

Computational and Robotic Models of the Hierarchical Organization of Behavior: An Overview

Gianluca Baldassarre and Marco Mirolli

Laboratory of Computational Embodied Neuroscience,
Institute of Cognitive Sciences and Technologies,
National Research Council (LOCEN-ISTC-CNR),
Via San Martino della Battaglia 44, I-00185 Rome, Italy,
{gianluca.baldassarre, marco.mirolli}@istc.cnr.it

Abstract. The hierarchical organisation of behaviour is a fundamental means through which robots and organisms can acquire and produce sophisticated and flexible behaviours that allow them to solve multiple tasks in multiple conditions. Recently, the research on this topic has been receiving increasing attention. On the one hand, machine learning and robotics are recognising the fundamental importance of the hierarchical organisation of behaviour for building robots that scale up to solve complex tasks, possibly in a cumulative fashion. On the other, research in psychology and neuroscience is finding increasing evidence that modularity and hierarchy are pivotal organisation principles of behaviour and of the brain. This book reviews the state of the art in computational and robotic models of the hierarchical organisation of behaviour. Each contribution reviews the main works of the authors on this subject, the open challenges and promising research directions. Together, the contributions give a good coverage of the most important models, findings, and challenges of the field. This introductory chapter presents the general aims and scope of the book and briefly summarises the contents of each chapter.

1 Modeling the Hierarchical Organization of Behavior

The performance of flexible behaviour to accomplish multiple goals requires a hierarchical organisation of actions. Each action consists in a sensorimotor mapping that associates a flow of motor commands to the flow of sensory inputs. The mappings related to different actions can be substantially different. When this happens, different actions have to be encoded in distinctive portions of the architecture of the control system so to avoid a cross-talk or *catastrophic interference* between them (French, 1999; McCloskey and Cohen, 1989). On the other hand, when sensorimotor mappings are very similar their encoding in common structures facilitate *generalisation* and the *reuse of knowledge* for the accomplishment of different purposes (Meunier et al., 2010; Singh, 1992). Hierarchical control architectures can allow both the avoidance of catastrophic interference and the

exploitation of previously acquired skills to accomplish new tasks. Furthermore, they can also allow the decomposition of complex control problems into smaller tractable problems (Hart and Grupen, 2011), and the *chunking* of simpler actions in higher level actions so that increasingly complex behaviours can be acquired cumulatively (Bakker and Schmidhuber, 2004; Balleine and Dickinson, 1998).

For all these reasons, hierarchical architectures are becoming increasingly important in robotics, in particular when robots are requested to solve not only one tasks but multiple tasks, and not only in one condition but in multiple conditions. Hierarchical architectures are now generally considered as the necessary condition to allow robots to undergo a prolonged autonomous development Baldassarre and Mirolli (2010), and to scale up robot behaviour to address real-life problems (e.g., Demiris and Khadhour, 2006; Yamashita and Tani, 2008). As shown in various chapters presented in this book, state-of-art robotics use hierarchical control architectures to solve multiple tasks, to facilitate the re-use of acquired knowledge to solve other tasks, to facilitate human-to-robot transfer of knowledge, to avoid interference, and so on.

If the adoption of hierarchical architectures is rather new in robotics, the recognition that animal behaviour is hierarchical organised is quite older, dating back at least in the early 1960s (Miller et al., 1960). Today, the hierarchical organisation of behaviour is given for granted in psychology, where it is generally conceived that humans cumulatively build a repertoire of skills that can be flexibly composed to form increasingly complex action programs. Empirical, theoretical, and computational research has been carried out in psychology to understand the details of such a hierarchical organisation of behavior (Botvinick and Plaut, 2004; Cooper and Shallice, 2000; Fischer, 1980; Schneider and Logan, 2006; Zacks et al., 2007).

Recent research has also been demonstrating that to the hierarchical organisation of behaviour corresponds a hierarchical organisation of the brain (Fuster, 2001; Meunier et al., 2010). In particular, animals' brains have been suggested to exploit hierarchy in order to chunk pieces of behaviour so to reuse them in new tasks (Graybiel, 1998), to easily recall behaviours in later times (e.g., to pursue goals associated with them Redgrave and Gurney, 2006), to avoid cross-talk, and to exploit the compositionality allowed by a modular organisation of information (Graziano, 2006). Furthermore, recent neuroscientific research is revealing that the brain is hierarchically organised at multiple levels both within cortical (Miller and Cohen, 2001) and sub-cortical regions (Yin and Knowlton, 2006). In this respect, it can be said that much of the behavioural flexibility exhibited by real organisms depends on the fine hierarchical organisation of the underlying brain structure (Meunier et al., 2010).

The aim of the present book is to review the state of the art in computational and robotic models of the hierarchical organisation of behaviour. The book covers the full spectrum of models: from (scientific) models that try to explain behavioural and/or neural phenomena found in real animals to (technological) models that aim to endow robots with increasingly powerful controllers (including models that try to do both things). Indeed, we are convinced that

the cross-fertilisation between different disciplines and approaches is of the most importance both if we want to better understand human behaviour and if we want to construct more useful artificial systems (in the case of hierarchies as in any other). And that computational modelling is the lingua franca that can help to bridge different disciplines that have different concepts, methods, and traditions.

In order to provide a wide overview of current computational research on the hierarchical organisation of behaviour we asked the authors of the various contributions, which represent some of the most active researchers in the field today, to: 1) offer a broad survey of their main works on hierarchical behaviour, 2) highlight the problems and open challenges they see in their area, and 3) when possible, point to the most promising research direction for tackling those challenges. We hope that, taken together, the resulting contributions, collected in one unique venue and following homogeneous guidelines, will aid the reader novel to the field to have a comprehensive panoramic on hierarchical behaviour, and the more expert reader to be informed on the last advancements in the field.

2 Book Overview

The chapters of the book have been organised as follows. Part I presents the contributions that address the issue of the hierarchical organisation of behaviour by mainly addressing the problem of how building skilled robots. Some of these contributions are ‘bio-inspired’ but their distinctive feature is to aim to build useful intelligent technological artefacts. Part II presents computational models directed to understand the hierarchical organisation of behaviour in real animals. The distinctive feature of the models of this part is to aim to answer scientific questions on the organisation of natural intelligence. Although some contributions are robotic and/or refer also to elements of the brain organisation, its major focus is however the study of animal behaviour. Finally, Part III presents computational models focussed on understanding the hierarchical organisation of animal’s brain. As those of the previous part, these contributions aim to understand natural intelligence. Although these models refer heavily to behaviour, given their system-level approach to the study of brain, their distinctive feature is their close reference to the empirical evidence on the organisation and functioning of real nervous systems. Together, the models illustrated in the chapters represent a solid basis from which to depart to build future robots solving multiple problems or to foster further investigations of the hierarchical organisation of action in the behaviour and brain of real animals. We now present brief summaries of the contents of each book chapter.

2.1 Part I: Hierarchical Organisation of Behaviour in Robots

Behavioral Hierarchy: Exploration and Representation. In chapter ??, **Andrew G. Barto**, **George Konidaris**, and **Christopher Vigorito** discuss the advantages of behavioural hierarchy from the point of view of hierarchical reinforcement learning. In particular, the authors explore two kinds of benefits that are

often overlooked: the benefits given by behavioural hierarchies with respect to the problem of how to efficiently *explore* the environment, and those related to the problem of how to appropriately *represent* the state space on which an agent works. Each of the two kinds of benefits is exemplified by reviewing two computational experiments: the first two experiments show how a behavioural hierarchy can improve exploration in structured environments so as to significantly improve learning speed; the other two experiments show how behavioural hierarchies can allow the use of low complexity function approximation methods for complex problems and how they can allow the selection of different state abstractions for different skills to be learned, thus facilitating learning in high dimensional state spaces. Finally, the authors discuss the generality of their findings and some promising directions for future research in hierarchical reinforcement learning.

Self-organized functional hierarchy through multiple timescale: neuro-dynamical accounts for behavioral compositionality. In chapter ??, **Yuichi Yamashita** and **Jun Tani** review their work on functional hierarchies of behaviors in distributed neural systems that work at multiple time-scales. They first discuss the problems encountered in scaling up a modular architecture to complex robotic systems with many degrees of freedom and then present two kinds of systems based on distributed representations. In the first system, the recurrent neural network with parametric biases (RNNPB), different sensory-motor sequences are encoded in the fixed activation patterns of the parametric bias units. In the more recent multiple timescale recurrent neural network model (MTRNN), different classes of context units have different time-constant thus working at different time-scales. Through a supervised learning processes, fast units learn to encode different primitives while slow units learn to encode appropriately switch between those primitives so to perform higher level behaviors. The authors conclude by discussing the challenge of overcoming the limits of both distributed and localist systems, and how these might be tackled in future work.

Autonomous Representation Learning in a Developing Agent. In chapter ??, **Jonathan Mugan** and **Benjamin Kuipers** present a hierarchical computational model, called *Qualitative Learner of Action and Perception* (QLAP), that has been built for developing high-level qualitative representations from low-level continuous sensor and actuator representations while autonomously interacting with its environment. The system has been built on the basis of four principles, devoted in particular to facing the problem of learning useful representations: (1) the exploitation of the synergy between created representations and the agent's development; (2) the generation of new qualitative representations that capture in an abstract form relevant 'phenomena' in the environment; (3) the decomposition of the environment in many sub-parts; (4) the creation of representations making learning more efficient. After explaining these principles, the authors present QLAP in detail and show how such system can effectively learn useful representations and hierarchical actions in a simulated robotic environment with realistic physics.

Hierarchies for Embodied Action Perception. In chapter ??, **Dimitri Ognibene, Yan Wu, Kyuhwa Lee, and Yiannis Demiris** present the *Hierarchical Attentive Multiple Models for Execution and Recognition* (HAMMER) architecture, a bio-inspired learning framework for endowing robots with the ability to understand and imitate human actions. The model, which is based on a repertoire of paired inverse and forward models, is based on three key principles: (1) human knowledge is hierarchically structured, both at the perceptual and at the execution level; (2) perception is active, and driven by task requirements and context; (3) learning is pivotal for action perception, permitting the continuous acquisition of new behaviours. After discussing these principles and describing the general architecture and functioning of HAMMER, the authors review some robotic experiments directed to test the architecture. The experiments demonstrate how the hierarchical organisation and the repertoire of inverse and forward models of the architecture facilitate its autonomous acquisition of complex behaviours and also the acquisition of sequences of actions on the basis of imitation processes.

Learning and Coordinating Repertoires of Behaviors with Common Reward: Credit Assignment and Module Activation. In chapter ??, **Constantin A. Rothkopf** and **Dana H. Ballard** present a reinforcement learning computational model that autonomously learns different simultaneous tasks on the basis of different computational modules and to appropriately mix them to maximise a common indistinct reward. The model is composed by many modules that have different state representations and that collectively control the behaviour of the system by sharing the same action space. The proposed algorithm learns to use the modules that are more appropriate for each task at hand by combining each module's reward estimate with an error signal depending on the difference between the unique reward estimate and the sum of the reward estimates of other co-active modules. The efficacy of the model is demonstrated through different experiments involving both simple abstract grid-world tasks and a more complex navigation task in a virtual 3D environment.

2.2 Part II: Hierarchical Organisation of Animal Behaviour

Modular, Multimodal Arm Control Models. In chapter ??, **Stephan Ehrenfeld, Oliver Herbort, and Martin V. Butz** discuss some important challenges faced by human motor control and present modular and hierarchical computational architectures that are able to meet those challenges. In particular, the authors identify three main challenges for flexible human motor control: (1) sensory redundancy, related to the fact that different sources of information about the state of the body and of the environment must be considered for motor control; (2) motor redundancy, related to the fact that different postures and trajectories can be used for the same goal, thus requiring the resolution of behavioural alternatives; (3) uncertainty, related to the fact that many features of the motor task can change from time to time and even during the execution of a movement while both movement execution and sensory processing are always noisy. After

identifying the modularity of representations and the hierarchical organisation of planning and control as the two main mechanisms that the brain may exploit to meet these challenges, the authors present a computational model, based on direct inverse modelling mechanisms, that has been used to account for the human behavioural flexibility in motor planning, the SURE-REACH model. Furthermore, since the representational scheme of SURE-REACH is computationally very expensive, a new model, called the Modular Modality Frame (MMF) model, is presented. This model aims to cope with the problem of scalability by further modularising the representational space, and by exploiting modular and hierarchical representations at several levels.

Generalization and Interference in Human Motor Control. In chapter ??, **Luca Lonini**, **Christos Dimitrakakis**, **Constantin Rothkopf**, and **Jochen Triesch** identify the problem of generalization without interference as a fundamental issue in modeling human motor control. In fact, a cumulative learning system that is learning a new motor skill must be able to efficiently exploit previously acquired abilities while preventing that catastrophic interference disrupt old knowledge, as is typically the case in simple neural systems. The authors first review the available empirical data on consolidation of procedural memories and then discuss the different computational models that have been proposed for learning multiple tasks in bio-inspired learning architectures.

A Developmental Framework for Cumulative Learning Robots. In chapter ??, **Mark Lee**, **James Law**, and **Martin Hülse** provide an extensive review of their work on the Lift-Constraint, Act, Saturate (LCAS) approach, a framework for building robot controllers inspired by developmental psychology (developmental robotics). The basic idea behind LCAS is that development proceeds in a staged fashion thanks to the presence of learning constraints of various sorts (anatomical, maturational, computational, environmental, etc.). These facilitate skill learning and are progressively released as soon as learning is saturated, thus allowing the cumulative development of increasingly complex abilities. The authors first present their approach to learn sensorimotor mappings: these constitute the fundamental building block of all their systems. Then the authors show how such sensorimotor mappings can be acquired from autonomous experience, and how the LCAS framework can lead to a staged development similar to those observed in children. Finally, the authors discuss the role of novelty in the LCAS framework, a possible developmental program for humanoid robots based on it, the research challenges facing the framework, and its relations with other work.

The Hierarchical Accumulation of Knowledge in the Distributed Adaptive Control Architecture. In chapter ??, **Encarni Marcos**, **Milanka Ringwald**, **Armin Duff**, **Martí Sánchez-Fibla**, and **Paul F.M.J. Verschure** present the Distributed Adaptive Control (DAC) architecture as a biologically-motivated cumulative learning system. In this hierarchical framework, a reactive layer stores built-in stimulus-response associations, an adaptive layer co-adapts behavioral

responses and perceptual representations according to reinforcements, and a contextual layer stores sensory-motor chains and uses them to perform memory-based goal-directed behaviors. The authors review several recent studies on this architecture focused on the interactions and arbitration between the higher contextual layer and the lower-level ones, and discuss how important challenges in cumulative learning can be tackled within the DAC framework.

2.3 Part III: Hierarchical Organisation of Animal Brain

The Hierarchical Organisation of Sensorimotor Cortical and Basal-Ganglia Brain Systems: A Computationally-Informed Review and Integrated Hypothesis. In chapter ??, **Gianluca Baldassarre**, **Daniele Caligiore** and **Francesco Mannella** discuss an important problem existing in the neuroscientific literature, related to the presence of two research frameworks focussed on either the hierarchical organisation of cortex or the hierarchical organisation of sub-cortical structures, in particular basal ganglia. The problem is that these two research threads proceed quite in isolation from each other, so missing account for the integrative nature of hierarchical brain and encountering important limitations in its explanation. To better illustrate this problem, the authors review in detail two of their computational models developed respectively in each of the two research frameworks. This allows them to exemplify in detail the problems, brain areas, experiments, main concepts and limitations of the two frameworks. On this basis, the authors then propose a theoretical integration of the two perspectives, and show how this leads to solve most problems found by the two accounts when taken in isolation. Overall, the paper shows that cortex and the basal ganglia form a whole highly-integrated system solving all the challenges of choice, selection, and behaviour implementation posed by adaptive behaviour on the basis of a sophisticated hierarchical organisation.

Divide and Conquer: Hierarchical Reinforcement Learning and Task Decomposition. In chapter ??, **Carlos Diuk**, **Anna Shapiro**, **Natalia Cordova**, **Jose Ribas-Fernandez**, **Yael Niv**, and **Matthew Botvinick** present recent empirical research that investigates the brain mechanisms underlying complex human behaviour requiring task decomposition by using hierarchical reinforcement learning computational models as interpretative tools. In particular, after briefly reviewing the field of hierarchical reinforcement learning, the authors summarise two recent experiments that demonstrate the existence of neural correlates of key predictions of hierarchical reinforcement learning, i.e. the presence of prediction error signals at different levels of abstraction, and the presence of pseudo-reward signals generated in presence of sub-goal accomplishment. The authors then focus on the important problem, currently still open in both robotics and neuroscience/psychology, of how a system can autonomously acquire goals/sub-goals that can guide the acquisition of repertoires of skills, essential for supporting hierarchical behaviour. Finally, the authors review other three behavioural and neuroimaging experiments devoted to investigate the brain mechanisms underlying the solution of this problem in animals.

Neural Network Modelling of Hierarchical Motor Function in the Brain. In chapter ??, **Juan M. Galeazzi** and **Simon M. Stringer** review their recent work on modelling of the hierarchical organisation of motor function in the brain. This work is based on biologically-plausible neural network architectures in which hierarchically organised dynamical and recurrent neural populations learn through biologically-plausible Hebbian synaptic learning rules. The authors review various simulations in which they demonstrate how their system can learn to perform both low-level motor primitives and high-level motor programs, how novel movement sequences can be learned through the modulation of high-level motor programs by context, and how motor sequences can be learned on the basis of a delayed reward signals thanks to an associative learning rule. Finally, the authors discuss what they consider the most important future challenges in the modelling of hierarchical motor function.

Restoring Purpose in Behavior. In chapter ??, **Henry H. Yin** provides an original perspective towards behaviour. Yin is not a computational modeller but a behavioural neuroscientist, and has in particular contributed to the understanding of the hierarchical organisation of the brain as expressed in goal-directed behaviour. In this contribution, the author criticises the dominant paradigm in the behavioural and brain sciences, which views behaviour as determined by its antecedent causes (i.e. external or internal stimuli), on the ground that organisms' behaviour is theleologic, and hence must be understood by considering the animal's *goals*. The author's proposal, which builds on the works of cybernetics and its concept of negative feedback, views behaviour as the manifestation of control in systems that act to make inputs match their goals. The author tries to demonstrate the flaws of previous theories in the explanation of behaviour and then discusses the kind of control allowed by a hierarchical organisation of behaviour. Finally, the contribution ends with a discussion of the possible experimental protocols that may be used to exploit the proposal for improving our understanding of animal behaviour.

3 Conclusions

The present book offers a broad overview on the major works and open problems in the field of the hierarchical organisation of behaviour in robots and organisms. The contributions presented here, coming from some of the most active and important researchers of the field, can surely give both the expert and the non-expert reader a panoramic knowledge on what can be found in this area, and can hopefully prompt new ideas. We thus hope that the book will attract new researchers and foster further investigations in this exciting front-edge field of research at the core of robotics and cognitive science.

Acknowledgments

This paper and a large part of the effort that led to this book have been supported by the Project 'IM-CLeVeR - Intrinsically Motivated Cumulative Learning Ver-

satellite Robots’ funded by the European Commission under the 7th Framework Programme (FP7/2007-2013), ‘Challenge 2 - Cognitive Systems, Interaction, Robotics’, Grant Agreement No. ICT-IP-231722. Support or co-support from other institutions, where present, is described in the “Acknowledgments” section of each chapter. The Editors of the book thank the EU reviewers (Benjamin Kuipers, Luc Berthouze, and Yasuo Kuniyoshi) and the EU Project Officer (Cecile Huet) for their valuable advices and their encouragement. For more information on the IM-CLeVeR project, and for additional multimedia material, see the project website: <http://www.im-clever.eu/>. We also thank Simona Bosco for her editorial help with some contributions.

Bibliography

- Bakker, B. and Schmidhuber, J. (2004). Hierarchical reinforcement learning based on subgoal discovery and subpolicy specialization. In Groen, F., Amato, N., Bonarini, A., Yoshida, E., and Kruse, B., editors, *Proceedings of the 8-th Conference on Intelligent Autonomous Systems (IAS-8)*, pages 438–445.
- Baldassarre, G. and Mirolli, M. (2010). What are the key open challenges for understanding the autonomous cumulative learning of skills? The Newsletter of the Autonomous Mental Development Technical committee (IEEE CIS AMD Newsletters). Vol. 7 N. 1, pp. 11.
- Balleine, B. W. and Dickinson, A. (1998). Goal-directed instrumental action: contingency and incentive learning and their cortical substrates. *Neuropharmacology*, 37(4-5):407–419.
- Botvinick, M. and Plaut, D. C. (2004). Doing without schema hierarchies: a recurrent connectionist approach to normal and impaired routine sequential action. *Psychological review*, 111(2):395–429.
- Cooper, R. and Shallice, T. (2000). Contention scheduling and the control of routine activities. *Cognitive neuropsychology*, 17(4):297–338.
- Demiris, Y. and Khadhour, B. (2006). Hierarchical attentive multiple models for execution and recognition of actions. *Robotics and autonomous systems*, 54(5):361–369.
- Fischer, K. W. (1980). A Theory of Cognitive Development: The Control and Construction of Hierarchies of Skills. *Psychological Review*.
- French (1999). Catastrophic forgetting in connectionist networks. *Trends Cogn Sci*, 3(4):128–135.
- Fuster, J. M. (2001). The prefrontal cortex—an update: time is of the essence. *Neuron*, 30:319–333.
- Graybiel, A. M. (1998). The basal ganglia and chunking of action repertoires. *Neurobiol Learn Mem*, 70(1-2):119–136.
- Graziano, M. (2006). The organization of behavioral repertoire in motor cortex. *Annu Rev Neurosci*, 29:105–134.
- Hart, S. and Grupen, R. (2011). Learning generalizable control programs. *IEEE Transactions on Autonomous Mental Development*, 3(1):216–231.
- McCloskey, M. and Cohen, N. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. In Bower, G. H., editor, *The psychology of learning and motivation*, volume 24, pages 109–165. Academic Press, San Diego, CA.

- Meunier, D., Lambiotte, R., and Bullmore, E. T. (2010). Modular and hierarchically modular organization of brain networks. *Front Neurosci*, 4:200.
- Miller, E. K. and Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annu Rev Neurosci*, 24:167–202.
- Miller, G. A., Galanter, E., and Pribram, K. H. (1960). *Plans and the structure of behavior*. Adams-Bannister-Cox, New York.
- Redgrave, P. and Gurney, K. (2006). The short-latency dopamine signal: a role in discovering novel actions? *Nature Reviews Neuroscience*, 7(12):967–975.
- Schneider, D. W. and Logan, G. D. (2006). Hierarchical control of cognitive processes: switching tasks in sequences. *Journal of experimental psychology. General*, 135(4):623–40.
- Singh, S. (1992). Transfer of learning by composing solutions of elemental sequential tasks. *Machine Learning*, 8(3):323–339.
- Yamashita, Y. and Tani, J. (2008). Emergence of functional hierarchy in a multiple timescale neural network model: a humanoid robot experiment. *PLoS computational biology*, 4(11):e1000220.
- Yin, H. H. and Knowlton, B. J. (2006). The role of the basal ganglia in habit formation. *Nat Rev Neurosci*, 7(6):464–476.
- Zacks, J. M., Speer, N. K., Swallow, K. M., Braver, T. S., and Reynolds, J. R. (2007). Event perception: a mind-brain perspective. *Psychological bulletin*, 133(2):273–93.