

2018

Analysis of Embodied and Situated Systems from an Antireductionist Perspective

Da Rold, Federico

<http://hdl.handle.net/10026.1/11302>

<http://dx.doi.org/10.24382/785>

University of Plymouth

All content in PEARL is protected by copyright law. Author manuscripts are made available in accordance with publisher policies. Please cite only the published version using the details provided on the item record or document. In the absence of an open licence (e.g. Creative Commons), permissions for further reuse of content should be sought from the publisher or author.

This copy of the thesis has been supplied on condition that anyone who consults it is understood to recognize that its copyright rests with its author and that no quotation from the thesis and no information derived from it may be published without the author's prior consent.

Analysis of Embodied and Situated Systems from an Antireductionist Perspective



Federico Da Rold

School of Computing Electronics and Mathematics

University of Plymouth

A thesis submitted to the University of Plymouth
in partial fulfilment of the requirements for the degree of

Doctor of Philosophy

August 2017

Acknowledgements

The author is extremely grateful to the members of supervisory team Tony Belpaeme, Lingfen Sun and in particular to his Director of Studies Angelo Cangelosi that kindly accepted me as a student when my former main supervisor secured a position in another country, easing the final part of my journey. During my doctoral studies I have been surrounded by great colleagues that significantly improved my life and work Frank Broz, Martin Peniak, Marek Rucinski, Anthony Morse, Samantha Adams, Michael Klein, George Terzakis, Christopher Ford, Daniela Conti. I would also like to thanks Carole Watson for her great support. I am deeply thankful to Alessandro Di Nuovo for his encouraging comments and for improving the quality of this work. I am also grateful to Thomas Wennekers for his insightful comments on the first version of this manuscript. During my short stay in Tokyo I met a truly great lab and a very motivated and innovative visiting scholar, Miguel Aguilera. Thank to all of you. My early training as predoctoral researcher at the Institute of Cognitive Sciences and Technologies in Rome was fundamental for finishing this dissertation. I would like to deeply thank Stefano Nolfi, Vieri Santucci, Francesco Rigoli, Onofrio Gigliotta, Giancarlo Petrosino, Dimitri Ognibene, Giovanni Pezzulo, Marco Mirolli, Elio Tuci, Valerio Sperati, Fabio Paglieri. Many thanks to three great friends and colleagues that supported me over these years, Francesco Capponi, Thomas Mindos and Peter Gibbons. You guys really made my life easier! Finally, the scientists that made this thesis possible, mentoring and shaping me as a researcher. Takashi Ikegami, as during the short stay in his lab I developed a great number of ideas that form the backbone of this work. Anna Borghi, for introducing me to this field, mentoring and assisting me throughout my career. Domenico Parisi, for training me during my first year as postgraduate researcher, constantly giving me new ideas and perspectives, and teaching me how to efficiently and enthusiastically handle this difficult job.

Abstract

The analysis of embodied and situated agents from a dynamical system perspective is often limited to a geometrical and qualitative description. However, a quantitative analysis is necessary to achieve a deep understanding of cognitive facts.

The field of embodied cognition is multifaceted, and the first part of this thesis is devoted to exploring the diverse meanings proposed in the existing literature. This is a preliminary fundamental step as the creation of synthetic models requires well-founded theoretical and foundational boundaries for operationalising the concept of embodied and situated cognition in a concrete neuro-robotic model. By accepting the dynamical system view the agent is conceived as highly integrated and strictly coupled with the surrounding environment. Therefore the antireductionist framework is followed during the analysis of such systems, using chaos theory to unveil global properties and information theory to describe the complex network of interactions among the heterogeneous sub-components.

In the experimental section, several evolutionary robotics experiments are discussed. This class of adaptive systems is consistent with the proposed definition of embodied and situated cognition. In fact, such neuro-robotics platforms autonomously develop a solution to a problem exploiting the continuous sensorimotor interaction with the environment.

The first experiment is a stress test for chaos theory, a mathematical framework that studies erratic behaviour in low-dimensional and deterministic dynamical systems. The recorded dataset consists of the robots' position in the environment during the execution of the task. Subsequently, the time series is projected onto a multidimensional phase space in order to study the underlying dynamic using chaotic numerical descriptors. Finally, such measures are correlated and confronted with the robots' behavioural strategy and the performance in novel and unpredictable environments.

The second experiment explores the possible applications of information-theoretic measures for the analysis of embodied and situated systems. Data is recorded from perceptual and motor neurons while robots are executing a wall-following task and pairwise estimations of the mutual information and the transfer entropy are calculated in order to create an exhaustive map of the nonlinear interactions among variables. Results show that the set of information-theoretic employed in this study unveils characteristics of the agent-environment interaction and the functional neural structure.

This work aims at testing the explanatory power and impotence of nonlinear time series analysis applied to observables recorded from neuro-robotics embodied and situated systems.

Authors declaration

At no time during the registration for the degree of Doctor of Philosophy has the author been registered for any other University award. Work submitted for this research degree at Plymouth University has not formed part of any other degree either at Plymouth University or at another establishment. Relevant scientific seminars and conferences were regularly attended at which work was often presented.

Word count for the main body of this thesis: **38083**

Signed: _____

Date: _____

Publications:

Da Rold, F. *Defining Embodied Cognition: The Problem of Situatedness*. New Ideas In Psychology 2018 - accepted

Da Rold, F. *Information-Theoretic Decomposition of Embodied and Situated Systems*. Neural Networks 2018, 105, 94-107

Da Rold, F. *Chaotic analysis of embodied and situated agents*. Robotics and Autonomous Systems 2017, 95, 143-159

Da Rold, F. *Deterministic chaos in mobile robots*. In International Joint Conference on Neural Networks (IJCNN) 2015, 1-7

Contents

- Acknowledgements** **iii**

- Abstract** **iv**

- Author’s declaration** **v**

- 1 Introduction** **9**
 - 1.1 Aims and objectives 13
 - 1.2 Contribution to knowledge 13
 - 1.3 Thesis outline 14

- 2 Background** **17**
 - 2.1 Embodied and situated cognition 17
 - 2.1.1 Computationalism and symbolic artificial intelligence 17
 - 2.1.2 Connectionism and the sub-symbolic representation of knowledge 19
 - Artificial neural networks and learning 20
 - Genetic algorithms 22
 - Evolutionary robotics 24
 - 2.1.3 A new perspective on intelligence: Embodiment and situatedness 26
 - 2.1.4 Robots: scientific tools for studying embodied cognition 27
 - 2.1.5 The problem of definition 30
 - 2.1.6 Different embodied views in neural modeling: developing the concept of time and language grounding 33
 - 2.1.7 Closing remarks on the concept of embodied cognition: the problem of situatedness 37
 - 2.2 Understanding an embodied and situated system 40
 - 2.2.1 Reductionism and antireductionism 41
 - 2.2.2 Holistic science: general system theory and cybernetics 43
 - 2.3 Nonlinear science 45
 - 2.3.1 Information theory: nonlinearities in communication channels 45

2.3.2	Application of information-theoretic measures to embodied and situated robotic systems	46
2.3.3	The revolution of chaos theory	47
2.3.4	Informal definition of chaos and the problem of empirical determinism	49
2.3.5	Application of chaotic measures to data recorded from humans and robots	50
3	Methods	53
3.1	Adaptive systems	53
3.1.1	Artificial neural networks	53
3.2	Nonlinear time series analysis	56
3.2.1	Phase space reconstruction	56
3.2.2	Chaos theory	58
3.2.3	Information theory	61
3.3	Evolutionary robotics simulator	65
3.4	Experiment: Chaotic dynamics in simulated robots	66
3.4.1	Evolutionary robotics experiment	66
3.5	Experiment: Information theoretic decomposition of embodied systems	69
3.5.1	Simulated robotic scenario	70
4	Chaotic dynamics in simulated robots	73
4.1	Introduction	73
4.2	Results	74
4.2.1	Chaotic analysis	75
4.2.2	Chaotic measures and performance	79
4.2.3	Chaotic dynamics and behavioral strategies	83
4.3	Discussion	93
4.4	Conclusion	96
5	Decomposing an embodied and situated system <i>via</i> information theoretic measures	98
5.1	Introduction	98
5.2	Results	100
5.2.1	Local mutual information	102
5.2.2	Local transfer entropy	109
5.2.3	Mutual information	114

CONTENTS

5.2.4	Transfer entropy	115
5.2.5	Decomposing the behaviour via information-theoretic measures	119
5.3	Discussion	123
5.4	Conclusions	126
6	Conclusions	128
6.1	Current limits and future works	133
	Information theory	135

List of Figures

3.1	Neural architectures. (a) Feedforward neural network (b), simple recurrent neural network, and (c) recurrent neural network.	53
3.2	The Lorenz system. (a) shows the chaotic evolution of the system defined by the equations. (b) depicts the same attractor reconstructed using the variable x	57
3.3	EvoRobot simulator. The picture shows a running experiment, depicting the environment and the robot on the left, the neural activity on the right.	65
3.4	The simulated environment. The squared arena contains a target area placed at the center. The target area changes color between green and blue at regular intervals. The robot, depicted as a red circle, achieves positive reward remaining on the green target area. When the target area becomes blue, the reward is negative and the robot is forced to escape from danger, entering in the empty space of the arena.	67
3.5	The experimental scenario. (a) depicts the maze and the robot (black circle), placed at the starting position. (b) shows a schematic representation of the robot's morphology with the displacement of the sensors. The left (IF-L) and right (IF-R) infrared sensors utilized in the experiment are painted in red. (c) the artificial neural network. Numbers refer to the identification numbers (ID- n) used in subsequent figures for referring to specific connections. The output layer includes two motor neurons that separately control the left (M-L) and right (M-R) motors.	70
4.1	Time series embedding. Reconstructed chaotic attractor from the vector storing the position of the best robot generated from seed 24, controlled by a SRN, and evolved in environments where the target area changes color with 61 different timescales (E3) in the range $T_{evo} = \{470, 471, \dots, 529, 530\}$. (a) shows the position of the robot, data used for the reconstruction of the attractor. (b) illustrates the robot's trajectory in Cartesian space and (c) displays the reconstructed attractor in phase space.	75
4.2	Parametrization of the embedded vector. A correct reconstruction of the attractor from a mono-dimensional time series depends upon appropriate values of embedding delay τ and dimension m . Data is collected from the best robot of seed 24, evolved in E3 with timescales $T_{evo} = \{470, 471, \dots, 529, 530\}$, and controlled by a SRN. (a) shows the estimated mutual information between the original time series and itself after applying temporal delays (Δt). The first minimum is reached with $\Delta t = 8$ (red dot), which represents the correct value for τ . (b) displays the percentage of false nearest neighbors for different embedding dimensions. To fully unfold the attractor the minimum m is 3, as it lowers the percentage below the threshold of 5% (red dashed line).	76

- 4.3 **Attractor dimension.** The fractal dimension is estimated with the correlation integral. The algorithm is applied to the best robot of replication 24, controlled by a SRN and evolved in E3 with timescales $T_{evo} = \{470, 471, \dots, 529, 530\}$. (a) shows the logarithmic plot of the correlation integral versus the distance threshold r for different embedding dimensions m . The linear fit (red dashed line) of the scaling region estimates the fractal dimension D_2 . (b) reports the estimated values of D_2 for different embedding dimensions, starting from the minimum m , for the chaotic robot (black dots), and a stochastic system generated with synthetic data (red stars). 77
- 4.4 **Level of chaos.** The level of chaos is quantified by λ_1 , calculated from the logarithm of the average local divergence of neighboring points over time. The linear fit of the scaling region (red dashed line) estimates the value of λ_1 of the recorded data. (a) shows convergence of the estimation of λ_1 at higher embedding dimensions for the best robot of seed 24, whereas (b) shows that the agent evolved from seed 27 fails to stabilize. Both datasets are recorded from robots controlled by a SRN and evolved in environments with timescales $T_{evo} = \{470, 471, \dots, 529, 530\}$. 78
- 4.5 **Mean difference bootstrapping of the fitness between chaotic robots controlled by FF or SRN neural network.** The statistical test evaluates the effect of the neural controller on the performance in the presence of chaos. Positive confidence intervals (red solid lines) indicate a better performance of agents with FF controllers. During the testing phase, robots act in environments never experienced during evolution, expanding the upper and lower bounds of the set of timescales by 50. (a), (b), and (c) refer to robots evolved in E1, E2, and E3, respectively. 81
- 4.6 **Mean difference bootstrapping of the fitness between chaotic and non-chaotic robots.** The statistical test evaluates the effect of chaotic behaviors on the robots' performance during the testing phase. Positive confidence intervals (red solid lines) indicate a better performance of chaotic agents. (a) shows that there is no significant difference in performance in E1. (b) and (c) show a larger fitness for chaotic robots evolved in E2 and E3, respectively. 82
- 4.7 **Mean difference bootstrapping of the fitness between chaotic and non-chaotic robots evolved in E1 and E2, tested in a larger number of novel environments.** Robots evolved in E1 and E2 are tested with 161 timescales, as robots of E3. Positive confidence intervals (red solid lines) indicate a better performance of chaotic agents. (a) shows no significant difference in E1. (b) show a better performance for chaotic robots evolved in E2. 83
- 4.8 **Correlations between chaotic measures and fitness.** The fitness value measures the average performance of chaotic robots during the testing phase, when the number of timescales is increased of ± 50 . Points of the scatter plot represents single robots. The gray dashed line indicates the Pearson correlations of λ_1 (top) and D_2 (bottom) with the fitness for robots evolved in E1(a), E2(b), and E3(c). 84

4.9 **Behavioral analysis of robots evolved in E1.** (a) shows the values of λ_1 plotted against D_2 for the whole set of chaotic robots evolved in E1. Markers indicate robots selected for the analysis. (b-e) depicts the robots' trajectories over 7,000 time steps and (f-i) the fitness achieved in every environment experienced by the agents during the testing phase. 86

4.10 **Behavioral analysis of robots evolved in E2.** (a) scatter plot of the level of chaos (λ_1) and complexity (D_2) of chaotic robots evolved in E2. Markers refer to robots selected for the analysis. (b-e) show the evolution of the trajectories for 7,000 time steps in Cartesian space. (f-i) show the fitness score in 121 environments. 88

4.11 **Behavioral analysis of robots evolved in E3.** (a) points represent chaotic robots with the associated level of chaos (λ_1) and fractal dimension (D_2). Markers show the selected robots. (b-e) depict the trajectories followed by the analyzed robots while executing the task for 7,000 time steps. (f-i) illustrates the performance of the 4 robots during the testing phase in 161 different environments. 90

4.12 **Behavioral analysis of robots that developed a similar strategy in different experiments.** Trajectories followed by robots with high fractal dimension evolved in E1(a), E2(b), E3(c), and the corresponding fitness achieved during the testing phase (d-f). Strategies adopted by robots with low fractal dimension (g-h), and the performance of robots trained in E1(j), E2(k), and E3(l). 92

5.1 **Time series recorded from the robots' perceptual and motor neurons.** The figure depicts the neural activity during two laps around the maze. (a) and (b) represent the left and right infrared sensor, respectively. (c) refers to the left motor actuator and (d) to the right output. Replications 1-8 navigates counter-clockwise, whereas the 9th and 10th follow the opposite direction. The neural activation is characterised by large positive spikes of the infrared sensor directed toward the outer walls and drops in activity to the inner motor. 100

5.2 **Average values of the local mutual information estimated on unpruned neural networks.** The outcome of the pairwise local mutual information is normalised separately for Gaussian and rectangular kernels in the range $[0, 1]$ to obtain a weighted ranking of interactions. The degree of exchange of information between variables unveils functional relationships of the underlying neural structure. 103

5.3 **Estimation of the local mutual information in pruned neural networks evolved in the 4th replication.** The local mutual information is calculated for all functional neural sub-architectures, creating separated weighted rankings. Synapses are labelled with identification numbers. This replication shows resiliency towards pruning, relying mainly on 3 connections that are characterised by high degrees of information exchange. 105

5.4 **Estimation of the local mutual information in pruned neural networks evolved in the 10th replication.** The local mutual information detects high degrees of dependence between variables that form the functional backbone of the neural network. Furthermore, the local mutual information individuates the presence of a synapse in the left sensorimotor chain, which stabilises the robot's trajectory. . . . 106

5.5 **Estimation of the local mutual information in pruned neural networks evolved in the 2nd replication.** The uniqueness of this neural structure is the necessity of a connection between the sensor directed towards the inner wall and the outer motor actuator. This characteristic is captured by the local mutual information that estimates high degree of information exchange. However, probabilities estimated with different kernels cause inconsistencies on the mapping of the functional neural structure. 107

5.6 **Effects of the smoothing parameter on the local mutual information estimated with Gaussian kernel.** The formulation of Gaussian kernel includes a parameter α that modifies the bandwidth h or values stored in the bandwidth matrix \mathbf{H} . Bar charts refer to the complete set of all functional neural architectures evolved in the 10th replication, demonstrating that different values of this parameter do not significantly affect the ranking of the local mutual information. 108

5.7 **Accuracy of the local transfer entropy.** The accuracy of the local transfer entropy for describing the information dynamics is tested applying lesions during temporal intervals characterised by information transfers from the motor to the inputs. The dynamic unveiled by the local transfer entropy is compared with surrogate datasets, thus applying lesions during random time steps. The prediction horizon is defined as the convergence of the average lifespan of robots pruned according to the information dynamic and random lesions. The local transfer entropy is estimated with different parametrizations, specifically x and y refer to recorded time series, X and Y are embedded vectors, \dot{x} is the difference of signal. (a) and (b) show the prediction horizon and the percentage of lesions applied during the test for rectangular and Gaussian kernel, respectively. (c) summarises the result with an accuracy measure, based on the number of lesions and the robots' survival capability. 111

5.8 **Dynamic of the information transfer.** The local transfer entropy estimated between the right infrared sensor and left motor actuator for the robot of the 1st replication during the execution of a lap. In (a) probabilities are estimated with Gaussian kernel on embedded time series, whereas in (b) a rectangular kernel is employed. (c) and (d) depict the local information transfer calculated on the recorded signals using Gaussian and rectangular kernel, respectively. 113

5.9 **Estimation of the mutual information.** The dependences between variables unveil characteristics of the agent-environment interaction. The systems display high degrees of dependence between the left infrared sensor and the motor actuators in robots that navigate counter-clockwise. On the contrary, agents that follow the opposite direction are characterised by high degrees of dependence between the right sensor and the outputs. 115

5.10 **Estimation of the transfer entropy.** The information flow unveils characteristics of the agent-environment interaction and the underlying functional neural structure. The information transfer is directed from the motors to the inputs as the behaviour is mainly driven passively by the biases. However, the transfer entropy detects the importance of fundamental synapses, which are characterised either by information flow directed towards the output, or the smallest degree of information transfer in the opposite direction. The transfer entropy is calculated on embedded vectors. Positive values on the recurrent connection in the output layer refer an information flow from left to right. 116

5.11 **Estimation of the transfer entropy in pruned neural networks evolved in the 10th replication.** Bar charts show the restructuring of the information flow caused by pruning. The transfer entropy detects an increased activity on the synapse connecting the left infrared sensor with the right motor. This result is consistent with an observation of the robot's trajectory that becomes more unstable, thus requiring active adjustments during the execution of the task. 117

5.12 **Estimated local transfer entropy and robot's trajectory.** Graphical representation of the information dynamics related to the robot's spatial position. Red circles indicate information transfers from the sensors to the motors, or a left to right direction for recurrent connections in the output layer. The size of the marker is proportional to the magnitude of the estimated local transfer entropy. Blackened environments (ID1, ID3 and ID4, bottom row) refer to lesioned synapses. 120

Chapter 1

Introduction

Understanding cognitive phenomena and intelligent behaviours is a challenging endeavour. Until recently, cognition was exclusively studied from an anthropocentric perspective, focusing mainly on high level cognitive functions typically found in humans. Traditional cognitive science conceives the mind as a computational machine that operates exclusively over symbols using discrete and rigid logical rules, abstracting away from the physicality of the agent and the environment (Fodor, 1975; Pylyshyn, 1986). Hence, cognition is essentially amodal and detached from experience. Typically, the mind is described as a computer, capable of performing algorithmic thinking with an extraordinary efficiency.

The field of embodied cognition, however, is opposing such perspective, underlining the importance of bodily experiences, physicality and biological grounds for constructing and structuring concepts and knowledge (Wilson, 2002). This alternative explanation strongly rejects the ontological separation between the abstract mind and the physical world defined by Descartes, thus considering the abstract mind and the concrete agent as an indivisible entity.

By embracing this theoretical framework, cognitive scientists radically reconsider the approach followed while studying cognitive facts (Parisi et al., 1990). The idea of a mind separated from contextual information and biological reality is meaningless as the cogitating agents are constantly immersed in a flow of sensorimotor interactions (Beer, 2008), where ecological properties and environmental features are fundamental for explaining cognition (Francisco Varela & Rosch, 1991). Clearly, the embodied view extends the boundaries of cognitive science, formerly an exclusive domain of psychologists, philosophers, linguists and computer scientists, to a larger number of scientific fields. In fact, the reformulation proposed by the novel perspective opens

the field of cognitive science to several disciplines such as neuroscience, biology, robotics, and physics, which are all involved in the attempt to explaining intelligence. However, transdisciplinarity naturally leads to fragmentations and numerous different ideas regarding the concept of embodiment. In fact, every field of study is biased toward a peculiar aspect of the integrated system brain-body-environment. This causes wide variations on the definition of embodied cognition without reaching a substantial agreement concerning the foundation of the theory. Surely, this dynamism within the field is positive, avoiding radicalisation and depicting the field as still at its infancy, and thus with great potential for further development. However, this issue is of particular importance for roboticists, computer scientists and computational modelers. Recreating a cognitive agent, either instantiated in a physical robot or a simulated scenario, requires a precise definition that draws clear boundaries. Modelers build cognitive agents from scratch, and thus the necessity of defining the subject of study is critical.

Embodied cognitive science is related to two other concepts, namely grounded and situated cognition. Although there is no general consensus regarding the meaning of these concepts, the existing literature allows an at least vague framing of the theory (Pezzulo et al., 2011; Fischer, 2012; Myachykov et al., 2014). Overall, embodiment refers to the agent's peculiar body that simultaneously shapes and constraints the cognitive abilities. The term embodied entails also facts concerning specific phylogenetic developments, thus referring to the stable characteristics that build a common ground among conspecifics, as well as idiosyncrasies that are the product of the development of the single individual. The constant and immutable characteristics of the physical world are instead the domain of grounded cognition. The surrounding environment is the material that grounds concepts and structures knowledge. Typically, the subjects of study are statistical regularities that are captured by the brain, somehow similarly to an associative engine. The term situated cognition, instead, conveys two different meanings. Firstly, situatedness is the unstable contextual knowledge, either available in specific spatiotemporal physical settings or defined by the society where the agent lives in. From this perspective, situatedness includes the concrete spatiotemporal instantiation of the physical world and social structures (Barsalou, 2008). Secondly, such concept is framed in the context of dynamical system theory, referring to the continuous sensorimotor interaction between the agent and the environment (Beer, 2000).

In this thesis, I embrace the second definition of situated cognition, stressing the importance of a dynamical unfolding of cognitive facts, rather than emergent statistical structures that structure our knowledge.

Embracing a dynamical system perspective requires a deep reformulation of the analytical methods employed. In fact, the brain-body-environment system is studied from an integrative perspective. The network of concepts and the intelligent behaviours are global features that emerge from nonlinear interactions among the heterogeneous sub-components of the system. The disentanglement of such a holistic system during the analysis of embodied and situated agents restricts the explanatory power of cognitive modeling, obscuring the underlying dynamics. Thus, a scientific method based on reductionism may not suffice. A natural phenomenon is not always fully understood by studying single variables in isolation. The modern formulation of the antireductionist philosophy of science developed from two distinct fields. In biology, Von Bertalanffy (1968) conceived a foundational framework named system science. Currently, science is highly specialised and different disciplines study natural phenomena in isolation, decontextualizing the subject of investigation. For example, biologists study the metabolic processes of a single neuronal cell without considering the brain, neuroscientists focus on biological neural networks excluding the metabolic processes of the organism, and so forth. According to the principles established in the context of system thinking, biological or physical entities are open systems capable of communicating with the external environment (Fang & Casadevall, 2011). Studying the sub-components in isolation is not a wrong approach. However, such scientific method ignores global behaviours and properties that emerge from the interactions among different variables that shape the system (Ahn et al., 2006b). Therefore, developing a method aimed at describing the global characteristics that emerge from the interaction of several different layers of heterogeneous components is essential. At the same time, cybernetics was defining a novel approach to system engineering based on the same idea of self-regulation and continuous and real time interaction with the contextual information incoming from the environment (Wiener, 1961). Early cyberneticists did not restrict the field to inanimate artefacts, claiming the applicability of the novel science to any system, including biological organisms (Ashby, 1956).

Recently, several fields are accepting the intrinsic limits of a purely antireductionist perspective (Ahn et al., 2006b,a). The more evident example is the mapping of the human genome. Genetics and biologists realised that a static mapping of the DNA does not suffice at predicting the ontogenetic development of a biological organism (Fang & Casadevall, 2011). Phenotypic gene expressions are explained by a complex network of interactions between genes that are contextualised in a specific environment, rather than considering the genome as an array of switches that activates-deactivates certain features of the organism. Similarly, functional connectivity in biological neuronal networks is understood in terms of complex nonlinear interactions among neurons, metabolic processes and environmental influences (Stam, 2005). Clearly, proposing a profound reformulation of the scientific methodology is extremely problematic. In fact, mathematical tools normally utilised the reductionist framework are not applicable. Reductionism avoids complexity simplifying the system by assuming linear relationships and homogeneity among variables. On the contrary, antireductionism stresses the importance of nonlinearities and interactions among heterogeneous components forming the sub-system, focusing also on the temporal evolution of the system. Typically, the mathematical tools developed in the context of nonlinear dynamical system theory are employed by scientists following the antireductionist philosophy of science. However, this field of applied mathematics is still at an early stage of development and many aspects are not fully understood. Moreover, the application of dynamical system theory to empirical data is often qualitative and geometrical, rather than quantitative (Beer, 1995; Carvalho & Nolfi, 2016). In this thesis, I propose a quantitative and analytical explanation of embodied and situated systems employing data-driven approaches developed in the field of nonlinear time series analysis (Kantz & Schreiber, 1997), strictly following the antireductionist perspective. The embodied and situated systems are modelled with evolutionary robotics (Nolfi & Floreano, 2000), an established method for the study of minimal cognitive agents (Harvey et al., 2005). The evolved robots are testbeds for chaos theory (Kaplan & Glass, 2012) and information theory (Shannon & Weaver, 1949), models of scientific explanations that capture global properties of the systems and complex relationships among the sub-components. The scientific tools utilised in this thesis are simulated robotic platforms, as synthetic models of cognitive agents abstract away from the complexity of a biological living system by providing a fully controllable experimental environment (Nolfi & Gigliotta,

2010). For example, a simulated robotic platform overcomes the problem of collecting spurious data resulting from noisy environments. Similarly, a robot develops concepts that are unbiased by previous knowledge or personal experiences. Furthermore, a robotic platform easily allows for invasive interventions on the agent, such as synaptic lesions or significant impairments on the bodily functionalities (Nolfi & Floreano, 2000). Therefore, cognitive neuro-robotics provides a valuable experimental practice for studying embodied and situated systems (Harvey et al., 2005).

1.1 Aims and objectives

In what follows, I summarise the motivations and theoretical foundations that guide this dissertation. During the design of the robotic experiments, as well as the subsequent data analysis, the author is inspired by the following principles. Firstly, embodied and situated cognition is conceived according to the dynamical system perspective, and thus the triadic system brain-body-environment is strictly coupled and indivisible. Secondly, the operationalisation of the definition of embodied cognition, as well as the analytical framework aimed at understanding such systems, are consistent with an antireductionist philosophy of science, thus avoiding disentanglements of the sub-components of the system. Thirdly, the mathematical tools employed during data analysis are developed in the field of nonlinear time series analysis, which provides a numerical description of the embodied and situated system, rather than a solely qualitative and geometrical characterisation concerning the temporal evolution of the dynamic. The aim of this work is to investigate nonlinearities in embodied and situated systems from an antireductionist perspective, addressing the following scientific questions:

- Are chaotic measures (i.e. level of chaos and fractal dimensions) correlated to the robots' performance?
- Do information-theoretic measures (i.e. mutual information and transfer entropy) detect functional relationships among artificial neurons and/or between the agent and the environment?

1.2 Contribution to knowledge

The contributions to knowledge consists of:

- A preliminary review of the current definitions of embodied cognition, underlining the different conceptualisations of the term situatedness. In particular, I individuate two main perspectives, where situated cognition refers either to opaque contextual information or the dynamical and real time interaction of the agent with the environment. Following the second approach, derived from a dynamical system view of cognition, I propose an antireductionist and systemic view of scientific investigation, which may overcome current limits that affect the analysis of embodied and situated systems.
- The first experimental part of the dissertation investigates the efficacy of chaos theory on the study of global properties that autonomously emerge in complex systems. The evolutionary robotic model is an operationalisation of the more general class of biological and artificial embodied and situated systems, and thus the study aims at showing the potentials and limits of a formal quantitative method that is applicable to a great number of different experimental scenarios. Furthermore, I contribute to debate surrounding the field of chaos theory, which is affected by issues concerning the applicability of chaotic measures to empirical data.
- The second experimental part focuses on the complex network of interactions among the sub-components of the system. The mathematical tools employed are developed in the field of information theory, showing that nonlinear time series analysis unveils characteristics of the underlying functional neural architecture, as well as characteristics of the agent-environment interaction.

1.3 Thesis outline

This section summarises the content of the chapters, following the outline of the thesis.

- **Chapter 2** provides the necessary background to motivating and understanding the experiments described in later sections. This chapter develops covering three main topics. Firstly, the history of cognitive science is briefly discussed, explaining the transition from the computational tradition to the embodied view. Several definitions of embodied cognition proposed in the existing literature are reviewed from a theoretical point of view, focusing in particular on the different meanings of term situatedness. The topic is further clarified

and supported with examples of cognitive neuro-robotics models. Secondly, the differences between the reductionist and antireductionist approach to scientific investigation are discussed, covering the topic of system science and cybernetics. Finally, this chapter underlines the lack of quantitative measures for the analysis of embodied and situated system from a dynamical system and holistic perspective. The solution proposed is nonlinear time series analysis, in particular chaos theory and information theory.

- **Chapter 3** describes the machine learning method applied to neuro-robotics systems and the set of nonlinear mathematical tools utilised during data analysis. Evolutionary robotics, a machine learning technique inspired by biological evolution, is employed for developing minimal cognitive abilities in mobile robots controlled by artificial neural networks. In the second part, concepts of nonlinear time series analysis are explained, including the set of chaotic tools that are applied to attractors reconstructed from recorded data and information-theoretic measures.
- **Chapter 4** investigates the possible applications of chaotic analysis in miniature mobile robots executing a temporal task. The aim of the experiment is twofold. Firstly, investigate the possibility of applying chaotic measures to embodied and situated systems. Secondly, the study provides evidence to the controversial topic concerning the application of chaos theory to empirical data by providing a large dataset. Compared with similar works, the experiment tests the capability of chaotic measures at unveiling minor differences that are the product of the evolutionary process in different populations of robots performing the same task, rather than manipulating macroscopic variables.
- **Chapter 5** describes a wall-following scenario where robots autonomously develop a behavioural strategy to navigating in a simple squared maze. Similarly to the experiment of the previous chapter, macroscopic variables, e.g. body morphology or environmental properties, are not manipulated. However, evolution produces significant differences in the resulting behavioural strategies and functional neural structures. Time series recorded from perceptual and motor neurons are analysed with mutual information and transfer entropy in order to unveil underlying characteristics of the artificial brain and peculiarities of the

agent-environment interaction in different populations of robots.

- **Chapter 6** summarises the results of the experiments reported in this thesis, drawing overall conclusions regarding the application of nonlinear time series analysis from an antireductionist perspective to embodied and situated agents. Furthermore, limits and future works are discussed.

Chapter 2

Background

2.1 Embodied and situated cognition

The aim of this section is to provide a historical and conceptual overview about the field of cognitive science and artificial intelligence. The endeavour of building intelligent machines challenged the scientific community since the development of computers (Turing, 1950). However, cognitive processing is conceived in two different senses. The computational tradition defines the mind as a computational amodal machine that manipulates abstract symbols according to precise logical rules (Newell & Simon, 1976). This perspective developed from the field of mathematical logic, linguistics and the tradition of analytical philosophy (Fodor, 1975). The consequence of such theory of cognition is the study of the mind detached from the physical instantiation of the agent's body and decontextualized from the external world (Pylyshyn, 1986). Conversely, the embodied and situated view stresses the importance of contextual information defined by the environment and the peculiarities determined by the body (Francisco Varela & Rosch, 1991). Thus, cognition develops from real time and modality-specific interactions with the world (Barsalou, 2008).

2.1.1 Computationalism and symbolic artificial intelligence

Through the centuries, countless philosophers and scientists have tried to answer fundamental questions concerning the way we think. What is intelligence? What is the relationship between the mind and the body? What is the relationship between the external world and the subjective experience? These deep and theoretically challenging questions still struggle to receive a unique answer leaving the debate still open. Crucial to understanding the perspective of embodied cognitive science is the mind-body problem, including its extension, which entails comprehending how perception and action relate to our intelligence.

Although the dispute about the relation between mind and body can be traced back to the ancient Greek philosophy and medieval scholastic tradition, Descartes is the first author that frames the problem in modern terms. He supports the idea of substance dualism, giving to mind and body two different ontological statuses: the material *res extensa*, a mechanical body subjected to the laws of physics, and the abstract *res cogitans*, realm of the laws of thought. Clearly, this position has deep epistemological implications. By defining *res cogitans* as persistent and indivisible the study of the empirical world is now freed of any religious implications, making the theory consistent with the scholastic tradition and the idea of a soul.

There are strong implications following the acceptance of an ontological dualism, and the most important for the field of cognitive science is an idealistic view of cognition. The central element for understanding intelligence is the mind, abstracted from any physical instantiation. The physical traits of the cogitating agent, including brain, body, and perceptual system, are marginal for the study of cognitive facts.

One of the most revolutionary inventions of the last century was the development of computers, which is rooted in the *a-machine* (Turing, 1936), also known as Turing machine. In 1928, David Hilbert challenged the community of mathematicians with the Entscheidungsproblem, asking whether a function on the natural numbers is computable. The solution of the problem addresses the notion of effective computability, which requires an algorithm for solving the given function, that is, a mechanical procedure based on a finite set of rules. The concept of algorithm was still lacking of a formal definition and Alan Turing proposed an anthropocentric explanation, finding the necessary and sufficient mental abilities that a human uses when solving a logical problem. The Turing machine has a tape of infinite length divided into discrete cells that can store one of the three symbols of the simple alphabet $\{0; 1; B\}$, where B is blank. The machine reads the content of a cell and according to a table that defines deterministic logical rules. The possible (computational) actions are writing a symbol, moving to another cell and halting. The latter command solves the algorithm without entering in an infinite loop. The architecture of modern general-purpose digital computers has been developed by John von Neumann at Princeton and is explicitly based on Turing's theoretical work. The machine, now physical, has a central pro-

cessing unit, a memory, and it is capable to receiving inputs from and showing some outputs to the user.

During the second half of the last century, classic cognitivism (Fodor, 1975) and computational theories of mind (Pylyshyn, 1986) developed a framework which bases the study of cognition and intelligence on the analogy between a computer and any form of mental activity, extending the original operationalisation of human algorithmic thinking to every domain of cognition. From this perspective, the mind operates over abstract symbols following syntactical rules and the experience of the bodily-self, as well as the external environment, is a mere source of inputs. Perception is passive and action is just an effect of the cognitive outcome our inner thoughts. Therefore, the concepts developed throughout the process of thinking are abstract and amodal. One of the most notable examples of computational theory of mind is the physical symbol system hypothesis proposed by Allen Newell and Herbert Simon (Newell & Simon, 1976). A physical symbol system is a formal language capable of combining meaningful symbols to generate expressions that, in turn, can be manipulated in a combinatorial way to generate new expressions. According to Newell and Simon, such system is necessary and sufficient for describing, and thus replicating, general intelligence. Defining intelligence as the ability of manipulating symbols implies that machines based on computers may rightfully be consider as intelligent agents, since they are able to process a logical formal system. However, robots controlled by a system of first-order logic acting in an environment is a contradiction. In fact, such system operates robustly only if the logical system includes a number of axioms that exclude any arbitrary change in the environment (McCarthy & Hayes, 1969).

2.1.2 Connectionism and the sub-symbolic representation of knowledge

The year 1943 saw the birth of artificial neural networks. In a seminal paper, McCulloch & Pitts (1943) developed a model of a biological brain, showing that neuronal ensembles are capable of solving propositional logic. The authors propose a mathematical formalisation that captures some essential characteristics of the existing physiological model of biological neurons. The building block of the computational model is an artificial neuron, and the artificial brain is a network of interconnected basic units. The information is processed hierarchically and serially, the lower

layers sending output to the higher layers. In those years, the neuron was conceived as an all-or-none functional unit, which is either active or inactive. Therefore, the artificial neuron is represented by a step function that returns 1 if the input is above a predefined threshold. Today, binary units are seldom employed and are substituted by nonlinear activation functions (e.g. logistic function). Further details on the artificial neural networks utilised in this thesis are provided in Section 3.1.1. Each unit in the higher layers receives inputs from several artificial neurons of the lower layers. The connections between neurons, which represent artificial synapses, are weighted and may strengthen or weaken the incoming signal. A neural network may also be provided with an internal memory by designing architectures with recurrent connections (Elman, 1990). A hidden or output neuron receives as input its activation during the previous time step, thus storing the internal states that characterised the past activity of the network.

The term *connectionism* was first introduced by Hebb (Elman, 1998b). However, the prevailing direction of research is based on the work of the parallel and distributed processing approach (McClelland & Rumelhart, 1987; Rumelhart & McClelland, 1988), which grows in opposition to the computational and symbolic view of cognition. The fundamental concepts that distinguish the connectionist view are the followings. Firstly, cognitive processing develops in *parallel*, there is no functional modularity as in the abstract model of the mind, with the analogous of a central computational unit and a memory reservoir; instead, nature provides a large number of simple interconnected units that work simultaneously. Secondly, knowledge representation is sub-symbolic and distributed within the network of synaptic weights. Thirdly, cognition is robust towards malfunction of the system, e.g. lesions of synaptic weights, and noise. The imprinting of a biologically inspired perspective on intelligence and cognition is evident, and the decades that follow the rebirth of neural computation lead to a critique towards the centrality of high-level cognitive functions and abstract reasoning, typically found in humans.

Artificial neural networks and learning

Like a biological neural system, artificial neural networks are capable to learn from experience (Hebb, 1949). In fact, the strength of the synaptic weights is subjected to modifications throughout the learning process. The learning methods to adjust the set of values are classified into three distinct

groups.

Supervised. The desired answer is known, and thus the mapping $f(x, y)$ between definite values of input x and output vectors y constructs a training set. An example is the gradient descent method of the backpropagation algorithm (Rumelhart et al., 1986). This learning method is often used in top-down modeling in cognitive robotics and connectionism (McClelland & Rumelhart, 1987; Rumelhart & McClelland, 1988; Flusberg et al., 2010; Stramandinoli et al., 2012), where input output patterns are linguistic labels.

Semi-supervised. Although there is general knowledge about the appropriate outputs, the specific set of output values is unknown. Algorithms within this class are reinforcement learning (Sutton & Barto, 1998) and artificial evolution (Nolfi & Floreano, 2000) that adjust synaptic weights according to a reward policy or a global utility function. These methods are widely employed in autonomous robotics where the researcher studies the necessary and sufficient condition for the emergence of a desired behaviour. For example, if a robotic platform has to approach a light source (Tuci et al., 2004), the mapping $f(x, y)$ of the sensorimotor flow at every time step is not known. By exploiting the mechanism of artificial evolution, a fitness function defines the desired behaviour and the space of the solution is explored in order to find a maximum.

Unsupervised. The network is trained to autonomously extract generalities in given dataset. A widely used training method is based on Kullback-Leibler divergence and Gibbs sampling, framing the neural network in the class of generative models. Recently, a neural system that gained popularity is the restricted Boltzmann machine, which forms the building block for deep architectures (Ackley et al., 1985). In its basic formulation, the neural network has one input layer and one hidden layer connected with bidirectional connections. The units are binary and stochastic, and the probability of being active is defined by a logistic activation function that processes the incoming signal. The input layer in a trained neural network reproduces the vector originally clamped at the beginning of the bidirectional and recursive propagation. In a restricted Boltzmann machine the output layer is absent (Hinton et al., 2006), and thus the neural network does not provide an explicit response to a given input. Rather, the hidden layer autonomously extracts common features of different classes of elements included in the training set.

Genetic algorithms

The genetic algorithms are a class of machine learning methods that explore the space defined by a complex problem, the so-called fitness landscape, searching for the optimal solution. The idea of Evolutionsstrategie (evolutionary strategies) was developed by Eigen (1973) during the 1960s, in order to create software capable of adjusting autonomously its own parameters. The author takes inspiration from the Darwinian evolutionary theory, creating a parent characterized by a genotype, which generates an offspring with random mutations. Simultaneously, John Holland begun his effort in creating a formal model aimed at studying the theory of evolution from a computational perspective. A decade later, the genetic algorithm was a fully developed and functional theory (Holland, 1992). Similarly to the early mathematical model of the brain, which started the field of the artificial neural networks, the genetic algorithm operationalises the concept of Darwinian evolution. From a structural point of view, a vector represents a haploid genome, where each element is an allele, that is, a specific instantiation among the several mutually exclusive variants of a single locus in the sequence. The functional aspects of the algorithm follow few simple steps. The biological principle known as survival of the fittest that underlies the evolution of biological species is founded on two mechanisms:

- *Vitality*: an individual is able to survive for a period of time sufficient to reach a reproductive stage.
- *Fertility* an individual replicates the genome which will populate future generations.

In natural environments, defective genomes do not live long enough to fully develop and reproduce. Furthermore, weak individuals are more likely to generate a small number of offspring. In the mathematical model, the concepts of vitality and fertility which guide the natural process of evolution are substituted with a precise fitness function that defines *ad hoc* objectives. Each solution produced by different genomes is numerically quantified according to the fitness function, and the best individuals are selected for generating offspring, and thus transmit the genetic heritage to future generations. There are several criteria for selecting the genotypes that will reproduce. The simplest and most commonly used are:

- *Rank based*: in a population of N individuals, the n genotypes that achieved the largest fitness score generate m offspring in order to create a new population of N individuals.
- *Roulette-wheel selection*: each genotype has a probability proportional to the associated fitness score to generate offspring.
- *Tournament selection*: the genotypes are randomly paired; a random number is associated to each pair and if the value is above an arbitrary threshold, the weakest individual is selected for reproduction.

The evolutionary process explores the fitness landscape by means of random genetic mutations. While generating offspring, single values in the artificial genomes have a small probability to change value, thus slowly moving on the space of the possible solutions. Another possibility for inducing mutation in the genome, possibly increasing the fitness score, is the cross-over, where non-overlapping portions of different genomes selected for reproduction are combined. The random mutation is a local operator that smoothly explores the search space. On the contrary, cross-over is a global operator that may cause major shifts on the fitness landscape. Following this approach, optimal sub-solutions to a global problem are merged together. The genetic algorithm heavily relies on random events such as random selection, mutations and cross-over, while optimising a solution to a problem, during an iterative process that unfolds through several generations. In order to avoid regressions, elitism is often employed: a clone of each parent is included in the population of the next generation.

An early application showing the efficiency of the genetic algorithm is the solution to the iterated prisoner dilemma (Axelrod, 2006). The problem is a fictional scenario, where two rational individuals are in solitary confinement and have the options to remain silent or betray, confessing the crime of the other inmate. The outcome is represented by a bi-dimensional square matrix of order 2: if they betray each other or remain silent, both prisoners spend two or one year in prison, respectively. On the other hand, if one individual confesses and the other stays silent, the betrayer is set free, whilst the other inmate is condemned to a charge of three years. In the iterated version of the problem, the question is asked several times and the prisoners remember a defined number of previous games. Axelrod tested the genetic algorithm with an iterated

prisoner dilemma where the individuals remember the three previous games, and thus there are $4^3 = 64$ possible combinations. The artificial genome is binary, encoding the two possible options available to the prisoners, and has a total of 70 loci, which define all possible combinations and 3 initial moves for each player. Therefore, the fitness landscape is defined by 2^{70} possible states. The population consists of 20 genotypes evolved for 50 generations. The genetic algorithm solved efficiently the prisoner dilemma with limited resources, considering that only $20 \times 50 = 1000$ states on the fitness landscape are explored.

What lays underneath such efficiency at exploring the fitness landscape? Why is a subset of few genotypes enough for finding an optimal solution? The efficiency of the heuristic search paradigm is justified by the building block hypothesis (Holland, 1992), which assumes that well-performing genomes are partially defined by chunks that are relevant for the solution of the problem. The idea is formalised with the notion of schemata (Holland, 1992; Mitchell, 1998). For the simple case of binary encoding, the schemata is defined by an alphabet $L = \{0; 1; *\}$, where 0 and 1 are precise instances of the allele and the symbol * is a wildcard referring to both 0 and 1. Each schemata is defined by an order, which represents the total number of 0s and 1s, and a length, that is, the number of symbols between the first and the last locus containing either 0 or 1. For example, a schemata of a genotype consisting of 8 loci of order 4 and length 5 may be written as $S = \{ *1 * *01 * 0 \}$. Each genome with l loci is an instance of several schemata, in particular adopting a binary encoding 2^l . Therefore, a population of n individuals is actually searching with a number between 2^l and $n2^l$ schemata. This hypothesis sheds light on the fast convergence shown by the application of the optimization algorithm. By using the genetic algorithm, the fitness landscape is explored with the explicit parallelism of concrete instantiations of different genotypes, but also with the implicit parallelism of the schemata.

Evolutionary robotics

Evolutionary robotics models adaptive and emergent autonomous behaviour mimicking the process of Darwinian evolution (Nolfi & Floreano, 2000). Although practical applications of this technique are multifaceted, the backbone is the application of genetic algorithms to adjust free parameters in a robotic system. In the experiments reported in this thesis, I utilise a genetic al-

gorithm to train artificial neural networks mounted in real or simulated miniature robots to solve tasks defined by a fitness function. By using artificial genomes, populations of simulated robots learn how to solve efficiently the given task, while interacting with the surrounding environment in real time. The process of evolution unfolds in few steps.

1. **Create a population.** An initial set of N vectors that represent the artificial genome are generated with random values.
2. **Phenotypic expression.** The genetic traits are mapped onto a matrix W , defining the strength of the synaptic weights.
3. **Fitness evaluation.** Each robot acts in the same environment for t time steps and the performance is evaluated according to a fitness function.
4. **Selection and reproduction.** Robots are ranked based to the fitness score and a subset is designated for reproduction according to a selection rule.
5. **Evolution.** Offspring encoding the parent's genotype with random mutations populate the next generation. Steps 2-4 are iterated for a fixed amount of generations.

Throughout generations, artificial evolution finds the correct values for synaptic weights in the parameter space. If evolution is successful, at the end of the last generation robots learn to act in a robust and consistent manner in the environment, solving complex tasks. The main difference with connectionism is the lack of predefined training sets that define sequences of input-output patterns and the presence of a physical or simulated agent. In fact, the training by artificial evolution is a semi-supervised method, which does not necessarily state the desired answer given an input vector. Instead, the fitness function frames a general problem, e.g. move towards a light source (Tuci et al., 2004), and populations of robots autonomously develop correct input-output patterns through an implicit process of trial and error. The word *implicit* underlines the difference with another semi-supervised learning algorithm, reinforcement learning (Sutton & Barto, 1998), where the agent is explicitly rewarded with positive or negative feedbacks.

Another direction of research follows the approach of morphological computation (Pfeifer & Bongard, 2006), where the artificial genome stores the phylogenetic expression of an agent's body.

The vector storing the genetic traits defines the set of sensors and actuators included in the agent's perceptual and motor system, their position in the body and the size of different body parts. The two methods are not mutually exclusive, and significant advancements in adaptive systems are achieved by co-evolving body morphology and controller (Bongard, 2011).

2.1.3 A new perspective on intelligence: Embodiment and situatedness

In the following years, the symbolic approach to cognition and artificial intelligence has been widely criticised (Brooks, 1990; Thelen & Smith, 1996; Beer, 1995; Francisco Varela & Rosch, 1991; van Gelder, 1998). One of the strongest and more robust arguments against computationalism is a thought experiment that underlines the absence of a fully developed semantic in a cognitive theory which is solely based on an abstract symbol system. The Chinese room argument (Searle, 1980) creates a fictional scenario where an English speaker is locked inside a room with a set of rules for responding in a meaningful way to messages written in Chinese. Although the outcome is surely consistent with the Chinese language, the imaginary human subject does not have a true understanding of the ideograms. The central point of this argument is that a seemingly intelligent syntactical manipulation of symbols, an appropriate output given a certain input, is not a sufficient condition for intelligent behaviours. According to Searle's position, an intelligent agent must also understand the meaning of symbols, that is, the link between symbols and the real world (Harnard, 1990; Cangelosi et al., 2002).

Criticisms towards the physical symbol system hypothesis and computationalism come also from applied fields when roboticists realised that controllers based on formal symbolic logic are extremely fragile and inefficient (Minsky, 1988; Moravec, 1988; Brooks, 1990). During the sixties and seventies there have been several attempts to building intelligent robots with cognitive architectures based on abstract symbolic systems to represent the world, and all of them had limited capacity to interact with the environment. In the middle eighties, Rodney Brooks took a completely different approach developing behaviour-based robotics (Brooks, 1999). The central idea in Brooks' work is that the "world is its own best model" (Brooks, 1999, p. 115) and using abstract symbols to represent the environment is deemed to fail. As underlined by the author, evolution required a considerable amount of time to displaying basic locomotion skills and low level forms

of intelligent behaviour in the simplest organisms. Conversely, high level cognitive functions are a relatively novel fact in the history of life. Thus, human cognitive skills may not be the pinnacle of the evolutionary process. This argument is surely debatable and not immune from criticisms, for example following the biological theory of punctuated equilibrium that postulates sudden significant changes during the evolutionary process, rather than smooth and progressive modifications (Gould & Eldredge, 1977). However, the argument establishes a radical reinterpretation of the concept of intelligence, challenging an anthropocentric view of cognition. To overcome the inherent limits of representational architectures, Brooks developed the subsumption architecture: the problem that the robot faces is divided into subtasks defined hierarchically. In the global architecture, low level layers correspond to primitive behaviours such as obstacle avoidance, whereas high level layers embody more abstract goals, for example exploring the world. The development of the whole cognitive architecture is built incrementally, following a bottom-up approach where lower layers are implemented and debugged before the more abstract cognitive modules are added to the cognitive architecture.

The core idea behind behaviour-based robotics is elegantly captured by Moravec's paradox (Moravec, 1988), the discovery in artificial intelligence and robotics that it is relatively easy for a machine to successfully accomplish high-level tasks, e.g. playing chess, rather than solve real time problems with a heavy load of sensorimotor activity and perceptual features. This view strides with the more traditional opinion, based on an anthropocentric preconception, that higher level cognitive functions are the apex of intelligent behaviour.

2.1.4 Robots: scientific tools for studying embodied cognition

Originally, the term robot appeared in the title of the play Rossum's Universal Robots of the Czech writer Karel Čapek. The literal meaning of the word robot was "serf labour". Today, this connotation partially captures the wider spectrum of meanings that it conveys. In fact, a machine conceived as automated workforce that substitutes humans in dangerous or repetitive tasks holds true only for industrial robots. However, robots have also a scientific utility as they provide synthetic model for studying human cognitive facts (Parisi, 2014). A physical or simulated robotic platform becomes a valuable tool aiding the development and testing of scientific theories and hypothesis.

Compared to experiments conducted with human or animal models, artificial systems provide a controlled experimental scenario.

In the experimental part of this thesis, I operationalize the concept of embodied and situated cognition with evolutionary robotics experiments, where synaptic weights of the neural controller are adjusted exploiting the process of synthetic evolution. Although simplified, the experimental design comprises all the properties that characterize a simple embodied and situated cognitive agent. To conduct the experiments, I utilise two-wheeled epuck miniature robots (Mondada et al., 2006) equipped with several input sensors. In fact, the robotic platform has a body with sensors and actuators, which allow the agent to perceive and act in the world throughout a continuous flow of sensorimotor loops. Furthermore, the agent is controlled by an artificial neural network, a simple model of a biological brain.

A necessary step to justify the decision of adopting evolutionary robotics for studying cognitive agents involves a discussion about the scientific foundations of this area of research. What is the purpose of evolutionary robotics? Is it aimed at solving practical engineering problems or answering scientific inquiries? In a recent article, Trianni (2014) elucidates and organizes different perspectives on the field, individuating two mainstreams directions of research. Firstly, evolutionary robotics is conceived as a design method, aimed at exploring the search space of a complex robotic problem. By applying artificial evolution, the designer avoids *a priori* assumptions that may lead to sub-optimal solution. Secondly, evolutionary robotics is utilised as a model of biological or cognitive facts. From this perspective, a robotic system is seen as the synthetic equivalent of a living organism, and thus the system becomes a tool for the scientific investigation of natural phenomena. According to Harvey et al. (2005), evolutionary robotics is a useful synthetic tool for studying cognition. By designing minimal models that entail the properties of the cognitive fact under scrutiny, scientists have a ready available testbed aimed at providing evidence to scientific hypothesis in a prejudice-free fashion. Overall, cognitive scientists that accept the validity of the synthetic approach to science, where real or simulated robots are regarded as valid model for rational agents, reject the centrality of human cognition. Therefore, a simple artificial organism suffices at capturing the necessary and sufficient conditions which shape the

basis for the autonomous development of low-level cognitive skills.

A recent model based on an evolutionary robotics experiments demonstrate that offloading the solution of a navigation task in a double T-maze environment on the agent-environment interaction does not prevent the development of cognitive strategies that rely on the internal states of a neural network (Carvalho & Nolfi, 2016). The authors demonstrate robots without hidden artificial neurons evolve a sub-optimal solution. On the contrary, agents controlled by multi-layered networks are capable of generalising and storing sensorimotor patterns, achieving a higher fitness score. To support this conclusion, populations of robots are also evolved in less structured and noisy environments (e.g. perturbing the orientation and position of the agent during half of the trials), showing that in absence of regularities in the agent-environment interaction limit the usefulness of hidden artificial neurons. Matsuda et al. (2014) replicate the vicarious trial-and-error, a seemingly hesitating behaviour typically observed in rat experiments. A widely accepted hypothesis related to such a behavioural pattern is that vicarious trial-and-error enhances learning efficiency compared to animals that exhibit a purely reactive decision-making process. This claim is however difficult to prove in animal models as the internal neural dynamics are inaccessible to the experimenter. The proposed neuro-robotic model aids the investigation of this cognitive phenomena by providing a simplified and controllable synthetic scenario. Results indicates that the presence of vicarious trial-and-error in the evolved behavioural strategy leads to more adaptability to the environment.

Similarly to the approach followed in this thesis, embodied and situated agents has been used as testbed to develop novel methodologies to analyse data collected during experiments. The advantages derived from creating simple cognitive models is evident during the attempt to create analytical frameworks compatible with a theory that stresses the primary role of sensorimotor coupling between the agent and the environment. For example, Beer (1995) applies the dynamical system framework to study the evolved behaviour of legged agent controlled by an artificial neural network. The aim of this work is to show the potential of such analytical framework for studying embodied and situated agents. Recently, an information theoretic approach has been employed to characterise the agent-environment interaction of a minimal cognitive agent invol-

ved in a relational categorisation task (Williams & Beer, 2011). An agent controlled by a recurrent neural network and equipped with a simple visual system moves left or right in a bi-dimensional environment. By exploiting the perceptual input and the internal dynamic of the neural controller the agent categorises several objects of different shape and size. Furthermore, the agent exploits the simple left-right movement to observe the items from different perspective. This model is used as testbed for demonstrating the benefits of an information-theoretic analysis that may be utilised with biological embodied systems.

The following sections discuss the problem of defining the embodied cognition view, demonstrating the benefit of utilising robotic systems for modelling cognitive agents. In particular, I will show how minimal differences on the definition of the field result in maximal divergences during the operationalisation in experimental robotics.

2.1.5 The problem of definition

After the brief historical sketch outlined in the previous sections is now time to provide an overview of the current position of the embodied and situated framework. After 30 years of scientific investigation, the field develops several ramifications with very different conceptualizations. Nowadays, answering the question "what is embodiment?" is not an easy task. A scientific field lacking of precise boundaries and well-founded theoretical boundaries is especially problematic for computational modelers and roboticists. In fact, psychologists and neuroscientists deal with existing organisms, manifestation of phylogenetic and ontogenetic developments. On the contrary, computer scientists have to design and create their own agents starting from an empty software project. Therefore, a well-founded theoretical ground is a necessary condition in order to operationalise, and thus investigate, a scientific theory. The debate around the definition of the umbrella term 'embodied cognition' has been progressing for at least the past decade, yet an agreement is still missing and far to be reached. Although an exhaustive review of the repertory of different conceptualisation of embodied cognition goes beyond the aim of this thesis, a general overview is necessary to motivate and understand the experiments reported in later chapters, as well as the analytical tools applied to data collected from the artificial agents. Several key words have been created to name the novel perspective on cognition, each term bringing its own

meaning of embodied cognition. In the existing literature, cognition is defined as grounded in perception and action (Barsalou, 2008), situated in a social and physical context (Maturana & Varela, 1991; Beer, 1995; Pfeifer & Bongard, 2006; Barsalou, 2008), and the mind is conceived as extended (Clark & Chalmers, 1998) and embodied (Francisco Varela & Rosch, 1991; Lakoff & Johnson, 1999; Pfeifer & Bongard, 2006). All such terms share the same common ground, which is the fundamental role of the bodily interaction with the environment, thus rejecting an amodal explanation of conceptual knowledge, but subtle differences drive the theory to substantially different paths.

Wilson (2002) made a first attempt aimed at organising different conceptualisations of the embodied cognition view. The author identifies six interpretations of the theory:

- Cognition is a situated activity. The cognizant acts within an external context which may be social or physical. The agent is constantly involved in a perceptual-action loop of relevant inputs and outputs.
- Cognition is time pressured. The interaction with the contextual information unfolds under the pressure of real-time sensorimotor loops.
- The cognitive load is off-loaded onto the environment. To overcome limits of our computational abilities, we exploit the environment to reduce the complexity of cognitive tasks.
- The environment is part of the cognitive system. The agent and the environment are strictly coupled and it is worthless studying the mind in isolation; the world is integrated in the cognitive system.
- Cognition is for action. Behavioural outputs are the basis of cognition, and thus cognitive facts must be analysed in terms of appropriate behavioural responses to a situation. Although the mind is not studied in isolation, this view weakens the role of the environment, both in terms of the more general contextual information and the physical world, as the focus is the cogitating agent in place of a strictly coupled agent-environment system.
- Off-line cognition is body based. Even in absence of perceptual features and motor actions, conceptual knowledge is grounded in bodily states.

Following the conceptual clarification made by Wilson, Ziemke (2002) investigated different kinds of body-environment interactions and bodily structures. From a computer scientist perspective, the six definitions enumerated in (Wilson, 2002) do not suffice for building intelligent systems. As noted by Ziemke, only the sixth notion of embodiment refers explicitly to the body, while the other five aim at illustrative different views concerning the connexion between mind and environment.

Recently, the problem of defining the theory has been tackled by framing different concepts of embodiment from a hierarchical point of view, in the attempt to creating a taxonomy of the terms utilised. The theoretical foundations are built upon the distinction between grounded, embodied and situated cognition first delineated by Pezzulo et al. (2011). Fischer (2012) proposes a classification which differentiates between top-down and bottom-up processing during knowledge retrieval. The author identifies the idea of grounded cognition as the most universal, as it refers to invariants of the physical world. The notion of embodiment, instead, is based on the ontogenetic development of a specific individual, subjected to objective physical properties as well as idiosyncratic past sensorimotor experiences, and thus it is more variable. Situatedness is the most dynamic concept, referring to the available contextual knowledge; at the same time, it is also the most flexible towards a constantly changing environment. Therefore, grounded cognition is a better framework to analysing top-down cognitive phenomena, whereas situated cognition perspective is more consistent with a bottom-up approach. Reaching similar conclusions, Myachkov et al. (2014) finds that the notions of grounded, embodied and situated cognition are used interchangeably. By developing a hierarchical definition structured on the idea of stability, that is, how much the ground for conceptual knowledge is constant through time and under different conditions, the authors attempt to set clear boundaries between several conceptualizations of the embodied view. The argument is centred on the notion of simulation, asking which features of the body and the world are exploited during knowledge retrieval. The main problem is the disentanglement of the two notions of grounded and situated cognition. According to the authors, the concept of grounded cognition refers to objective and universal physical properties of the world, and it is the most stable. Furthermore, the term should be substituted with tropism, borrowing the idea from plant biology of an organism self-directing towards a stimulus. On the contrary, situatedness is defined as contextual information in a general sense, including introspective states and

social backgrounds, which leads to an unstable and opaque source of knowledge grounding. The concept of embodiment is the less problematic, referring to the domain of sensorimotor experiences. The degree of stability of embodied cognitive phenomena lies between the objectiveness of the grounded cognition and the unstable dynamicity which characterizes the situated framework.

Overall, there is a general agreement on the notions of grounded and embodied cognition, which clearly refer respectively to the physical world and the peculiar body which characterizes the artificial or living organism. On the other hand, the idea of situated cognition conveys at least two different meanings.

This issue is evident when considering the embodied view from a dynamical system perspective (Smith, 2005; Beer, 2008). By studying the temporal evolution of embodied systems, the concept of situatedness focuses on the real time interaction of an agent with the surrounding environment throughout a flow of perceptual-motor loops. For example, Cangelosi et al. (2015) clearly interpret the idea of situated cognition as a real time interaction of the agent with the environment, where cognitively sound actions modify the environment; these alterations of the world are immediately perceived by the cognizant, which instantly reacts. The static definition of situatedness as an unstable contextual information shares similarities with the ongoing process of changes produced by a dynamic and real time interaction of an agent with the environment. However, the dynamical system perspective underlines the importance of the strict coupling between the agent and the surrounding environment, which leads to significant differences on the concept of situatedness.

2.1.6 Different embodied views in neural modeling: developing the concept of time and language grounding

In this section I describe two cognitive facts that have been modelled with agents controlled by artificial neural networks, showing how different theoretical views of embodied cognition lead to substantially different operationalisations during the experimental design. The first cognitive phenomenon is the development of time perception and conceptualisation in biological neuronal systems (Karmarkar & Buonomano, 2007), which has been replicated and studied in embodied and situated artificial systems (Yamauchi & Beer, 1994; Tuci et al., 2004; Maniadakis et al., 2009;

Flusberg et al., 2010). Recently, a connectionist model has been proposed by Flusberg et al. (2010), which takes inspiration from the grounding of higher level concepts by exploiting the mechanism of conceptual metaphor (Lakoff & Johnson, 1999). In the embodied view developed within the linguistic tradition, metaphors are mechanisms that permit the conceptualisation of knowledge not directly grounded in action and perception. For example, in social organisations the word "power" is understood in analogy with upward spatial direction. The aim of the connectionist model is to investigate the relation between time and space, in the attempt of understanding which concept grounds the other (Boroditsky, 2000). Although the words embodied and situated are used in the paper, in this model the network is not controlling a physical agent, and there is not a real or simulated environment. The inputs of the artificial brain are linguistic labels, such as days of the week or cardinal directions, hardwired by the experimenter taking inspiration from the real world. The neural network includes two sets of inputs, one for items (e.g. "Monday", "Yellow"), the other for relations (e.g. "later", "west of"). Each input layer is connected to a dedicated representational hidden layer, which are in turn connected to a single deeper hidden layer which integrates the information incoming from the two distinct lower layers. The output layer consists of a set of labels which are the logical consequence of items and relations presented to the neural network in the input layers. For example, given the inputs "Monday" in the item input layer and "later" in the relation input layer, a possible correct output may be "Friday". The model shows that temporal relations are grounded in spatial relations. From an embodied and connectionist perspective, "Monday" and "Friday" are consistent with the spatial relation "move forward". The authors clarify that the labels of the input layers are arbitrary abstractions of a hypothetical lower level, modality-specific modules, i.e. true perceptual states of an embodied and situated agent. However, the aim of the model is to investigate neural structures developed at the end of the learning process, rather than real-time conceptualisations achieved through genuine sensorimotor interactions with the environment.

A very different approach is followed Tuci et al. (2004) while modeling the evolution of a sense of time in an embodied and situated mobile robotic platform. The authors use evolutionary robotics to train a two-wheeled cylindrical miniature robot controlled by a continuous time recurrent neural network. By building an architecture with recurrent connections, the agent is provided with a short

time memory, which tracks events occurred in the near past. The free parameters of the neural controller are adjusted with artificial evolution, a learning method discussed in detail in the method section of this thesis. Two light sensors placed on the opposite sides of the body and a ground sensor form the robot's perceptual system. A light bulb is placed in the arena with a white floor, except for a circular shaded grey area surrounding, but not including, the light source. The robot learns to approach the light and navigate on the grey area without touching an inner black line, which represents a wall. There are two possible environments to which the robot may be exposed. In the first environment, the light is completely surrounded by the grey area and the black line, whereas in the second scenario there is a discontinuity. A second objective given to the robots is to categorise the environment by autonomously evolving a signalling strategy. The categorisation is a product of the internal dynamic of the neural networks, which develops a feeling for the flow of time by experiencing the persistence of a stimulus over time. In fact, the robotic platform perceives the colour of the arena with the ground sensor, developing a sense of time by experiencing the presence of the grey area over time. In the environment with a discontinuity in the target area, the robot does not emit any signal in the white space of the arena while approaching the light source. By entering in the shaded portion of the environment, the neural network feels the time spent in the coloured area and is capable of emitting a signal when the external stimulus, recorded by the ground sensor, detects the white portion of the environment.

Clearly, the two approaches are quite different. The connectionist model is designed by the experiment following an *a priori* analysis of situated social contextual features of the human linguistic experience. The features relevant for the study are formalised in a set of static and descriptive inputs which form the input layers of the artificial neural network, e.g. linguistic conventions such as "west of" and "move forward". The same principle guides the selection of sensorimotor and embodied traits selected for the model, as well as physical properties of the environment. On the contrary, the evolutionary robotic experiment does not follow any ontological decomposition of the real world. The model, instead, entails general principles of an embodied and situated system. In fact, the robot is provided with a sensorimotor system, and the agent is situated within the context of a real time dynamical interaction with the environment. The aim of the robotic experiment is to understand a cognitive fact with a minimal model of a living system, which develops

its own cognitive characteristics in a fully autonomous manner.

The second example aimed at showing the differences in the embodied cognition framework is the development of higher level linguistic skills. The grounding of abstract concepts (Arbib, 2008; Borghi & Binkofski, 2014) and linguistic compositionality (van Gelder, 1990) are two cognitive phenomena difficult to explain from the point of view of a theory that underlines the primary role of direct perception and motor activity during the development of conceptual knowledge.

In a neuro-robotic scenario, Stramandinoli et al. (2012) proposes a model aimed at unveiling the mechanism for grounding higher level concepts and words. The input layer consists of linguistic binary units that represent embodied and situated contextual features, as for example "grasp" and "smile". The linguistic layer is fully connected to a set of recurrent hidden units, which are in turn linked to an output layer. The output units are binary and there is a one-to-one mapping with external procedural top-down modules controlling the joint space of the arm of a humanoid robot. This scenario is methodologically similar to the model described in (Flusberg et al., 2010), as the neural network is not directly controlling an agent. In fact, the outputs of the network are labels themselves, and the motor control is off-loaded to a dedicated software. To form higher-order concepts, the network learns sequences of actions. For example, "give" is an ordered combination of the action primitives "grasp", "push" and "release". Although the temporal sequence is unfolded by the recurrent neural network, the problem is not framed from a dynamical system perspective, where the robotic platform interacts in real time with the environment throughout a flow of continuous sensorimotor loops. The experimental design, as well as the subsequent analysis of the hidden artificial neurons, is centred on discrete sequential processing of linguistic labels to unveil sub-symbolic structures for storing and retrieving conceptual knowledge.

Sugita & Tani (2005) study the problem of developing semantic compositionality in an embodied and situated neuro-robotic model from a dynamical system perspective. In the experimental scenario, a two-wheeled robot is equipped with a camera for discriminating three objects of different colour placed in the environment. In addition to the motor outputs that regulate the speed of the motors, a simple robotic arm is mounted on the front side of the body. The robot is controlled by a modular recurrent neural network, consisting of two sub-networks for learning sequences

of perception-action loops and linguistic labels. The two modules are connected with a set of units called parametric biases, which function as modulators for the dynamic (for a detailed explanation see (Tani, 2003)). The robot learns three actions, "point", "hit" and "push", the different colours of the objects, and three spatial positions relative to the robot's subjective experience. During the supervised training process, the robot learns to understand simple sentences, e.g. "point red" or "hit left", producing an appropriate behavioural response. The training set consists of a subset of all the possible linguistic combinations. At the end of the learning phase, the robot is exposed to sentences not included in the training set, showing that the agent is capable of responding with the correct action. By confronting the two operationalisations of the embodied view, the differences are evident. The robotic model described in (Stramandinoli et al., 2012) abstracts away from a real time interaction with the environment, whereas Sugita & Tani (2005) explicitly exploit the continuous sensorimotor flow for grounding linguistic labels.

2.1.7 Closing remarks on the concept of embodied cognition: the problem of situatedness

The transdisciplinary nature of the embodied and situated view of cognition inexorably leads to different conceptualisations, mainly biased by the field of study. For neuroscientists, the brain and neural structures acquire centrality while explaining cognitive or computational functions (Rizzolatti & Arbib, 1998; Gallese & Lakoff, 2005); on the other hand, an experimental psychologist or a linguist shifts the focus on the subjective experience of a cogitating agent, which grounds the processes generated by the mind onto sensorimotor experiences (Lakoff & Johnson, 1999; Barsalou, 2008; Borghi & Binkofski, 2014). Furthermore, by following a dynamical system approach the boundaries dividing inner and outer elements, as well as brain, body and environment, are dissolved into the unity of a systemic view (Beer, 1995; Thelen & Smith, 1996). Although the boundaries of embodied and grounded cognition are somewhat stable, referring to the subjective ontogenetic development, body morphology and physical properties of the environment, a common vision about the idea of situatedness is still a problem. Following the discussion introduced in the previous sections of this chapter, I propose two distinct definitions of the term:

- **Situated grounded cognition:** the agent is situated in an opaque social context which is

complementary to the physical world; the subjective experience of development, the stability of objective physical properties of the environment, and the unstable contextual information structure knowledge and ground concepts.

- **Situated dynamical cognition:** the agent is situated in a physical and social environment; cognition emerges from a real time, continuous and strictly coupled sensorimotor interaction between an unstable subjective experience and a unstable objective world.

The problem of conceiving a widely accepted operationalisation of the concept of situatedness was an issue implicitly discussed in the early years of connectionism. The theoretical limits of an approach based solely on brain modeling is underlined by Parisi et al. (1990), stating that several experimental biases may follow from the absence of a physical or simulated environment. In fact, the numbers of input units that represent a natural or social domain, such as spatial or temporal relations, as well as the frequency and the variety of inputs, are arbitrarily decided by the experimenter. Thus, there still is an ontological decomposition of the phenomena and the model is designed in a top-down fashion. Furthermore, there is no real interaction with the world, as the state of the neural network at time t is not influenced by past sensorimotor iterative loops at time $t - n$. To overcome these limits, they propose an ecological approach to connectionist models, where the artificial neural networks control a body placed in an environment, although the three elements may be extremely simplified. A conceptual clarification was also proposed by Elman (1998a). He distinguishes between connectionism, artificial life and dynamical system approach to cognition. In his view, the three approaches share the common ground of a biologically oriented stance by mimicking brain functions, biology, or the physics of nature. However, they differ for the aim and scope. Connectionism primarily concerns emergent structures, as well as storage and retrieval, of high-level cognitive functions in the brain. Therefore, the main area of investigation is oriented at understanding how sub-symbolic computational and distributed machines handle concepts and knowledge. Artificial life rejects the centrality of high-level functions, focussing also on low-level cognitive facts and adaptive behaviours of early and less developed forms of life. The dynamical system approach studies the emergent properties of adaptive systems, including neuronal ensembles and body-environment interactions, from a mathematical

perspective.

In this thesis, I support the position stated in (Parisi et al., 1990), where the importance of a real time and dynamical interaction of the triadic systemic entity brain-body-environment is rigorously operationalised. Therefore, the conceptualisation of embodied and situated cognition corresponds to the definition as situated dynamical cognition proposed in this section. Connectionist models are surely useful tools for investigating how knowledge is structured in the brain. However, current models are often undermined by an important methodological flaw, that is, the arbitrary and top-down attitude during the design of the model. To overcome this issue, the neural architecture and the training set require support from an accurate analysis of the cognitive fact investigated by the model. For example, while designing a situated grounded cognition model of language grounding, a frequentist analysis of words included in linguistic corpora may guide the construction of the input vector and the training set. Otherwise, the creation of the model follows principles not far apart from the computational tradition, and a real agent may show incapacity of handling the noise and variety of a continuously mutating environment. On the other hand, a situated dynamical cognition proposes models weakened by two issues. Firstly, the construction of simulated or physical experimental environments is extremely simplified, and thus cognitive agents develop solutions for toy-problems. Secondly, dynamical system theory is still a scientific framework at its infancy. In fact, during the analysis of complex interactions between the agents with the surrounding environment either few macroscopic variables of minimal models are manipulated (Harvey et al., 2005), or a holistic but qualitative description is employed (Yeragani et al., 2002). Clearly, the definitions and meanings of embodied cognition reported in this section form a continuum, that is, there are not clear boundaries and the different sub-fields partially overlap. In what follows, I propose an analytical framework aimed at understanding such complex systems. The embodied view, considering all the multifaceted meanings, promotes the importance of an integrative stance towards the relationship between the agent and the environment. By accepting this view, scientists are forced to deal with systems consisting of heterogeneous parts, as well as a mind built upon multimodal information. Thus, the necessity of mathematical tools that capture the intricacy of complex networks of interaction among heterogeneous components is crucial.

2.2 Understanding an embodied and situated system

In this thesis we stress the importance of the real time agent-environment interaction to fully develop the embodied and situated view of cognition. In this respect, the theory is framed within the context of dynamical system theory, that is, brain-body-environment are deeply intertwined and the resulting system is conceived as a unity in a strict sense. Thus, understanding embodied and situated systems requires appropriate analytical tools to avoid a disentanglement of the three elements. In the field of synthetic modeling, the experimenter does not divide the system into sub-components in order to achieve a meaningful behaviour. Instead, the solution of a predefined task emerges thorough the coupled interaction of the robots' artificial brain during the bodily interaction with the world. By avoiding *a priori* decomposition of the problem, the framework is clearly holistic and antireductionist: intelligent behaviour emerges from complex and nonlinear interactions of simple, and non-necessarily homogeneous, parts of the system. For instance, the solution to a problem found by a two-wheeled robot arises, according to the theoretical principles of embodied cognition, from the coupled interaction of controller, body, and environment. But there is also a more microscopic level of description, where groups of artificial neurons, sensors, and objects in the environment are the atomic parts of the system. In this regard, evolutionary robotics does not impose any assumption, as the only factor decided by the experimenter is a global utility function.

The first step concerns the design, creation and training of adaptive systems. After the training process ends, and the robot has learned how to execute a task, the behaviour of the agent undergoes an analysis aimed at understanding the factors that convey to a stable and robust solution. In this respect, a reductionist approach is still a widely used framework. Consequently, the principles that guide the creation of the model are completely discarded. The focus is now the single part of the system, rather than global properties or local interactions among sub-parts. Reductionism is not the only approach used for the analysis of embodied adaptive systems, especially by who is keener on drawing the analogy between cognition and dynamical system theory. For example, Carvalho & Nolfi (2016) investigate the interplay between reactive intelligence and internal states of a miniature mobile robot using evolutionary robotics technique. The authors give a qualitative description of the robots' behaviour from a dynamical system perspective, wit-

hout using specific measures. A similar approach is followed by Montebelli et al. (2008) while describing behaviours evolved by populations of robots in terms of visually different attractors.

In the following sections, I propose an epistemological and methodological framework for the analysis of embodied and situated cognitive agents, arguing that a systemic and antireductionist perspective may guide the development of quantitative tools aimed at understanding synthetic and biological models.

2.2.1 Reductionism and antireductionism

Descartes was the first author that framed the mind-body problem in modern terms. But the French scientist and philosopher developed also a ground-breaking advancement in the field of scientific methodology. He conceived the idea of reductionism, where a complex problem is reduced into smaller and simpler parts. Therefore, a natural phenomenon is subjected to *a priori* decomposition, in order to isolate single elements that are more tractable compared to investigating the system as a unity. A reductionist approach achieved numerous advancements in science throughout the centuries. However, reductionism has intrinsic limits. In fact, decomposing a system into simpler sub-components terminates any attempt aimed at understanding global properties of the whole system. Furthermore, reductionism assumes linear models and the consequent superposition principle: two simultaneous inputs in a system cause the same effect resulting from the sum of the two inputs given individually. In simpler terms, an effect is proportional to the cause.

While a large number of scientific successes are rooted into reductionist ground, global properties as well as complex interactions between variables are completely neglected. However, complexity and nonlinear dynamics are central in the conceptual foundation of embodied and situated cognitive science. A methodological aid may come from general system theory (Von Bertalanffy, 1968) and the related discipline of cybernetics (Ashby, 1956). From a system perspective, there are two ways to approach the study of a natural phenomenon as a whole. Firstly, scientists may model biological facts using nonlinear dynamical systems, or infer characteristic of the underlying dynamic from recorded data embedding the time series in order to reconstruct the attractor in phase space. By estimating dynamical measures, for example level of chaos or fractal dimen-

sion for chaotic systems, global properties are unveiled. Secondly, the focus may move towards the relationships between the components that structure the system. In this respect, analytical tools developed in the field of information theory provide useful methods to map and characterize the information flow and nonlinear relations between variables. Furthermore, information theory is also strictly related to cybernetics, where the chain of sensors and effectors are interpreted as channels of communication.

Despite of subtle differences, the conceptual grounding of the two approaches is analogous, as they conceive a system as a whole, stressing the importance of contextual information coming from the environment. In recent years, these theories inspired system biology (Fang & Casadevall, 2011), an integrative approach to science that combines empirical experiments with mathematical modeling (Ahn et al., 2006b). Originally, system biology was created to address the complexity of living systems. One of the major forces that increase the interest of system thinking is the human genome project. The aim of the large scale international scientific collaboration, completed in 2003, was the mapping of all the genes that form the human genome. A striking result derived from the completion of the project is indeed negative, underlining the lack of explanatory power of a mere structural description of the DNA. The complexity of a fully developed living system is not captured by the genes considered in isolation, where the genome is a simple array of switches that activates certain characteristics of a biological entity. Rather than structural, the human genome requires a functional examination, as genes interact in a complex and nonlinear way, and are constantly affected by the environment (Basso et al., 2005; Margolin et al., 2006). The analogy with embodied and situated cognition, where complex sensorimotor chains build our cognition, is evident, and a reformulation of the methods employed may advance this novel theoretical framework.

A related problem concerns the identification of functional connectivity in biological neuronal ensembles, that is, assessing dependencies between two or more patterns of neural activity (Seth & Edelman, 2004; Vicente et al., 2011; Friston et al., 2013). In other words, the aim is to create graphs for understanding the coupling between different brain areas, without imposing any assumption about effective connections (Friston et al., 2013). Functional brain mapping relies

on nonlinear time series analysis of datasets recorded from fMRI or EEG, which capture the complexity of the ongoing dynamic caused by intricate networks of neural activity.

Recently, antireductionism has also been criticized in the field of clinical medical science (Ahn et al., 2006b,a). Current diagnoses of pathological states, as well as the design of treatments, are rooted into a reductionist framework. Existing medical concepts consider healthy states as normalcy, and risk prevention is the key element to preserve this condition. Furthermore, biological systems are seen as the product of linear combinations of the constituent parts, and abnormal states are justified by a small subset of factors, often by a single cause. Therefore, predictions are made with static mathematical models based on frequentist statistics. A very different point of view guides system science: the functions of a biological system are emergent properties caused by nonlinear interactions of the constituent parts that are understood by analysing the system as a whole. Moreover, contextual information and the evolution of the organism in space and time acquire centrality. Thus, the analysis of biological systems requires either stochastic or chaotic models. Although reductionism is efficient on detecting and solving pathologies where the relation between cause and effect are evident, as for example for an inflammation of the appendicitis, intrinsic limits of the theoretical framework fail at explaining biological phenomena where several factors are intertwined in a complex way. In fact, current medical science struggles to provide a fully developed explanation for diabetes (Ahn et al., 2006a) or psychiatric disorders.

2.2.2 Holistic science: general system theory and cybernetics

During the first half of the 20th century, scholars of two different disciplines challenged the well-founded and widely accepted philosophical framework known as reductionism. In theoretical biology, an antireductionist philosophy of science was under development. Von Bertalanffy (1968) proposed general system theory, a foundational framework aimed at guiding the scientific investigation in natural and human science. The author highlights a major problem affecting models created for the study of natural phenomena. In fact, scientists always assume that systems are close. A model of a neuron decontextualizes the cell from the surrounding environment: other neurons, cells forming the body, the influence of external forces, and so on. Living organisms are instead open systems, which necessitates of oxygen for surviving, food for developing and

other conspecifics to mate and reproduce, if the system reached an advanced phylogenetic development. The observer creates boundaries. While studying a single neuron, the subject of study becomes the single cell. Neuronal ensembles are only an objective external context, unnecessary to understanding the functional principles and the properties of single neurons. The traditional analytical and reductionist approach to science provides a full description of the single parts of the system. By embracing the superposition principle, thus imposing a linear interaction between cause and effect, a complete understanding of the single parts suffices for explaining the behaviour of the whole system. On the contrary, system science starts from the whole, a super-system characterised by global properties that emerge from a complex interaction of several, potentially heterogeneous, sub-components. A central element in Bertalanffy's position is the generality, as the final objective is the study of principles and mechanisms that characterise any kind of system. This is an important point, transforming general system theory from a methodological aid for biology into a fully developed philosophy of science, open to transdisciplinarity.

The second area of research that embraces a systemic and holistic stance is the so-called cybernetics. The term cybernetics was introduced by Wiener (1961), defining a novel field aimed at studying human thinking, as well as control and communication in machines. The etymology is rooted into the ancient Greek word *kubernetes*, which means steersman, or the art of governing. According to Wiener, the proposed approach to science is expected to unveil common principles governing self-regulatory machines and neural systems alike. Ashby (1956) broadened the scope of the discipline by defining cybernetics as the theory of machines. The study of biological or artificial systems is grounded on the idea of information and control, from which concepts of recursiveness, circularity and feedback (restricted to negative in those years) derives. These foundational assumptions shape this new field as the study of behaviour rather than objects. Such perspective is consistent with the definition of embodiment and situatedness followed in this thesis, where the subjective experience is submerged in an environmental context, producing continuous communication between perceptual and motor channels. The theory is also close to the first era of connectionism. The early model of artificial neural networks based on a binary formalisation of the neurons, similar to relay switches, fitted perfectly with a mechanistic view of biological organism. Therefore, cybernetics is a major candidate as the perfect bridge

between different disciplines, sharing principles with general system theory, and thus from foundational biological science, psychology, neuroscience and the formal methods developed within the context of information theory (Shannon & Weaver, 1949).

The two disciplines are often conceived as similar and complementary, and surely for the broad scientific endeavour the differences are marginal. General system theory and cybernetics originated in different contexts, namely basic science and technological application. Furthermore the two fields differ for the nature of the models employed, as the biological tradition is oriented towards open systems and dynamic interaction, whereas cybernetics is focussed on homeostasis and feedback (Drack & Pouvreau, 2015). However both fields have the common interest for system organisation and teleological behaviour, taking a holistic stance when studying biological or artificial machines.

2.3 Nonlinear science

2.3.1 Information theory: nonlinearities in communication channels

Information theory is a branch of applied mathematics that quantifies the information content in communication channels. The principles were first developed by Shannon & Weaver (1949), to investigate characteristics of signal processing such as reliability, storing, efficiency in communication and data compression. Although the term information is central, the notion itself is left undefined. Therefore, the theory develops as a practical quantification of the efficiency of communication in the wider sense, assuming an incomplete knowledge about the future and a past corrupted by noise. The basic measure of the theory is entropy, which quantifies the amount of uncertainty of a random variable. The interpretation is thus consistent with the analogous quantity established in the field of thermodynamics, utilised to estimate the degree of disorder in a system. From the interpretation of entropy in information-theoretic terms, communication is practically conceived as the reduction of uncertainty.

Information theory is not intrinsically holistic as the various measures that extend entropy are aimed at quantifying the efficiency of communication channels. However, multivariate measures derived from entropy captures nonlinear interactions among variables, and importantly, homoge-

neity is not a fundamental assumption. Therefore, these quantities are useful tools for studying variables that represent observables recorded from heterogeneous sub-parts of a super-system. Furthermore, the mathematical tools that information theory provides to the scientific communities are nonlinear. According to the embodied view of cognition, especially framed in the situated and dynamical system perspective, the importance of nonlinear bodily interactions with the surrounding environment, as well as nonlinearities in brain activation, are a pillar. Thus, an analytical framework able of capturing nonlinear dependencies is fundamental.

Finally, information theory is model free. The estimation of measures is based on probabilities estimated from empirical data employing binning or kernel density that are essentially data-driven. Therefore, *a priori* assumptions about the nature and type of the underlying model are unnecessary, which render the mathematical framework consistent with an emergentist perspective on the development of cognitive skills.

2.3.2 Application of information-theoretic measures to embodied and situated robotic systems

Information theoretic metrics have already been applied to neuro-robotic platforms in order to unveil how embodied and situated agents structure the information. Tarapore et al. (2004) simulate an Extended Braitenberg Vehicle which approaches and explores cylindrical objects placed in an arena. The simulated agent is equipped with proximity sensors, camera, and two wheeled differential motor system. Time series are recorded from the activation of the perceptual system. The robot-environment interaction is characterized by calculating entropy, pairwise linear correlations, and pairwise mutual information between sensors. Results show different level of disorder in the system, quantified with entropy, for different behaviors exhibited by the robot, e.g. exploring the arena versus navigating around the cylinder. Furthermore, mutual information captures nonlinearities in the time series, detecting sensory interactions at longer time scales compared to linear correlation. In a related work (Gomez et al., 2005) the same simulated robot is extended with a Kohonen feature map (Kohonen, 1998) for implementing a neural visual system. The environment is populated with cylinders and cubes which are approached by the agent. The authors estimate pairwise mutual information between proximity sensors as well as receptors in the

visual field, showing that the structure of the information exchange between time series depends on the shape and size of the objects that are perceived by the agent. The analysis of embodied systems with information theoretic tools is extended to robots characterized by different body morphologies (Sporns et al., 2006). Lungarella & Sporns (2006) utilize entropy, mutual information, multi-information, complexity, and transfer entropy in order to generate a map of the perceptual and motor system of different robots, while they are interacting with the surrounding environment. The experimental setting is similar for each robot, as agents mount cameras to detect simple objects placed in the environment. The input flow incoming from cameras is processed by a neural network, which functions as a saliency map. The neural architecture is common to all robotics platforms. The study of the information structure is thus focused on differences caused by different embodiments. In fact, the three robots used in the experiments are a simple humanoid, a four legged mobile platform, and a simulated two wheeled vehicle which have a different motor systems and body shape. By applying information theoretic measures, results show that body morphology, as well as the learning process, affect the structures of the information flow in the robots. The efficacy of applying information theoretic measures to agents controlled by artificial neural system is confirmed by a model aimed at simulating a minimal klinotaxis neural circuit of the nematode *c. elegans* (Izquierdo et al., 2015). By calculating mutual information and transfer entropy, the authors describe the complete loop of information flow from environmental stimuli, to sensory and hidden neurons, to motor neurons and actuators. The application of mutual information to embodied and situated systems is extended to a human-robot interaction scenario (Stoelen et al., 2015) in order to develop a set of measures aimed at characterizing real-world applications of adaptive systems. The information exchange between a neuro-robotic platform and human participants during the interaction is quantified by values of mutual information.

2.3.3 The revolution of chaos theory

Chaos is a relatively new discovery, developed within the context of the more general field of nonlinear dynamical systems. The subject originated in the 17th century when Newton invented differential equations to solve the two-body problem, calculating the position of the earth during the orbital trajectory around the sun. In the years following the publication of Newton's laws of

motion, physicists tried to solve the three-body problem with no success. In the late 1800s, Poincare revolutionised the way to studying dynamical systems, giving more importance to qualitative rather than quantitative questions. For example, instead of determining the exact position of every planet in the solar system in a specific time, the new approach asks whether the solar system is stable forever or some planet will escape to infinity. In order to tackle complex problems from the novel perspective, the French mathematician developed a new set of analytical tools, based on a geometric methodology. Poincare is also the first scientist who saw the possibility of chaotic behaviour in a system. Notwithstanding the intuition of the brilliant French mathematician, to have a fully developed chaos theory a key element that was missing: the modern development of computers during the second half of the last century. The new technology changed deeply the way of studying dynamics, allowing mathematicians to have completely different insights on the subject. Indeed, the possibility of running intensive numerical simulations led Lorenz to discover his famous chaotic dynamics. In the early 1960s, he was studying a model of convection rolls in the atmosphere using a system of three differential equations (Lorenz, 1963):

$$\begin{aligned}\dot{x} &= \sigma(y - x) \\ \dot{y} &= -xz + \rho x - y \\ \dot{z} &= xy - \beta z\end{aligned}\tag{2.1}$$

For certain values of the parameters σ , ρ and β , the system exhibits an erratic behaviour, following repetitive trajectories in a bounded region of the phase space, but never entering in the same exact orbit. Another important characteristic of the system is its unpredictable behaviour. If the initial values of any of the three variables differs slightly, after several time steps the evolution takes a very different path. There is an important implication surrounding the discovery of chaos in dynamical systems. A natural process that appears to be random might not always be the product of noise; instead, the laws governing a phenomenon might be fully deterministic. In other words, a deterministic nonlinear dynamical system, with certain parameters, may produce a seemingly random and unpredictable behaviour. Thus, the meaning of the word *chaos* in dynamical system theory does not convey the idea of random disorder. Although the temporal evolution is erratic, a chaotic system is structured and fully deterministic.

2.3.4 Informal definition of chaos and the problem of empirical determinism

The evolution of a chaotic dynamical system is seemingly random, but it is governed by deterministic laws. Such complex dynamics are good models for some natural phenomena, including physiological data. A rigorous and widely accepted definition of chaos is still subject of vibrant debate (Poon et al., 2010). Typically, low-dimensional chaos refers to an erratic behavior of an autonomous nonlinear dynamical system, that is, a system without noisy or deterministic inputs. In the present study we follow the informal definition proposed by Kaplan & Glass (2012), widely used in literature: chaos is an aperiodic, bounded, and deterministic dynamic with sensitive dependence to initial conditions. Thus, there are four important characteristics in a dynamical system that exhibits chaotic dynamics:

- Aperiodicity: the trajectory of the dynamic never repeats itself, meaning that the system does not enter in the same state more than once.
- Bounded: the evolution of the dynamic is confined in a region of the phase space and the states will never approach infinity.
- Determinism: there are no random terms in the equations governing the system, which implies that every state is described by definite rules.
- Sensitivity to initial conditions: two starting points arbitrarily close will diverge exponentially fast over time. This is a crucial feature of chaos because it implies the impossibility of long term predictions.

Indubitably, sensitivity to initial condition is the only signature for chaos that has reached consensus in literature. Recently, Poon et al. (2010) suggest a less stringent definition of chaos as a low-dimensional nonlinear dynamic with long-term unpredictability. Therefore, non-autonomous systems with deterministic or stochastic inputs can generate chaotic dynamics. In this thesis, we estimate two measures that capture different properties of chaotic attractors. The first measure is the fractal dimension of the attractor, which relates to the complexity of the dynamics, quantifying the number of degrees of freedom of the system. The estimated dimension is a static metric and

it does not inform about the evolution of the dynamic. There exist several algorithms to calculate the degrees of freedom of a chaotic attractor, and in the present work I employ the correlation dimension (D_2) (Grassberger & Procaccia, 1983). The second measure applied to the analysis of data recorded from embodied and situated agents quantifies the level of chaos in the system, and it is a dynamic measure. Lyapunov exponents relates to the local exponential divergence of nearby trajectories on the attractor. A dynamical system has a Lyapunov exponent for each dimension of the phase space and in a chaotic system at least one is positive. The largest Lyapunov exponent (λ_1) quantifies the level of chaos in the system as it estimates the exponential growth of small differences on the initial conditions. A large λ_1 means that the system is extremely sensitive to small perturbations on the initial conditions, limiting the prediction horizon. The explanatory power of nonlinear dynamical systems in a chaotic regime is surely interesting for modeling the complexity of nature with precise rules defined by deterministic equations. But there is also an opposite approach, using empirical data to unveil the characteristics of the underlying dynamic. Starting from measurements, for example the position of the robot in the environment, I use tools of nonlinear time series analysis to reconstruct the attractor in phase space in order to quantify the level of chaos and the degree of complexity of the system. However, detection of chaos in recorded time series is problematic due to instrumental or dynamic noise, which may produce seemingly chaotic behavior without satisfying the criteria of determinism. Existing methods, i.e. numerical titration (Barahona & Poon, 1996; Poon & Barahona, 2001), are not reliable at discriminating between noisy and fully deterministic time series, especially in the presence of brown noise (Freitas et al., 2009). Therefore, the utility of chaos theory is a matter of debate, the major critic being the lack of applicability to empirical datasets.

2.3.5 Application of chaotic measures to data recorded from humans and robots

After the initial enthusiasm triggered by pioneering works on chaotic measures applied to empirical data, in particular in the field of nonlinear EEG analysis (Babloyantz & Destexhe, 1986; Skarda & Freeman, 1987), the validity of these results was critically re-examined (Theiler, 1994; Theiler & Rapp, 1996). Although controversial (Glass, 2009), chaos theory is still successfully used with experimental physiological data. Nonlinear time series analysis and chaotic measu-

res have been applied to time series recorded with EEG or MEG and they have been reliable estimators to detect epileptic seizures and schizophrenia (Faure & Korn, 2001; Sarbadhikari & Chakrabarty, 2001; Korn & Faure, 2003; Le Van Quyen et al., 2003; Stam, 2005). Bob et al. (2009) show that λ_1 calculated from electrodermal activity differs in schizophrenic and control subjects during a word association task. λ_1 applied to electrocardiogram data (Perc, 2005a) discriminates patients with major depressive disorder and normal control subjects in supine and standing posture (Yeragani et al., 2002). The level of chaos calculated from the heart rate is a reliable estimator to detect the anaerobic threshold in a not invasive way (Silva et al., 2012). A recent study (Mangin et al., 2011) shows that λ_1 calculated from time series recorded from ventilatory flow is lower in patients with carotid atherosclerosis. λ_1 has also been applied to the analysis of dynamic gait (Dingwell & Cusumano, 2000; Dingwell et al., 2001; Perc, 2005b; Dingwell et al., 2008; Bruijn et al., 2009; Kang & Dingwell, 2009; Cignetti et al., 2012; Rispens et al., 2014), showing that different walking speeds (Bruijn et al., 2009) and age (Cignetti et al., 2012) are characterized by different levels of chaos. Fractal dimension of reconstructed attractors from EEG data quantify the degree of complexity during different sleeping stages in young adults (Shen et al., 2003) and newborns (Janjarasjitt et al., 2008). The fractal dimension D_2 is also a valuable tool for detecting color vision deficiencies from visual evoked potentials (Boon et al., 2008).

In the field of autonomous mobile robotics, chaos theory has been successfully used for describing quantitatively agents' behaviour (for a recent review see (Zang et al., 2016)). Smithers is the first roboticist that applies nonlinear time series analysis to datasets produced by embodied and situated robotic systems (Smithers, 1995). By recording the values of the robots' infrared sensors while performing a wall-following task, fractal dimension is calculated on the reconstructed attractor. In another early attempt which exploits chaotic measures for explaining robots' behavior, the position of the agent in the environment over time is recorded in order to estimate the level of chaos in the system (Nehmzow & Walker, 2003; Nehmzow, 2012). These preliminary works are aimed at demonstrating the applicability of chaos theory to mobile robotics. Monirul Islam & Muraase (2005) attain a significant advancement in the field of nonlinear time series analysis applied to robotics correlating systematically chaotic measures to different kinds of behaviour. They use evolutionary robotics to train Kephra robots controlled by neural networks at solving navigation

tasks in different environments. At the end of the evolutionary process, data is collected from infrared sensors, which is used to calculate both level of chaos and complexity of the system. Interestingly, fractal dimension is significantly different depending on the environment where the robot is acting. Furthermore, the degrees of freedom of the system increase throughout generations. Clearly, dissimilar environments, as well as different stages of the learning process, affect the robots' behaviour. Therefore, there is a correlation between fractal dimension and the behavioural strategy employed by the agents. The correlation between behavioural strategies and chaotic dynamics is confirmed by an evolutionary robotics experiment, where simulated e-puck robots learn to escape from a target area at regular time intervals (Da Rold, 2017). Results show a correlation between the robots' performance and chaotic measures. In a robotic experiment aimed at replicating vicarious trial-and-error, a behavioural response used by rats which suggest self-conflict during a decision process, Matsuda et al. (2014) estimated the largest Lyapunov exponent from neural activation and perceptual information. They show a connection between chaotic states and the manifestation of vicarious trial-and-error in the robotic model.

Chapter 3

Methods

3.1 Adaptive systems

This section describes in detail using mathematical language the concept of neural networks employed in this thesis: feedforward multilayer perceptron, Elman recurrent networks (Elman, 1990), and recurrent networks. Different learning algorithms are discussed, with a particular emphasis on the genetic algorithm.

3.1.1 Artificial neural networks

This section provides a rigorous mathematical description of the neural networks utilised in this dissertation. The history and informal explanation of the computational model have been discussed in Section 2.1.2. An artificial neural network is a biologically plausible machine learning algorithm, which operationalises in mathematical terms the functional properties of a neuronal ensemble. The constituent basic unit is the artificial neuron, which is characterised by a nonlinear




Figure 3.1: **Neural architectures.** (a) Feedforward neural network (b), simple recurrent neural network, and (c) recurrent neural network.

activation function:

$$\Phi(x) = \frac{1}{1 + e^{-\beta x}} \quad (3.1)$$

The above equation defines the logistic squash function, where the parameter β define the slope of the sigmoidal function, which is set to 1 in all the experiments described in the following chapters. Any nonlinear function may be employed and an alternative often chosen in neural modelling is the hyperbolic tangent. The artificial neurons are grouped into layers and represented as vectors. The basic model of artificial neural network is the multilayer perceptron (Figure 3.1(a)), also called feed forward architecture (FF). The model consists of a vector of input units $X = \{x_1, \dots, x_n\}$ connected with directed forward synapses to a vector of $Y = \{y_1, \dots, y_n\}$ of hidden units. The multilayer perceptron is equipped with a vector $O = \{o_1, \dots, o_n\}$ of output units. A neural architecture may consist of several forwardly connected hidden layers, or designed with branches to create different neural pathways. The set of synapses connecting two layers of artificial neurons is defined as follow:

$$W = \begin{bmatrix} w_{1,1} & w_{1,2} & \cdots & w_{1,n} \\ w_{2,1} & w_{2,2} & \cdots & w_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ w_{m,1} & w_{m,2} & \cdots & w_{m,n} \end{bmatrix} \quad (3.2)$$

where $w_{m,n}$ is the strength of the synaptic weight connecting the unit m to the unit n . The threshold that characterised early development of neural modelling is substituted with a bias unit, which has a constant activation of -1. The bias units are connected to hidden and output neurons with a matrix B of synaptic weights that are modified during the learning process. The input signal normalised in the range $[0, 1]$ and clamped into the input layer. The information is propagated to the higher level layers, and the activation of a single hidden or output neuron is updated with the following formula:

$$y_i = \Phi \left(\sum_{j=0}^N w_{ji}x_j - b_i \right) \quad (3.3)$$

where y_i is the output of the i^{th} unit, x_j is the input from the j^{th} unit of the previous layer, w_{ji} is the strength of the synaptic weight connecting the two units, b_i is the bias and Φ is the logistic activation function. Similarly, the activation of an output unit o_i is updated according to the signal incoming from the connected units y_j

The FF model is a static neural network that processes every input individually and serially by propagating the signal from the input towards the output layer. At the end of the computation the networks produces an answer and the neural activation is set to 0 before presenting the next input vector. To generalise over events that unfold in time, artificial neural networks with recurrent connections have been developed. Elman (1990) proposed the so-called simple recurrent neural network (SRN), an architecture equipped with an additional vector of context units that store the activation of the hidden units at time t (Figure 3.1(b)). During the next time step $t + 1$, the context layer functions as an additional set of input units connected to the hidden layer with a matrix C of synaptic weights. This simple form of short-term working memory stores information about several time steps due to recursion. The SRN updates at every time step according to the following formula:

$$y_i = \Phi \left(\sum_{j=0}^N w_{ji}x_j - b_i + w_{ki}c_k \right) \quad (3.4)$$

where the additional term c_k is the context unit that saves the activation of hidden unit y_i during the previous time step, and w_{ki} is the synaptic weight connecting the context unit with the associated hidden unit.

The two architectures are similar, except for the additional layer of context units in the SRN, which allows the neural network to have memory about past events. The presence of an internal dynamic transforms a FF from a statistical classifier into a dynamical recognizer (Pollack, 1991), that is, a dynamical system. At every time step, the output of the hidden layer is stored without any modification in the context units that function as an additional set of inputs during the next time step. Therefore, the neural network is informed about its past internal states. Compared with a FF, which can classify only statistical regularity, a SRN is capable to generalize upon events

that unfold in time. Another possibility aimed at providing a form of memory about past events to the neural network is the design of architectures with recurrent connections:

$$y_i^t = \Phi \left(\sum_{j=0}^N w_{ji} x_j^t + \sum_{k=0}^N c_{ki} y_k^{t-1} - b_i \right) \quad (3.5)$$

where y_k^{t-1} is the activation of the hidden unit y_k during the previous time step and c_{ki} is the recurrent connection between the k^{th} and i^{th} units of the same layer (Figure 3.1(c)). The elements of the matrices W, B and C define the space of the free parameters of the system.

3.2 Nonlinear time series analysis

In this section, I describe the nonlinear mathematical tools applied to datasets recorded from the robotic platforms. In particular, concepts and formalisations concerning the attractor reconstruction with the embedded delay vector technique, chaotic analysis and the information-theoretic measures employed in later chapters are explained.

3.2.1 Phase space reconstruction

According to the delay embedding theorem (Takens, 1981) the attractor described by a dynamical system can be reconstructed from a sequence of observables. The mono-dimensional time series is projected in a higher dimensional phase space using a delay vector:

$$\mathbf{x}_\tau^m = (x_i, x_{i+\tau}, x_{i+2\tau}, \dots, x_{i+(m-1)\tau}) \quad (3.6)$$

where m is the embedding dimension, τ is the embedding lag, and x_i is the i^{th} element of the vector containing the position of the robot. The idea behind the phase space reconstruction is to apply time shifts on the mono-dimensional time series to create a multidimensional geometrical object.

A correct reconstruction of the chaotic attractor depends on two parameters. Firstly, an appropriate embedding lag τ is estimated. While the robot acts in the world, points that are successive in the time series may not denote a significant change in the behaviour of the system. An appropriate lag corresponds to a temporal separation where new and relevant information is brought




Figure 3.2: **The Lorenz system.** (a) shows the chaotic evolution of the system defined by the equations. (b) depicts the same attractor reconstructed using the variable x .

into the system. A suitable mathematical tool is the mutual information, an information theoretic measure used to estimate the exchange of information between two time series:

$$MI(X; Y) = \sum_{x \in X} \sum_{y \in Y} p(x, y) \log \left(\frac{p(x, y)}{p(x)p(y)} \right) \quad (3.7)$$

A more detail explanation of the mutual information, as well as methods for estimating the probabilities are discussed in detail in the next section.

Clearly, having only one time series, a second dataset is generated applying temporal shifts. Two vectors $X_a = \{x_1, \dots, x_{n-\tau}\}$ and $X_b = \{x_{1+\tau}, \dots, x_n\}$ are created from the original vector X that contains the set of observable. Mutual information is calculated for several values of τ with unitary increments, creating a plot of τ versus $MI(X_a; X_b)$. The first minimum of the slope is the appropriate value of τ . An alternative method is based on the autocorrelation function, but it has limits when applied to intrinsically nonlinear systems.

The second parameter is the embedding dimension m . As demonstrated by Takens (1981), a sufficient large embedding dimension fully unfolds the attractor, preserving the same dynamical properties of the phase space described by the original variables of the system. The minimum

embedding dimension is calculated with the false nearest neighbor algorithm (Kennel, M. B., Brown, R., Abarbanel, 1992):

$$R_i = \Phi \left(\frac{|s_i^{m+1} - s_k^{m+1}|}{|s_i^m - s_k^m|} - r \right) \quad (3.8)$$

where s_k^m is the nearest neighbour to point s_i^m , Φ is a step function and r is a tolerance value. According to Kennel, M. B., Brown, R., Abarbanel (1992) $r = 10$ performs well with most datasets. When the attractor is completely unfolded, points that are close in phase space of dimension m are close also in dimension $m+1$. The algorithm is iterated in several phase spaces of increasing dimension until the percentage of false neighbors is acceptable. Figure 3.2 shows an example of phase space reconstruction of the Lorenz system defined with Equation 2.1 with parameters $\sigma = 10$, $\rho = 28$ and $\beta = 8/3$.

3.2.2 Chaos theory

On the reconstructed attractor, nonlinear chaotic measures are applied to uncover the dynamical properties of the system. In particular, I employ two classical measures: level of chaos λ_1 and fractal dimension D_2 . The dimension of the attractor gives an insight about the complexity of the system, unveiling the number of variables that define the dynamic. There are several methods to estimate the fractal dimension of a chaotic attractor. We use the correlation dimension D_2 (Grassberger & Procaccia, 1983), which is a widely tested algorithm, robust with finite time series. D_2 is based on the correlation integral, a method that counts the number of points inside a sphere of radius r , centred in a reference point x_i :

$$C^m(r) = \frac{1}{N^2} \sum_{\substack{i,j=1 \\ i \neq j}}^N \Theta(r - |x_i - x_j|) \quad (3.9)$$

where m is the embedding dimension, N is the total number of points in the attractor and Θ is the Heaviside step function, whose values is 0 for negative arguments and 1 otherwise. To estimating the fractal dimension, D_2 can be approximated from the slope of the straight scaling region of the plot $\log(C^m(r))$ versus $\log(r)$ (see for example Figure 4.3(a) in a following chapter). The slope of the scaling region is calculated for several embedding dimensions. If the system is

deterministic, the estimated dimension D_2 remains constant. Conversely, in random processes D_2 constantly increases when the attractor is embedded in higher dimensions.

The second chaotic measure used in this study is λ_1 , which quantifies one aspect of the definition of chaos, that is, sensitivity to small differences on initial conditions. There exist several methods to estimate λ_1 from recorded data. The algorithm proposed by Rosenstein et al. (1993) has become popular due to its robustness to noise and its efficiency with short time series (for a benchmark see (Cignetti et al., 2012; Rispens et al., 2014)). Moreover, the steps of the algorithm partially overlap with the calculation of the correlation integral, thus lowering the computational load. The estimation of the rate of divergence of two adjacent trajectories in phase space quantifies the level of chaos in the system. After taking a reference point p and the nearest point k that lies on a nearby trajectory of the attractor, the Euclidean distance $d_j = (p, k)$ is calculated. An important constrain is that p and k are not in the same trajectory. To exclude points that are close in time, the mean period is estimated using fast Fourier transform and neighbouring points within this temporal window are discarded. In order to quantify the rate of divergence of two trajectories, the distance between the pair of points $d_j(t)$ is calculated for t time steps. In practice, the evolution of two points that are spatially close in the phase space is followed to verify whether they tend to locally diverge or converge. The process is repeated for every point of the attractor and for several time steps t . Finally, an approximation of the true value of λ_1 is derived from the linear fit of the slope defined by the following formula:

$$y(t) = \frac{1}{\Delta t} \overline{\ln d_j(t)} \quad (3.10)$$

where $\langle \ln d_j(t) \rangle$ is the mean divergence of all pair of points.

The presence of a positive λ_1 and a fractal value of D_2 does not guarantee that the dynamic is generated by a nonlinear, low-dimensional, and deterministic system. Thus, further analyses are necessary to exclude the possibility that underlying dynamic is the product of a linear system with noise. In a chaotic deterministic system, the evolution of the attractor in phase space is regular. Therefore, the evolution for n time steps of a reference point p_t is approximated by the average trajectory of k_t neighbors. A simple method to unveil deterministic structures in a system is the

nonlinear prediction error (Kaplan & Glass, 2012):

$$E = \frac{1}{N} \sum_{n=1}^N (p_{t+n} - \bar{k}_{t+n})^2 \quad (3.11)$$

Eq (3.11) alone is not sufficient to evaluate the quality of the prediction. A convenient way to estimate the magnitude of the error is calculated dividing the prediction error E by the variance σ^2 of the time series. If the value is close to zero, the prediction error is small and the dynamic is deterministic. On the contrary, a value above one indicates that neighbours are a bad model for predicting the future trajectory of p_t , as the system is contaminated with high level of noise.

One of the major drawbacks of nonlinear time series analysis is the lack of mathematical tools that give a definitive proof of the presence of chaos. When dealing with empirical data, the true dynamic is not explicit in form of equations. Therefore, asserting that the time series is characterized by a chaotic structure requires statistical validation: The assumption that the dynamic is the product of a nonlinear deterministic system is tested against the null hypothesis of a linear system with noise. One way of verifying the presence of low-dimensional nonlinear dynamics is the surrogate data testing. The original time series is scrambled in order to disrupt the hypothetical nonlinear component. After calculating the fast Fourier transform, the phase is replaced with white noise in the range 0 to 2π . A surrogate time series is created computing the inverse fast Fourier transform of the original amplitude and the randomized phase. To evaluate the null hypothesis, a measure that is consistent only with nonlinear chaotic systems is estimated on the original time series and the surrogates. If the values of the scrambled data are significantly different the null hypothesis is rejected.

Another necessary constrain for applying nonlinear tools to empirical data is the stationarity of the time series. The behaviour of a dynamical system depends on the values of its parameters: different parameters may lead the same set of functions to a stable fixed point, limit cycles or chaotic attractors. Calculating chaotic measures of a system whose dynamic changes on different time windows is meaningless, because a single value characterizes dynamics of qualitatively different attractors. Most of the existing tools in the field of linear and nonlinear time series analysis requires stationarity. Therefore, several statistical methods have been developed to

asses the stationarity of recorded datasets. However, they have been mainly developed in the context of linear time series analysis and their applicability in a nonlinear system suffers the same limitations of the autocorrelation function for the estimation of the embedded lag. To overcome this limitation, Schreiber (1997) proposes a nonlinear method called cross-prediction error, which is a second order statistics that verifies the stationarity directly on phase space. A chaotic attractor is characterized by regularities in the evolution of the system and cross prediction error verifies whether this is true or not for the system under scrutiny. The points of the attractor are divided into n non-overlapping segments of equal length, and different sections of the attractor are used to predict the position of points on other regions of the phase space. Because of the characteristic self-similarity and regular oscillations of a chaotic attractor, segments that are close in phase space are expected to behave in a similar way. Consequently, points in two adjacent segments are good predictors. Clearly, the opposite is true for segments far apart in the attractor, as they are expected to provide a bad prediction. The expectation is to having a mixture of good and bad predictors with a regular alternation. Formally, given a point $x_t \in X$ future states x_{t+n} are predicted with the following formula:

$$\delta_{XY} = \frac{1}{N_X} \sum_{n=1}^{N_X} (x_{t+n} - \bar{y}_{t+n})^2 \quad (3.12)$$

where N_X is the number of points in segment X , \bar{y}_{t+n} is the center of mass of points in segment Y , and n is the number of forward time steps used for the prediction, which is kept constant to 1. As in (Perc, 2006), points $y \in Y$ that are farther from x than a threshold $\theta = \sigma/4$ are excluded, where σ is the standard deviation of the recorded time series. If the set of proximal points θ_y has less than ten elements, the threshold is increased up to a value equal to the standard deviation of the signal.

3.2.3 Information theory

This section summarizes the information theoretic measures employed for the analysis of the evolved embodied and situated system. The building block of more complex metrics is the Shannon entropy (Shannon & Weaver, 1949), which measures the degree of uncertainty, and thus the amount of disorder, in a flow of information. In a neuro-robotic system, entropy estimates

the average information content at each time step of any time series recorded from the neural activation of the controller, as well as other variables such as the agent's position, while the robot is executing a task. The entropy is calculated according to the following formula:

$$H(X) = - \sum_{x \in X} p(x) \log p(x) \quad (3.13)$$

The entropy of a recorded dataset is a global measure, as it evaluates the average degree of uncertainty. A related measure is the local entropy, which characterizes the dynamics of the information flow of time series at specific temporal instants:

$$h(x) = - \log p(x) \quad (3.14)$$

The local entropy estimate the information content of a single event, i.e. a specific measure, in a recorded signal. A global estimate of the amount of disorder in a recorded time series from a dynamical perspective is derived by averaging the set of local values.

The amount of disorder that characterizes a single variable is not sufficient for mapping the information flow in an embodied and situated agent controlled by a neural network. Therefore, the mutual information (Equation 3.7), a measure related to the entropy that quantifies the degree of dependence of two time series, is estimated on recorded neural activity. This information-theoretic measure unveils nonlinear correlations between variables in the system. The information gain of two variables may be estimated locally, according to the following formula:

$$mi(x; y) = \log \left(\frac{p(x, y)}{p(x)p(y)} \right) \quad (3.15)$$

The outcome of the local form consists of a time series describing the exchange of information thorough time between two channels of communication. Importantly, local mutual information may assume negative values. The interpretation of negative values is that a variable is misinformative, increasing the uncertainty of the other variable (Fano & Hawkins, 1961).

Estimating the amount of information shared between two variables in a system may unveil non-

trivial dependencies in an embodied and situated system. However, mutual information does not reveal the directionality of the information flow. To overcome this issue, Schreiber (2000) developed a specific measure, the so-called transfer entropy, determining which variable sends messages to the other:

$$TE_{X \rightarrow Y} = \sum_{x \in X} \sum_{y \in Y} p(y_{t+n}, \mathbf{y}_t^{\text{dy}}, \mathbf{x}_t^{\text{dx}}) \log \left(\frac{p(y_{t+n} | \mathbf{y}_t^{\text{dy}}, \mathbf{x}_t^{\text{dx}})}{p(y_{t+n} | \mathbf{y}_t^{\text{dy}})} \right) \quad (3.16)$$

where y_{t+n} is a measurement of the random variable Y at time $t+n$, \mathbf{y}_t^{dy} and \mathbf{x}_t^{dx} are embedded vectors. By embedding time series, hidden dynamics are unveiled exploiting one observable and past events are explicitly included in the estimation. In practice, transfer entropy quantifies the amount of information transmitted from the dynamic of X , evaluating also the dynamic of Y , to future states of Y . The time horizon of the prediction is defined by the parameter n , which holds a constant value of 1 in the experiments described in later chapters.

Crucially, transfer entropy is asymmetric, thus the magnitude and direction of the information transfer from a random variable X towards Y is estimated by subtracting the outcome of $T_{X \rightarrow Y}$ and $T_{Y \rightarrow X}$. Furthermore, transfer entropy captures the amount of information transmitted from past states of the destination variable to the source, which is consistent with the interpretation of causality introduced by Wiener (Wiener, 1959; Wibral et al., 2014). The temporal evolution of the information transfer between two variables is calculated by applying the local transfer entropy to single realizations of the stochastic processes (Lizier et al., 2008):

$$te_{x \rightarrow y} = \log \left(\frac{p(y_{t+n} | \mathbf{y}_t^{\text{dy}}, \mathbf{x}_t^{\text{dx}})}{p(y_{t+n} | \mathbf{y}_t^{\text{dy}})} \right) \quad (3.17)$$

The local transfer entropy assess which variables are acting as emitter and receiver during a specific time step, while the robot is executing the task.

These information theoretic measures rely on a correct estimation of the probabilities from recorded data points. Although for entropy and mutual information, where one or two dimensions are involved, a naïve estimator such as fixed binning usually suffices, a correct evaluation of transfer

entropy is more problematic due to high-dimensional spaces.

In fact, the recorded signal is projected onto a multi-dimensional phase space by embedding the time series. Thus, simple estimators are not reliable (Lee et al., 2012; Lizier, 2014) and probability distributions are reconstructed employing more sophisticated techniques such as kernel density estimation. In this work, the rectangular kernel is employed to infer the distribution from the recorded finite dataset, which is the estimator utilised in the original formulation of the transfer entropy (Schreiber, 2000):

$$K(x) = \frac{1}{N_X} \sum_{i=1}^{N_X} |x - x_i| \leq \frac{h}{2} \quad (3.18)$$

This kernel draws a n -dimensional box with sides of fixed size around each point, defined with a constant value of $h = 1$ in this article, counting the number of element x_i of the dataset that are close to the reference point x . The number of dimensions n depends on the number of variables that define the probability distribution. Further analysis are based on the Gaussian kernel, defined as follows for the univariate case:

$$K(x) = \frac{1}{N_X} \sum_{i=1}^{N_X} \frac{1}{\sqrt{2\pi\alpha h^2}} e^{-0.5\left(\frac{x-x_i}{\alpha h}\right)^2} \quad (3.19)$$

Where h is the bandwidth. The choice of a correct bandwidth is essential as it defines the smoothness of the distribution. If h is chosen too small, spurious results are included in the reconstructed probability distribution. On the other hand, a large bandwidth obscures fine structures that may characterize the recorded dataset. The values of the parameter α is manipulated in order to fine tune the smoothing of the distribution. An efficient and widely tested method to assess the correct value of h is the Silverman's rule of thumb (Silverman, 1986):

$$h = \sigma \left(\frac{4}{3n} \right)^{\frac{1}{5}} \quad (3.20)$$

where n is the number of points of the recorded time series and σ is the standard deviation.

The calculation of mutual information and transfer entropy includes joint distributions and em-

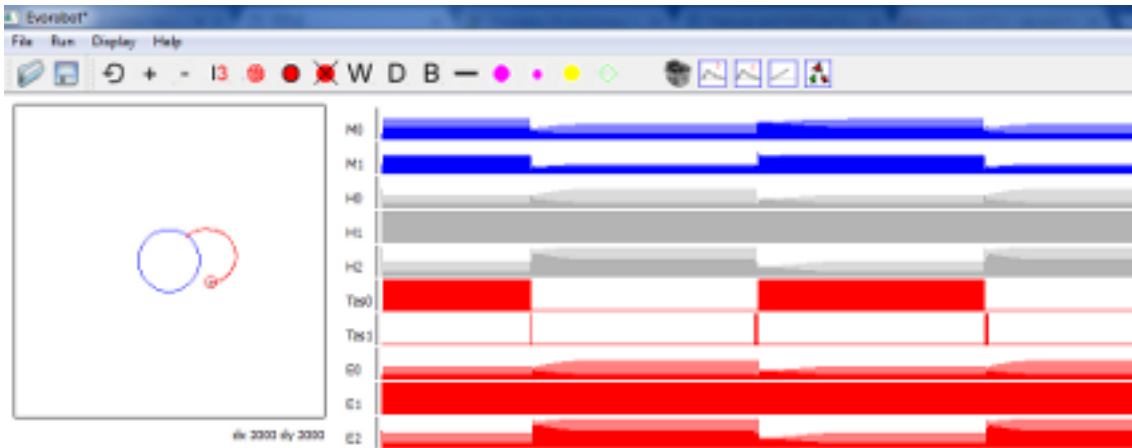


Figure 3.3: **Evorobot simulator.** The picture shows a running experiment, depicting the environment and the robot on the left, the neural activity on the right.

bedded multi-dimensional vectors. The extension of Gaussian kernel to multivariate distributions is defined as follows:

$$K(x) = \frac{1}{N_X} \sum_{i=1}^{N_X} \frac{1}{(2\pi)^{D/2} |\alpha \mathbf{H}|^{1/2}} e^{-0.5(\mathbf{x}-\mathbf{x}_i)^T \alpha \mathbf{H}^{-1} (\mathbf{x}-\mathbf{x}_i)} \quad (3.21)$$

where D is the number of dimensions and \mathbf{H} is the bandwidth matrix. A correct estimation of multivariate probability densities relies on the correct choice of \mathbf{H} , which is estimated with the Silverman's rule of thumb generalized to D dimensions as follows:

$$h = \Sigma \left(\frac{4}{(D+2)n} \right)^{\frac{1}{D+4}} \quad (3.22)$$

where Σ is the covariance matrix of the multivariate dataset, n is the number of recorded points and D is the number of dimensions.

3.3 Evolutionary robotics simulator

In this thesis I utilise Evorobot (Nolfi & Gigliotta, 2010), an open source software tool for running evolutionary robotics experiments with epuck robots (Mondada et al., 2006), which are two-wheeled cylindrical platforms of 3.5 cm radius equipped with various sensors (e.g. infrared sensors) and motor actuators. An example of a running simulation is depicted in Figure 3.3. The simulator is written in C/C++ and includes the following features. Firstly, an evolutionary robotics

algorithm allows the experimenter to define a write fitness functions to evaluate individual robot's performance. The algorithm includes mechanism for selection and reproduction available for the experimenter. Secondly, a neural network simulator permits the user to define the architecture and the activation function that characterises the artificial neurons. The parameters of the neural network and genetic algorithm are defined in a configuration file. Thirdly, the software includes a simulator for the robot's body, including sensors and actuators, as well as the environment and a set of objects (e.g. target areas, walls). To conduct the experiment presented in this thesis, I designed and implemented the fitness functions. Furthermore, the software has been modified to collect data necessary to perform chaotic and information-theoretic analyses.

3.4 Experiment: Chaotic dynamics in simulated robots

In what follows, I describe the experimental scenario and the mathematical tools utilized for nonlinear time series analysis. The evolutionary robotics experiment is conducted using Evorobot (Nolfi & Gigliotta, 2010). The software employed during the estimation of chaotic measures has been developed by the author using Matlab, and tested replicating results reported in (Rosenstein et al., 1993).

3.4.1 Evolutionary robotics experiment

The robotic platform is a miniature vehicle equipped with several input and proprioceptive sensors. For the purpose of the experiment, I limit the perceptual information to a ground sensor placed on the front-bottom part of the robot. The wheels of the robot are controlled by two distinct step motors, and thus the platform is capable of navigating the environment either following a straight line, or performing turns of various angle. The simulations are replicated with robots controlled by different artificial neural networks, either a FF, or a SRN (Elman, 1990). The signal incoming from the ground sensor is preprocessed in order to discriminate two colors, green and blue, which are the input features of two binary units forming the input layer of the neural controller. The output layer consists of two units, each defining the angular velocity of one the step motor. The environment is a squared arena of 200 cm side surrounded by walls, containing a target area with a radius of 20 cm placed at the center (Figure 3.4). The target area has a simple intrinsic temporal dynamic, changing color at regular intervals between green and blue. The only

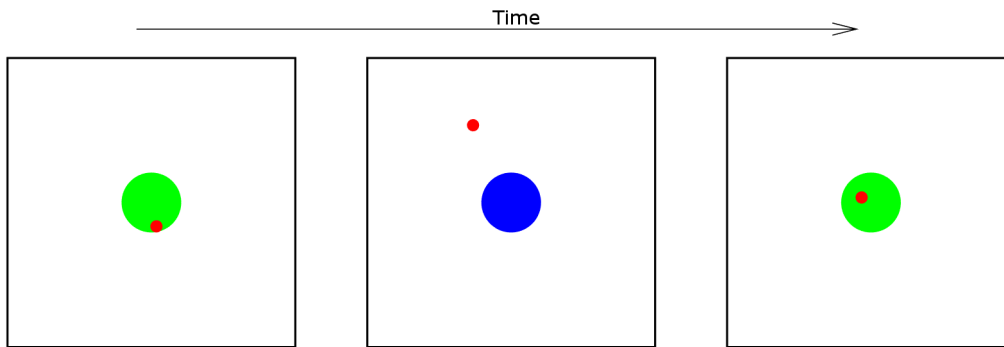


Figure 3.4: **The simulated environment.** The squared arena contains a target area placed at the center. The target area changes color between green and blue at regular intervals. The robot, depicted as a red circle, achieves positive reward remaining on the green target area. When the target area becomes blue, the reward is negative and the robot is forced to escape from danger, entering in the empty space of the arena.

environmental variable manipulated in the experiment is the timescale that defines the number of time steps elapsing before a change of color in the target area. In order to create variations on the evolutionary constrains, three main experimental conditions are devise by manipulating the amount of timescales that a population of robot experience during the learning process. In fact, the evolutionary niche is characterized by either 1(E1), 21(E2), or 61(E3) timescales. Therefore, different populations of robots experience a different number of temporal dynamics during the learning process. Each experimental condition is replicated 5 times, using timescales of different length (Table 3.1), for a total of 15 simulated environments.

The robot gets a positive reward navigating inside the target area when its color is green, whilst negative score is gained if the agent stays on the blue colored target area. The portion of the environment outside the target area is a safe zone, where the robot has neither positive nor negative reward. Formally, the fitness function is defined as follows:

$$\text{Fitness} = \begin{cases} +1 & \text{green} \\ -1 & \text{blue} \\ 0 & \text{otherwise} \end{cases} \quad (3.23)$$

The robot lifespan during a trial is limited to 10,000 time steps, where a time step corresponds to

3.4. EXPERIMENT: CHAOTIC DYNAMICS IN SIMULATED ROBOTS

	Timescales				
E1	{100}	{200}	{300}	{400}	{500}
E2	{90, ..., 110}	{190, ..., 210}	{290, ..., 310}	{390, ..., 410}	{490, ..., 510}
E3	{70, ..., 130}	{170, ..., 230}	{270, ..., 330}	{370, ..., 430}	{470, ..., 530}

Table 3.1: Timescales used in the three experimental conditions. The values refer to the number of time steps elapsing between changes of color in the target area. In environments characterized by a set of multiple timescales, one value is randomly chosen when the robot is placed in the environment. Each experimental condition is replicated 5 times with variations of the temporal dynamic, modifying the length of the timescales.

200 ms. If the robot collides against a wall the trial terminates. At the beginning of the trial, the robot is placed at the center of the target area, which always starts with green color. For each experimental setup, a population of 100 genotypes is randomly generated with 8 bit encoding, which define the values of the free parameters of the neural network. The values contained in the genotype vector are normalized in the range $[-5,+5]$ for the synaptic weights and $[-1,+1]$ for biases. The size of the vector describing the genotype is 17 for the FF and 26 for the SRN network. For each experiment, the population of robot is evolved for 200 generations, which are enough to reach a steady state in the evolutionary process. At the end of each generation, the 5 best robots generate 20 offspring. In order to smoothly explore the space of the possible solutions, the evolutionary process is driven by mutation, which is a local search operator. During reproduction, there is a 2% probability of flipping a bit in the artificial genome. Although debatable, cross-over is not employed as the operator may cause instabilities in the learning process, generating major transitions on the fitness landscape (Nolfi & Floreano, 2000). In order to avoid regressions in the evolutionary path, elitism is employed, that is, each robot selected for reproduction generates a clone. The number of trials available to the robots during each generation is proportional to the number of time intervals the robot may be exposed throughout the evolutionary process. Thus, robots trained in environments characterized by 21 or 61 timescales have a high probability to experience every time interval during phylogenetic development. Every experiment is repeated 40 times for each type of neural controller, generating new initial genotypes using different random seeds. The parameters are summarized in Table 3.2.

In the input layer, FF and SRN neural controllers have two units that encode the incoming information from the ground sensor. To build a system as simple as possible, continuous input value

Parameter	Value
Population	100
Replications	10
Parents	20
Offspring	5
Encoding	8 bits
Mutation	2%
Elitism	1
Weights	[-5,5]
Biases	[-1,1]
Parameters FF	17
Parameters SRN	26
Generations	200

Table 3.2: Genetic algorithm parameter table for experiment 1

of the single ground sensor are pre-processed to feed two input units with binary activation, where each neuron reacts to one of the two colors of the target area. With this simple perceptual system, robots do not gain any sensory information when they navigate the environment outside the target area. The input signal is propagated to a three neurons hidden layer, which is fully connected to two output neurons that regulates the speed of left and right motors. Both hidden and output neurons are connected to a bias unit, and each unit is updated according to Equation 3.8. The SRN has an additional layer of context units that store the output values of the hidden layer at the previous time step. This additional layer is connected to the hidden units and serves as additional input to the network (Equation 3.4). The artificial neurons are formalised with a logistic function (Equation 3.1).

3.5 Experiment: Information theoretic decomposition of embodied systems

In this section I report the details of the experimental setup and the measures utilized during data analysis. Evolutionary robotics is a semi-supervised learning method that autonomously explores the solution space defined by a global utility function. By mimicking the evolutionary process, populations of robots develop a solution finding the optimal free parameters of the system, adjusting the values of synaptic weights and motor biases. At the end of evolution, data is recorded from the 4 neurons of the neural network during the execution of the task. The recorded time series are subsequently analysed with information-theoretic measures.




Figure 3.5: The experimental scenario. (a) depicts the maze and the robot (black circle), placed at the starting position. (b) shows a schematic representation of the robot's morphology with the displacement of the sensors. The left (IF-L) and right (IF-R) infrared sensors utilized in the experiment are painted in red. (c) the artificial neural network. Numbers refer to the identification numbers (ID- n) used in subsequent figures for referring to specific connections. The output layer includes two motor neurons that separately control the left (M-L) and right (M-R) motors.

3.5.1 Simulated robotic scenario

Populations of simulated e-puck robots are evolved using Evorobot, a scientific software developed for evolutionary robotics experiments (Nolfi & Gigliotta, 2010). The simulated agents are miniature cylindrical robots with two differential wheels and 8 infrared sensors placed around the body (Mondada et al., 2006). In this experiment, only the front middle left and front middle right sensors are activated and connected to the controller (Figure 3.5(b)).

The environment is a simple squared maze with inner walls of 30 cm length and outer walls of 54 cm, shaping 4 corners of 90° . At the beginning of the trial robots are placed after the bottom-right corner, as shown in Figure 3.5(a), with the front part of the agent directed towards the upper-left corner. The aim of the simulation is to mapping the relations between variables forming the system, rather than the robots' ability to learn multiple tasks or develop robustness toward noise. Therefore, the environment where robots evolve is not manipulated in order to avoid a combinatorial explosion of the factors involved during the information-theoretic analysis. The length and displacement of the walls, as well as the starting position, are constant. During the learning stage, the life span of a robot is limited to 2,000 time steps and a collision with the walls

terminates the trial. A population of 100 robots characterised by different genotypes with 8 bits encoding is tested and the performance is measured according to the following fitness function:

$$Fitness = (m_1 + m_2) - 2|m_1 - m_2| - 0.5(i_1 + i_2) \quad (3.24)$$

where m_1 and m_2 are the speed of the left and right motor actuators, i_1 and i_2 the activation of the left and right infrared sensors. The fitness function selects robots that move in the arena at the highest speed. The first two components of the fitness function affect the motor behavior, rewarding a high forward angular velocity while penalizing rotations of the robot's body. The third component punishes the activation of the two infrared sensors, forcing wall avoidance. The selection of the robots that generate the population of the next generation is rank based. The 20 robots which achieve the largest fitness generate 5 offspring with a 2% probability of mutating a single bit of the vector encoding the genotype. Reproduction does not follow a cross-over scheme, where non overlapping parts of two different artificial genomes are mixed. Although the utility of this operator is controversial in the field of evolutionary robotics, mixing different genotypes may lead to instabilities during the evolutionary process (Nolfi & Floreano, 2000). One offspring of each parent is replicated with elitism, that is, the genotype is cloned without any mutation. Therefore, the best solution is preserved to guarantee stability during the evolutionary process, avoiding descends on the fitness landscape. The evolution iterates for 200 generations and the simulation is replicated 10 times with different random seeds to generating the initial set of genotypes. The parameters are summarized in Table 3.3.

The neural network architecture is recurrent (Equation 3.5) and designed avoiding complexity, consisting of two input neurons that encode the activation of the left and right proximity sensors normalised in the range $[0, 1]$. The input layer is fully connected to two output units (Figure 3.5(c)). Each artificial neuron of the output layer regulates the speed of a motor actuator and the two units are connected with lateral connections. The motor neurons are connected to a bias unit and the activation of the incoming signal is updated with following equation:

$$m_t^a = \Phi(m_{t-1}^b + i_1 + i_2 - c_a) \quad (3.25)$$

Parameter	Value
Population	100
Replications	10
Parents	20
Offspring	5
Encoding	8 bits
Mutation	2%
Elitism	1
Weights	[-5,5]
Parameters	8
Generations	200

Table 3.3: Genetic algorithm parameter table for experiment 2

where m_{t-1}^b is the activation of the output b at the previous time step, c_a is the bias unit, i_1 and i_2 are the input neurons that encode the activation of left and right infrared sensor, respectively; output neurons are update with the logistic squash function (Equation 3.1). The number of free parameter encoded in the genotype is 8. The binary values stored in the genome are normalized in the range $[-5,+5]$ for synaptic and bias connections.

The set of information-theoretic measure and the kernel density estimators for reconstructing the probability distributions from the recorded data are described in Section 3.2.3. Equation 3.16 and 3.17 that define the transfer entropy include embedded vectors. The embedded lag τ is estimated with the mutual information or the average value of the local form, as discussed in the result section. The minimum embedding dimension d is estimated using the false nearest neighbour algorithm (Equation 3.8). The threshold of false neighbour is set to 20%, as the underlying dynamics may not be fully deterministic.

Chapter 4

Chaotic dynamics in simulated robots

4.1 Introduction

This experimental chapter aims at investigating the possibility of understanding the global properties of an embodied and situated system, which is one aspect of the antireductionist and holistic perspective. A mathematical theory that easily map onto such perspective on scientific investigation is chaos theory, where a system is described in terms of a set of coupled nonlinear equations. The analogy with an integrative view of science is evident as the subcomponents cannot be separated and studied in isolation. The interest of such models is focused on global properties that characterize the system conceived as a whole. Chaos theory is surely a fascinating mathematical fact, where deterministic systems exhibit erratic behaviour. However, the model does not include stochastic components and the seemingly random evolution of the dynamic is the product of nonlinear and complex interactions among the variables of the system. What makes nonlinear chaotic systems so appealing is the possibility of creating mathematical models that are fully understood, without incorporating stochastic elements to formalise the inexplicable. However, major issues undermine chaos theory if the focus moves from top-down abstract models to the domain of empirical data. In fact, models may be developed in a bottom-up fashion, starting from observables. Following a data-driven approach, chaotic measures are estimated on attractors reconstructed in phase space from recorded time series. However, discriminating a low-dimensional deterministic system from linear and noisy dynamics is still an unsolved problem. Notwithstanding the impossibility of providing a final proof that demonstrates the presence of chaos in recorded data, I test the theory on a large number of robots in order to assess the practical utility and concrete applicability of fractal dimension D_2 and level of chaos λ_2 to real-

world phenomena. The analysis of recorded time series is guided by a precise operationalisation of the definition of chaos presented in Section 2.3.4 with appropriate nonlinear tools. Thus, acknowledging the impossibility of definitive hallmarks for chaos, the risk of misinterpreting the results is reduced to an acceptable minimum. Relevant previously published works are reviewed in Section 2.3.5, demonstrating that chaotic measures are capable of quantifying the dynamic underlying robots' behaviour with the phase space reconstruction technique. Another direction of research in mobile robotics focuses in top-down modelling, where robots' behaviour is guided by chaotic controllers. Results show that two wheeled robots efficiently explore unpredictable environments Nakamura & Sekiguchi (2001); Martins-Filho et al. (2004); Volos et al. (2013), outperforming agents controlled by stochastic controllers.

In what follows, I report a series of evolutionary robotics experiments where simulated e-puck robots develop a behavioural response to a target area that changes colour and properties over time. At regular intervals, a target area placed at the centre of the arena switches between positive and negative reward. The dynamic governing the agent-environment interaction is analysed estimating chaotic measures on the reconstructed attractor. By using a large number of simulated robots, I investigate the applicability of chaos theory to the more general class of embodied and situated agents.

4.2 Results

At the end of the evolutionary process, the robot of each replication that achieved the largest fitness score during the last generation is selected for data collection. Therefore, a total of 1,200 different trajectories are used for chaotic analysis. The best robots of each replication do not show a significant difference in performance, thus the selected agents are tested in environments never experienced during evolution. To evaluate the robots' adaptiveness towards novel situations, the lower and upper bounds of the set of timescales used during training is expanded of ± 50 . For instance, if the robots have been trained in a world characterized by a set of timescales $T_{evo} = \{90, 91, \dots, 109, 110\}$, the performance is evaluated using intervals in the range $T_{test} = \{40, 41, \dots, 159, 160\}$. During evolution the temporal dynamic of the target area is chosen randomly. On the contrary, for collecting data the best robots act for 10,000 time steps for

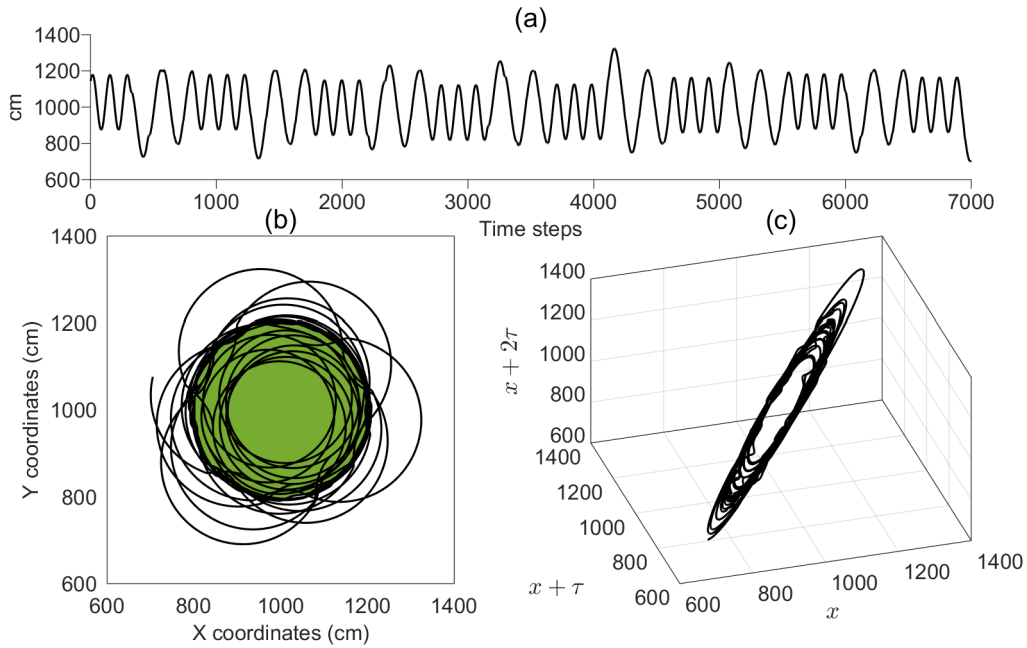


Figure 4.1: Time series embedding. Reconstructed chaotic attractor from the vector storing the position of the best robot generated from seed 24, controlled by a SRN, and evolved in environments where the target area changes color with 61 different timescales (E3) in the range $T_{evo} = \{470, 471, \dots, 529, 530\}$. (a) shows the position of the robot, data used for the reconstruction of the attractor. (b) illustrates the robot's trajectory in Cartesian space and (c) displays the reconstructed attractor in phase space.

each timescale. Robots are exposed to several environments and performing chaotic analysis for every trajectory is computationally expensive. In order to restrict the number of time series, trajectories are recorded only from the environment where robots score the largest fitness. The first 3,000 time steps of the time series are discarded in order to avoid artifacts generated by the initial transient phase of the robot's trajectory.

4.2.1 Chaotic analysis

The first step to take in order to conduct nonlinear time series analysis is the reconstruction of the attractor in phase space (Equation 3.6). From an embodied and situated perspective neural controller, robot, and environment are parts of a system governed by some unknown set of laws. Therefore, any time series conveys information about the entire system, in which all the subcomponents directly or indirectly interact. In this study we arbitrarily analyze the time series recorded from the position of the robot in the environment (Figure 4.1(a)). Specifically, we

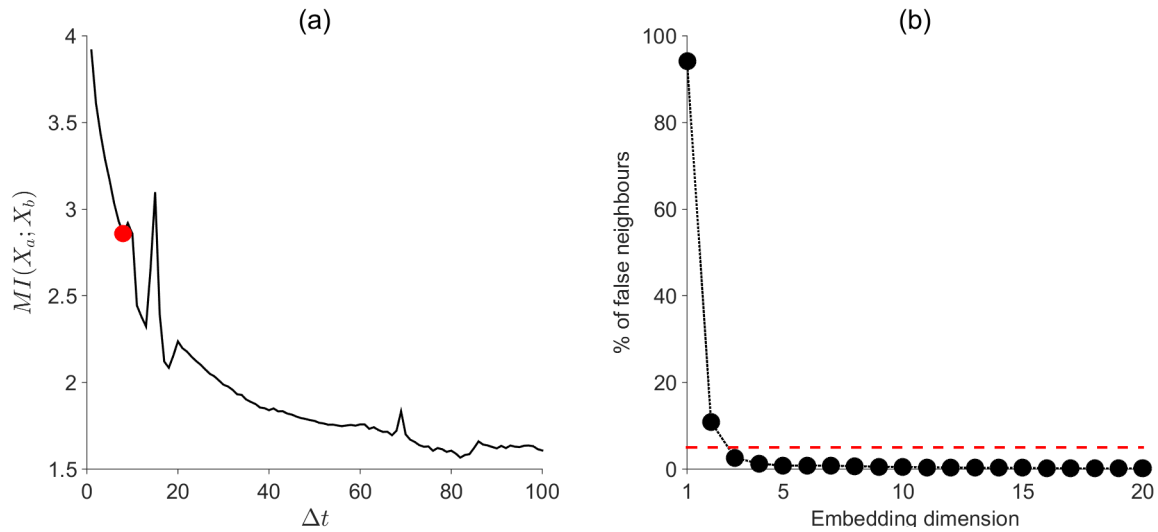


Figure 4.2: Parametrization of the embedded vector. A correct reconstruction of the attractor from a mono-dimensional time series depends upon appropriate values of embedding delay τ and dimension m . Data is collected from the best robot of seed 24, evolved in E3 with timescales $T_{evo} = \{470, 471, \dots, 529, 530\}$, and controlled by a SRN. (a) shows the estimated mutual information between the original time series and itself after applying temporal delays (Δt). The first minimum is reached with $\Delta t = 8$ (red dot), which represents the correct value for τ . (b) displays the percentage of false nearest neighbors for different embedding dimensions. To fully unfold the attractor the minimum m is 3, as it lowers the percentage below the threshold of 5% (red dashed line).

consider the environment as a Cartesian space (Figure 4.1(b)) and the time series is generated recording the position of the robot in the abscissa.

As discussed in the method section, a correct reconstruction of the attractor depends on a correct parametrization of the parameter τ and d that define the embedding lag and dimension, respectively. The embedding lag τ is approximated by the first minimum of the estimated mutual information (Equation 3.7), as discussed in Section 3.2.1. In practice, the information-theoretic measures captures the nonlinear correlations in the system, capturing the average temporal interval elapsing between crucial sub-behaviours. For example, the navigation in a straight corridor exhibited by robotic platform exploring a maze is irrelevant from a dynamical perspective. On the contrary, turning behaviour is relevant, as new information is brought into the evolution of the system. In this experiment, marginal and joint probabilities are calculated naïvely with bins of fixed size. The only parameter, the bin size, is optimized using the technique proposed by Shimazaki & Shinomoto (2007). Following the estimation of the embedding lag, the attractor is reconstructed




Figure 4.3: Attractor dimension. The fractal dimension is estimated with the correlation integral. The algorithm is applied to the best robot of replication 24, controlled by a SRN and evolved in E3 with timescales $T_{evo} = \{470, 471, \dots, 529, 530\}$. (a) shows the logarithmic plot of the correlation integral versus the distance threshold r for different embedding dimensions m . The linear fit (red dashed line) of the scaling region estimates the fractal dimension D_2 . (b) reports the estimated values of D_2 for different embedding dimensions, starting from the minimum m , for the chaotic robot (black dots), and a stochastic system generated with synthetic data (red stars).

in several increasing dimensions to unveil the minimum values of d utilizing the false nearest neighbour algorithm (Equation 3.8). In the experiments the threshold is set to a value of 5% and m is increased to a maximum value of 20. Figure 4.2 shows an example of parametrisation of τ and d based on a specific robot, whereas Figure 4.1(c) depicts the fully unfolded chaotic attractor.

In order to quantify the underlying dynamics, D_2 is estimated on the phase space using the correlation integral (Grassberger & Procaccia, 1983) defined by Equation 3.9. The algorithm counts the number of points surrounding a reference point p_k , iterating the process for the whole trajectory of the attractor. The correlation integral includes the parameter r that defines the length of the boundaries for searching neighbouring points. The range of values of r is calculated taking the smallest distance r_{min} and the largest distance r_{max} between all pairs of points. I use 30 different spheres of radius $r_n = r_{min} + n(r_{max} - r_{min})/30$, with n incrementing from 1 to 30. To estimating the fractal dimension, D_2 can be approximated from the slope of the straight scaling region of the plot $\log(C^m(r))$ versus $\log(r)$, as shown in Figure 4.3(a). The slope of the scaling




Figure 4.4: Level of chaos. The level of chaos is quantified by λ_1 , calculated from the logarithm of the average local divergence of neighboring points over time. The linear fit of the scaling region (red dashed line) estimates the value of λ_1 of the recorded data. (a) shows convergence of the estimation of λ_1 at higher embedding dimensions for the best robot of seed 24, whereas (b) shows that the agent evolved from seed 27 fails to stabilize. Both datasets are recorded from robots controlled by a SRN and evolved in environments with timescales $T_{evo} = \{470, 471, \dots, 529, 530\}$.

region is calculated for embedding dimensions $d \in D$ in the range $D = \{m, \dots, m + 10\}$, where m is the the minimum embedding dimension (Figure 4.3(b)).

The second chaotic measure calculated on the reconstructed attractor is λ_1 , which quantifies the level of chaos in the system. In this thesis I employ the method developed by Rosenstein et al. (1993) formalised with Equation 3.10. The algorithm estimates the local exponential divergence of nearby trajectories. λ_1 is approximated by the linear fit applied on the scaling region of the slope obtained from the plot of $\langle \ln(div) \rangle$, which is the logarithmic of the average divergence, versus the time horizon (Figure 4.4).

As discussed in the method section of this thesis, appropriate nonlinear statistics is applied on the reconstructed attractors to verify the presence of a low-dimensional deterministic system. Thus, nonlinear prediction error is estimated to test the determinism of the underlying dynamics according to Equation 3.11. In particular, the number of neighbors is varied in the range 1 to 300, whereas prediction horizon is left constant with $N = 1$. Furthermore, the stationarity of the embedded time series is assessed with the cross-prediction error (Equation 3.12). Finally, the

presence of a linear system with injected noise is investigated with the surrogate data testing on scrambled time series, that is, synthetic data generated by disrupting hypothetical nonlinearities on the recorded signal. In this experiment, nonlinear prediction error is calculated on the surrogate time series. The process is repeated for 40 scrambled time series and the nonlinear prediction error of the original time series is compared with a Wilcoxon rank sum test.

The first step to take before estimating λ_1 and D_2 is to verify whether the time series is a valid candidate for chaotic analysis. For evaluating the efficiency of the evolved behavioral strategy in the presence of chaos, I divide the robots into two groups: chaotic and non-chaotic. Following the steps described in detail in the previous section, trajectories that fail to lower the amount false nearest neighbors below 5% at the tenth dimension are labeled as non-chaotic. Furthermore, datasets which fail to exhibit a deterministic dynamic due to a large nonlinear prediction error are included in the same set. Moreover, trajectories unable to reject the null hypothesis, postulated with the surrogate data testing, of a linear dynamical system with noise are also classified as non-chaotic. Finally, a constant increase of the fractal dimension D_2 (Figure 4.3) or a flattening of the scaling region for λ_1 (Figure 4.4), when the time series is embedded in higher dimension is also a hallmark for stochastic system. Finite dimensional systems converge once the embedding dimension is large enough and the attractor is properly unfolded. Stochastic systems are infinite dimensional and the estimated fractal dimension never settles to a constant value. At the same time, λ_1 decreases because higher embedding dimensions are large enough to accommodate the dynamic. Therefore embedded time series that fail the convergence of λ_1 or D_2 are also categorized as non-chaotic. Nonstationary time series are not included in any of the two groups as a dynamic that qualitatively changes over time undermines the reliability of λ_1 and D_2 . Any robot that is not included either in the non-chaotic group, or discarded after the cross-prediction error for stationarity, is classified as chaotic. Average values of λ_1 and D_2 , as well as the number of chaotic robot for each experimental scenario, are reported in Table 4.1.

4.2.2 Chaotic measures and performance

A first question I want to answer is whether robots that exhibit a chaotic behavior have advantages over robots in the non-chaotic group. To compare different groups of robots non-parametric

4.2. RESULTS

	FF			SRN		
	N	λ_1	D_2	N	λ_1	D_2
E1	38	0.057(± 0.052)	1.66(± 0.67)	59	0.065(± 0.057)	1.64(± 0.67)
E2	52	0.060(± 0.051)	1.61(± 0.63)	51	0.068(± 0.064)	1.53(± 0.70)
E3	49	0.046(± 0.043)	1.55(± 0.70)	58	0.070(± 0.041)	1.32(± 0.60)

Table 4.1: **Average values of chaotic measures.** Number of chaotic robots (N), level of chaos (λ_1), and fractal dimension (D_2).

bootstrap is used, which is a simple form of Monte Carlo estimation. In frequentist statistics, the sample is representative with some degree of approximation of the unknown true distribution. The idea behind bootstrap is to create fake samples $\hat{\theta}(D)$ exploiting the values of the original dataset D . In practice, to comparing two empirical distributions the algorithm involves few simple steps. Firstly, a new sample $\hat{\theta}(D_1)$ of the same size n of the original data D_1 is created choosing random subjects with repetition and substitution. Similarly, a new dataset $\hat{\theta}(D_2)$ of size m is randomly generated for distribution D_2 and the difference d_k between the average value of resampled populations $\hat{\theta}(D_1)$ and $\hat{\theta}(D_2)$ is calculated. The process is iterated k times and the final result is a bell shaped distribution consisting of the differences of the means between samples bootstrapped from the recorded data D_1 and D_2 . Finally, upper and lower 95% confidence intervals are derived from the 5th and 95th percentiles of the distribution. The only parameters is the number of bootstrap, which is set to $k = 1,000$ as the estimation of the percentiles stabilizes. This simple technique can be easily used to evaluate significant statistical differences between two populations. If 0 is included between the upper and lower confidence intervals, the two populations are equal. On the contrary, two negative or positive confidence intervals indicate a higher score for D_2 and D_1 , respectively. The motivations that lead to applying bootstrapping for comparing two populations are multifaceted. From a methodological perspective, non-parametric bootstrapping avoids assumptions on the true distribution. The available data is exploited to generate a more tractable normal curve. Moreover, there is a practical reason that arises from the different size of the samples, which undermines the statistical power of traditional statistical tests. Before testing the difference in performance between robots that exhibit a chaotic behavior and agents that are not governed by such dynamics, preliminary analyses are conducted on chaotic robots controlled by different kinds of neural networks. The motivation is to verifying whether the presence




Figure 4.5: Mean difference bootstrapping of the fitness between chaotic robots controlled by FF or SRN neural network. The statistical test evaluates the effect of the neural controller on the performance in the presence of chaos. Positive confidence intervals (red solid lines) indicate a better performance of agents with FF controllers. During the testing phase, robots act in environments never experienced during evolution, expanding the upper and lower bounds of the set of timescales by 50. (a), (b), and (c) refer to robots evolved in E1, E2, and E3, respectively.

of deterministic chaos affects the quality of the solution despite the neural controller is static or dynamic. Results reported in Figure 4.5 show that there is no significant difference of fitness between FF and SRN in all three kind of environments, that is, experiments where robots are trained with 1, 21 or 61 timescales. Therefore, robots controlled both by FF and SRN are grouped together and tested against agents that do not show presence of chaos. However, there is no significant difference between the performance of chaotic and non-chaotic robots evolved in E1 (Figure 4.5(a)).

Interestingly, non-parametric bootstrapping demonstrates that robots characterized by chaotic dynamics have a significant higher fitness score in environments where they have been exposed to a variety timescales during evolution, as shown in Figures 4.6(b-c).

To strengthen the result, robots trained using a single interval are tested with ± 80 timescales, that is, the same number of environments experienced by robots evolved with 61 environments during the testing phase. The same test is performed on the population of robots evolved with 21 timescales, expanding the original set of time intervals of ± 70 . Results confirm the hypothesis




Figure 4.6: Mean difference bootstrapping of the fitness between chaotic and non-chaotic robots. The statistical test evaluates the effect of chaotic behaviors on the robots' performance during the testing phase. Positive confidence intervals (red solid lines) indicate a better performance of chaotic agents. (a) shows that there is no significant difference in performance in E1. (b) and (c) show a larger fitness for chaotic robots evolved in E2 and E3, respectively.

of a higher ability of robots driven by a chaotic dynamics to adapt to novel environments. Comparing Figure 4.6(a) with Figure 4.7(a), there is a noticeable increment on the difference of the performance between chaotic and non-chaotic robots. Consistently, the increasing divergence on performance when robots trained with 21 timescales are tested in a wider variety of intervals is also evident, as shown in Figure 4.6(b) and Figure 4.7(b).

So far I have evidence that robots governed by chaotic dynamics have a better behavioral response to environments never experienced during the evolutionary process. A further question concerns whether λ_1 or D_2 are related to the robots' performance. To assess the effect of λ_1 and D_2 on the robots' behavior, Pearson correlation coefficient is estimated to quantify the degree of linear dependence between the average fitness score on the testing scenario and chaotic measures. Results summarized in Figure 4.8 show a moderate negative correlation between D_2 and fitness score, whereas λ_1 has a weak positive correlation.

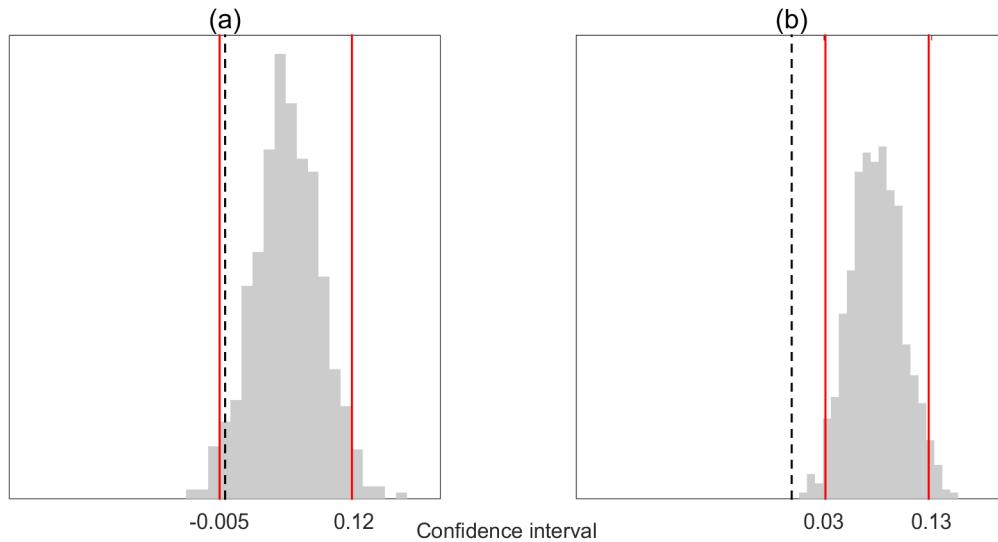


Figure 4.7: Mean difference bootstrapping of the fitness between chaotic and non-chaotic robots evolved in E1 and E2, tested in a larger number of novel environments. Robots evolved in E1 and E2 are tested with 161 timescales, as robots of E3. Positive confidence intervals (red solid lines) indicate a better performance of chaotic agents. (a) shows no significant difference in E1. (b) show a better performance for chaotic robots evolved in E2.

4.2.3 Chaotic dynamics and behavioral strategies

In this section I verify whether the level of chaos, quantified by λ_1 , or the number of degrees of freedom, estimated with fractal dimension D_2 , are related to different behavioral strategies. These dynamic descriptors are calculated in an abstract phase space, however, sensitivity to initial conditions and complexity of the system affect the actual behavior of the robot in physical space.

Robots' behavior is evaluated by distal analysis (Nolfi & Floreano, 2000), that is, a behavioral description from the point of view of the observer. By inspecting the trajectories of the robots, I group agents that display similar solutions, and for each experiment two prototypical behaviors are selected. The choice of the selected behaviors is arbitrary, but it follows a rationale. I analyze behaviors that are characteristic of environments where robots are exposed either to 1, 21, or 61 timescales during the evolutionary process, comparing the effect of chaotic measures on the trajectories followed by the robots and the performance. Furthermore, I select a strategy which emerges both in environments where robots are trained with 1 timescale and scenarios that force

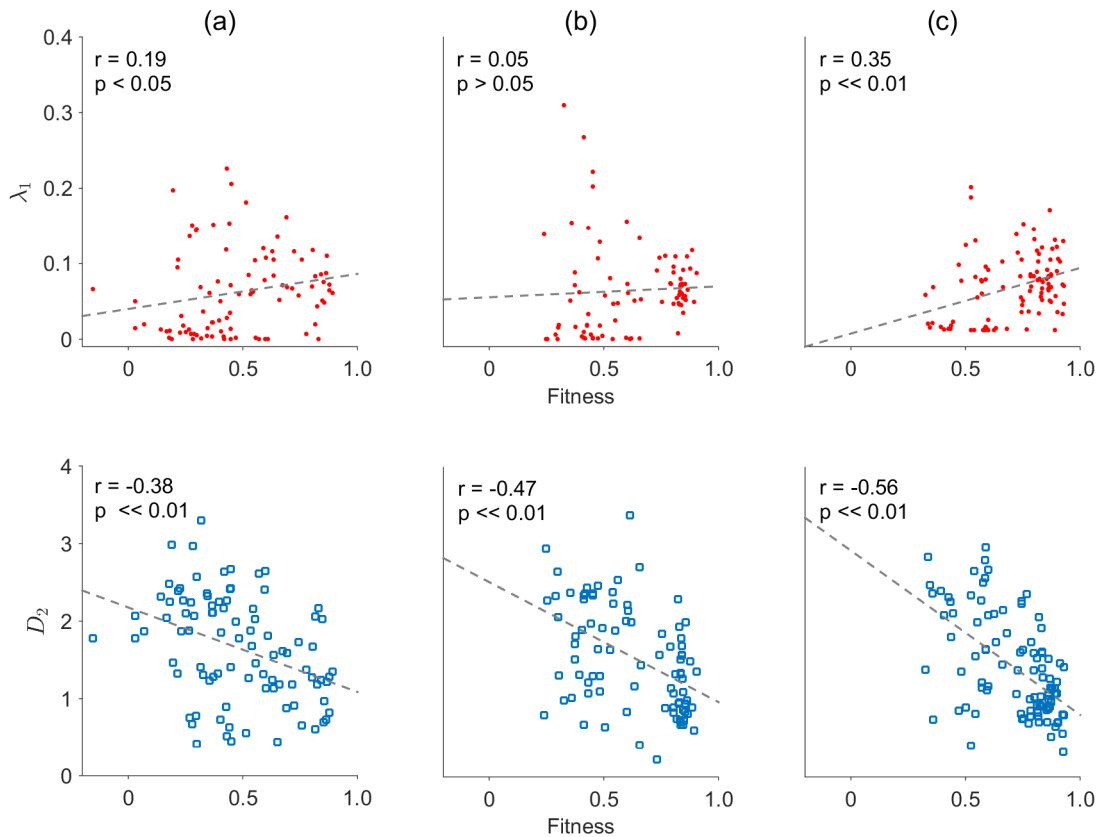


Figure 4.8: Correlations between chaotic measures and fitness. The fitness value measures the average performance of chaotic robots during the testing phase, when the number of timescales is increased of ± 50 . Points of the scatter plot represents single robots. The gray dashed line indicates the Pearson correlations of λ_1 (top) and D_2 (bottom) with the fitness for robots evolved in E1(a), E2(b), and E3(c).

the agent to adapt to several timescales during evolution.

Experiment 1

In the group of chaotic robots evolved in environments characterized by a single timescale, 4 robots develop a behavioral strategy which is not present in scenarios where agents are exposed to several intervals. Figure 4.9(b) shows a prototypical trajectory drawn by Robot 1 Experiment 1 (R1E1) controlled by a FF neural network. A similar behavior (Figure 4.9(c)) is produced by Robot 2 Experiment 1 (R2E1), which mounts a SRN. Both agents are evolved in an environment where the target area changes color every 200 time steps. As shown in Figure 4.9(a), estimated chaotic measures are very similar for robots controlled either by a FF ($D_2 = 2.11$, $\lambda_1 = 0.022$) or a

SRN ($D_2 = 1.99$, $\lambda_1 = 0.002$). Furthermore, the performance of the two robots during the testing phase is comparable. As shown in Figure 4.9(f) and Figure 4.9(g), both agents have a heterogeneous efficiency towards different timescales. Overall, the evolved behavioral strategy is resilient to disturbances given by novel timescales, scoring an average fitness of 0.37 ± 0.21 (R1E1) and 0.47 ± 0.18 (R2E1). In comparing the first two agents, R1E1 and R2E1, I show that robots exhibit a similar behavior in Cartesian space, and are characterized by similar chaotic measures estimated in phase space, react likewise to environment never experience during learning. To further investigate the effect of dynamic descriptors on the quality of the solution, I examine a third robot that exhibits a different behavior, but characterized by equivalent complexity and level of chaos. The selected Robot 3 Experiment 1 (R3E1) is controlled by a SRN and is evolved in an environment where the target area switches color every 200 time steps. Although chaotic measures of the reconstructed attractor are similar ($D_2 = 2.07$, $\lambda_1 = 0.015$), the trajectory plotted in Figure 4.9(d) differs from the behavior displayed by the other two robots, namely, R1E1 and R2E1. The strategy developed by R3E1 consists of an elliptical orbit outside the blue target area, which leads the robot to enter with a trajectory tangent to the circumference of the target area when the color switch happens. While on the green target area, the robot draws circles of small radius gaining positive reward. This solution is less robust compared with the prototypical behavior of R1E1 and R2E1, which directs the robots inside the green target area with a perpendicular trajectory. Figure 4.9(h) shows a high fitness when R3E1 is tested in the environment experienced during the evolutionary process, but the performance clearly drops with different timescales, lowering the average fitness score to 0.03 ± 0.08 . The last behavioral analysis of robots evolved in environments with a single timescale is focused on agents with opposite values of chaotic measures, that is, low fractal dimensions and higher levels of chaos. A fourth agent, Robot 4 Experiment 1 (R4E1), is randomly chosen among the group of robots with $\lambda_1 = [0.15; 0.25]$ and $D_2 = [0; 1]$. The robot's controller is a SRN, evolved in an environment characterized by a timescale of 200 time steps. The strategy evolved by R4E1 is shown in Figure 4.9(e). According to the results reported in the previous section, λ_1 has a weak positive correlation with the robots' ability to adapt to novel environments, but there is a moderate negative correlation between low fractal dimensions D_2 and the performance. Therefore, there is a high probability of selecting a robot that evolved

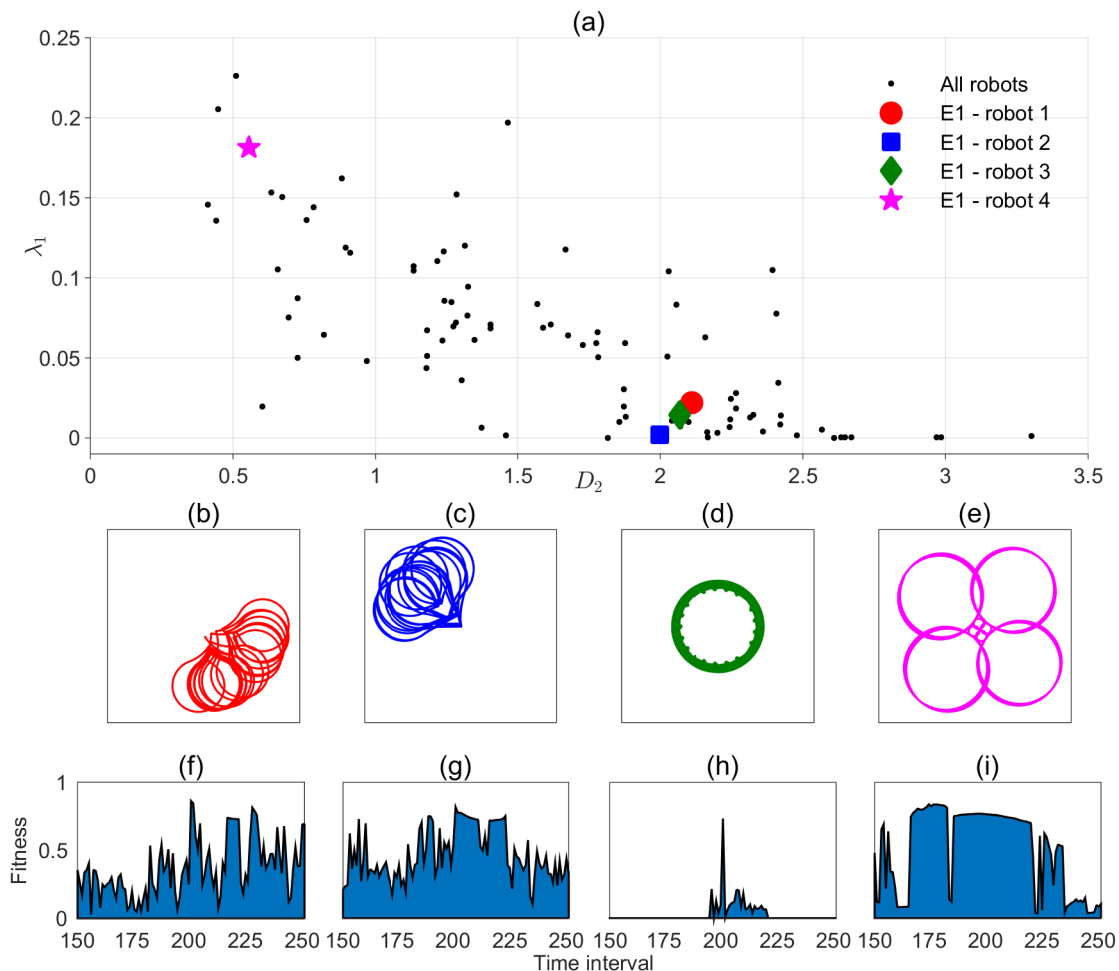


Figure 4.9: **Behavioral analysis of robots evolved in E1.** (a) shows the values of λ_1 plotted against D_2 for the whole set of chaotic robots evolved in E1. Markers indicate robots selected for the analysis. (b-e) depicts the robots' trajectories over 7,000 time steps and (f-i) the fitness achieved in every environment experienced by the agents during the testing phase.

a strategy robust to environmental noise produced by a change in the timescale. Figure 4.9(a) confirms our prediction showing normalized fitness scores of 0.51 ± 0.30 when the robot acts in the 101 environments experienced during the testing phase.

From the analysis of robots' trajectories I can formulate the hypothesis that there is a correlation between chaotic measures and robustness of the solutions toward novel environments, with the constraint that the agents exhibit the same behavioral strategy. In other terms, dynamic quantities

are useful tools for predicting the quality of the behavioral strategy if the underlying dynamical system is the same, but variables representing parameters have different values.

Experiment 2

In order to provide evidence for the formulated hypothesis, 4 more robots (Figure 4.10(a)) are examined within the group of chaotic agents trained in environments with 21 different temporal dynamics. Among the various solutions found by the evolved robots, I select the trajectories followed by Robot 1 of Experiment 2 (R1E2) and Robot 2 of Experiment 2 (R2E2), shown in Figure 4.10(b) and Figure 4.10(c), respectively. The first robot is controlled by a FF and the evolutionary niche has timescales in the range $T = \{390, 391, \dots, 419, 420\}$. R2E2 evolves in environments where the color switch happens with intervals $T = \{190, 191, \dots, 219, 220\}$, and the agents mounts a SRN. This kind of behavioral pattern is exhibited only by robots trained in environment with 21 timescales and it is adopted by a total number of 8 robots. The avoidance strategy of the blue target area is based on circular trajectories that bring the robot in the safety zone, re-entering at regular intervals to check the state of the target area. If the target area is green, the robot stays inside, navigating on the central portion of the green target area. Interestingly, the two robots have a similar behavior and level of chaos, $\lambda_1 = 0.088$ for R1E2 and $\lambda_1 = 0.057$ for R2E2, but the estimated fractal dimension is substantially different. In fact, the underlying dynamic of R1E2 has a dimension $D_2 = 1.72$, whereas the number of degrees of freedom involved in the behavior of R2E2 is much lower, with $D_2 = 0.65$. Both robots achieve a large fitness in a wide number of scenarios never experience during the evolutionary process (Figures 4.10(f-g)). As expected, the robot with lower fractal dimension has a slightly better performance, scoring an average fitness of 0.41 ± 0.29 versus a fitness equal to 0.37 ± 0.26 scored by the robot with a dynamic where a higher degree of complexity is involved. To randomly choose Robot 3 experiment 2 (R3E2), I follow the same logic used with the third robot of the E1, that is, an agent that displays a similar level of chaos ($\lambda_1 = 0.081$) and a comparable fractal dimension ($D_2 = 1.64$) with R1E2. The robot's artificial brain is recurrent and synaptic weights adapt to environments with $T = \{490, 491, \dots, 519, 520\}$. The trajectory of the robot is shown in Figure 4.10(d) and the fitness score recorded in the 121 intervals of the testing phase are plotted in Figure 4.10(h). Results show a considerably better

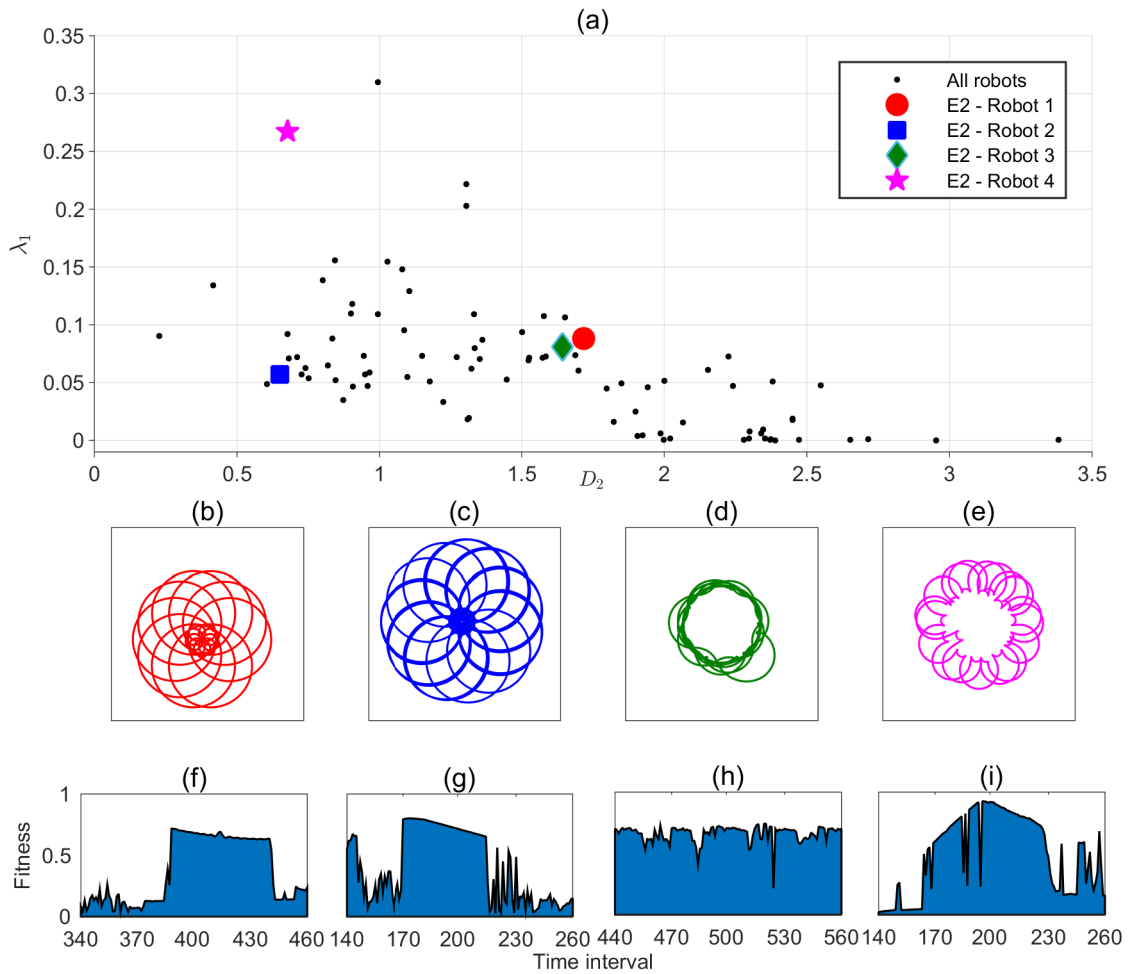


Figure 4.10: **Behavioral analysis of robots evolved in E2.** (a) scatter plot of the level of chaos (λ_1) and complexity (D_2) of chaotic robots evolved in E2. Markers refer to robots selected for the analysis. (b-e) show the evolution of the trajectories for 7,000 time steps in Cartesian space. (f-i) show the fitness score in 121 environments.

performance compared to R1E2 and R2E2, with an average fitness of 0.66 ± 0.07 . Finally, I analyze a robot with high level of chaos ($\lambda_1 = 0.267$) and a low fractal dimension ($D_2 = 0.68$). The trajectory of the selected agent, Robot 4 Experiment 2 (R4E2), and the overall performance is shown in Figure 4.10(e) and Figure 4.10(i), scoring an average fitness of fitness of 0.49 ± 0.32 . The agents has the same controller and evolves in the same environments of R2E2. Contrary to the prediction that can be infer from the Pearson correlation shown in Figure 4.8, the average performance of the robot in environments not experienced during evolution is worse than R3E2.

This result supports the hypothesis stated above concerning the role of chaotic characteristic measures on the performance of the robots. In fact, level of chaos and degree of complexity are not absolute invariant measures that correlate with the efficiency of the solution despite of the behavioral strategy employed by the agents. Instead, they are useful tools for explaining the quality of the solution if the robots' behaviors are governed by the same dynamic.

Experiment 3

To strengthen the preliminary conclusions achieved through behavioral analysis of selected robots evolved in E1 and E2, 4 more robots of E3 are analyzed (Figure 4.11(a)). For the behavioral analysis of the last experiment, where robots experience 61 intervals during evolution, I first examine two robots that exhibit the same kind of behavior, but the underlying dynamics have different chaotic measures. The motivation is to validate the hypothesis which follows from the comparison between the average fitness of R1E2 and R2E2, that is, a low fractal dimension have an effect on the performance if robots have a similar behavior. The strategy adopted by 6 robots evolved in environments with 61 timescales consists in elliptical orbits that move the agents outside the blue target area, followed by regular circles of wide radius drawn inside the green target area. Within the group of robots that display such a solution, I compare Robot 1 Experiment 3 (R1E3), whose trajectory is shown in Figure 4.11(b), and Robot 2 Experiment 3 (R2E3), which follows more regular orbits, as shown in Figure 4.11(c). The two robots clearly have different values of chaotic measures, with R1E3 having a $D_2 = 1.65$ and $\lambda_1 = 0.086$, against $D_2 = 0.91$ and $\lambda_1 = 0.165$, which characterize the dynamic of R2E3. In terms of fitness attained during testing stage, R1E3 reaches a mean value of 0.66 ± 0.03 (Figure 4.11(f)), while R2E3 scores 0.80 ± 0.03 (Figure 4.11(e)). As expected, lower fractal dimension and higher level of chaos have a positive impact on the overall performance of the robot if the two agents have a similar behavior. Remarkably, the trajectory followed by R2E3 is visually more regular than the path exhibited by R1E3. Moreover, 7 robots of E2 behave in a similar manner, including R3E2, which extends the analysis to robots evolved with a different number of timescales. Although the behavioral strategy is similar, Figure 4.10(d) and Figure 4.11(b) show a clear difference between trajectories of R3E2 and R1E3. As reported previously, fractal dimension, level of chaos, and average fitness have

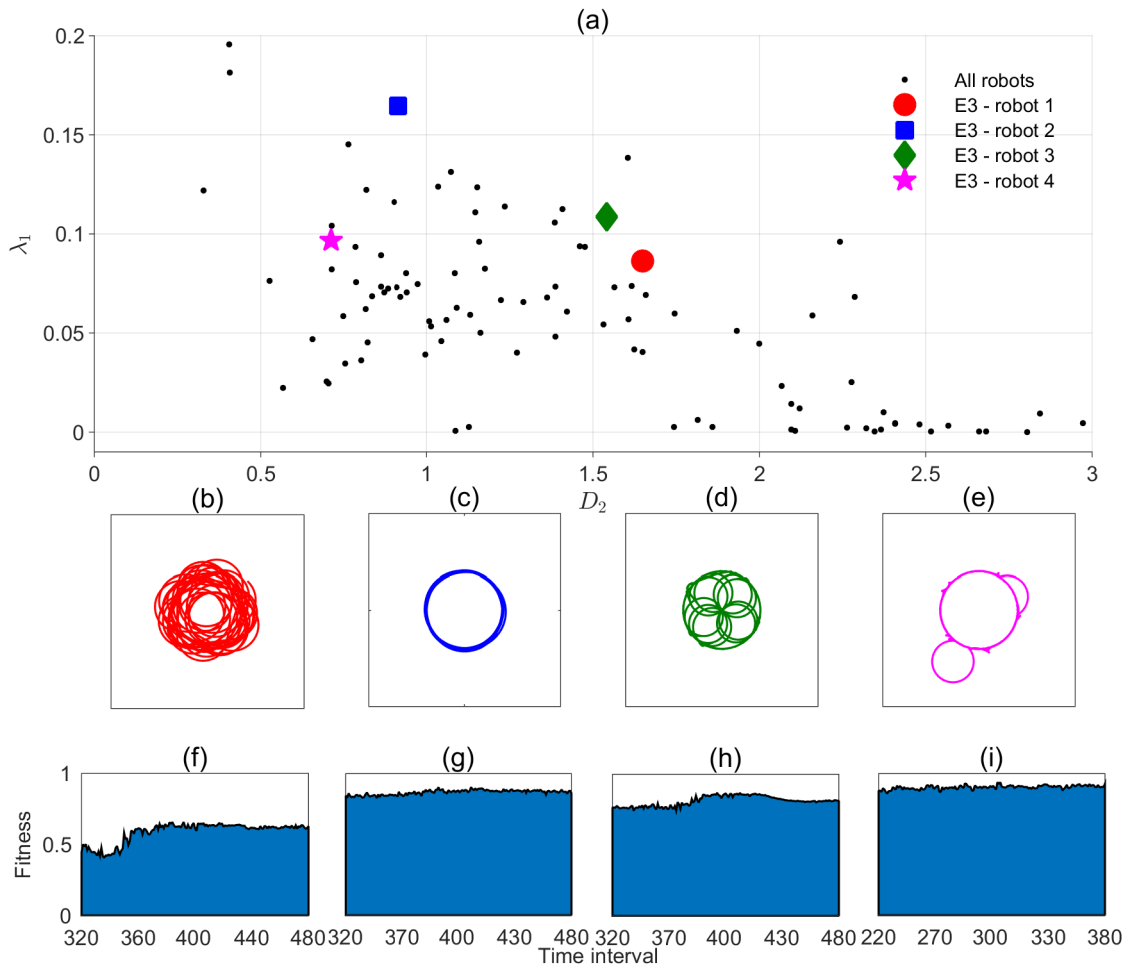


Figure 4.11: **Behavioral analysis of robots evolved in E3.** (a) points represent chaotic robots with the associated level of chaos (λ_1) and fractal dimension (D_2). Markers show the selected robots. (b-e) depict the trajectories followed by the analyzed robots while executing the task for 7,000 time steps. (f-i) illustrates the performance of the 4 robots during the testing phase in 161 different environments.

similar values for both robots. This is consistent with the hypothesis of a correlation between chaotic characteristic measures and fitness, with the restriction that agents employ a similar behavioral strategy. Furthermore, the evolutionary niche that shapes the solution is subordinated to dynamics resulting from the evolutionary process. In fact, similar behaviors characterized by comparable level of chaos and complexity, but evolved with a different number of timescales, have the same capability to adapt to novel environments. As for the previous experiments, an agent with similar level of chaos and fractal dimension of R1E3 is selected. Figure 4.11(d) shows

the trajectory drawn by Robot 3 Experiment 3 (R3E3) and the fitness achieved in the 161 environments of the testing phase is reported in Figure 4.11(h). Although the dynamic of R3E3 is characterized by chaotic measures similar to R1E3, with $D_2 = 1.54$ and $\lambda_1 = 0.109$, the solution found by R3E3 is more efficient than the strategy of R1E3. In fact, the average fitness is equal to 0.81 ± 0.03 , which is closer to R2E3. This confirms our claim that D_2 and λ_1 are not absolute invariant measures that correlates with the performance, but they are reliable estimators only when robots are governed by the same dynamic. The first three agents are controlled by a SRN and are trained using the set of timescales $T = \{370, 371, \dots, 429, 430\}$. The dynamic of the last agent, Robot 4 Experiment 3 (R4E3), is characterized by an intermediate level of chaos ($\lambda_1 = 0.097$) and a low fractal dimension ($D_2 = 0.71$). The trajectory of R4E3 is shown in Figure 4.11(e) and the fitness gained in the test environments is reported in Figure 4.11(i). The behavioral strategy is the product of a SRN, evolved with a set of timescales $T = \{290, 291, \dots, 319, 320\}$. The solution is robust towards environments experienced during training and novel scenarios, achieving an average fitness of 0.90 ± 0.01 .

Results concerning the behavioral analysis of robots evolved in E3 are consistent with deductions that follow from the examination of robots' trajectories evolved in experiments characterized by 1 and 21 timescales. Therefore, the hypothesis that D_2 and λ_1 are correlated to the fitness score when robots behave in a similar manner is confirmed. Moreover, I argue that D_2 and λ_1 are invariant with respect to the evolutionary niche.

Common behaviour for experiment 1,2 and 3

In order to provide further evidence to this claim, a behavioral strategy that emerges in every experiment is examined. Estimated chaotic measures show a variety of qualitatively different attractors, with level of chaos bounded in a small interval ($\lambda_1 = [0.0003; 0.13]$), but with noticeable differences for the number of degrees of freedom ($D_2 = [0.68; 2.38]$). Robots that evolve this strategy navigate inside the green target area following the perimeter. When the target area is subjected to the color switch, agents steer into the empty portion of the environment. While the color of the target area is blue, robots move following a circular trajectory of small radius. By employing this behavioral strategy, the agents enter inside the target area at regular intervals in order

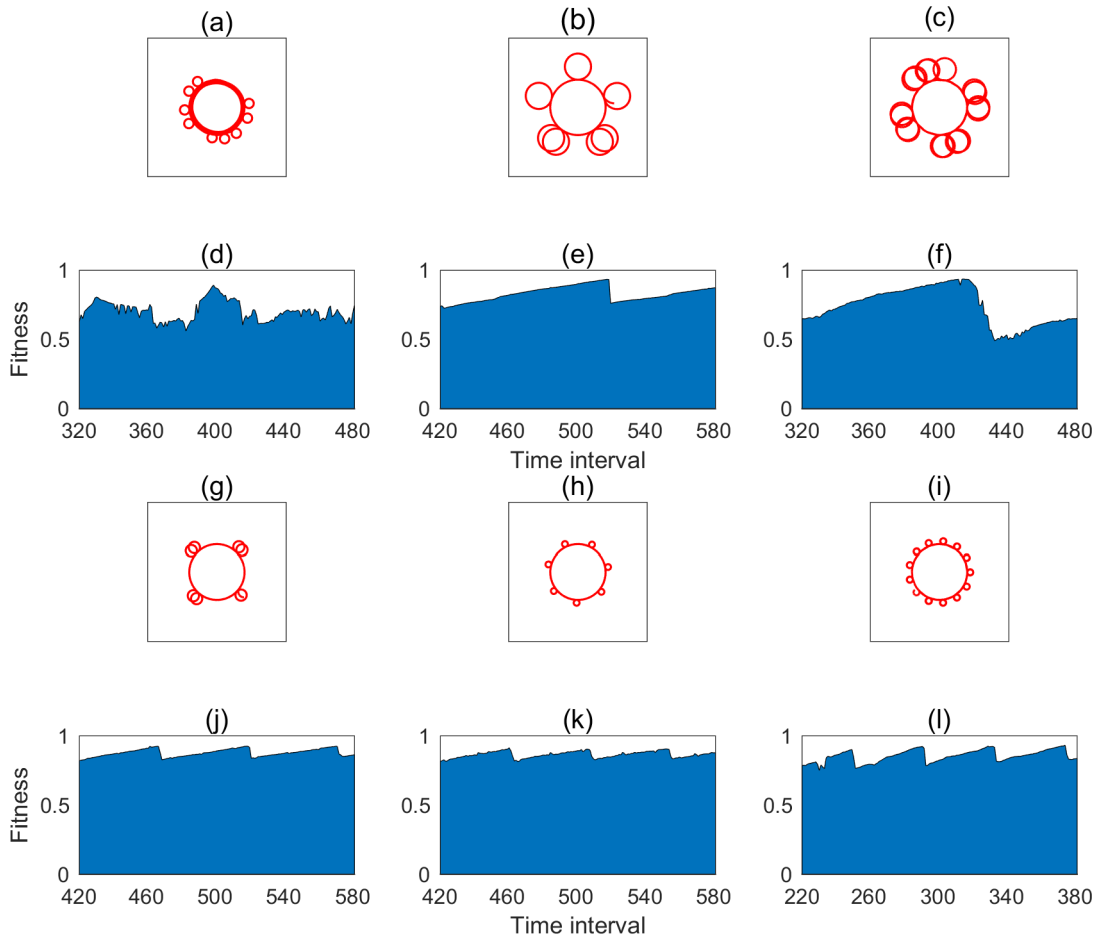


Figure 4.12: Behavioral analysis of robots that developed a similar strategy in different experiments. Trajectories followed by robots with high fractal dimension evolved in E1(a), E2(b), E3(c), and the corresponding fitness achieved during the testing phase (d-f). Strategies adopted by robots with low fractal dimension (g-h), and the performance of robots trained in E1(j), E2(k), and E3(l).

to check the current color. In order to have an unbiased comparison between robots evolved in different experiments, that is, trained in environment with 1, 21 or 61 timescales, agents pertaining to E1 and E2 are tested in 161 environments. Although the strategy evolved by this pool of agents is efficient and almost every robot achieves a high average fitness score (0.80 ± 0.01), chaotic measures still correlate with efficiency of the solution. The robot that uses the weakest behavioral strategy has also the highest fractal dimension $D_2 = 2.38$, scoring a noticeable lower average fitness of 0.60 ± 0.79 . I compare two clusters of three robots, grouped together according

to the values of D_2 and λ_1 . The criteria for selecting the trajectories is based on robots of E1, as only a limited number of agents evolve the analyzed prototypical behavior. For each robot of E1 that undergoes a detailed behavioral analysis, two robots of E2 and E3 are randomly selected within an interval of ± 0.1 for fractal dimension and ± 0.01 for level of chaos. Therefore, robots included in the same group have similar D_2 and λ_1 values. There is a single robot evolved with 1 timescale which dynamic is characterized by a significant high fractal dimension, thus the choice is constrained to one agent. Instead, the agent governed by a dynamic with low complexity is randomly chosen among attractors with $D_2 < 1$. From a visual inspection of the robots' trajectories, and evaluating the fitness achieved in different environments, the effect of the degree of complexity is evident. Firstly, robots with a low-dimensional chaotic dynamic (Figures 4.12(g-i)) avoid the blue target area drawing small circles in the empty zone of the environment. Instead, dynamics with a higher degree of complexity lead the robots either to follow a circular trajectory of wider radius in the empty part of the environment (Figures 4.12(b-c)), or navigate inside the green target area following a slightly spiral orbit (Figures 4.12(a)). Secondly, Robots with a higher fractal dimension have a slightly worst performance during the testing phase (Figures 4.12(d-e)) compared with agents' behavior characterized by a lower degree of complexity (Figures 4.12(j-l)), as reported in Table 4.2.

	High dimensional			Low dimensional		
	<i>Fitness</i>	D_2	<i>Controller</i>	<i>Fitness</i>	D_2	<i>Controller</i>
E1	0.70(± 0.07)	1.59	FF	0.877(± 0.03)	0.82	FF
E2	0.83(± 0.05)	1.57	FF	0.865(± 0.02)	0.82	FF
E3	0.74(± 0.13)	1.62	SRN	0.854(± 0.04)	0.87	SRN

Table 4.2: Fitness and fractal dimension of robots that exhibit a similar behavior in different evolutionary niches. Robots are randomly selected from a group of agents that evolved the same behavioral strategy. Agents are evolved in environments characterized by a change in color of the target area with different numbers and lengths of time intervals. Data show that the performance is unrelated to the evolutionary niche and the neural controller. Fractal dimension, instead, is a numerical descriptor that predicts the robustness of the solution towards novel environments.

4.3 Discussion

In this chapter I estimate chaotic measures reconstructing the attractor from the behavior of embodied and situated neuro-robotic systems, trained with genetic algorithm. The procedure for

calculating level of chaos and degrees of freedom follows a widely accepted definition of chaos as a deterministic aperiodic dynamic, bounded in a region of the phase space, and sensitive on small variations on the initial conditions (Kaplan & Glass, 2012). The definition is operationalized using appropriate nonlinear mathematical tools applied on the reconstructed attractor in phase space. Suitable tests are employed to verify the presence of a low-dimensional and deterministic chaotic dynamic. Nevertheless, current analytical methods in nonlinear science do not provide a definitive signature for the presence chaos in empirical data. Although a recent framework relaxes the definition of chaotic system, including non-autonomous dynamical systems with stochastic inputs (Poon et al., 2010), the presence of dynamic or instrumental noise invalidates estimates aimed at quantifying level of chaos or fractal dimension. In the presence of stochastic components, measures developed for low-dimensional systems are applied to high-dimensional cases. Therefore, the reliability of such estimators becomes questionable (Kantz & Schreiber, 1998; Faure & Korn, 2001; Sarbadhikari & Chakrabarty, 2001; Korn & Faure, 2003; Le Van Quyen et al., 2003; Stam, 2005; Bob et al., 2009; Freitas et al., 2009; Glass, 2009). Considering this issue, I utilize chaotic measures on a large number of simulated autonomous robots. The main question concerns whether chaotic measures are useful estimators for explaining the solution developed through the real time interaction of agents with the surrounding environment. In this regard, I agree with the interpretation suggested by Bob et al. (2009). The authors find a correlation between schizophrenic associations and increased level of chaos calculated from recorded electrodermal activity, but they are aware of the current limitations for discriminating stochastic and fully deterministic systems. A large λ_1 is thought as a metaphor for pathologically increasing complexity and less predictability in the neural activity, rather than a certain hallmark for an underlying chaotic dynamic. Lacking of a reliable test for determinism in nonlinear science, the contribution of this chapter to chaos theory is to provide evidence about the utility of chaotic measures applied to empirical data, while avoiding strong and definitive claims regarding the underlying dynamics.

The transdisciplinary effort that follows from the embodied and situated view of cognition requires a shared mathematical framework. Moreover, chaos theory has the advantage of creating models both from equations and recorded data. Therefore, nonlinear analysis of time series

recorded from behavioral patterns or neural activity are useful for creating a common ground which links different disciplines. In a recent article, Reynolds et al. (2016) show that mud snails exhibit chaotic patterns while searching for food. This result may inspire roboticists while building a robotic system that executes efficiently a garbage collection task or explores an unknown area (Nakamura & Sekiguchi, 2001; Martins-Filho et al., 2004; Volos et al., 2013). On the other hand, ecologists that follow a dynamical system approach have an artificial model which can be used for confirming hypotheses developed with animal models (Floreano et al., 2007). Results reported in this paper show that robots involved in a temporal task, where agents develop a correct timing in order to avoid a target area that changes properties at regular intervals, are more efficient if the evolved behavioral strategy is chaotic. By taking into consideration the described robotic model, ecologists may be inspired while studying animal behaviors produced in similar tasks.

Furthermore, the neuro-robotic model confirms theoretical frameworks postulated by the embodied and situated view of cognition. In fact, robots' performance is unrelated to the kind of neural controller. Each experiment is replicated with FF and SRN neural networks, and in the presence of chaotic dynamics there is no significant difference in the performance. This result is somewhat unexpected as recurrent connections give a short term memory to the robots. An internal dynamic may cause changes of direction without perceptual information, condition that is fulfilled while robots are avoiding the blue target area. Although some behavioral strategies emerge only in populations of robots controlled by SRN, there is no effect on the fitness. Therefore, the emergence of an effective solution is offloaded in the coupled interaction of body, artificial brain, and environment during the evolutionary process. Results show that the entangled interaction among sub-components of the system is captured by chaotic measures, as they are correlated with the performance, rather than the type of controller. Similarly, same behaviors evolve in environments characterized by different amounts and lengths of timescales.

Leaving aside the interdisciplinary effort that links human and animal science with artificial systems, I now focus on the contribution given by our model to the field of mobile robotics. Firstly, results reported in this chapter are consistent with other evolutionary robotics scenarios, where

chaotic measures are applied to the behavior of two wheeled robots (Monirul Islam & Murase, 2005). By applying chaotic measures to the embedded time series recorded from the activity of the infrared sensors, there is an evident correlation between the robots' performance and fractal dimension. Level of chaos, instead, is not informative about the quality of the solution evolved by the robots. Moreover, different architectures of the neural controllers do not strongly affect either the level of chaos, or the degree of complexity of the system, as well as the fitness score. It is worth noting that λ_1 is calculated with a different algorithm (Wolf et al., 1985) and data is collected from perceptual states of the robots, rather than their position in the environment, but results are consistent with findings reported in this chapter.

The novel contribution provided by our experiments concerns the interplay between fractal dimension and the type of behaviors exhibited by the robots. The simple environment where evolution takes place is not constrained by narrow corridors, and robots of different populations adopt radically different behavioral strategies. By inspecting the robots' trajectories in Cartesian space, I conclude that there is not an absolute correlation between fractal dimension and performance. Rather, the degree of complexity of the system correlates with the robustness of the solution developed by the robots only if they behave in a similar manner. Secondly, robots that exhibit chaotic behaviors are more resilient towards environmental changes, scoring high fitness in environments either encountered through evolution, or experienced only during the testing phase. This result is aligned with experiments where robots' behavior is controlled by a chaotic dynamic resulting from a set of nonlinear Equations (Nakamura & Sekiguchi, 2001; Martins-Filho et al., 2004; Volos et al., 2013). For example, Nakamura & Sekiguchi (2001) show the advantage of a chaotic trajectory over a random walk when a two wheeled robot is exploring the environment.

4.4 Conclusion

In this chapter, I show that chaos theory can be successfully applied to datasets recorded from the position in the environment of embodied and situated agents. The definition of chaos as an aperiodic, bounded, and deterministic dynamic, with sensitive dependence on initial condition is operationalized with appropriate nonlinear analytical tools. Although the application of chaotic measures to empirical data is extremely controversial, results show that robots governed by a

4.4. CONCLUSION

chaotic dynamic are more resilient to environmental changes compared with non-chaotic agents. Furthermore, data support a correlation between the robustness of the solution evolved by the robots and the level of chaos, as well as the degree of complexity of the underlying dynamic. Results based on behavioral analyses show that chaotic measures are not absolute invariants. In fact, level of chaos and degree of complexity of the system explain robots' capability of adapting to novel environments only if agents have the same underlying dynamic, which leads to similar behavioral patterns. The existing set of mathematical tools for nonlinear time series analysis still lacks of a definitive test for discriminating fully deterministic and stochastic systems. However, chaos theory can be applied *de facto* to observables producing consistent results.

Chapter 5

Decomposing an embodied and situated system *via* information theoretic measures

5.1 Introduction

The embodied and situated view of cognition stresses the importance of real-time and nonlinear bodily interaction with the environment for developing concepts and structuring knowledge. From this perspective, the agent forms an integrated system with the surrounding world and thus understanding cognitive processing requires appropriate analytical tools. In accordance to an antireductionist stance, the experiment described in this chapter moves from the study of global properties towards the analysis of relationships among the sub-components of the system. By estimated Information-theoretic measures on pairings of variables to unveiling nonlinear interactions that structure the agent-environment system are unveiled.

Current approaches are mainly based on a qualitative geometrical description of the systems' behaviour from a dynamical system perspective (see for example (Beer, 1995; Tani & Nolfi, 1999; Montebelli et al., 2008; Carvalho & Nolfi, 2016)). Although visual descriptions have the merit of explaining the evolution of agent-environment interaction thorough time, a precise numerical explanation of observed phenomena is an essential prerequisite to advancing a scientific field. Furthermore, little effort is directed towards the study of relationships that characterise the subcomponents of the system. To overcome this issue, I analyse an evolutionary robotic model (Nolfi & Floreano, 2000) performing a wall-following task employing nonlinear mathematical tools developed in the field of information theory. The choice of such cognitive model is motivated by its consistency with an operationalisation of the embodied and situated view. In

the experimental scenario, a robot, capable of receiving perceptual inputs and producing motor responses, is controlled by an artificial brain and the agent autonomously develops a solution to a problem defined by the surrounding environment. To avoid a disentanglement of the sub-components of the system, this work is framed within the context of an antireductionist philosophy of science (Ahn et al., 2006b,a; Fang & Casadevall, 2011) and system science (Von Bertalanffy, 1968), using concepts recently developed in the field of system biology (Basso et al., 2005; Margolin et al., 2006). Specifically, we utilize mathematical tools developed in the field known as information theory. Information-theoretic measures do not directly accommodate into an antireductionist scientific framework, as for example chaos theory (Da Rold, 2017; Kesić, 2016). However, we apply such measures to all possible pairings of connected variables, thus unveiling the intricate web of dynamic relationships and structures of the indivisible system brain-body-environment. Thus, the idea of information-theoretic decomposition refers to a mapping of the relationships among sub-components of systems, to understanding how these interact within the system.

Recently, Beer and Williams (Beer & Williams, 2015) analyse a minimal cognitive agent involved in a relational categorization task, outlining a proposal for relating dynamical systems and information theory, showing the complementarity and consistency of the two styles of explanations. In this work, we employ local forms of information-theoretic measures, which directly capture characteristics of the ongoing dynamics. However, the existing literature reviewed in Section 2.3.2 is mainly focused on detecting modifications of the information structure when macroscopic alterations occur, as for example substantial modifications of the environments, or differences in the robots' body. On the contrary, in this experiment environmental variables and body morphologies are not manipulated to testing the robustness of the information-theoretic measures in detecting minor differences that emerge in different populations during the evolutionary process. Furthermore, neural networks in the evolved agents are subjected to synaptic lesions in order to verify and strengthen the conclusions that follow from the analysis. In what follows, I demonstrate that the set of information-theoretic measures provides a decomposition of the system, capturing the intricacy of nonlinear relationships that characterise robots' behaviour and neural dynamics.

5.2. RESULTS




Figure 5.1: Time series recorded from the robots' perceptual and motor neurons. The figure depicts the neural activity during two laps around the maze. (a) and (b) represent the left and right infrared sensor, respectively. (c) refers to the left motor actuator and (d) to the right output. Replications 1-8 navigates counter-clockwise, whereas the 9th and 10th follow the opposite direction. The neural activation is characterised by large positive spikes of the infrared sensor directed toward the outer walls and drops in activity to the inner motor.

5.2 Results

At the end of the evolutionary process every population evolve an efficient strategy for navigating around the squared maze. During testing and data collection the lifespan of robots is extended to 10,000 time steps in order to improve the accuracy of information-theoretic measures. The analysis is restricted to the robot that achieved the highest fitness score in each replication. The starting position directs the robots counter clockwise, but agents evolved in the 9th and 10th replication perform a 180° turn after approaching the first corner and circulate in the opposite direction. To unveil the minimal architecture necessary for the robots to execute the task, as well as every functional sub-architecture, all possible combinations of synaptic lesions are applied to the neural controllers. The pruned neural networks are evaluated counting the number of laps achieved by the robots. The robots controlled by a full architecture achieves on average 28.3 ± 6.2 laps. We utilize this benchmark as the fitness score may give spurious results if lesioned robots spin, thus scoring low amounts of reward, instead of touching a wall after few time steps terminating the

trial. Although the experimental setting is extremely simple and evolution is not characterized by noise, e.g. different starting positions or wall lengths, the resulting embodiment is very diverse in different replications. In fact, robots evolved from different random seeds presents noticeable dissimilarities concerning the possible sub-architectures that accomplish a number of laps equal, or very close to, the baseline test where no lesions are applied. Specifically, the average number of viable neural structures across the 10 replications is 8.1 ± 4.8 , ranging from a minimum of 2 for the best robot of the 3rd population to a maximum of 16 for the robots evolved from the 5th and 7th seeds. Despite of dissimilarities in the evolved neural structures, there are common features shared by all neural architectures. The synapses connecting the infrared sensor oriented toward the outer wall and the motor close to the inner wall is always necessary for a correct functioning of the robots. Furthermore, every evolved artificial brain requires at least one of the two connections between the motor outputs.

By inspecting the recorded neural activity shown in Figure 5.1 the only clear difference that emerges is between robots that navigate clockwise or counter clockwise. The pattern of the time series is similar for every robot, with positive spikes of the infrared sensor oriented towards the outer wall, followed by a drop of the motor neuron close to the inner wall. However, diversities in the neural structure and minor behavioural differences are not captured from an observation of the recorded time series.

In the following sections, mutual information and transfer entropy, including their local form, are estimated on signals recorded from the perceptual and motor neurons while robots are navigating in the arena. The initial 500 time steps of the recorded signals are discarded in order to avoid artifacts which may follow from a transient phase where robots are stabilizing the repetitive trajectory. Instabilities during the beginning of the trial are obvious for the two robots that follow a clockwise direction, as the agents invert the direction with a 180° turn at the first corner of the maze. However, the initial part of the time series recorded from agents that navigate counter clockwise is also characterized by a considerable amount of outliers. The information-theoretic measures are estimated on the recorded time series, and their first difference, for pairs of variables connected with a synapse. Datasets are normalized with 0 mean and standard deviation

1. Probabilities are calculated with univariate or multivariate Gaussian kernel, according to the dimensionality of the sample, as well as rectangular kernel. The aim of the analysis is to infer functional relations between variables of the same robot, rather than comparing the magnitude of information-theoretic measures that emerge in different populations. Therefore, the outcome of the various information-theoretic measures employed in this work are normalised, as we investigate differences that emerge in the same system, rather than focusing on intra-seeds evaluations.

In the next section I first report the results of the local forms of the local mutual information and transfer entropy, which provide a dynamic description of the exchange of information and causal relationships. Subsequently, the global versions of such estimates are applied to the embodied and situated systems for unveiling mutual dependencies and the structure of the information flow. Finally, inferences concerning the agent-environment interaction and the functional neural architecture based on different information-theoretic measures are combined in order to draw a behavioural description without observing the robot during the execution of the task.

5.2.1 Local mutual information

The averages value of the pairwise local mutual information between variables of the systems do not show substantial differences. Furthermore, the time series representing the dynamic of the exchange of information are not characterised by significant fluctuations. However, a normalization in the range $[0, 1]$ of the outcome of the local mutual information calculated on the 5 possible pairings between neurons of the same robot offers insights about the underlying functional neural structures. By normalising, minor differences on the estimations are magnified, thus creating a weighted ranking of interactions between variables of the same system.

Figure 5.2 shows the rankings of the averaged local mutual information estimated on the recorded time series using both rectangular and Gaussian kernels. Results demonstrate that information dynamics and nonlinear interactions that are not captured by observing the recorded activity of the perceptual and motor neurons are unveiled. On the contrary, the estimation of the information-theoretic measure on the first difference does not show any pattern that either maps the functional connectivity, or describes peculiarities of the behavioural strategy (data not shown).

The outcome of the local mutual information reveals functional dependencies between variables

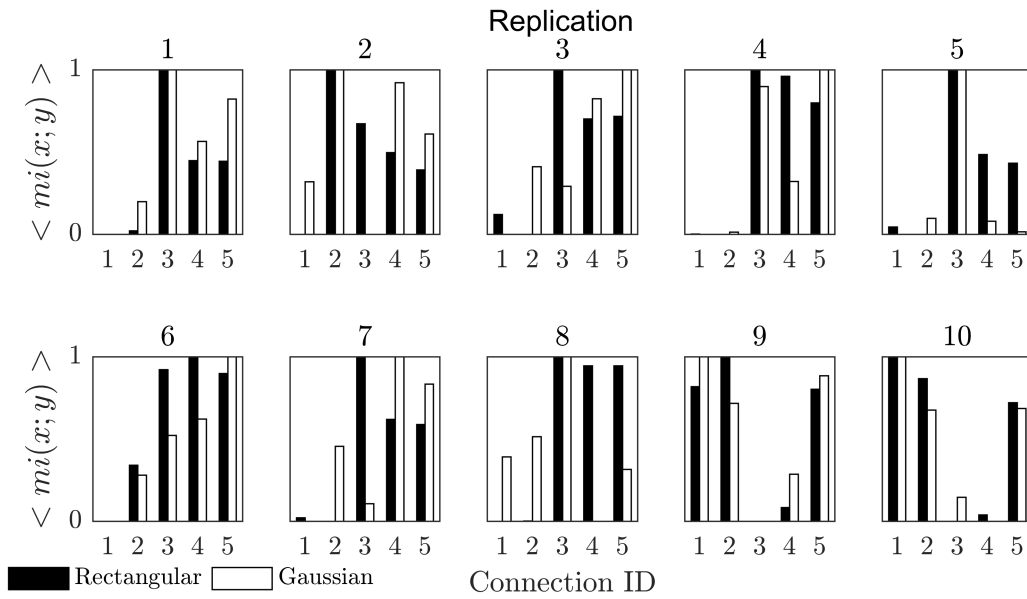


Figure 5.2: Average values of the local mutual information estimated on unpruned neural networks. The outcome of the pairwise local mutual information is normalised separately for Gaussian and rectangular kernels in the range $[0, 1]$ to obtain a weighted ranking of interactions. The degree of exchange of information between variables unveils functional relationships of the underlying neural structure.

of the system. For example, there is an overall high mutual dependence between the external infrared sensor and the inner motor, which is consistent with the incapability manifested by the robots of exploring the maze while this neural connection is lesioned. Interestingly, the 2nd replication is the only population of robots that navigates counter clockwise requiring the connection between the inner infrared sensor and the outer motor. The average local mutual information shows a high value of the aforementioned connection, thus detecting an association between the sensory information and the motor output. The local mutual information detects also peculiarities on the embodiment of the 5th replication. In fact, the neural architecture is resilient towards several combinations of synaptic lesions. The only synapse always present in all the 16 functional sub-architectures connects the outer sensor with the opposite motor and the information exchange between these two variables is significantly higher if compared to other replications. Therefore, the ranked values of the local mutual information are useful for detecting the importance of direct connections between variables that form the functional backbone of the neural system.

We observe noticeable differences on the estimation depending on the kernel employed for the reconstruction of the probability distributions. These variations are justified, as the local form of the mutual information does not include a linear component in the formula and thus low probabilities cause an exponential growth in the entropy level. In particular, the noticeable differences are visible in the 2nd, 3rd and 7th replication, where the Gaussian kernel is underestimating the connection between the outer infrared sensor and the inner motor. Other dissimilarity are evident in the 8th population, where the local mutual information estimated with rectangular kernel accentuates the importance of the sensorimotor chain in the right-hand side of the body and the recurrent connection of the output layer, whereas Gaussian kernel detects important exchanges of information on the opposite side of the robot. However, this agent functions with several sub-architectures that involves both sides of the body. Finally, both kernels estimates a low average value of information exchange between the inner infrared sensor and the inner motor of the 3rd replication, whereas the synapse connecting the neurons is fundamental for the execution of the task.

To further investigate the effect of different kernels on the estimation of information-theoretic measures, local mutual information is calculated for every possible neural architecture that allows a correct functioning of the system. In this chapter, we report results of 3 arbitrarily chosen replications. The rationale for selecting robots for the analysis with different sub-architectures aims at providing a variety of behavioural strategies that emerge in different brain-body-environment systems.

The first selected robot evolved from the 4th population (Figure 5.3). The agent neural structure is particularly resilient towards synaptic lesions, without visible effects on the recorded time series. The estimation of the local mutual information between the outer infrared sensor and the outer motor is significantly lower when probability distributions are reconstructed employing a Gaussian kernel. Despite of negligible differences, the average value of the local mutual information is significantly high between all sensors necessary for a correct execution of the task.

The second agent is the best robot of the 10th population. The agent navigates the arena clockwise and exhibits slightly different trajectories depending on the synaptic lesions applied to the

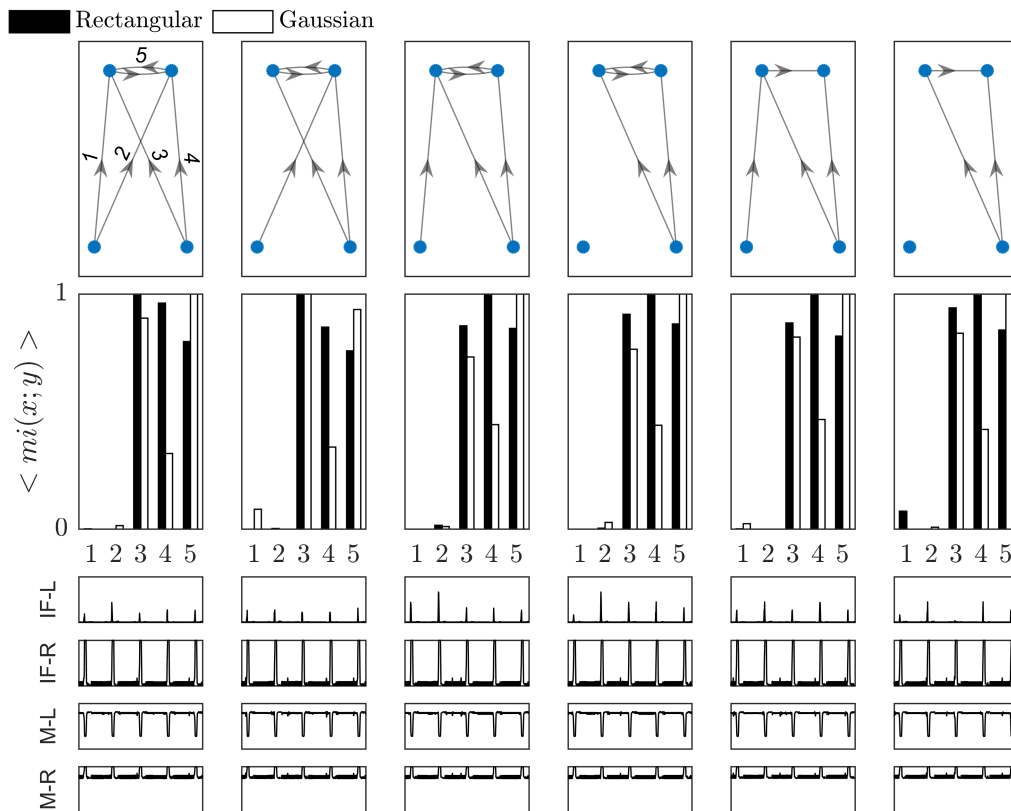


Figure 5.3: Estimation of the local mutual information in pruned neural networks evolved in the 4th replication. The local mutual information is calculated for all functional neural sub-architectures, creating separated weighted rankings. Synapses are labelled with identification numbers. This replication shows resiliency towards pruning, relying mainly on 3 connections that are characterised by high degrees of information exchange.

neural controller. Figure 5.4 shows a clear effect caused by the presence of the connection between the left infrared sensor and the left motor actuator, which stabilises the trajectory, and thus activation of the sensor. Although the visible effect is on the time series recorded on the opposite motor, the local mutual information is capable of detecting minor changes on the information exchange within the sensorimotor chains of the left-hand side of the robot. Overall, results are consistent with the sub-architectures identified through systematic lesions applied to the neural controller. Firstly, there is a high degree of dependence between the outer sensor and the inner motor, with the exception of one sub-architecture, where the Gaussian kernel estimates a low value. Secondly, the recurrent dynamic connecting the output layer is recognised as important

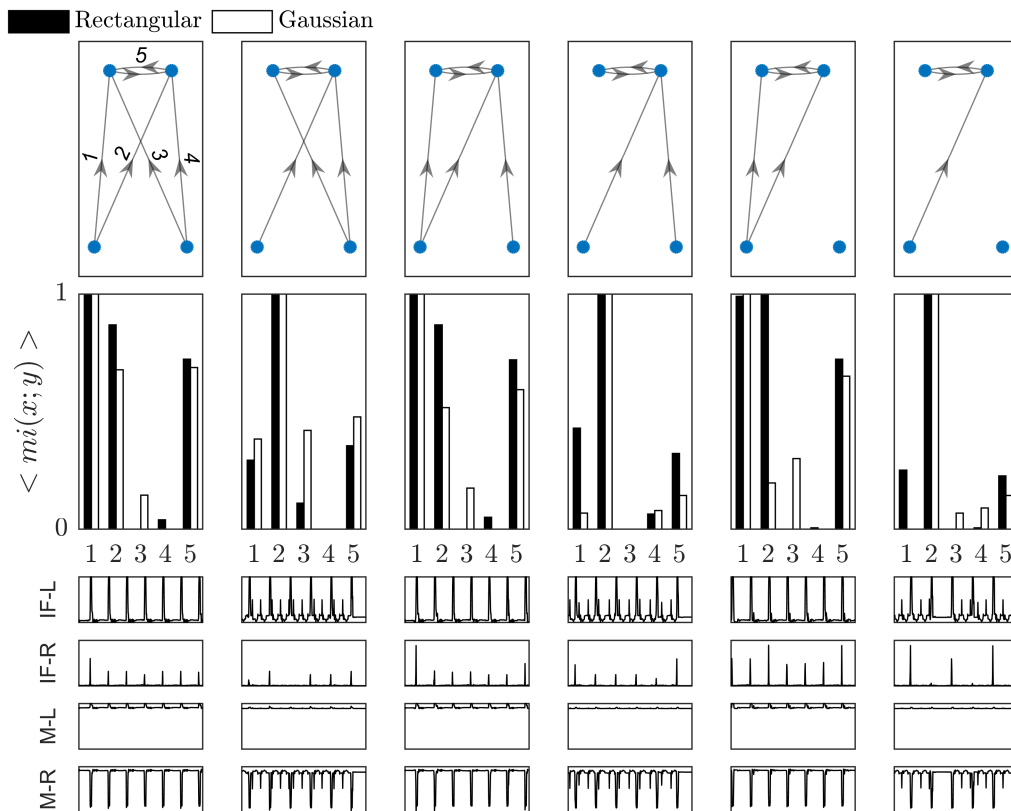


Figure 5.4: Estimation of the local mutual information in pruned neural networks evolved in the 10th replication. The local mutual information detects high degrees of dependence between variables that form the functional backbone of the neural network. Furthermore, the local mutual information individuates the presence of a synapse in the left sensorimotor chain, which stabilises the robot’s trajectory.

for 4 architectures. Finally, the absence of the connection between the left sensor and motor is detected as the measure of dependence drops to very low levels.

The last robot included in the analysis evolved from the 2nd replication, which is the only neural controller that follows a counter clockwise direction and requires the connection between the inner infrared sensor and the outer motor (Figure 5.5). When the robot navigates in the environment there are fast bursts in activity of the outer infrared, synchronised with low peaks of the activation of the inner motor. The inner infrared sensor have a complex pattern of activity, characterised by smooth changes within the interval $[0.0, 0.5]$. The outer motor is always producing high outputs with small fluctuations. Interestingly, both rectangular and Gaussian estimators de-

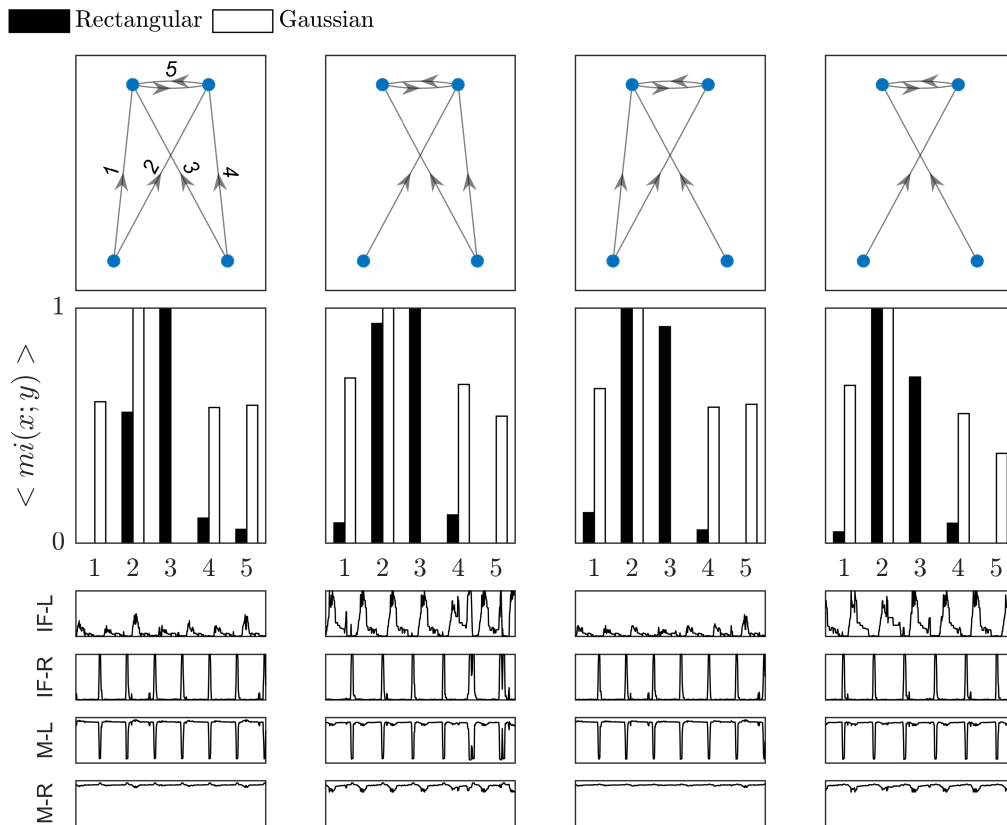


Figure 5.5: Estimation of the local mutual information in pruned neural networks evolved in the 2nd replication. The uniqueness of this neural structure is the necessity of a connection between the sensor directed towards the inner wall and the outer motor actuator. This characteristic is captured by the local mutual information that estimates high degree of information exchange. However, probabilities estimated with different kernels cause inconsistencies on the mapping of the functional neural structure.

to detect a high degree of dependence between the inner sensor and the outer motor, a peculiarity of this replication. The synaptic connection between the left infrared sensor and the left motor stabilises the system, as for the best robot of the 10th replication. In fact, a lesion produces visible fluctuations on the time series recorded from the left infrared sensor and the right motor. By analysing this replication, manifest discrepancies on the probability estimation clearly emerge. Firstly, Gaussian kernel does not detect an association between the outer infrared sensor and the inner motor actuator. Although the synapse is necessary for avoiding collisions, Gaussian kernel is estimating low average values of local mutual information. Secondly, the outcome of the local mutual information calculated with Gaussian kernel shows high values of information exchange




Figure 5.6: Effects of the smoothing parameter on the local mutual information estimated with Gaussian kernel. The formulation of Gaussian kernel includes a parameter α that modifies the bandwidth h or values stored in the bandwidth matrix \mathbf{H} . Bar charts refer to the complete set of all functional neural architectures evolved in the 10th replication, demonstrating that different values of this parameter do not significantly affect the ranking of the local mutual information.

between all connected variables, except for the pair consisting of the right infrared sensor and the left motor, whereas the rectangular kernel detects mutual association between the right sensor and the left motor.

This issue is clarified by inspecting the probability densities reconstructed with different kernels. We limit the description to the univariate case, which is sufficient for explaining substantial inconsistencies in the estimation of the local mutual information. The datasets recorded from observables characterized by spikes of activity or inactivity include a large number of values close to 0 after normalization and long one-sided tails either in the negative range for motors or positive for infrared sensors. Employing kernel density estimation on such datasets leads to significantly different probabilities distributions that depend on the estimator. In fact, Gaussian kernel produces a slope with a narrow and smooth spike centered around 0, whereas rectangular kernel estima-

tes probabilities equal to 0 as values becomes higher or lower than the bandwidth, shaping the distribution as a step function.

A possible solution to limit the differences on the estimation of the local mutual information is the manipulation of the parameter α included in the equation describing the Gaussian kernel, which adjusts the bandwidth h , or the values of \mathbf{H} for multivariate datasets, thus affecting the smoothing of the reconstructed probabilities. However, Figure 5.6 shows that the size of the bandwidth does not significantly affect the ranking of the averaged local mutual information estimated on every functional architecture evolved from the 10th replication. Similar results are obtained on the other nine replications. The question concerning which kernel is the right choice for mapping the information structure does not have a definitive answer. In general, rectangular kernel provides a better mapping of the underlying neural structure. However, Gaussian kernel is more sensitive towards the restructuring of the dynamics of the information exchange between variables caused by synaptic lesions. Furthermore, the choice of the kernel is also task-dependent, as datasets recorded from diverse experimental scenarios may present peculiar patterns of activation whose properties are unveiled by different probability estimators.

5.2.2 Local transfer entropy

In the estimation of the local transfer entropy, several parameters are involved. Firstly, the estimation may be based on the recorded time series or the derivatives. Secondly, different kernels may produce different results in detecting the direction and magnitude of the dynamics that characterise the information transfer. Thirdly, estimating the information transfer on the single-dimensional time series by setting parameters $\tau = 1$ and $d = 1$, instead of embedding the signal recorded from the variable sending the signal, or both time series, may have a significant impact on the reliability and precision of the calculation. Finally, the estimation of the embedding lag τ differs if the average of the local mutual information is employed in place of the mutual information while embedding the recorded dataset. In order to test the best parametrisation of the local transfer entropy, the information-theoretic measure is calculated for all possible combinations of parameters. To have a benchmark we use as case study the connection between the outer infrared sensor and the inner motor, which have been detected as fundamental in every replication by the top-down

lesions applied during testing. Furthermore, the selected connection is generally characterised by high levels of local mutual information when the full architecture is employed, lowering the possibility of non-homogeneous scenarios as robots of different populations, and thus with diverse neural structures, are included in the dataset. The ranking derived from the normalisation of the estimated local transfer entropy do not show patterns related to the underlying functional neural structures or the robots' behaviour. However, the dynamic of the information transfer individuates time steps during which a lesion on the synapse does not cause major failures of the system.

In order to find the best parametrisation, we devise a test aimed at evaluating the reliability of the local transfer entropy in detecting temporal instants where the forward connection linking the infrared sensor to the motor is unnecessary. At the beginning of the test, the robot is placed in the same starting position of the evolutionary process and the architecture without any synaptic lesions is downloaded in the system. The robot navigates the arena and the neural connection is deactivated during time steps characterised by an information transfer from the motor output toward the infrared sensor, which is the direction opposite to the true connectivity. Before lesioning the synapse, the current configuration of the system is saved, recording the robot's position in the maze and the activation of the neurons. After the lesion is applied, the robot acts for a fixed amount of time steps or until the agent collides against a wall. The robots' ability at surviving is tested with several temporal horizons $T = \{1, \dots, 100\}$, that define the maximum lifespan of the robot during the trial. At the end of each trial, the robot is repositioned in the environment and the activation of the neural system is restored according to the values previously recorded. The test is performed for each time step characterized by negative information transfer. In order to assess the exactness of the information-theoretic measure, the test is repeated 50 times with random sequences synthetically generated. The values stored in the vector describing the temporal sequence of the estimated local transfer entropy is shuffled using the Fisher-Yates algorithm. Therefore, the total number of time steps is preserved, while the temporal order is scrambled. The precision of the local transfer entropy is evaluated by comparing the total lifespan of a robot with a controller lesioned according to the dynamic predicted with the local transfer entropy, against the average number of time steps resulting from randomly applied lesions. The limit of the prediction horizon of the local transfer entropy is defined by time $t \in T$ that represent the convergence of




Figure 5.7: Accuracy of the local transfer entropy. The accuracy of the local transfer entropy for describing the information dynamics is tested applying lesions during temporal intervals characterised by information transfers from the motor to the inputs. The dynamic unveiled by the local transfer entropy is compared with surrogate datasets, thus applying lesions during random time steps. The prediction horizon is defined as the convergence of the average lifespan of robots pruned according to the information dynamic and random lesions. The local transfer entropy is estimated with different parametrizations, specifically x and y refer to recorded time series, X and Y are embedded vectors, \dot{x} is the difference of signal. (a) and (b) show the prediction horizon and the percentage of lesions applied during the test for rectangular and Gaussian kernel, respectively. (c) summarises the result with an accuracy measure, based on the number of lesions and the robots' survival capability.

the total lifespan of robots pruned according to the local transfer entropy and random lesions.

To evaluate the accuracy rate of different parametrisations, we create a simple measure $A = (l/t) + (l \times t)$, where l is the percentage of time steps with a lesioned synapse and t is the average time that a robot controlled by a pruned neural network survives before colliding against a wall. The first component of the equation is a ratio that rewards the reliability of the local transfer en-

tropy. The second component accounts for the total number of lesions during the robot's lifespan that the information-theoretic measure detects. The expectation is that a good parametrisation detects the largest number of lesions applied with a correct timing. Figure 5.7 shows that several parametrisation are effective for detecting the absence of information transfer and thus the possibility of lesioning the synaptic connection, during several time steps. Interestingly, rectangular kernel is a better estimator for low dimensional cases. In fact, it outperforms the Gaussian kernel if the time series is not embedded, or the mutual information is utilised for estimating the parameter τ , as the false nearest neighbour algorithm returns on average lower minimum embedding dimension d . In fact, the mutual information estimates the embedding delay $\tau = 1.05 \pm 0.22$ and dimension $d = 3.56 \pm 1.09$ for rectangular kernel and $\tau = 1.57 \pm 1.45$, $d = 3.99 \pm 1.52$ for Gaussian kernel. Embedding parameters based on the averaged value of the local mutual information are $\tau = 3.84 \pm 2.00$ and $d = 6.26 \pm 1.56$ for probabilities reconstructed with rectangular kernel and $\tau = 5.62 \pm 4.53$, $d = 5.45 \pm 1.38$ if Gaussian kernel is employed. Results demonstrates that the most reliable estimation is achieved embedding the recorded signals using the local mutual information for evaluating the parameter τ , and reconstructing the probabilities with Gaussian kernel.

The effect of different parameters on the estimation of the local transfer entropy is shown in Figure 5.8. Although the pattern of the time series describing the local transfer entropy is similar, exhibiting positive spikes corresponding to the corners of the maze, there are clear differences on the representation of the dynamic of the information transfer. The local transfer entropy calculated from distributions reconstructed by a Gaussian kernel on the embedded time series show clear regular spikes, followed by small negative fluctuations with occasional forward transfers of moderate intensity 5.8(a). A similar model of information transfer is reconstructed by employing a rectangular kernel, but a significant lower number of time steps between high peaks of positive information transfer are characterised by activity from the infrared sensor directed towards the motor output 5.8(b). Overall, the evaluation of the dynamic describing the information transfer is significantly less efficient when multivariate probabilities are reconstructed with a rectangular kernel. In fact, small positive values are seldom detected while the agent is navigating following a straight line. Furthermore, there are temporal offsets in the detection of high degrees of posi-

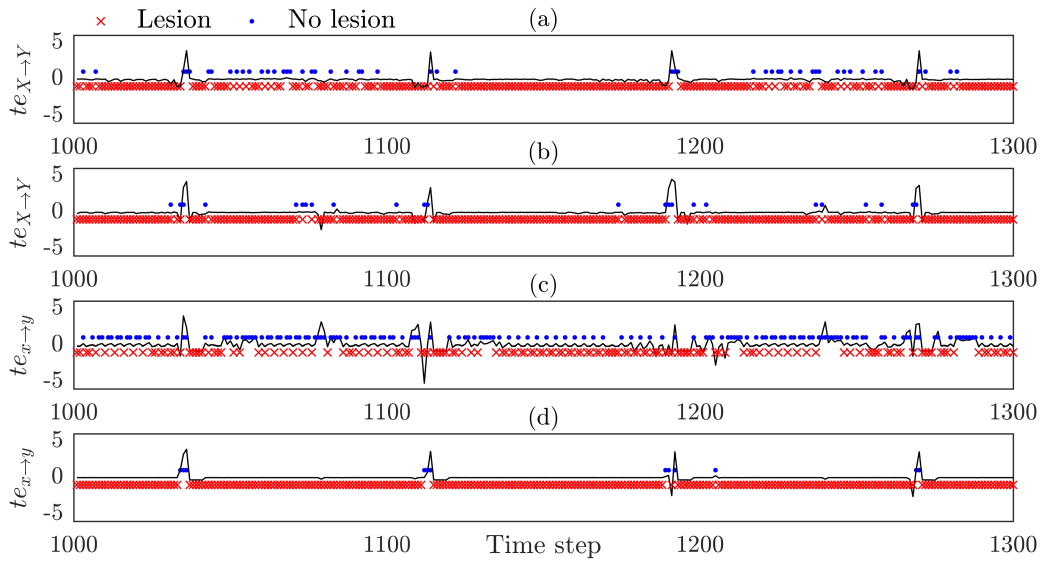


Figure 5.8: Dynamic of the information transfer. The local transfer entropy estimated between the right infrared sensor and left motor actuator for the robot of the 1st replication during the execution of a lap. In (a) probabilities are estimated with Gaussian kernel on embedded time series, whereas in (b) a rectangular kernel is employed. (c) and (d) depict the local information transfer calculated on the recorded signals using Gaussian and rectangular kernel, respectively.

tive information transfer when the robot is approaching the corner. The latter is surely the most important factor, as the cornering behaviour is a more difficult sub-task. For example, failing to activating the synaptic connection in the straight corridor may lead to a collision with the walls after several time steps. Instead, the 90° corner requires an immediate reaction by the robot with a hasty reconfiguration of the motor activation.

The differences in the estimation of the local transfer entropy in relation to the kernel employed are more evident in the low dimensional case, where the time series is not projected onto a multidimensional space. Specifically, Gaussian kernel produces wider fluctuations and the synaptic connection is active for about 50% of the time 5.8(c). However, overestimating the number of time steps defined by a positive information transfer is not useful as the most important factor for lesioning the synaptic connection in real time, while avoiding collisions, is a correct information-theoretic description of the directed temporal dynamic. On the other hand, probabilities estimated with a rectangular kernel are less sensitive towards small variations of the information transfer, detecting only spikes that correspond to the corners of the arena 5.8(d).

As a consequence of the time series embedding, past events of the dynamic stored in the activity of the infrared sensor and the motor output are unveiled. Therefore, the estimation of the local transfer entropy entails a fine grained description of the temporal structure that leads to an activation of the connection with a consistent timing, thus avoiding malfunctions in the system. By inspecting the time series produced by the local transfer entropy, the effects produced by embedding the time series, as well as the kernel used for determining the probability distribution, are clearer. We argue that in low dimensions the precision of the Gaussian kernel does not compensate the lack of knowledge about past events of the dynamic, which is achieved by embedding the time series. Therefore, the local transfer entropy detects fluctuations but places positive values at the wrong time. However, the weakness in the low dimensional case becomes a strength in higher dimensions, where knowledge about the past states of the system are included in the multidimensional embedded time series, which unfolds the dynamics. In fact, the Gaussian kernel is more resilient towards the curse of dimensionality and captures finer structures in the multidimensional data. Therefore, past events that characterised the complex dynamic of the information transfer between sensory input and motor output are detected and exploited.

5.2.3 Mutual information

The mutual information applied on recorded time series identifies whether robots execute the wall following clockwise or counter clockwise, detecting strong mutual dependences between the inner infrared sensor and both motor outputs (Figure 5.9). Similar results are obtained when the mutual information is estimated on time series recorded from robots controlled by lesioned architectures. This result is somewhat unexpected, as signals of the infrared sensor oriented toward the outer wall is characterized by large spikes during the execution of the task, whereas the inner infrared sensor exhibits small fluctuations (Figure 5.1). Furthermore, tests on pruned architectures demonstrate that in every population the most important synapse connects the infrared sensor facing the outer wall of the arena with the opposite motor. Therefore, the mutual information captures characteristics of the robot-environment interaction, rather than the information structure between sensors and motor actuators of the system. The mutual information estimated from the first difference of the recorded time series does not show any clear pattern

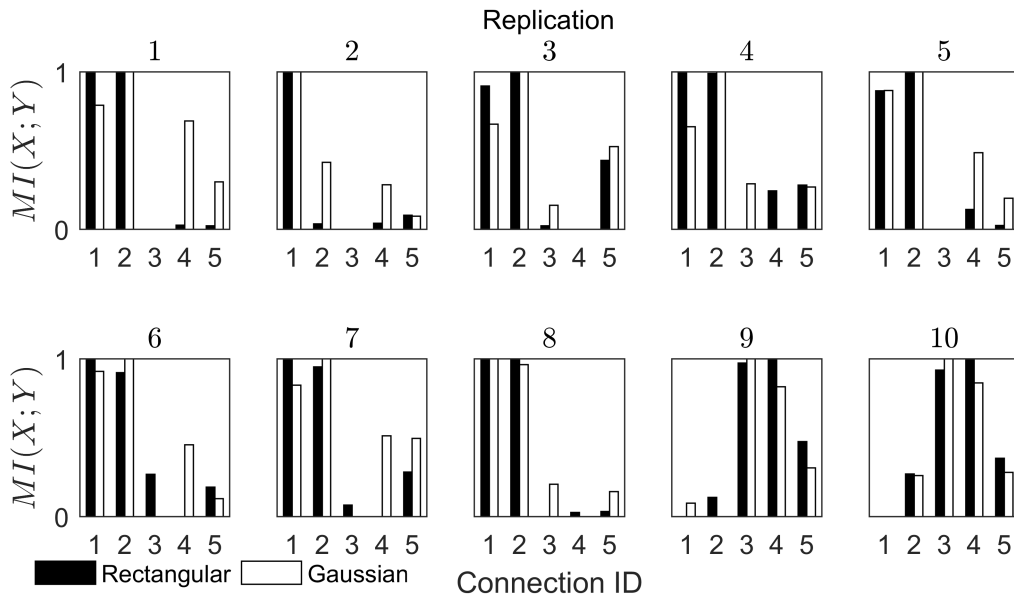


Figure 5.9: Estimation of the mutual information. The dependences between variables unveil characteristics of the agent-environment interaction. The systems display high degrees of dependence between the left infrared sensor and the motor actuators in robots that navigate counter-clockwise. On the contrary, agents that follow the opposite direction are characterised by high degrees of dependence between the right sensor and the outputs.

(data not shown).

5.2.4 Transfer entropy

As expected, the estimation of the information transfer differs significantly depending on the embedding parameters, both in terms of directionality and magnitude. However, the parametrisation that best matches the true functional connectivity estimates the pairwise transfer entropy on the recorded time series, assessing the value of the embedding lag τ using the mutual information (Figure 5.10). Overall, the direction of the information transfer is directed from the motors toward the sensors. However, this result is consistent with the task as in the straight corridors the robot is not actively using the neural network. Indeed, previous analysis based on the local form of the transfer entropy demonstrates that during a significant amount of time steps the robot is not using the neural connections (Figure 5.8). In fact, the navigation is mainly driven passively by the biases, which are always set with positive values by evolution, and small adjustment of the trajectory induced by low activation of the infrared sensors. For every replication, excluding the 2nd




Figure 5.10: Estimation of the transfer entropy. The information flow unveils characteristics of the agent-environment interaction and the underlying functional neural structure. The information transfer is directed from the motors to the inputs as the behaviour is mainly driven passively by the biases. However, the transfer entropy detects the importance of fundamental synapses, which are characterised either by information flow directed towards the output, or the smallest degree of information transfer in the opposite direction. The transfer entropy is calculated on embedded vectors. Positive values on the recurrent connection in the output layer refer an information flow from left to right.

that is discussed later in this section, and with different embedding parameters, the connection linking the outer infrared sensor and the inner motor is characterized by the highest forward or lowest backward information transfer from the input to the output. Furthermore, transfer entropy calculated with rectangular kernel shows an information transfer from the right to left motor for robots that explore the maze counter clockwise and the opposite direction for the two agents that follow a clockwise direction. Interestingly, the transfer entropy detects a high degree of information from the inner sensor towards the outer motor in the robot of the 2nd replication, which is the

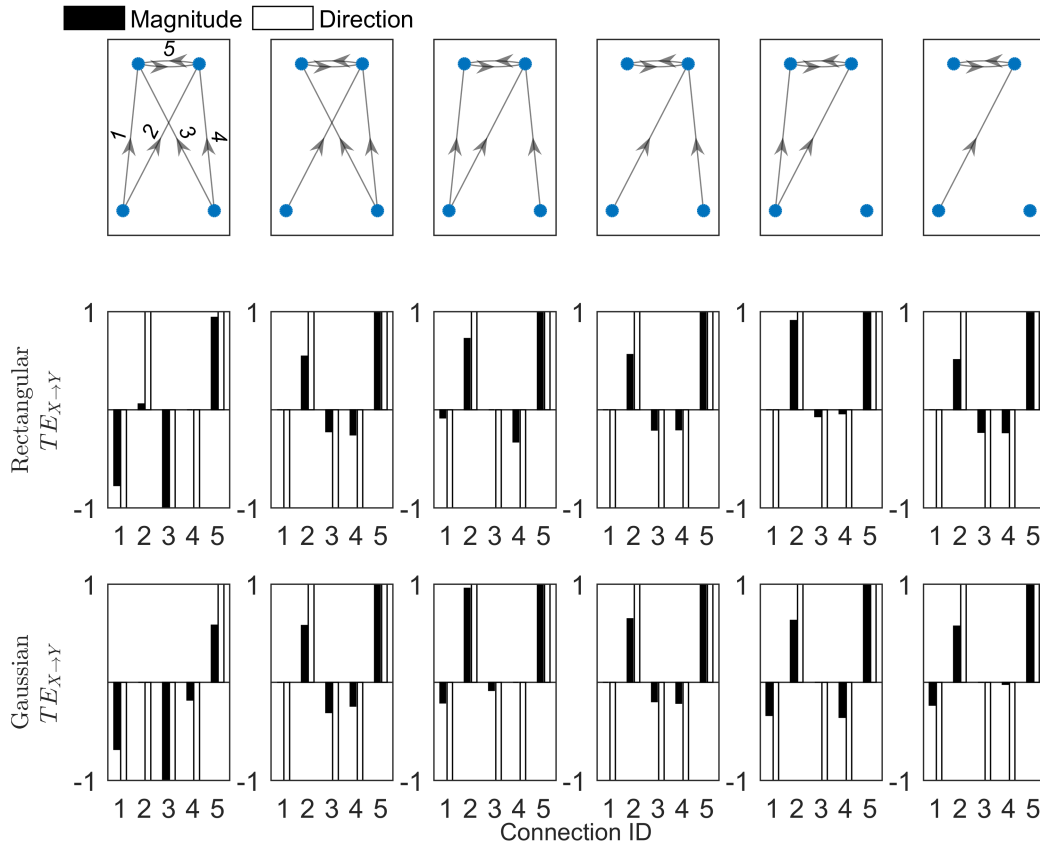


Figure 5.11: Estimation of the transfer entropy in pruned neural networks evolved in the 10th replication. Bar charts show the restructuring of the information flow caused by pruning. The transfer entropy detects an increased activity on the synapse connecting the left infrared sensor with the right motor. This result is consistent with an observation of the robot's trajectory that becomes more unstable, thus requiring active adjustments during the execution of the task.

only agent that explores the maze counter clockwise necessitating this connection. Moreover, the 2nd replication exhibits the most unstable trajectory and the transfer entropy detects an overall higher degree of information transfer directed from the sensors to the motors.

The transfer entropy is also calculated on robots controlled by pruned sub-architectures, as shown in Figure 5.11 for the best robot of the 10th replication. Results further clarify the role of the transfer entropy in describing the behaviour of an embodied and situated system. Firstly, lesioned neural networks require higher degrees of forward information transfer from the left infrared sensor to the right motor, which is the functional backbone of the architecture. This outcome is justified by a more unstable trajectory produced by a pruned architecture, and thus an expected increment

of the activity to compensate erratic behaviours. Another indispensable synapse, the recurrent connection between the two motor actuators, is always transferring high degrees of information from left to right. Secondly, the connection between the left infrared sensor and the left motor is characterised by a high degree of information transfer from the output to the input. Conversely, result from the averaged local mutual information (Figure 5.2), as well as synaptic lesions, demonstrate the importance of a sensorimotor chain on the left-hand side of the robot's body. The outcome of the two information-theoretic measures may appear contradictory, the transfer entropy falsifying the mapping of the information dynamic unveiled by the local mutual information. However, these results underline the fact that the two measures capture different properties of the embodied system. We argue that the transfer entropy unveils both functional relations between variables, similarly to the local mutual information, as well as aspects of the agent-environment interaction. To clarify, we sketch an explanation of the behavioural dynamic. A robot controlled by the full architecture mainly navigates passively, exploiting the inertial movement generated by the positive biases. In the presence of critical input flows, the left infrared sensor triggers a sensorimotor chain that involves both motors. Thus, the robot produces a behavioural response, which is promptly detected by the left infrared sensor. However, the configuration of the system brain-body-environment is actively modified by the agent, explaining the high degree of information transfer directed toward the sensors, which is affected by the robot's movement. The transfer entropy reveals that there is a significant relationship between the two variables when the connection is active, but the strongest driving force of the sensorimotor loop is, in this case, the agent-environment interaction. In fact, when the synapse is lesioned the transfer entropy detects low values of information transfer between the sensor and motor placed on the left side of the body. On the contrary, sub-architectures with an active forward connection from the left infrared sensor to the right motor actuator show higher degrees of information transfer, but lower than the full architecture. However, low values of transfer entropy in the synapse connecting the left-hand side sensorimotor chain are explained by the high levels of information transfer from the outer sensor to the inner motor in robots mounting pruned neural networks, to compensating the instability of the systems.

5.2.5 Decomposing the behaviour via information-theoretic measures

In this section we develop a behavioural description of selected agents from an information-theoretic perspective. In what follows, we perform a thought experiment, assuming that the robot's behaviour has not been observed, thus attempting to infer the characteristics of the system from the results of the information-theoretic analysis. The rankings describing the interaction between variables of the system defined with the pairwise estimation of the information-theoretic measures unveil meaningful functional dynamics of the sensorimotor network, as well as characteristics of the agent-environment interaction. In fact, from the results of the information-theoretic analysis, we may hazard an *ansatz* about the functional connectivity of the underlying neural structure and the behaviour exhibited by the embodied and situated systems. Figure 5.12 shows the dynamics of the local information transfer while robots are completing a lap of the maze. In order to compare the magnitude of the information transfer between different pairs of variables, the values of the local transfer entropy are normalised in the range $[0; 1]$ across the five time series. The differences on the dynamics of the information transfer are evident and are consistent with different behavioural strategies. Furthermore, the best robot mounting the minimal functional neural architecture evolved in the 10th replication is included in the analysis, in order to demonstrate the capability of information-theoretic measures of unveiling minor modifications in the system.

The direction followed by the robots is deduced from the mutual information. Figure 5.9 shows high degrees of mutual dependence between the left infrared sensor and the motors for the 2nd robot, an information structure characteristic of robots that navigate counter-clockwise. The outcome of the transfer entropy and the average values of the local mutual information offer insights on the underlying functional neural architecture. Both Gaussian and rectangular kernels identify strong information exchange between the left infrared sensor and the right motor (Figure 5.5). However, the local mutual information presents inconsistencies on the relationship between the right infrared sensor and the left motor that depends on the estimator employed. Therefore, we inspect the results of the transfer entropy (Figure 5.10), which individuates a strong forward information transfer from the left infrared sensor towards the right motor, as well as a weak forward interaction between the right sensor and the left motor. The transfer entropy calculated with




Figure 5.12: Estimated local transfer entropy and robot's trajectory. Graphical representation of the information dynamics related to the robot's spatial position. Red circles indicate information transfers from the sensors to the motors, or a left to right direction for recurrent connections in the output layer. The size of the marker is proportional to the magnitude of the estimated local transfer entropy. Blackened environments (ID1, ID3 and ID4, bottom row) refer to lesioned synapses.

rectangular kernel detects strong positive information transfer on the left-hand side sensorimotor chain that is not confirmed neither by the Gaussian estimator, nor the averaged local mutual information. Thus, we advance a hypothesis about the functional connectivity of the neural network, consisting of 2 synapses connecting the infrared sensors with the motors placed on the opposite side of the body and the necessity of a connection in the output layer.

Besides relationships between variable of the system, the transfer entropy discloses characteristics of the agent-environment interaction. In the previous section, we notice a higher degree of information transfer from the inputs to the outputs in robots controlled by pruned architecture

that compensates instabilities of the behaviour. Therefore, we infer that the 2nd replication evolved a behavioural strategy characterised by numerous adjustment of the trajectory during the execution of the task, hypothesis confirmed by inspecting the robot's positions in the environment (Figure 5.12). In fact, the robot steers towards the inner wall after completing the 90° left turn, which leads to fluctuations on the information transfer of moderate intensity. Furthermore, 3 forward connections linking the input layer with the motor actuators display fluctuations synchronised with the critical sub-task that consist of a sharp turn. The remaining forward synapse, which connects the right infrared sensor and the right motor, shows spikes of information transfer while the robot is exiting the bottom-right corner and during the execution of the left turn in the upper-right corner. Interestingly, the local transfer entropy between the two motors displays low values. However, this result is consistent with the data analysis reported in the previous section that demonstrates the importance of the directionality of the information transfer during the temporal evolution of the dynamic, rather than high magnitudes (Figure 5.8). Moreover, the agent is controlled by the full architecture and thus adjustment in the trajectories may be offloaded to other synaptic connections.

The unimportance of estimating large values of local information transfer is confirmed with the analysis of the 10th replication, either controlled by the neural controller without synaptic lesions, or the minimal functional architecture consisting of 3 connections. In fact, Figure ?? demonstrates that the dynamic of the information flow between the left infrared sensor and the right motor is characterised by small fluctuations of information transfer. Nevertheless, results from the averaged local mutual information and the global transfer entropy identify the importance of a synapse linking this sensorimotor chain. Considering the other synaptic connections, the outcome of the transfer entropy (Figure 5.11) is aligned with the functional connectivity deduced from the degrees of information exchange (Figure 5.4). The structure of the information flow in the robot controlled by the full architecture is characterised by significant amounts of global information transfer directed from the outputs to the inputs, which is a fingerprint for a regular trajectory determined by the biases and small behavioural adjustments. However, the global transfer entropy detects the importance of the synapse connecting the left infrared sensor with the right motor. In fact, Gaussian kernel estimates negative information transfer but this connection is characterised

by the weakest magnitude. Moreover, the global transfer entropy calculated with rectangular kernel describes this synapse as the only forward connection characterised by an information flow directed towards the motor neuron.

The hypothesis regarding the functional neural structure and the dynamics of the agent-environment interaction are confirmed by the outcome of the transfer entropy in pruned neural networks. Firstly, the system displays high degrees of information transfer from the left sensor to the right motor, underlining the importance of a direct channel of communication (Figure 5.11). Secondly, the agent controlled by the minimal sub-architecture, which comprises only the artificial synapse connecting the left infrared sensor with the right motor output, as well as the two lateral connections of the output layer, shows small fluctuations on the local transfer entropy between the left input sensor and the right motor (Figure 5.12, bottom row). Clearly, describing a correct timing for the information directed toward the output is more important than estimating large fluctuations in the dynamics of the information transfer. In fact, this neural connection is the only active between the inputs and the motors.

Interestingly, significant differences emerge from the information-theoretic decomposition of the behavioural strategies evolved in different populations. In fact, the 10th replication mounting the full architecture is characterised by large spikes of local information transfer from the left motor to the right motor while the robots is turning right in the 90° corners of the maze (Figure ??, middle row). Furthermore, the dynamic of the local transfer entropy between the left infrared sensor and the left motor actuator exhibits a large positive spike followed by high degrees of negative local information transfer while entering in the the corners. The clockwise direction followed by the robot is deduced from the results of the mutual information, which correlates the left infrared sensor with the motor actuators (Figure 5.9). The local transfer entropy between the left infrared sensor and the right motor exhibits small fluctuations while the robot is crossing the straight corridor and small spikes in the corners of the maze. Undoubtedly, the structures and dynamics of the information differ significantly compared with the robot evolved in the 2nd replication.

The capability of capturing minor behavioural differences exhibited by robots is evident by comparing the best agent of the 10th replication mounting the full architecture or the minimal functional

neural network. Results of the global transfer entropy between variables of the system is consistent with the mapping of the information exchange unveiled with the local mutual information. Furthermore, results reported in Figure 5.4 depict high degree of dependence between the sensor and motor placed on the left-hand side of the robot if the connection is active. In particular, functional neural sub-architectures that include this synapse show the absence of fluctuations on the time series recorded from the left infrared sensor and the right motor. The role of a sensori-motor chain on the left part of the body is now explained by the dynamic of the information flow. By inspecting the time series representing the values of the local transfer entropy, the presence of a connection from the left infrared sensor to the right motor stabilizes the trajectory on the initial part of the corner. A final remark concerns the lesioned synapses. Although these connections exhibit complex dynamics of the local information flow, characterised by significant amounts of large spikes, the averaged values of the local mutual information (Figure 5.4) and the outcome of the global transfer entropy (Figure 5.11) detect their irrelevancy.

5.3 Discussion

The endeavour attempted in this chapter is to examine the explanatory power of different information-theoretic measures during the analysis of embodied and situated systems. The testbed is a neuro-robotic model that autonomously adapts to the environment mimicking the process of evolution. The decision of employing this class of adaptive systems is justified as follows. Firstly, a major principle guiding the evolutionary robotics framework is an antireductionist approach, where the solution to a problem defined by a global utility function is achieved in a bottom-up fashion from the interaction of the robot with the surrounding environment, thus avoiding *a priori* analytical decomposition of the system to achieving a top-down solution. Thus, the experimenter has no or little knowledge about the functioning of the systems, which may bias the subsequent information-theoretic analysis. Secondly, neuro-robotic systems trained with semi-supervised learning algorithms are a sound operationalisation of embodied and situated view of cognition. In fact, the system brain-body-environment is conceived as a holistic system, where the agent's body and brain are strictly coupled with the surrounding environment. Boundaries that divide concepts such as inner and outer, neurons and sensors, as well as agent and physical world

are blurred and tend to dissolve into the unity of a system perspective, where sub-components interact in a dynamical and nonlinear way during the learning process. Finally, a neuro-robotic system is characterised by concrete instantiations of communication channels, the artificial synapses, which are manipulated in order to verify the hypotheses formulated about the map of interactions between variables.

The application of information-theoretic measures leads to a decomposition of the embodied and situated system, without rejecting the antireductionist perspective. In fact, the system is analysed without disentangling the coupled sub-components and the network of relationships are unveiled by confronting and relating the whole set of possible interactions among variables that form the system. The outcome of the different information-theoretic measures estimated on all possible couples of variables that are connected by a synapse is normalised, creating rankings of interaction. Similarly, Bauer et al. (2007) summarises the results of the transfer entropy creating causal maps of industrial chemical processes. To create such map, variables are ordered in a binary matrix, where rows indicate causal variables and columns effect variables. Pairings that show a positive information transfers above a certain threshold are marked with 1, thus creating a matrix of causal interactions based on pairwise estimations of the transfer entropy. However, we do not set minimal values for assessing the significance of specific information transfers. Rather, normalisation creates a weighted ranking according to the magnitude of the information-theoretic measure, where 0 corresponds to the weakest and 1 to strongest interaction, thus mapping the entire network of relationships among variables of the system.

Results demonstrate that each information-theoretic measure unveils different characteristics of the underlying functional neural structure and the agent-environment interaction. In particular, mutual information is a useful behavioural descriptor, discriminating between the two classes of robots that navigate clockwise or counter-clockwise. On the other hand, the local formulation of the mutual information, which describes the dynamic of the information exchange between two random variables, provides an evaluation of functional relationships within the network of perceptual and motor variables. In the neuro-robotic systems described in this article, such relationships have a concrete instantiation with the artificial synapses, and thus communication between chan-

nel may be interrupted. Therefore, we move from a passive description of the system to an active verification of the results obtained with the information-theoretic analysis.

The transfer entropy is a time-asymmetric measure that quantifies the directed information transfer between two variables. This measure is particularly interesting if applied to embodied in situated systems, characterised by a continuous flow of sensorimotor loops. In fact, results reported in this article demonstrate that transfer entropy serves two purposes. Firstly, this information-theoretic measure unveils functional relationships concerning the variable of the system by detecting synapses that are fundamental for the execution of the task. Secondly, the transfer entropy captures dynamics of the agent-environment interaction. Thus, the mapping of the information flow among perceptual and motor variables describes functional characteristics of communication channels, similarly to the dynamics captured by of the local mutual information, conditioned however by the continuous interaction of the robot's body with the physical world. In fact, during the straight navigation in the corridors of the maze, the robot's body is moved forward by the positive values of the biases connected to the motor actuators. Consistently, the information flow is mainly directed from the actuators to the perceptual system, depicting the agent as passive. However, the transfer entropy detects synapses linking perceptual and motor neurons that are fundamental for the functioning of the systems. Furthermore, agents that exhibit unstable trajectories caused by a pruning of the neural network or sub-optimal evolution are characterised by larger values of information transfer directed from sensors to motor actuators if the channel of communication is utilised by the robots during the execution of the task. The interpretation of this result is that in the presence of an ordered system, characterised by stable trajectories, relevant connections intervene for steering in the corners of the maze and are mainly unused in the corridors, explaining the overall low magnitudes of information transfer. On the contrary, disordered systems requires a flow of information directed from the sensors towards the motors, in order to actively adjust the trajectory. Consistently, irrelevant synapses are described with large amounts of negative information transfer, as these connections do not actively influence the behaviour either in the straight part of the maze or during the turns.

The local transfer entropy is the only information-theoretic measure that does not unveil functional

relationship among sub-component of the system if the outcome is evaluated with the normalised ranking. However, the local transfer entropy detects the necessity of an active neural connection between two variables during specific time steps. This measure identifies time steps characterised by an information transfer from specific perceptual inputs towards a motor output, that is, the necessity of a synapse during specific temporal intervals. The effect of lesions applied on these synapses according to time steps characterised by output-input dynamics of information transfer are compared with randomly generated sequences of pruning, demonstrating that the local transfer entropy detects temporal instants where the presence of a specific neural connection is fundamental for exploring the environment avoiding collisions with the walls.

By combining the results of different information-theoretic measures, the experimenter may infer characteristics of the robot's behaviour and peculiarities of the neural controller without observing the system during the execution of the task. Furthermore, modifications of the information structures and dynamics caused by synaptic lesions are mapped and subsequently confronted with the description of fully functional agents, consolidating the hypotheses.

5.4 Conclusions

In this article, we investigate the explanatory power and impotence of information-theoretic measures applied to embodied and situated systems, operationalised with an evolutionary robotics model executing a wall-following task. Information theory provides a set of nonparametric and model-free statistical tools that unveils nonlinear relationships between time series recorded from perceptual and motor neurons of the agent. The aim is to assist and complementing with numerical characterisations the purely geometrical description of such systems offered by a dynamical system approach. The analysis strictly follows an antireductionist perspective, avoiding a disentanglement of the integrate system consisting of an agent controlled by an artificial brain that interacts in real time with the environment. In accordance with recent foundational developments in system biology, the whole network of interactions among variables of the system is analysed to capturing functional relationships. We adopt two measures, namely mutual information and transfer entropy, which investigate the interactions between pairings of variables connected by an artificial synapse. The mutual information measure the degree of dependence between

observables. This measure captures characteristics of the agent-environment interaction, discriminating between robots that navigate in the maze clockwise or counter-clockwise. The transfer entropy, instead, is time-asymmetric and extends the concept of mutual information determining the magnitude and direction of information transfer between two variables, thus revealing the complex network of causal interactions among the various sub-components of the system. Results demonstrate that the transfer entropy describes features of the robot's interaction with the surrounding environment, as well as the information flow within the neural system. These estimators describe the structure of mutual dependences and information flows over the entire robots' lifespan. To investigating the information dynamics, the local form of such measures is estimated on recorded time series, thus unveiling characteristics regarding the temporal evolution of the system from an information-theoretic perspective. The local mutual information quantifies the information exchange, determining to what extent two variables inform or misinform each other during the execution of the task. Tests conducted on robots with pruned neural architectures demonstrate that the estimation of high degrees of dependence between neurons correlates to the underlying functional neural structure, detecting synapses that are fundamental for executing the wall following. The local transfer entropy uncovers the importance of synaptic activity during specific temporal intervals, discriminating between active perceptual-motor loops and passive inertial motion produced by other parts of the system. Notwithstanding the limits of estimating such measures on pairings of variables rather than using multivariate extensions of mutual information and transfer entropy, as well as inconsistencies caused by different probability estimators, the information-theoretic tools are capable of decomposing the integrated embodied and situated agent, determining the functional relationships among the sub-components of the system.

Chapter 6

Conclusions

This thesis investigates the general problem of understanding embodied and situated systems from a dynamical system perspective. The argument develops following three levels of explanation. Firstly, providing a theoretical clarification concerning the concept of embodied cognition. Secondly, supporting an antireductionist philosophy of science for analysing such cognitive systems. Finally, applying appropriate nonlinear integrative mathematical tools during the analysis of data collected from neuro-robotic models.

The scientific questions proposed in the introductory chapter of the thesis have been consistently addressed, demonstrating the potential of a rigorous quantitative, data-driven and dynamical approach to the study of embodied and situated systems. In chapter 4, a large number of trajectories are recorded from different robots that react to environmental temporal dynamics. The chaotic measures calculated on attractors reconstructed by embedding the vectors storing the robots' position are useful numerical descriptors that correlate with the capability of adapting to unpredictable environments. The experiment is a significant step forward on the analysis of autonomous systems, demonstrating the explanatory power of chaos theory applied to empirical data recorded from robotic systems. In chapter 5, miniature mobile robots execute a wall-following and time series record the neural activity. During data analysis, two information-theoretic measures are applied to unveil aspects of the information structure that characterise the embodied and situated systems. Firstly, mutual information is estimated on each pair of neurons connected by a synapse to evaluate the degree of mutual dependence between neurons. Secondly, direction and magnitude of the information flow is reconstructed employing the transfer entropy for revealing the complex networks of underlying causal structures. The novelty introduced in this chapter is the estimation of the local form of such measures to assess the dynamic of the information

structure and flow while the robots are executing the task. Furthermore, the hypotheses and inferences derived from the analysis are verified by lesioning the artificial neural systems controlling the robots. Results show that each measure captures different relational aspects of the agent-environment interaction and functional neural connections.

A major issue affecting the field of embodied cognition is the lack of a widely accepted definition (Cangelosi et al., 2015), which causes a broad spectrum of experimental operationalisations. This problem is particularly critical for cognitive modellers and roboticists that study artificial cognitive agents that are built by the experimenter. Thus, a precise definition guiding the design of the model is crucial. Although the foundational traits of the field are still vague, the different explanations of the related concepts of embodied, grounded and situated cognition proposed in literature share some similarities and common meanings (Fischer, 2012; Myachykov et al., 2014). Embodied cognition is centred on the subjective experience of a physical brain placed in a physical body. Particular attention is directed towards the agent, member of a population of conspecifics. In fact, the body, a product of phylogenetic evolution throughout generations, shapes and constrains the cognitive abilities of the agents. Furthermore, the history of the ontogenetic development that characterises a specific individual affects the network of concepts stored in the brain. The term grounded cognition, instead, refers to immutable physical properties of the environment, whose regularities are perceived by the agent, structuring concepts and knowledge. From this perspective, the sensorimotor experience of the world is mainly a source of statistical invariants that affect the neuronal ensembles in the brain. Synthetic models are often aligned with the connectionist tradition, where the input-output mappings are concretised with linguistic labels, often abstracting away from the real time and dynamic of the agent-environment interaction. The last concept, situatedness, is the more problematic and entails at least two meanings. Firstly, situated cognition refers to the contextual information available in the world, including the physical environment and the social experience, in a precise spatiotemporal setting (Barsalou, 2008). Secondly, situatedness is defined as the continuous flow of sensorimotor loops experienced by the agent while interacting with the environment, a position aligned with the dynamical system approach to cognition (Beer, 2008). The experiments presented in this dissertation are based on the latter definition, thus stressing the importance of real-time and dynamical exchanges between

the robot and the external world.

By delimiting clear theoretical boundaries, the idea of embodied and situated cognition is operationalised employing evolutionary robotics (Nolfi & Floreano, 2000). This technique is based on neuro-robotics platforms trained with a learning process that models biological evolution. The agents autonomously adapt in bottom-up and semi-supervised fashion to the environment and the continuous sensorimotor interaction is a central element for developing cognitive abilities.

From this perspective, an analysis of embodied and situated systems solely based on a reductionist approach does not suffice. In fact, embracing a dynamical system perspective leads inexorably to conceiving the agent, brain and environment as a highly integrated and indivisible system. Cognitive facts are the product of a complex and nonlinear network of interactions among subsystems and a radical reformulation of the analytical methodology is a necessary condition for understanding embodied and situated cognitive agents. Recently, several scientific disciplines are reconsidering the centrality of reductionism, proposing a holistic and antireductionist philosophy of science (Ahn et al., 2006b,a; Fang & Casadevall, 2011). Hence, the focus is shifting from approaches based on a *divide et impera* stance, which examines specific parts of a biological phenomenon in isolation, towards the study of global properties of the system and the nonlinear interactions among its variables. The experiments reported in this thesis attempt to applying these concepts to evolved robotic platforms. An important principle guiding the experimental design is to avoiding the manipulation of macroscopic variables, such as robots' morphologies or major changes in the environment, in order to test the capability of nonlinear tools to detect minor modifications in the trained systems

The first work describes a scenario where robots are involved in a temporal task, performing a reaching-avoiding behaviour depending on the colour of a target area that switches between blue and green at regular intervals. The only environmental variables controlled by the experimenter are the lengths and the number of different temporal intervals elapsing between changes of colour in the target area experienced during the training process. A second variable that is modified concerns the neural controller, as the experiment is replicated with static and recurrent neural networks. At the end of the evolutionary process, the best robots are tested on environ-

ment characterised by target areas characterised by novel temporal dynamics never experienced during learning. The systems are analysed with mathematical tools developed in the field of chaos theory (Kaplan & Glass, 2012). By following this approach, the experimenter avoids an *a priori* decomposition of the system into sub-components, which are subsequently analysed in isolation. Rather, chaotic measures reveal and quantify global properties of the integrated system brain-body-environment. The robot's trajectory is recorded and utilized for reconstructing the attractor describing the underlying dynamics with the technique of delayed embedded vectors. Subsequently, chaotic measures are estimated on the reconstructed attractors that exhibit hallmarks of low-dimensional deterministic chaos. In particular, two dynamical measures are estimated on suitable robots: the level of chaos for estimating the long term unpredictability of the behaviour and the fractal dimension, which determines the number of degrees of freedom in the system. The results consistently demonstrate the utility of chaotic measures for unveiling properties and characteristics of the evolved robots, showing a correlation between the robots' ability at adapting to environments never experienced during the learning process and chaotic measures. Furthermore, robots that exhibit a chaotic behaviour are more resilient towards unpredictable environments if compared to agents that do not exhibit a chaotic behaviour. Therefore, the experiment supports the applicability of chaos theory to empirical data, including the class of embodied and situated systems.

The second experiment investigates the nonlinear relationships among variables of the system by applying information-theoretic measures (Shannon & Weaver, 1949) on robots performing a wall-following task inside a squared maze. The neural controller is based on a minimal architecture, where two perceptual neurons send signal to two motor actuators. The output units are connected with two lateral connections, thus creating a neural network that consists of six synapses. Similarly to the previous experiment, there is little intervention from the experimenter in the manipulation of variables. In fact, the environment and the neural architecture are invariant, and the dissimilarities that characterise different populations of robots are epigenetic phenomena. The activity of the perceptual and motor neurons is recorded while robots are exploring the maze and pairwise estimations of mutual information and transfer entropy are calculated between observable connected by a synapse. Although measures derived from Shannon entropy are not directly

related to an antireductionist framework, bivariate estimates are calculated on all the pairings of variables in order to numerically characterise the whole network of interactions. During the analysis, the degree of dependence between pairings of variables is estimated with the mutual information. The direction and magnitude of the information flow is instead captured by the transfer entropy. The analysis is further extended employing the local form of the mutual information and the transfer entropy, which capture the dynamic of the information during the execution of the task. The hypotheses derived from the information-theoretic analysis concerning the underlying functional neural structures are verified by pruning the evolved neural controllers. Interestingly, the local mutual information, as well as the global and local transfer entropy, discriminate synapses that are fundamental to avoid malfunctions in the system. Furthermore, the mutual information and the transfer entropy captures peculiarities of the robot-environment interactions, such as the direction followed by the agent during the exploration of the environment, or discriminating between agents that actively modify the trajectory exploiting the perceptual inputs from robots that are passively driven by the activation of the biases. The local transfer entropy unfolds the dynamics of the information flow between sensors and motors, thus discriminating temporal intervals that requires an active processing of the input signal to adjust the robots' trajectory.

Overall, the questions proposed in the initial part of this dissertation find an answer. Surely, the application of nonlinear quantitative tools to embodied and situated systems is an alternative and complementary framework to a static and structural explanation of the neural networks, or a qualitative and geometrical description of the evolution of the whole system. In particular, chaos theory applied to empirical data is a valuable numerical descriptor that captures characteristics of the underlying dynamics. Furthermore, chaotic systems may be created with coupled equations, thus complementing the bottom-up approach with a top-down modelling while studying a cognitive system.

The application of information theory covers the second fundamental aspect of an antireductionist approach to science, that is, the study of nonlinear relationships among heterogeneous sub-components that constitute a system. The information-theoretic measures are also less problematic, as strong assumptions regarding the dimensionality or the type of dynamics are

unnecessary. In fact, the approach is totally model-free and data-driven, based solely on probabilities reconstructed from recorded data. The analysis conducted on a minimal cognitive model demonstrates that information theory is a robust tool for unveiling functional neural structures and for explaining characteristics of the agent-environment interactions. In the following section I underline the current limits of the proposed approach, suggesting possible extensions to the experimental work described in this thesis.

6.1 Current limits and future works

The experiments described in this thesis are a testbed based on minimal cognitive models for the possible applications of nonlinear time series analysis to embodied and situated systems. Minimalism in the experimental design is necessary while assessing the explanatory power of an innovative approach for the study of embodied and situated systems. However, such simplicity permits stress tests on novel, and sometimes criticised, mathematical tools and approaches to scientific investigation. Thus, the experiments conducted in this thesis are explorative and commences a direction of research that surely requires more complexity to prove and confirm the reliability of nonlinear models of scientific explanation. Some limits that are affecting the state of the art in nonlinear time series analysis are evident, for example the controversial topic concerning the assumption of determinism that affects the field of chaos theory, or the absence of robust probability estimators in high dimensional spaces that are necessary for estimating information-theoretic measures. In this regard, embodied and situated neuro-robotics systems may provide interesting case studies for developing and testing future advancements in the field of nonlinear time series analysis. In fact, such models extend the existing methods, complementing the abstraction synthetically generated datasets and the impossibility of manipulating biological observables.

Therefore, this work has to be considered as the starting point for future research, listed below:

- Chaos theory and information theory are somewhat related. The level of chaos quantifies the sensitivity of the system to the initial conditions, limiting the long term prediction of the evolution of the dynamics. Similarly, the entropy captures the disorder of the system. The

two metrics and approaches clearly share similarities and a system analysis based on both frameworks, deterministic and stochastic, may create synergies to further explain properties and characteristics of embodied and situated systems.

- Another possibility is the creation of chaotic robots controlled by chaotic nonlinear dynamical system, thus following a top-down approach. In parallel, the same cognitive phenomena may be studied with autonomous systems, followed by a data-driven chaotic analysis. This approach combines, in the context of chaos theory, an *a priori* analytical decomposition and an emergentist framework, which may lead to interesting insights on the study of cognitive phenomena.
- Information-theoretic measures may find an application to systems characterised by higher degree of complexity, for example humanoid robots or swarms. These scenarios are useful for developing and testing multivariate information-theoretic measures.
- The nonlinear mathematical tools employed in the experiments reported in this thesis may apply also to biological data. An interesting perspective may stem from interdisciplinary experiments, where a cognitive fact is studied in natural and artificial systems, proposing a common background for data analysis.

The aim of the suggested future directions is twofold. Firstly, the proposed approach may advance the field of the embodied and situated view by providing a solid shared mathematical background for studying the complexity of cognition. From this perspective, novel experiments based on an integration of biological and artificial agents studied with a common methodological and analytical framework, capable of capturing the intrinsic nonlinearities of highly integrated systems, may cause significant advancements in the understanding of cognition. In fact, the possibility of sharing, comparing and communicating results amongst different disciplines is of primary importance for a scientific field heavily based on an interdisciplinary effort. Secondly, cognitive phenomena are investigated by several fields of research, ranging from psychology to neuroscience to robotics, thus providing a common ground for extending the existing set of nonlinear mathematical tools starting from an empirical perspective.

Information theory

Results reported in this thesis underline the importance of the kernel utilized during the estimation of probabilities from the recorded dataset. This issue is particularly evident in a simulated wall-following scenario, where time series are characterized by bursts of activity or inactivity, followed by relatively long periods characterised by small fluctuations. In this work, we restrict our investigation to the kernel density estimation approach, testing the effect of rectangular and Gaussian kernels, two widely employed distributions in the existing literature (see for example (Kaiser & Schreiber, 2002; Bauer et al., 2007; Wibral et al., 2014)). However, other kernels, e.g. triangular or Epanechnikov, may be valid candidates for the reconstruction of the probability distributions. These methods are model-free and data-driven and thus each kernel may disclose different properties of the recorded datasets. Furthermore, alternative methods to the kernel density estimation have been proposed. Kraskov et al. (2004) developed an adaptive partitioning method based on the estimation of k -nearest neighbours. The peculiarity of the algorithm is that the neighbouring space surrounding the points is defined by entropies calculated with Kozachenko-Leonenko technique. Although the procedure is originally developed for the evaluation of the mutual information, and thus employed for evaluating a maximum of 2 joint random variables, the authors discuss the robustness of the estimator in higher dimensional spaces. In fact, the method is utilized for the estimation of multivariate probabilities distributions employed in the formulation of the conditional mutual information (Frenzel & Pompe, 2007; Gómez-Herrero et al., 2015) and transfer entropy (Gómez-Herrero et al., 2015). A similar approach is followed by Lee et al. (2012), extending Darbellay-Vajda algorithm to 3 dimensions for achieving an adaptive partitioning of the sample space. Results demonstrate that method outperforms Gaussian kernel density estimation and fixed binning in mapping the information transfer with a benchmark test based on a biomedical dataset. However, the major drawback of this technique is the restriction to 3 dimensions, and thus this probability estimator is inapplicable to formulations of the transfer entropy based on embedded time series.

Another issue that limits the data analysis presented in this chapter is the absence of a direct statistical test to assessing true relationships between variables. These methods are mainly

based on surrogate data testing, where scrambled time series are generated from, and compared with, the recorded signal. The surrogate time series preserve the same statistical properties of the recorded observable, but correlations between the variables X and Y are disrupted (see for example (Vicente et al., 2011; Wibrál et al., 2014; Lizier et al., 2011)). However, the possibility of performing lesions in the neural architecture renders the detection of spurious correlations less stringent, justifying the decision of omitting surrogate data testing.

Moreover, the experiment is centred on direct relationships between pairs of variables and the maps of interactions are constructed with a normalised ranking. An extension of the experiments described in this article may focus on the evaluation of multivariate versions of the information-theoretic measures for describing complex sensorimotor and neural chains, revealing how ensembles of time series are interconnected (Verdes, 2005; Frenzel & Pompe, 2007; Flecker et al., 2011; Lizier et al., 2011; Williams & Beer, 2011; Runge et al., 2012b,a; Gómez-Herrero et al., 2015; Runge et al., 2017). However, this experiment aims at investigating the descriptive power of different information-theoretic measures estimated on neuro-robotics autonomous systems, rather than capturing the fine grained networks of complex causal chains of interactions. Therefore, limiting the information-theoretic measures to pairwise estimations is, to some extent, justified.

Bibliography

- Ackley, D. H., Hinton, G. E. & Sejnowski, T. J. (1985). A learning algorithm for Boltzmann machines. *Cognitive Science*, 9(1), 147–169.
- Ahn, A. C., Tewari, M., Poon, C.-S. & Phillips, R. S. (2006a). The clinical applications of a systems approach. *PLoS Med*, 3(7), e209.
- Ahn, A. C., Tewari, M., Poon, C.-S. & Phillips, R. S. (2006b). The limits of reductionism in medicine: could systems biology offer an alternative? *PLoS Med*, 3(6), e208.
- Arbib, M. A. (2008). From grasp to language: Embodied concepts and the challenge of abstraction. *Journal of Physiology Paris*, 102(1-3), 4–20.
- Ashby, W. R. (1956). *An introduction to cybernetics*. Chapman & Hail Ltd., London.
- Axelrod, R. M. (2006). *The evolution of cooperation*. Basic books.
- Babloyantz, A. & Destexhe, A. (1986). Low-dimensional chaos in an instance of epilepsy. *Proceedings of the National Academy of Sciences of the United States of America*, 83(10), 3513–3517.
- Barahona, M. & Poon, C. S. (1996). Detection of nonlinear dynamics in short, noisy time series. *Nature*, 381(May), 5–7.
- Barsalou, L. W. (2008). Grounded Cognition. *Annual Review of Psychology*, 59(1), 617–645.
- Basso, K., Margolin, A. A., Stolovitzky, G., Klein, U., Dalla-Favera, R. & Califano, A. (2005). Reverse engineering of regulatory networks in human B cells. *Nature Genetics*, 37(4), 382–390.
- Bauer, M., Cox, J. W., Caveness, M. H., Downs, J. J. & Thornhill, N. F. (2007). Finding the direction of disturbance propagation in a chemical process using transfer entropy. *IEEE Transactions on Control Systems Technology*, 15(1), 12–21.

- Beer, R. (1995). A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence*, 72(1-2), 173–215.
- Beer, R. D. (2000). Dynamical approaches to cognitive science. *Trends in Cognitive Sciences*, 4(3), 91–99.
- Beer, R. D. (2008). The Dynamics of Brain- Body-Environment Systems: A Status Report. *Handbook of Cognitive Science- An Embodied Approach* (pp. 99–120).
- Beer, R. D. & Williams, P. L. (2015). Information processing and dynamics in minimally cognitive agents. *Cognitive Science*, 39(1), 1–38.
- Bob, P., Susta, M., Chladek, J., Glaslova, K. & Palus, M. (2009). Chaos in schizophrenia associations, reality or metaphor? *International Journal of Psychophysiology: Official Journal of the International Organization of Psychophysiology*, 73(3), 179–185.
- Bongard, J. (2011). Morphological change in machines accelerates the evolution of robust behavior. *Proceedings of the National Academy of Sciences*, 108(4), 1234–1239.
- Boon, M. Y., Henry, B. I., Suttle, C. M. & Dain, S. J. (2008). The correlation dimension: A useful objective measure of the transient visual evoked potential? *Journal of Vision*, 8, 1–21.
- Borghini, A. M. & Binkofski, F. (2014). *Words as Social Tools: An Embodied View on Abstract Concepts*. Springer.
- Boroditsky, L. (2000). Metaphoric structuring: understanding time through spatial metaphors. *Cognition*, 75(1), 1–28.
- Brooks, R. A. (1990). Elephants Don't Play Chess. *Robotics and Autonomous Systems*, 6(1-2), 3–15.
- Brooks, R. A. (1999). *Cambrian intelligence: The early history of the new AI*, Volume 44. Cambridge, MA: MIT Press.
- Brujin, S. M., van Dieën, J. H., Meijer, O. G. & Beek, P. J. (2009). Statistical precision and

- sensitivity of measures of dynamic gait stability. *Journal of Neuroscience Methods*, 178(2), 327–333.
- Cangelosi, A., Bongard, J., Fischer, M. H. & Nolfi, S. (2015). Embodied intelligence. In *Springer Handbook of Computational Intelligence* (pp. 697–714). Springer.
- Cangelosi, A., Greco, A. & Harnad, S. (2002). Symbol Grounding and the Symbolic Theft Hypothesis. In *Simulating the Evolution of Language* (pp. 191–210). Springer.
- Carvalho, J. T. & Nolfi, S. (2016). Cognitive Offloading Does Not Prevent but Rather Promotes Cognitive Development. *PloS One*, 11(8), e0160679.
- Cignetti, F., Decker, L. M. & Stergiou, N. (2012). Sensitivity of the Wolf's and Rosenstein's algorithms to evaluate local dynamic stability from small gait data sets. *Annals of Biomedical Engineering*, 40(5), 1122–1130.
- Clark, A. & Chalmers, D. (1998). The Extended Mind. *Analysis*, 58(1), 7–19.
- Da Rold, F. (2017). Chaotic analysis of embodied and situated agents. *Robotics and Autonomous Systems*, 95, 143–159.
- Dingwell, J. B. & Cusumano, J. P. (2000). Nonlinear time series analysis of normal and pathological human walking. *Chaos (Woodbury, N.Y.)*, 10(4), 848–863.
- Dingwell, J. B., Cusumano, J. P., Cavanagh, P. R. & Sternad, D. (2001). Local Dynamic Stability Versus Kinematic Variability of Continuous Overground and Treadmill Walking. *Journal of Biomechanical Engineering*, 123(1), 27–32.
- Dingwell, J. B., Robb, R. T., Troy, K. L. & Grabiner, M. D. (2008). Effects of an attention demanding task on dynamic stability during treadmill walking. *Journal of Neuroengineering and Rehabilitation*, 5, 12.
- Drack, M. & Pouvreau, D. (2015). On the history of Ludwig von Bertalanffy's "General Systemology", and on its relationship to cybernetics-part III: convergences and divergences. *International Journal of General Systems*, 44(5), 523–571.

- Eigen, M. (1973). *Ingo Rechenberg Evolutionsstrategie Optimierung technischer Systeme nach Prinzipien der biologischen Evolution*. mit einem Nachwort von Manfred Eigen, Friedrich Frommann Verlag.
- Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, 14(2), 179–211.
- Elman, J. L. (1998a). Connectionism, artificial life, and dynamical systems: New approaches to old questions. In *A companion to cognitive science*. (pp. 233–227). Blackwell Publishing.
- Elman, J. L. (1998b). *Rethinking innateness: A connectionist perspective on development*. Cambridge, MA: MIT Press.
- Fang, F. C. & Casadevall, A. (2011). Reductionistic and holistic science. *Infection and Immunity*, 79(4), 1401–1404.
- Fano, R. M. & Hawkins, D. (1961). Transmission of information: A statistical theory of communications. *American Journal of Physics*, 29(11), 793–794.
- Faure, P. & Korn, H. (2001). Is there chaos in the brain? I. Concepts of nonlinear dynamics and methods of investigation. *Comptes Rendus de l'Academie des Sciences - Serie III*, 324(9), 773–793.
- Fischer, M. H. (2012). A hierarchical view of grounded, embodied, and situated numerical cognition. *Cognitive Processing*, 13(1 SUPPL), 161–164.
- Flecker, B., Alford, W., Beggs, J. M., Williams, P. L. & Beer, R. D. (2011). Partial information decomposition as a spatiotemporal filter. *Chaos*, 21(3), 37104.
- Floreano, D., Mitri, S., Magnenat, S. & Keller, L. (2007). Evolutionary conditions for the emergence of communication in robots. *Current Biology*, 17(6), 514–519.
- Flusberg, S. J., Thibodeau, P. H., Sternberg, D. A. & Glick, J. J. (2010). A connectionist approach to embodied conceptual metaphor. *Frontiers in Psychology*, 1(NOV), 142.
- Fodor, J. A. (1975). *The language of thought*. Harvard University Press.

- Francisco Varela, E. T. T. & Rosch, E. (1991). *The Embodied Mind: Cognitive Science and Human Experience*. Cambridge, MA: MIT Press.
- Freitas, U. S., Letellier, C. & Aguirre, L. A. (2009). Failure in distinguishing colored noise from chaos using the "noise titration" technique. *Physical Review E*, 79(3), 35201.
- Frenzel, S. & Pompe, B. (2007). Partial mutual information for coupling analysis of multivariate time series. *Physical Review Letters*, 204101(20), 204101.
- Friston, K., Moran, R. & Seth, A. K. (2013). Analysing connectivity with Granger causality and dynamic causal modelling. *Current Opinion in Neurobiology*, 23(2), 172–178.
- Gallese, V. & Lakoff, G. (2005). The Role of the Sensory-Motor System in Conceptual Knowledge: The Brain's Concepts. *Cognitive neuropsychology*, 22(3-4), 455–479.
- van Gelder, T. (1990). Compositionality: A connectionist variation on a classical theme. *Cognitive Science*, 14(3), 355–384.
- van Gelder, T. (1998). The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences*, 21(5), 615–628.
- Glass, L. (2009). Introduction to controversial topics in nonlinear science: is the normal heart rate chaotic? *Chaos (Woodbury, N.Y.)*, 19(2), 28501.
- Gomez, G., Lungarella, M. & Tarapore, D. (2005). Information-theoretic approach to embodied category learning. *Proceedings of the 10th International Symposium on Artificial Life and Robotics* (pp. 332–337).
- Gómez-Herrero, G., Wu, W., Rutanen, K., Soriano, M. C., Pipa, G. & Vicente, R. (2015). Assessing coupling dynamics from an ensemble of time series. *Entropy*, 17(4), 1958–1970.
- Gould, S. J. & Eldredge, N. (1977). Punctuated equilibria; the tempo and mode of evolution reconsidered. *Paleobiology*, 3(2), 115–151.
- Grassberger, P. & Procaccia, I. (1983). Measuring the strangeness of strange attractors. *Physica D: Nonlinear Phenomena*, 9, 189–208.

- Harnard, S. (1990). The Symbol Grounding Problem. *Physica D*, 42, 335–346.
- Harvey, I., Paolo, E. D., Wood, R., Quinn, M. & Tuci, E. (2005). Evolutionary Robotics: A New Scientific Tool for Studying Cognition. *Artificial Life*, 11(1-2), 79–98.
- Hebb, D. O. (1949). *The organization of behavior: A neuropsychological approach*. John Wiley & Sons.
- Hinton, G. E., Osindero, S. & Teh, Y.-W. (2006). A fast learning algorithm for deep belief nets. *Neural computation*, 18(7), 1527–1554.
- Holland, J. H. (1992). *Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence*. Cambridge, MA: MIT Press.
- Izquierdo, E. J., Williams, P. L. & Beer, R. D. (2015). Information Flow through a Model of the *C. elegans* Klinotaxis Circuit. *PloS one*, 10(10), e0140397.
- Janjarasjitt, S., Scher, M. S. & Loparo, K. a. (2008). Nonlinear dynamical analysis of the neonatal EEG time series: the relationship between sleep state and complexity. *Clinical Neurophysiology: Official Journal of the International Federation of Clinical Neurophysiology*, 119(8), 1812–1823.
- Kaiser, A. & Schreiber, T. (2002). Information transfer in continuous processes. *Physica D: Nonlinear Phenomena*, 166(1), 43–62.
- Kang, H. G. & Dingwell, J. B. (2009). Dynamics and stability of muscle activations during walking in healthy young and older adults. *Journal of Biomechanics*, 42(14), 2231–2237.
- Kantz, H. & Schreiber, T. (1997). *Nonlinear Time Series Analysis*. Cambridge University Press.
- Kantz, H. & Schreiber, T. (1998). Human ECG: nonlinear deterministic versus stochastic aspects. *IEE Proceedings - Science, Measurement and Technology*, 145(6), 279–284.
- Kaplan, D. & Glass, L. (2012). *Understanding nonlinear dynamics*. Springer.
- Karmarkar, U. R. & Buonomano, D. V. (2007). Timing in the Absence of Clocks: Encoding Time in Neural Network States. *Neuron*, 53(3), 427–438.

- Kennel, M. B., Brown, R., Abarbanel, H. D. I. (1992). Determining Embedding Dimensions for Phase-Space Reconstruction Using a Geometrical Construction. *Physical Review A*, 45(6).
- Kesić, S. (2016). Systems biology, emergence and antireductionism. *Saudi Journal of Biological Sciences*, 23(5), 584–591.
- Kohonen, T. (1998). The self-organizing map. *Neurocomputing*, 21(1), 1–6.
- Korn, H. & Faure, P. (2003). Is there chaos in the brain? II. Experimental evidence and related models. *Comptes Rendus - Biologies*, 326(9), 787–840.
- Kraskov, A., Stögbauer, H. & Grassberger, P. (2004). Estimating mutual information. *Physical Review E*, 69(6), 066138.
- Lakoff, G. & Johnson, M. (1999). *Philosophy in the Flesh*, Volume 4. Basic Books.
- Le Van Quyen, M., Chavez, M., Rudrauf, D. & Martinerie, J. (2003). Exploring the nonlinear dynamics of the brain. *Journal of Physiology Paris*, 97(4-6), 629–639.
- Lee, J., Nemati, S., Silva, I., Edwards, B. A., Butler, J. P. & Malhotra, A. (2012). Transfer Entropy Estimation and Directional Coupling Change Detection in Biomedical Time Series. *Biomedical Engineering Online*, 11(1), 19.
- Lizier, J. T. (2014). JIDT: An Information-Theoretic Toolkit for Studying the Dynamics of Complex Systems. *Frontiers in Robotics and AI*, 1, 11.
- Lizier, J. T., Heinzle, J., Horstmann, A., Haynes, J.-D. & Prokopenko, M. (2011). Multivariate information-theoretic measures reveal directed information structure and task relevant changes in fMRI connectivity. *Journal of Computational Neuroscience*, 30(1), 85–107.
- Lizier, J. T., Prokopenko, M. & Zomaya, A. Y. (2008). Local information transfer as a spatiotemporal filter for complex systems. *Physical Review E*, 77(2), 26110.
- Lorenz, E. N. (1963). Deterministic Nonperiodic Flow. *Journal of the Atmospheric Sciences*, 20(2), 130–141.

- Lungarella, M. & Sporns, O. (2006). Mapping information flow in sensorimotor networks. *PLoS Computational Biology*, 2(10), 1301–1312.
- Mangin, L., Lesèche, G., Duprey, A. & Clerici, C. (2011). Ventilatory Chaos Is Impaired in Carotid Atherosclerosis. *PLoS ONE*, 6(1), 1–10.
- Maniadakis, M., Tani, J. & Trahanias, P. (2009). Time perception in shaping cognitive neuro-dynamics of artificial agents. In *Proceedings of the International Joint Conference on Neural Networks* (pp. 1993–2000). IEEE Press.
- Margolin, A. A., Wang, K., Lim, W. K., Kustagi, M., Nemenman, I. & Califano, A. (2006). Reverse engineering cellular networks. *Nature Protocols*, 1(2), 662–671.
- Martins-Filho, L. S., Machado, R. F., Rocha, R. & Vale, V. S. (2004). Commanding mobile robots with chaos. In *ABCM Symposium Series in Mechatronics*, Volume 1 (pp. 40–46).
- Matsuda, E., Hubert, J. & Ikegami, T. (2014). A robotic approach to understanding the role and the mechanism of vicarious trial-and-error in a T-maze task. *PLoS ONE*, 9(7), e102708.
- Maturana, H. R. & Varela, F. J. (1991). *Autopoiesis and cognition: The realization of the living*. Springer.
- McCarthy, J. & Hayes, P. J. (1969). Some philosophical problems from the standpoint of artificial intelligence. In *Machine Intelligence* (pp. 463–502). Edinburgh University Press.
- McClelland, J. L. & Rumelhart, D. E. (1987). *Parallel distributed processing*, Volume 2. Cambridge, MA: MIT Press.
- McCulloch, W. S. & Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics*, 5(4), 115–133.
- Minsky, M. (1988). *Society of mind*. Simon and Schuster.
- Mitchell, M. (1998). *An introduction to genetic algorithms*. Cambridge, MA: MIT Press.

- Mondada, F., Bonani, M., Raemy, X., Pugh, J., Cianci, C., Klapotocz, A., Zufferey, J.-c., Floreano, D. & Martinoli, A. (2006). The e-puck , a Robot Designed for Education in Engineering. In *Proceedings of the 9th Conference on Autonomous Robot Systems and Competitions*, Volume 1 (pp. 59–65). IPCB: Instituto Politécnico de Castelo Branco.
- Monirul Islam, M. & Murase, K. (2005). Chaotic dynamics of a behavior-based miniature mobile robot: effects of environment and control structure. *Neural Networks*, 18(2), 123–144.
- Montebelli, A., Herrera, C. & Ziemke, T. (2008). On cognition as dynamical coupling: An analysis of behavioral attractor dynamics. *Adaptive Behavior*, 16(2-3), 182–195.
- Moravec, H. (1988). *Mind Children: The Future of Robot and Human Intelligence*. Harvard University Press.
- Myachykov, A., Scheepers, C., Fischer, M. H. & Kessler, K. (2014). TEST: A tropic, embodied, and situated theory of cognition. *Topics in Cognitive Science*, 6(3), 442–460.
- Nakamura, Y. & Sekiguchi, A. (2001). The chaotic mobile robot. *IEEE Transactions on Robotics and Automation*, 17(6), 898–904.
- Nehmzow, U. (2012). *Mobile robotics: a practical introduction*. Springer.
- Nehmzow, U. & Walker, K. (2003). The behaviour of a mobile robot is chaotic. *AISB Journal*, 1(4).
- Newell, A. & Simon, H. A. (1976). Computer science as empirical inquiry: symbols and search. *Communications of the ACM*, 19(3), 113–126.
- Nolfi, S. & Floreano, D. (2000). *Evolutionary Robotics: The Biology, Intelligence, and Technology*. Cambridge, MA, USA: Cambridge, MA: MIT Press.
- Nolfi, S. & Gigliotta, O. (2010). Evorobot. In *Evolution of communication and language in embodied agents* (pp. 297–301). Springer.
- Parisi, D. (2014). *Future robots: towards a robotic science of human beings*. John Benjamins Publishing Company.

- Parisi, D., Cecconi, F. & Nolfi, S. (1990). Econets: Neural networks that learn in an environment. *Network: Computation in Neural Systems*, 1(2), 149–168.
- Perc, M. (2005a). Nonlinear time series analysis of the human electrocardiogram. *European Journal of Physics*, 26(5), 757–768.
- Perc, M. (2005b). The dynamics of human gait. *European Journal of Physics*, 26(3), 525–534.
- Perc, M. (2006). Introducing nonlinear time series analysis in undergraduate courses. *Fizika a-Zagreb-*, 15(1/4), 91.
- Pezzulo, G., Barsalou, L. W., Cangelosi, A., Fischer, M. H., Spivey, M. & McRae, K. (2011). The mechanics of embodiment: A dialogue on embodiment and computational modeling. *Frontiers in Psychology*, 2(5), 1–21.
- Pfeifer, R. & Bongard, J. (2006). *How the body shapes the way we think: a new view of intelligence*. Cambridge, MA: MIT Press.
- Pollack, J. B. (1991). The induction of dynamical recognizers. *Machine Learning*, 7(2-3), 227–252.
- Poon, C. S. & Barahona, M. (2001). Titration of chaos with added noise. *Proceedings of the National Academy of Sciences of the United States of America*, 98(13), 7107–7112.
- Poon, C. S., Li, C. & Wu, G. Q. (2010). A unified theory of chaos linking nonlinear dynamics and statistical physics. *arXiv preprint arXiv:1004.1427*.
- Pylyshyn, Z. W. (1986). *Computation and cognition: Toward a foundation for cognitive science*. Cambridge, MA: MIT Press.
- Reynolds, A. M., Bartumeus, F., Kolzsch, A. & Van De Koppel, J. (2016). Signatures of chaos in animal search patterns. *Scientific Reports*, 6, 23492.
- Rispens, S. M., Pijnappels, M., van Dieën, J. H., van Schooten, K. S., Beek, P. J. & Daffertshofer, A. (2014). A benchmark test of accuracy and precision in estimating dynamical systems characteristics from a time series. *Journal of Biomechanics*, 47(2), 470–475.

- Rizzolatti, G. & Arbib, M. (1998). Language within our grasp. *Trends in Neurosciences*, 21, 188–194. *Trends in Neurosciences*, 21(1988), 188–194.
- Rosenstein, M. T., Collins, J. J. & De Luca, C. J. (1993). A practical method for calculating largest Lyapunov exponents from small data sets. *Physica D: Nonlinear Phenomena*, 65(1-2), 117–134.
- Rumelhart, D., Hinton, G. & Williams, R. (1986). Learning representations by Back-propagation Errors. *Nature*, 323(6088), 533–536.
- Rumelhart, D. E. & McClelland, J. L. (1988). *Parallel distributed processing*, Volume 1. Cambridge, MA: MIT Press.
- Runge, J., Heitzig, J., Marwan, N. & Kurths, J. (2012a). Quantifying causal coupling strength: A lag-specific measure for multivariate time series related to transfer entropy. *Physical Review E*, 86(6), 61121.
- Runge, J., Heitzig, J., Petoukhov, V. & Kurths, J. (2012b). Escaping the curse of dimensionality in estimating multivariate transfer entropy. *Physical Review Letters*, 108(25), 258701.
- Runge, J., Sejdinovic, D. & Flaxman, S. (2017). Detecting causal associations in large nonlinear time series datasets. *arXiv preprint arXiv:1702.07007*.
- Sarbadhikari, S. N. & Chakrabarty, K. (2001). Chaos in the brain: a short review alluding to epilepsy, depression, exercise and lateralization. *Medical Engineering & Physics*, 23(7), 445–455.
- Schreiber, T. (1997). Detecting and analyzing nonstationarity in a time series using nonlinear cross predictions. *Physical Review Letters*, 78(5), 843.
- Schreiber, T. (2000). Measuring information transfer. *Physical Review Letters*, 85(2), 461–4.
- Searle, J. R. (1980). Minds, brains, and programs. *The Behavioral and Brain Sciences*, 3(3), 417–457.

- Seth, A. K. & Edelman, G. M. (2004). Environment and behavior influence the complexity of evolved neural networks. *Adaptive Behavior*, 12(1), 5–20.
- Shannon, C. E. & Weaver, W. (1949). *The mathematical theory of communication*. Urbana: University of Illinois Press.
- Shen, Y., Olbrich, E., Achermann, P. & Meier, P. F. (2003). Dimensional complexity and spectral properties of the human sleep EEG. *Clinical Neurophysiology*, 114, 199–209.
- Shimazaki, H. & Shinomoto, S. (2007). A Method for Selecting the Bin Size of a Time Histogram. *Neural Computation*, 19(6), 1503–1527.
- Silva, F. M. H. S. P., Silva Filho, A. C., Crescencio, J. C. & Gallo Junior, L. (2012). Applying Lyapunov Exponents in Heart Rate Time Series to Identify the Anaerobic Threshold in Healthy Men. *2012 Computing in Cardiology (Cinc)*, 39, 629–632.
- Silverman, B. W. (1986). *Density estimation for statistics and data analysis*. CRC Press.
- Skarda, C. A. & Freeman, W. J. (1987). How brains make chaos in order to make sense of the world. *Behavioral and Brain Sciences*, 10(02), 161–173.
- Smith, L. B. (2005). Cognition as a dynamic system: Principles from embodiment. *Developmental Review*, 25(3-4), 278–298.
- Smithers, T. (1995). On quantitative performance measures of robot behaviour. *Robotics and Autonomous Systems*, 15(1-2), 107–133.
- Sporns, O., Karnowski, J. & Lungarella, M. (2006). Mapping causal relations in sensorimotor networks. In *Proceedings of the 5th International Workshop on Epigenetic Robotics*.
- Stam, C. J. (2005). Nonlinear dynamical analysis of EEG and MEG: Review of an emerging field. *Clinical Neurophysiology*, 116(10), 2266–2301.
- Stoelen, M. F., de Tejada, V. F., Huete, A. J., Balaguer, C. & Bonsignorio, F. P. (2015). Distributed and Adaptive Shared Control Systems: Methodology for the Replication of Experiments. *IEEE Robotics & Automation Magazine*, 22(4), 137–146.

- Stramandinoli, F., Marocco, D. & Cangelosi, A. (2012). The grounding of higher order concepts in action and language: A cognitive robotics model. *Neural Networks*, 32, 165–173.
- Sugita, Y. & Tani, J. (2005). Learning semantic combinatoriality from the interaction between linguistic and behavioral processes. *Adaptive Behavior*, 13(1), 33–52.
- Sutton, R. S. & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- Takens, F. (1981). Detecting strange attractors in turbulence. In *Dynamical systems and turbulence* (pp. 366–381). Springer.
- Tani, J. (2003). Learning to generate articulated behavior through the bottom-up and the top-down interaction processes. *Neural Networks*, 16(1), 11–23.
- Tani, J. & Nolfi, S. (1999). Learning to perceive the world as articulated: an approach for hierarchical learning in sensory-motor systems. *Neural Networks*, 12(7), 1131–1141.
- Tarapore, D., Lungarella, M. & Gómez, G. (2004). Fingerprinting agent-environment interaction via information theory. In *The Eighth International Conference on Intelligent Autonomous Systems (IAS-8)*, number LIS-CONF-2006-008 (pp. 512–520).
- Theiler, J. (1994). On the evidence for low-dimensional chaos in an epileptic electroencephalogram. *Physics Letters A*, 196(1-2), 335–341.
- Theiler, J. & Rapp, P. E. (1996). Re-examination of the evidence for low-dimensional, nonlinear structure in the human electroencephalogram. *Electroencephalography and Clinical Neurophysiology*, 98(3), 213–222.
- Thelen, E. & Smith, L. B. (1996). *A dynamic systems approach to the development of cognition and action*. Cambridge, MA: MIT Press.
- Trianni, V. (2014). Evolutionary robotics: model or design? *Frontiers in Robotics and AI*, 1, 13.
- Tuci, E., Trianni, V. & Dorigo, M. (2004). 'Feeling' the Flow of Time through Sensory-Motor Coordination. *Connection Science*, 4(16), 301–324.

- Turing, A. (1936). On Computable Problems with an Application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society*, 2(345-363), 42.
- Turing, A. (1950). Turing. Computing machinery and intelligence. *Mind*, 59(236), 433–460.
- Verdes, P. F. (2005). Assessing causality from multivariate time series. *Physical Review E*, 72(2), 26222.
- Vicente, R., Wibral, M., Lindner, M. & Pipa, G. (2011). Transfer entropy—a model-free measure of effective connectivity for the neurosciences. *Journal of Computational Neuroscience*, 30(1), 45–67.
- Volos, C. K., Kyprianidis, I. M. & Stouboulos, I. N. (2013). Experimental investigation on coverage performance of a chaotic autonomous mobile robot. *Robotics and Autonomous Systems*, 61(12), 1314–1322.
- Von Bertalanffy, L. (1968). *General system theory*. George Braziller.
- Wibral, M., Vicente, R. & Lizier, J. T. (2014). *Directed Information Measures in Neuroscience*. Springer.
- Wiener, N. (1959). The Theory of Prediction. *Modern Mathematics for Engineers*, 1, 125–139.
- Wiener, N. (1961). *Cybernetics or Control and Communication in the Animal and the Machine*. Cambridge, MA: MIT Press.
- Williams, P. L. & Beer, R. D. (2011). Generalized measures of information transfer. *arXiv preprint arXiv:1102.1507*.
- Wilson, M. (2002). Six views of embodied cognition. *Psychonomic Bulletin and Review*, 9(4), 625–636.
- Wolf, A., Swift, J., Swinney, H. & Vastano, J. (1985). Determining Lyapunov exponents from a time series Citado por mí. *Physica D: Nonlinear Phenomena* (pp. 285–317).
- Yamauchi, B. M. & Beer, R. D. (1994). Sequential behavior and learning in evolved dynamical neural networks. *Adaptive Behavior*, 2(3), 219–246.

BIBLIOGRAPHY

Yeragani, V. K., Rao, K. K., Smitha, M., Pohl, R. B., Balon, R. & Srinivasan, K. (2002). Diminished chaos of heart rate time series in patients with major depression. *Biological Psychiatry*, *51*(9), 733–744.

Zang, X., Iqbal, S., Zhu, Y., Liu, X. & Zhao, J. (2016). Applications of chaotic dynamics in robotics. *International Journal of Advanced Robotic Systems*, *13*(2), 60.

Ziemke, T. (2002). What's that Thing Called Embodiment ? In *Proceedings of the 25th Annual meeting of the Cognitive Science Society* (pp. 1305–1310).