green line in the plot) - there is almost no influence from the sensory readings. The $2K_c$ agent, however,

has good performance in both sensory modes, centring on the object in slightly different ways as the task progresses, depending on whether it has normal or inverted vision, making fine adjustments to the trajectory towards the end. A similar strategy is employed by the very strongly coupled agent ($10K_c$), however the accuracy with which the agent centres on the object towards the end of the task is poor and this divergence explains the lower fitnesses obtained. Nevertheless, both $2K_c$ and $10K_c$ agents display different behaviours from their autonomous trajectories, in contrast with the uncoupled network, indicating that they make use of the sensory information to solve the task.

An inspection of the dynamics of some of the system variables reveals that, as expected, the uncoupled network has some node activity varying directly with the sensory readings, with nodes without inputs keeping their natural frequency throughout the task (see Figure 9(a)). The synchronisation index remains low but has a significant variance. The lack of internal coupling results in the absence of any internal communication or flow of information as well as any influence caused by acting on the environment. As the task is intrinsically ambiguous, the agent cannot perform well.

The $2K_c$ network (Figure 9(b)), in contrast, remains synchronised for most of the time and the nodes can only "es-

cape" this state when sensory readings rapidly change their natural frequency. The agent's trajectories in both vision conditions are similar, but the synchronisation index plot reveals a diverse phase dynamics. Therefore, even though the network has a fixed, all-to-all pre-established coupling configuration, with no explicit plasticity mechanisms, it can respond and adapt to conflicting sensory conditions and obtain a good performance in the task. This behaviour is also observed in the $10K_c$ network (Figure 9(c)), but as the task approaches its end and the magnitude of the sensory readings increase, the frequency behaviour changes drastically, partly because of the high inputs, partly because of some instabilities inherent in the model (high coupling values, associated with properties of the nodes' frequencies distribution, may lead to unstable behaviour in the Kuramoto Model (Acebrón et al., 2005)). Taken together, these results suggest that more synchronised networks are more flexible in adapting to sensory ambiguity, but the variables' dynamics alone cannot provide a broader insight into the processes underlying this adaptation. In order to do that, we conduct an information flow analysis of different parts of the system.

A Transfer Entropy analysis between the agent's sensors and motors (environment and body) across all possible starting positions used during evolution for the $0K_c$ network (Figure 8(d)) reveals almost no information flow between the agent's body and the environment. As noted before, the lack of internal coupling prevents the agent from properly responding to the ambiguity of the task. The same analysis, carried out on the more strongly coupled networks (Figures 8(e) and 8(f)), shows that the normal and inverted vision conditions elicit different information flows at specific points of the agent's trajectory, reflecting its ability to successfully catch circles in both scenarios. For the $2K_c$ network, especially, starting positions that require a very different behaviour from that observed in the autonomous trajectory (e.g. agent-object separation between $[-50, 0]$, where the agent has to move against the straight autonomous trajectory) correspond to a higher (smaller) flow from sensors to motors in the normal (inverted) vision scenario in some parts of the task (between iterations 30 and 40), with the opposite happening from iteration 50 onwards. The same happens for the $10K_c$ network, although in this case there is a higher flow from motors to sensors in the inverted vision case between iterations 40 and 60 (again, corresponding to points where the agent has to actively move against its autonomous trajectory), and a higher flow from sensors to motors towards the end in both vision configurations. In summary, the point to stress here is the asymmetry of information flow and its behavioural correspondence obtained in fully connected oscillatory networks only by changing their level of synchronisation.

Looking at the average node activity for these coupling configurations (Figure 10(a)), one can see that the informa-
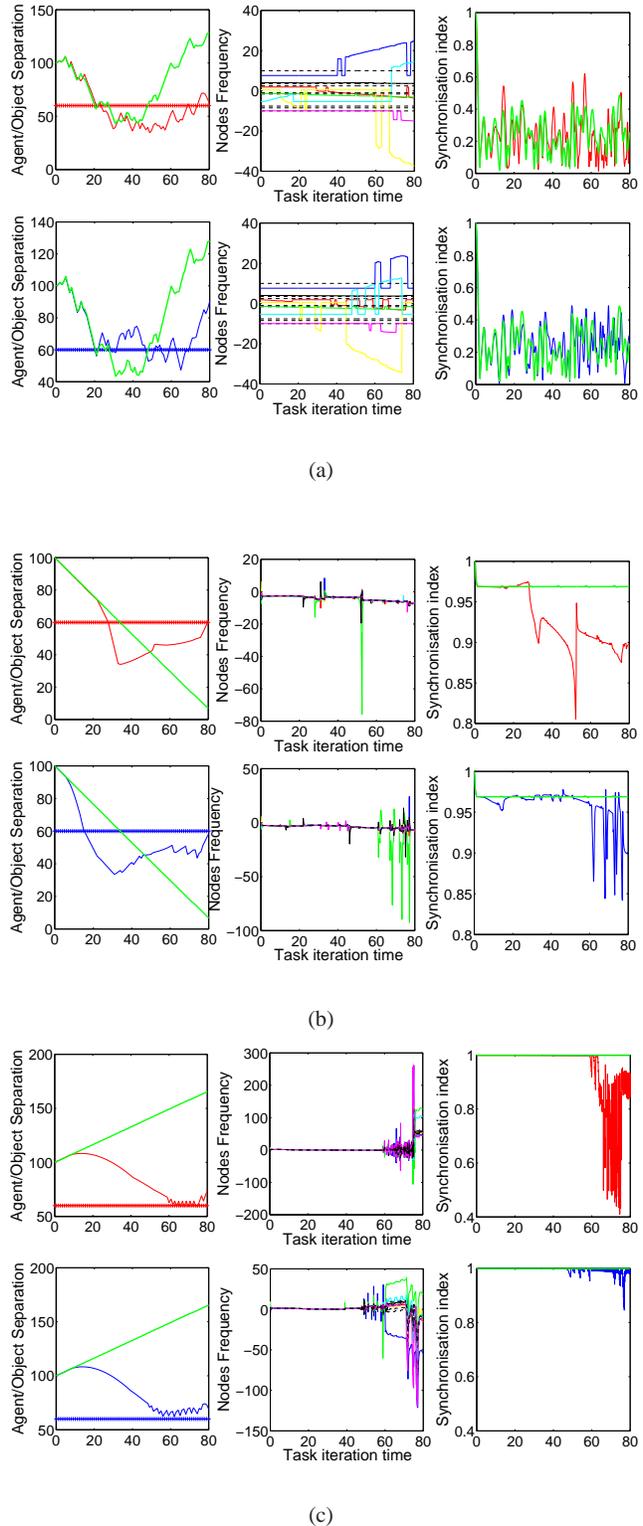

(a)

(b)

(c)

Fig. 9 Detailed behaviour of the internal and external dynamics for the orientation task agents. The top (bottom) three graphics of each figure refer to the circle catching under normal (inverted) vision behaviour. The leftmost column illustrates the horizontal coordinate of the agent and the object (straight horizontal lines), the middle one shows the frequency of each node of the network as the task progresses (colour lines for nodes with inputs, dashed lines otherwise) and the rightmost ones present the synchronisation index. Green lines reflect the autonomous trajectory. (a) $0K_c$, (b) $2K_c$ and (c) $10K_c$.
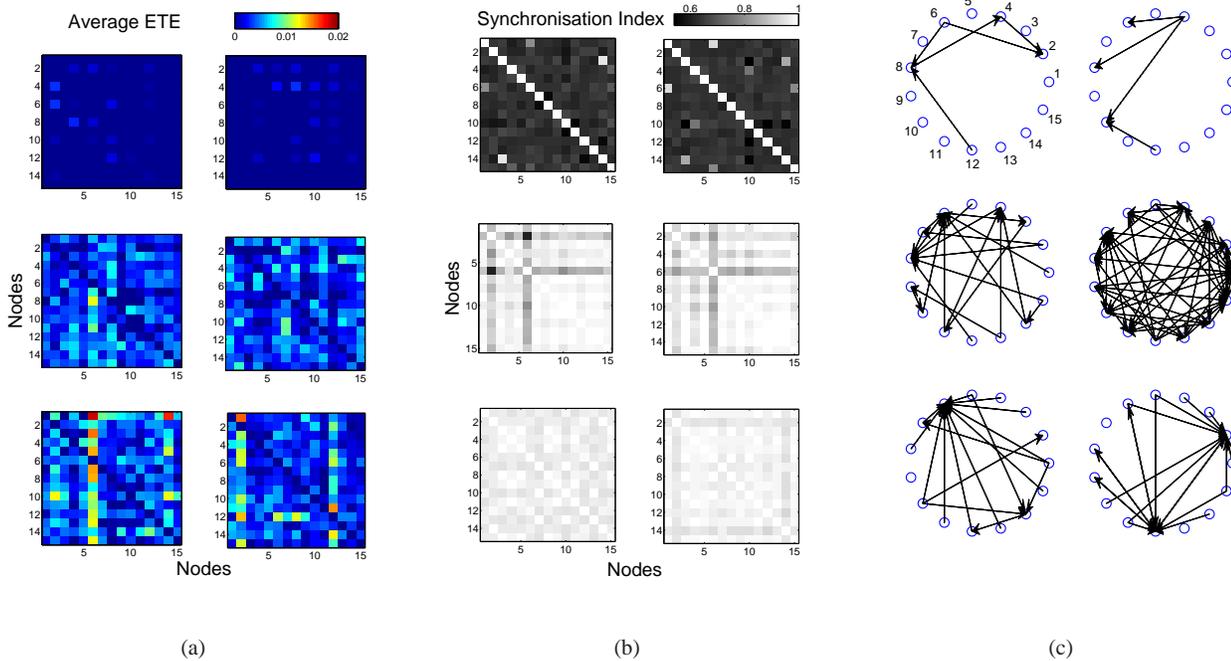
Fig. 10 (a) Averaged information flow between the oscillatory nodes of the network (Equation 11). The autonomous flow is subtracted in all plots to highlight the task-specific activity. (b) Same as (a) but for the synchronisation index. (c) functional network - a directed arrow is drawn whenever the Transfer Entropy between two given nodes exceeds 50% of the maximum ETE value obtained in the whole task. The rows correspond (top to bottom) to the $0K_c$, $2K_c$ and $10K_c$ agents. The left side of each column refers to orientation with normal vision, the right to inverted vision.

tion flow under normal and inverted vision is similar, with very little activity for the uncoupled network, whereas the more strongly coupled agents ($2K_c$ and $10K_c$), although having a nearly identical trajectory when catching the circles with normal and inverted vision in each corresponding scenario, show variation in their internal nodes' information flow. Having only the environmental feedback to shape its behaviour, the $0K_c$ network fails to detect and respond properly to this conflicting situation because the phase dynamics of each node can only be altered by the external stimulus, which in this case is made ambiguous, while the internal activity of the $2K_c$ and the $10K_c$ networks is modulated by the environmental context - the conflicting readings result in different phase dynamics that are exploited to generate fit behaviour. The synchronisation index adds to that perspective by showing that although highly synchronised, the $2K_c$ network is still flexible enough to escape the nearly entrained state of the $10K_c$ network or the incoherent behaviour of the uncoupled framework, producing the fittest agent.

With these two plots in mind, observe the picture of the equivalent functional network for the chosen agents in this task (Figure 10(c)). There are richer functional connections between nodes in the $2K_c$ network not only when comparing across different coupling configurations but also between the two vision conditions (inverted vision uses a

denser set of functional connections than normal vision). This higher performing network is more able to reconfigure its internal functional architecture than its counterparts, allowing it to cope well with both conditions.

Figure 11 shows the temporal evolution of the net Transfer Entropy between nodes (Equation 12). It reveals the absence of significant readings in the uncoupled network but shows much higher levels in the $2K_c$ network as the task progresses, with the $10K_c$ network having alternating moments of high and low readings. This plot provides a different angle of analysis on our system by making evident the differences in the flow of information within the different networks and hence in the relationships between the nodes in the various contexts. Given that the phase difference dynamics captures all the effective dynamics of the Kuramoto Model (Maistrenko et al., 2005), it's possible to see how different functional networks emerge from the interaction with the environment within the same coupling condition, stressing adaptability, or across different coupling conditions, revealing aspects of functional structure dependent on degree of coupling.

Comparing Figures 8(e) and 8(f) with Figure 11, notice how the magnitude of the ETE_{sm} and ETE_{ms} flows are high at particular moments of the task in the $2K_c$ agent, whilst the ETE between nodes has noticeable deviations

from zero across the trial. There is, therefore, no clear similarity between the sensorimotor and the network's nodes phase activity information flow. In the $10K_c$ network, however, the higher variability of the flow between nodes have a correspondence in the ETE_{sm} and ETE_{ms} flow. This correspondence can be explained by the higher value of coupling between nodes in the sense that the sensorimotor activity perturbs a much more rigid framework, driving all nodes' behaviour. The $2K_c$ network is more loosely coupled so that the perturbations modulate and influence the nodes' behaviour, but do not ultimately determine it.

In summary, the previous experiments with categorical perception revealed a limitation of both weakly and strongly coupled networks in achieving good performance in discriminating objects with different shapes. The second experiment with orientation in normal and inverted vision conditions shows that the weakly coupled network continues to fail in a conflicting scenario, but the more strongly coupled networks succeed. The conclusion is that, for our framework, networks with high synchronisation levels are able to detect and respond to visual stimuli but struggle to exploit the environmental context. In other words, the highly synchronised structures evolved here can detect (even with altered vision) but not discriminate among different external stimuli. In both tasks, networks with 'intermediate' levels of coupling (and hence propensity for synchronisation) perform best as this provides a more fluid, flexible network able to reconfigure and adapt to different behavioural conditions. However, as we have seen, this optimal 'intermediate' level is significantly different for the two tasks.

5 Discussion

Current research on neurophysiological processes shows that oscillatory neural activity is closely related to cognition and behaviour, with synchronisation mechanisms playing a key role in the integration and functional organization of different cortical areas (Varela et al., 2001; Engel et al., 2001; Womelsdorf et al., 2007). Nevertheless, its informational content and relationship with behaviour - and hence cognition - are still to be fully understood (Rieke et al., 1997; Mazzoni et al., 2008; Deco et al., 2011; Flanders, 2011).

In this context, following an Evolutionary Robotics approach, we evolved simulated robotic agents controlled by a spontaneously rhythmic network of coupled phase oscillators, showing that their performances in variations of a categorical perception task depend on the synchronisation regime of the network. This was also shown to be the case for an orientation task under normal and inverted vision, but the dependence on the synchronisation regime was quite different from in the first task. The analysis focused both on a behavioural level description, investigating the agent/object trajectories, and on a mechanism level description, exploring
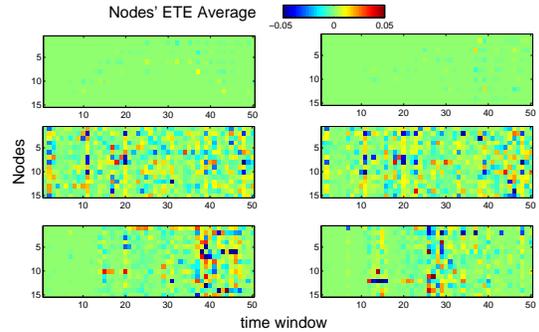

Fig. 11 Net Transfer Entropy from a given node to all remaining nodes of the network in 50 equally spaced time windows (Equation 12). The autonomous flow is subtracted in all plots to highlight the task-specific activity. The rows correspond (top to bottom) to K_c, $2K_c$ and $10K_c$ agents. The left column is for normal vision and the right for inverted vision.

the dynamics and the information transfer properties within and between the agent and the environment.

The design of the experimental framework attempted to capture relevant properties of its biological counterparts that would suit the aims of our study. The network model is inspired by the Kuramoto Model of coupled phase oscillators whilst the robotics approach explores the enactive nature of sensory and cognitive processes. Our focus is on three aspects: the information flow between sensors and motors, capturing the body/environment interaction; the average flow between nodes of the network, stressing the internal communication; and the evolution in time of the average net flow from a given node to all its counterparts, revealing at various moments of the task a higher/lower participation in the overall information dynamics.

The first experiment explored the fitness variation of evolved individuals under different network couplings K for a categorical perception task. The results showed that there is a statistically significant difference in performance for different strengths of coupling, with the best fitness evolution performance obtained when $K = K_c$, the critical value. Although the term "critical" here is not directly related to its meaning for self-organised criticality (Beggs, 2008), previous work (Kitzbichler et al., 2009) showed that the Kuramoto Model presents properties of critical systems (e.g. greater information transmission, storage, sensitivity to external stimulus) when its coupling is tuned to the critical value (Equation 4). Our results confirm this for the categorical perception task.

For the agent with best performance ($K = K_c$), the dynamical analysis shows that there is a very rich phase behaviour in the critical coupled network, with nodes oscillating at a variety of frequencies, occasionally synchronising and forming independent functionally connected clusters. In the Effective Transfer Entropy (ETE) analysis, we notice a higher flow magnitude in the sensorimotor loop time se-

ries (ETE_{sm} and ETE_{ms}), which has a behavioural correspondence and relates to sharper adjustments in the agent's trajectory, suggesting that the agent is more able to convert sensory information into motor responses to maximize its performance; we also observe a much broader internal communication between oscillators, with a great variety of functional connections in comparison with other coupling configurations. Remarkably, even though the network is fully connected, the flow is not symmetric, and the higher performance of the critical coupled network can be linked to its flexible ongoing dynamics which can be modulated by the environmental context.

The second experiment studied orientation to falling circles under normal and inverted vision, particularly interesting for presenting to the agent an ambiguous and conflicting scenario.

We found that the performance increases with the degree of coupling, but this time reaches its peak between 2 and $4K_c$, decreasing for the very strong coupling values greater than that. We compared three evolved individuals, one with no coupling ($0K_c$), one with strong coupling ($2K_c$) and the other with very strong coupling ($10K_c$) between oscillators. The first is unable to succeed on the task whilst the 2 and $10K_c$ agents produce good performance in both sensory modes, with the $2K_c$ agent outperforming the very strongly coupled one.

Dynamical analysis shows that the $2K_c$ network rapidly synchronises, with some nodes escaping the entrained state only when the corresponding sensory input is large. However, there is a diverse phase dynamics, with conflicting values resulting in different phase dynamics that are exploited to generate fit behaviour that can cope with the different visual conditions. Therefore, even though the network has a fixed, all-to-all pre-established coupling configuration, with no explicit plasticity mechanisms, it can respond and adapt to conflicting sensory conditions and obtain good performance on the task. This is further stressed by the Transfer Entropy analysis which shows that the normal and inverted vision conditions elicit a difference in the information flow at specific points of the agent's trajectory. The results highlight the asymmetry of information flow and its behavioural correspondence obtained in fully connected oscillatory networks only by changing their level of synchronisation. Hence, the results suggest that synchronised networks with an 'intermediate' level of coupling are more able to reconfigure functional connections not only when comparing across different coupling configurations but also between the two vision conditions, reflecting its superior performance on the task.

Concluding our analysis for both experiments, consider Figures 12(a) and 12(b), which display the formation of groups of oscillators during the execution of the task for the selected coupling conditions for Experiments 1 and 2, respectively.

The network has initially 15 single-node "clusters", each corresponding to an oscillator. A group (cluster) is merged whenever the phase difference between two given nodes is below a certain threshold (chosen to be 1 here) and the frequencies are similar (less than $1Hz$ apart). Hence values vary from 1 (single, giant cluster containing all the oscillators) to 15 (single-node clusters).

In the first experiment, the uncoupled network is associated with a large number of single-node groups, whereas the strongly coupled network has only one big group through nearly all of the task, indicating complete synchronisation. Curiously, the critically coupled network alternates between these two extremes. The same behaviour is observed for Experiment 2, with the exception that the very strong coupled network desynchronises at the end of the task, as revealed by earlier analysis. These plots reinforce an important aspect raised in the Introduction and mentioned throughout the paper: synchronisation dynamics modulate the assembly, collapse and reconfiguration of functional sub-networks. In our experiments, a higher diversity of such assemblies are present in the top performing agents, linking the adaptability of the agent to the flexibility of its functional connections.

The previous analysis is complemented by Figures 12(c) and 12(d), which present the dynamics of a metastability index λ and the Largest Lyapunov Exponent (LLE) for the chosen scenarios. λ (Equation 13) is calculated as the estimated variance of the synchronisation index r (Eq. 3) over all time points $t = 1 \ldots T$ and reflect the diversity of synchronisation states obtained for each coupling configuration of the network (Shanahan, 2010).

$$\lambda = \frac{1}{T-1} \sum_{t=1}^{T} (r(t) - <r>_T)^2 \tag{13}$$

The Largest Lyapunov Exponent, calculated using the phase variables θ_i, $i = 1 \ldots N$ (Maistrenko et al., 2005), indicate how chaotic the phase dynamics - and hence our system - are. To obtain the LLE, we followed the method described in Rosenstein et al. (1993) which is specially designed to be robust and fast in small data sets.

In both experiments, we notice that the metastability decreases (as expected) as the coupling increases, showing that there is less variation in the synchronisation index. The dynamics of the number of groups relate to the metastability index, with the critically coupled network also appearing in an intermediate position. An interesting point for the critical coupling condition is to notice how the agent/environment interaction in the categorical perception task reduces the spontaneous chaoticity level of the network (calculated for the autonomous network dynamics), whereas in the inverted vision scenario, near the critical coupling, chaotic behaviour is more prominent when the agent is engaged in the task than when under the influence of the autonomous network
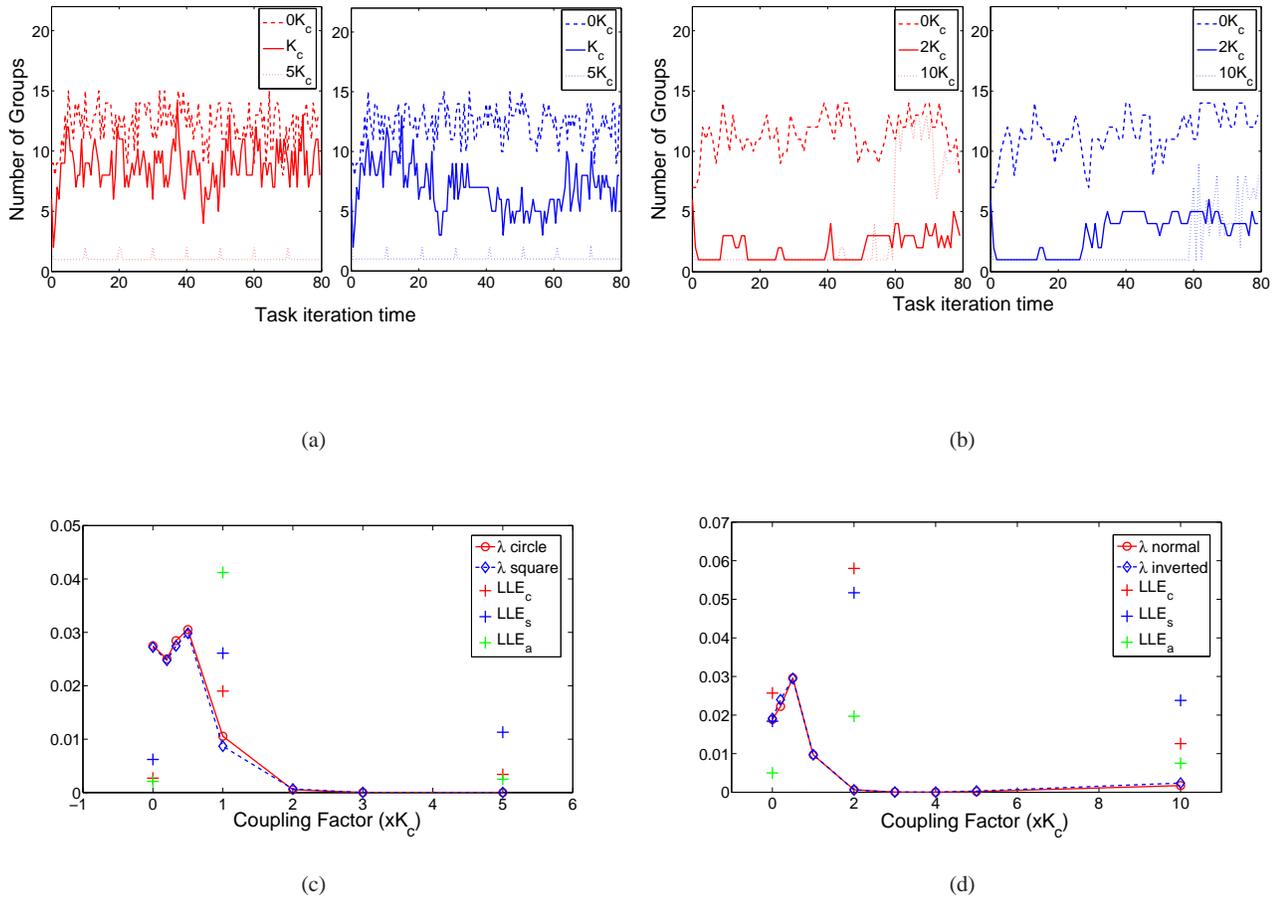
(a)   (b)

(c)   (d)

Fig. 12 Clustering of oscillators and metastable states for different coupling conditions as the task progresses. (a) and (b): nodes that have similar frequencies and a phase shift below a small threshold are grouped. Values vary from 1 (single, giant cluster containing all the oscillators) to 15 (single-node clusters). (c): metastability index λ (Equation 13) superposed with the Largest Lyapunov Exponent (LLE) for the circle (LLE_c), square (LLE_s) and autonomous (LLE_a) configurations. Larger λ indicates a great diversity of synchronisation states in the network as the task progresses; LLE greater than 0 suggests chaotic dynamics. (d): same as (c), but for the normal vision, inverted vision and autonomous configurations. Figures (a) and (c) refer to Experiment 1 and Figures (b) and (d) refer to Experiment 2.

dynamics. Also, note that the Largest Lyapunov Exponent is close to zero in the uncoupled autonomous scenario for both tasks ($0K_c$). This is anticipated and due to the lack of external stimulus, which makes each node of the network oscillate independently in a quasiperiodic way. The sensorimotor coupling changes this behaviour (even without connections between nodes, coupling can be seen as occurring "indirectly" through the environment), and the network activity becomes more chaotic. Therefore, the above observations reflect the adaptability and modulation of the network activity previously noticed in the information flow analysis. It also suggests that although neural synchronisation clearly plays an important role in the generation of behaviour in the fittest agents, other kinds of transient dynamics are also exploited (Santos et al., 2011).

There are some limitations for the framework presented here. Starting with the network model, the original formulation of the Kuramoto Model doesn't take into account important properties of real nervous systems (e.g. spatial distribution of units, transmission delays, asymmetrical interactions), therefore an important step towards a more biologically plausible architectures would be to implement extended versions of the model which tackle some of these constraints (Tass, 2006; Cumin & Unsworth, 2007; Breakspear et al., 2010). These extensions may also unfold more complex, metastable dynamics which are intrinsically connected to dynamic pattern formation in brain activity and hence are fundamental to adaptive behaviour (Omelćhenko et al., 2008; Tognoli & Kelso, 2009; Chialvo, 2010; Shanahan, 2010). As for the task scenarios, the experiments were conducted following an evolutionary robotics approach which,

despite all the advantages mentioned throughout the paper, suffers from high computational cost for fitness assessments which can be a limitation in scaling up to more complex tasks (see Bongard (2011) for an encouraging technique to tackle this problem). Also, simulated experiments have to be carefully set-up to avoid undesirable gaps when transferring to real world scenarios. In this sense, the addition of noise during simulations is essential to obtain good cross-platform performance (Jakobi et al., 1995). Another useful approach would be to incorporate ontogenetic plasticity mechanisms to the robot's controller, which are shown to greatly increase adaptability to sensory disruptions and dynamic environments (Urzelai & Floreano, 2001). It is, nevertheless, a promising approach to testing hypotheses which have not been extensively explored yet in the context of the relationship between neural synchronisation and behaviour. Finally, the Transfer Entropy analysis, in common with other information theoretic tools, relies on the proper choice of parameters, on accurate probability density estimations and require large number of data points for robust estimations. Using the extensions to the original Transfer Entropy measure suggested in Marschinski & Kantz (2002), we reduced some of the undesirable effects caused by small data sets. Likewise, further work should consider varying the orders of the Markov processes (parameters m and n in Eq. 5) for they influence the estimated value of transfer entropy and may be used to reveal time dependences since motor, sensors and nodes may operate at different time scales (Lungarella & Sporns, 2006; Gourévitch & Eggermont, 2007).

The Matlab source code for our numerical simulations is available at http://www.informatics.sussex.ac.uk/users/rm262 or from the authors on request.

References

Acebrón, J. A., Bonilla, L. L., Pérez Vicente, C. J., Ritort, F., & Spigler, R. (2005). The Kuramoto model: A simple paradigm for synchronization phenomena. *Reviews of Modern Physics*, 77(1), 137–185.

Arthuis, M., Valton, L., R´gis, J., Chauvel, P., Wendling, F., Naccache, L., Bernard, C., & Bartolomei, F. (2009). Impaired consciousness during temporal lobe seizures is related to increased long-distance cortical-subcortical synchronization. *Brain*, 132, 2091-2101.

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

Beggs, J. (2008). The criticality hypothesis: how local cortical networks might optimize information processing. *Phil. Trans. of the Royal Society, Series A*, 366(1864), 329-43.

Boden, M. A. (2006). *Mind as machine: a history of Cognitive Science*. Vol. I and II, Oxford Univ. Press., UK.

Bongard, J. C. (2011). Innocent until proven guilty: Reducing robot shaping from polynomial to linear time. *IEEE Trans. Evolutionary Computation*, 15(4), 571–585.

Borst, A. & Theunissen, F. E. (1999). Information theory and neural coding. *Nature Neurosci.*, 2(11), 947–57.

Breakspear, M., Heitmann, S., & Daffertshofer, A. (2010). Generative models of cortical oscillations: neurobiological implications of the Kuramoto model. *Front. Hum. Neurosci.*, 4.

Brown, P. (2003). Oscillatory nature of human basal ganglia activity: relationship to the pathophysiology of Parkinson's disease. *Mov. Disord.*, 18, 357–63.

Buehlmann, A. & Deco, G. (2010). Optimal information transfer in the cortex through synchronization. *PLoS Comput. Biol.*, 6(9).

Buzsaki, G. (2006). *Rhythms of the Brain*. Oxford University Press.

Chialvo, D. R. (2010). Emergent complex neural dynamics. *Nature Physics*, 6(10), 744–750.

Cumin, D. & Unsworth, C. (2007). Generalising the Kuramoto model for the study of neuronal synchronisation in the brain. *Physica D*, 226(2), 181–196.

Dale, K. & Husbands, P. (2010). The evolution of reaction-diffusion controllers for minimally cognitive agents. *Artificial Life*, 16(1), 1–19.

Deco, G., Buehlmann, A., Masquelier, T., & Hugues, E. (2011). The role of rhythmic neural synchronization in rest and task conditions. *Front. Hum. Neurosci.*, 5:4.

Di Paolo, E. A. (2000). Homeostatic adaptation to inversion of the visual field and other sensorimotor disruptions. *From Animals to Animals, Proc. Sixth Int. Conf. Sim. of Adaptive Behavior, Paris, J-A. Meyer, A. Berthoz, D. Floreano, H. Roitblat and S W. Wilson (eds)*, (pp. 440–449).

Engel, A. (2010). Directive minds: How dynamics shapes cognition. *in Stewart, Gapenne and Di Paolo (Eds.). Enaction: Toward a New Paradigm for Cognitive Science*.

Engel, A., Fries, P., & Singer, W. (2001). Dynamic predictions: oscillations and synchrony in top-down processing. *Nat. Rev. Neurosci.*, 2(10), 704–716.

Ermentrout, G. & Kleinfeld, D. (2001). Traveling electrical waves in cortex: Insights from phase dynamics and speculation on a computational role. *Neuron*, 29, 33–44.

Flanders, M. (2011). What is the biological basis of sensorimotor integration? *Biol. Cybern.*, 104, 1–8.

Floreano, D., Husbands, P., & Nolfi, S. (2008). Evolutionary robotics. *In Siciliano and Khatib (Eds.) Springer Handbook of Robotics*, (pp. 1423–1451).

Floreano, D. & Keller, L. (2010). Evolution of adaptive behaviour in robots by means of Darwinian selection. *PLoS Biology*, 8(1).

Floreano, D. & Mattiussi, C. (2008). *Bio-Inspired Artificial Intelligence*. MIT Press.

Glass, L. (2001). Rhythmic processes in physiology. *Nature 410 277-284*.

Gourévitch, B. & Eggermont, J. (2007). Evaluating information transfer between auditory cortical neurons. *J. Neurophysiol.*

Harnad, S. e. (1987). *Categorical Perception: The Groundwork of Cognition.* Cambridge University Press.

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

Hatsopoulos, N., Ojakangas, C., Paninski, L., & Donoghue, J. (1998). Information about movement direction obtained from synchronous activity of motor cortical neurons. *Proc. Natl. Acad. Sci. USA, Neurobiology*, 95, 1570615711.

Husbands, P. (2009). Never mind the iguana, what about the tortoise? Models in adaptive behaviour. *Adaptive Behavior*, 17(4), 320–324.

Husbands, P., Smith, T., Jakobi, N., & O Shea, M. (1998). Better living through chemistry: Evolving GasNets for robot control. *Connection Science*, 10, 185–210.

Ijspeert, A., Crespi, A., Ryczko, D., & Cabelguen, J. (2005). From swimming to walking with a salamander robot driven by a spinal cord model. *Science*, 315(5817), 1416–1420.

Izhikevich, E. (1999). Weakly pulse-coupled oscillators, FM interactions, synchronization, and oscillatory associative memory. *IEEE Trans. Neural Networks*, 10(3), 508–526.

Izquierdo, E. (2008). The dynamics of learning behaviour: a situated, embodied, and dynamical systems approach. *PhD thesis, CCNR, University of Sussex*.

Jackson, A., Gee, V., Baker, S., & Lemon, R. (2003). Synchrony between neurons with similar muscle fields in monkey motor cortex. *Neuron*, 38, 115–125.

Jakobi, N., Husbands, p., & Harvey, I. (1995). Noise and the reality gap: the use of simulations in evolutionary robotics. *In Proc. 3rd European Conference on Artificial Life. F. Moran et al. (Eds.)*, (pp. 704–720).

Jolliffe, I. T. (2002). *Principal Component Analysis*. Springer-Verlag New York.

Katz, P. E. (1999). *Beyond Neurotransmission: Neuromodulation and its Importance for Information Processing*. Oxford University Press.

Kitzbichler, M., Smith, M., Christensen, S., & Bullmore, E. (2009). Broadband criticality of human brain network synchronization. *PLoS Comput. Biol.*, 5(3).

Kullback, S. (1959). *Information Theory and Statistics*. Wiley, New York.

Kunyosi, M. & Monteiro, L. (2009). Recognition of noisy images by PLL networks. *Signal Processing*, 89.

Kuramoto, Y. (1984). *Chemical Oscillation, Waves, and Turbulence*. Springer, New York.

Li, Z. & Hopfield, J. (1989). Modeling the olfactory bulb and its neural oscillatory processings. *Biol. Cybern.*, 61(5).

Lungarella, M., Ishiguro, K., Kuniyoshi, Y., & Otsu, N. (2007a). Methods for quantifying the causal structure of bivariate time series. *Int. J. of Bifurcation and Chaos*, 17(3):903-921.

Lungarella, M., Pegors, T., Bulwinkle, D., & Sporns, O. (2005). Methods for quantifying the informational structure of sensory and motor data. *Neuroinformatics*, 3 (3), pp. 243-262.

Lungarella, M., Pitti, A., & Kuniyoshi, Y. (2007b). Information transfer at multiple scales. *Physical Review E*, 056117.

Lungarella, M. & Sporns, O. (2006). Mapping information flow in sensorimotor networks. *PLoS Comput. Biol.*, 2(10):e144.

Maistrenko, Y. L., Popovych, O. V., & Tass, P. A. (2005). Chaotic attractor in the Kuramoto Model. *Int. J. of Bifurcation and Chaos*, 15(11), 34573466.

Marschinski, R. & Kantz, H. (2002). Analysing the information flow between financial time series: An improved estimator for transfer entropy. *The European Physical Journal B*, 30 (2), p.275-281.

Mazzoni, A., Panzeri, S., Logothetis, N. K., & Brunel, N. (2008). Encoding of naturalistic stimuli by local field potential spectra in networks of excitatory and inhibitory neurons. *PLoS Comput. Biol.*, 4(12).

McDonnell, M., Ikeda, S., & Manton, J. (2011). An introductory review of information theory in the context of computational neuroscience. *Biol. Cybern.*, 105, 55–70.

Murray, J. D. (1989). *Mathematical Biology: An Introduction*. Springer-Verlag Berlin Heidelber.

Noe, A. (2004). *Action in perception*. Cambridge, MA: MIT Press.

Omelćhenko, O., Maistrenko, Y., & Tass, P. (2008). Chimera states: The natural link between coherence and incoherence. *Phys. Rev. Lett.*, 100, 044105.

Pfeifer, R., Lungarella, M., & Iida, F. (2007). Self-organization, embodiment, and biologically inspired robotics. *Science 318 1088*.

Pikovsky, A., Rosenblum, M., & Kurths, J. (2001). *Synchronization: A Universal Concept in Nonlinear Sciences*. Cambridge University Press.

Pitti, A., Lungarella, M., & Kuniyoshi, Y. (2009). Generating spatiotemporal joint torque patterns from dynamical synchronization of distributed pattern generators. *Front. NeuroRobotics*, 3:2, 1–14.

Pitti, A., Niiyama, R., & Kuniyoshi, Y. (2010). Creating and modulating rhythms by controlling the physics of the body. *Autonomous Robots*, 28:3, pp. 317-329.

Rieke, F., Warland, D., van Steveninck, R. R., & Bialek, W. (1997). *Spikes: Exploring the Neural Code*. Bradford Book - MIT Press, Cambridge MA.

Rolls, E. T. & Treves, A. (1998). *Neural networks and brain function*. Oxford University Press.

Rosenstein, M., Collins, J., & Deluca, C. (1993). A practical method for calculating largest Lyapunov exponents from small data sets. *Physica D: Nonlinear Phenomena*, 65(1-2), 117–134.

Santos, B., Barandiaran, X., & Husbands, P. (2011). Metastable dynamical regimes in an oscillatory network modulated by an agent's sensorimotor loop. *Proc. IEEE Symposium on Artificial Life (SSCI)*.

Schreiber, T. (2000). Measuring information transfer. *Phys. Rev. Lett.*, 85, 461464.

Shanahan, M. P. (2010). Metastable chimera states in community-structured oscillator networks. *Chaos*, 20, 013108.

Silverman, B. (1998). *Density Estimation for Statistics and Data Analysis*. London: Chapman and Hall.

Singer, W. (1993). Synchronization of cortical activity and its putative role in information processing and learning. *Annu. Rev. Physiology*, 55, 349–374.

Singer, W. (1999). Neuronal synchrony: a versatile code for the definition of relations? *Neuron*, 24(1), 49–65.

Staniek, M. & Lehnertz, K. (2008). Symbolic transfer entropy. *Phys. Rev. Lett.*, 100 (15).

Stewart, J., Gapenne, O., & Di Paolo, E. E. (2010). *Enaction: Toward a New Paradigm for Cognitive Science*. MIT Press.

Strogatz, S. (2000). From Kuramoto to Crawford: exploring the onset of synchronization in populations of coupled oscillators. *Physica D*, 143.

Szczepanski, J., Arnold, M., Wajnryb, E., Amigo, J. M., & Sanchez-Vives, M. V. (2011). Mutual information and redundancy in spontaneous communication between cortical neurons. *Biol. Cybern.*, 104, 161–174.

Taga, G. (1994). Emergence of bipedal locomotion through entrainment among the neuro-musculo-skeletal system and the environment. *Physica D*, 75(1-3), 190–208.

Tass, P. (2006). *Phase resetting in medicine and biology: stochastic modelling and data analysis*. Springer.

Tognoli, E. & Kelso, J. A. S. (2009). Brain coordination dynamics: True and false faces of phase synchrony and metastability. *Prog. Neurobiol.*, 87(1), 31–40.

Urzelai, J. & Floreano, D. (2001). Evolution of adaptive synapses: Robots with fast adaptive behavior in new environments. *Evol. Comput.*, 9(4), 495–524.

Varela, F., Lachaux, J., Rodriguez, E., & Martinerie, J. (2001). The brainweb: phase synchronization and large-scale integration. *Nat. Rev. Neurosci.*, 2(4).

Vicente, R., Wibral, M., Lindner, M., & Pipa, G. (2011). Transfer entropy: a model-free measure of effective connectivity for the neurosciences. *J. Comput. Neurosci.*, 30, 4567.

von der Malsburg, C. (1981). The correlation theory of the brain. *Internal report. Max-Planck-Institute for Biophysical Chemistry, Gottingen, Germany*.

Williams, P. & Beer, R. (2010). Information dynamics of evolved agents. *From Animals to Animals, Proc. 11th Int. Conf. Sim. of Adaptive Behavior, S. Doncieux, B. Girard, A. Guillot, J. Hallam, J.-A. Meyer and J-B. Mouret (eds)*, (pp. 38–49).

Winfree, A. (1980). *The geometry of biological time*. Springer-Verlag.

Womelsdorf, T., Schoffelen, J., Oostenveld, R., Singer, W., Desimone, R., Engel, A., & Fries, P. (2007). Modulation of neuronal interactions through neuronal synchronization. *Science*, 316(5831), 1609–1612.