

Intrinsically Motivated Learning Systems: An Overview

Gianluca Baldassarre and Marco Mirolli

Institute of Cognitive Sciences and Technologies, CNR, Roma, Italy
{gianluca.baldassarre,marco.mirolli}@istc.cnr.it

Abstract. This chapter introduces the field of intrinsically motivated learning systems and illustrates the content, objectives, and organization of the book. The chapter first expands the concept of intrinsic motivations, then introduces a taxonomy of three classes of intrinsic-motivation mechanisms (based on predictors, on novelty detection, and on competence), and finally introduces and reviews the various contributions of the book. The contributions are organized in six parts. The contributions of the first part provide general overviews on the concept of intrinsic motivations, the possible mechanisms that may implement them, and the functions that they can play. The contributions of the second, third, and fourth part focus on the three classes of the aforementioned intrinsic-motivation mechanisms. The contributions of the fifth part discuss mechanisms that are complementary to intrinsic motivations. The contributions of the sixth part introduce tools and experimental frameworks that can be used to investigate intrinsic motivations.

1 Intrinsically motivated learning

The capacity to learn autonomously and in a cumulative fashion is one of the hallmarks of intelligence. Higher mammals, especially when young, engage in a number of exploratory activities with the environment (think about children at play, for example; [von Hofsten, 2007](#)). These activities are not directed to pursue goals that are directly relevant for the the survival and reproduction of the organism, but are driven by *intrinsic motivations*, that is motivations such as curiosity, the interest for novel stimuli or surprising events, and the interest for learning new behaviours. The adaptive value of intrinsically motivated activities lies in that they allow the cumulative acquisition of knowledge and skills that can be later used (e.g., in adulthood) to accomplish fitness enhancing goals ([Baldassarre, 2011](#); [Singh et al., 2010](#)). Intrinsic motivations continue to operate also during adulthood and indeed in humans they underlie life-long learning and other typically human activities such as art and scientific discovery ([Schmidhuber, 2010](#)). Interestingly, intrinsic motivations are also at the basis of processes that strongly affect human well-being, such as the sense of competence, self-determination, and self-esteem ([Ryan and Deci, 2000](#)).

The concept of intrinsic motivation has been introduced in the 1950s in animal psychology ([Harlow, 1950](#)) and has been further elaborated in human psychology ([Deci, 1975](#); [Deci and Ryan, 1985](#)), where it is now widely applied,

especially in the fields of educational psychology (e.g. [Kohn, 1993](#)), developmental psychology (e.g. [Harter, 1981](#)), and organizational psychology ([Houkes et al., 2001](#)).

Furthermore, recent neuroscientific research on the learning signals used by real brains, in particular with respect to the learning of novel actions (e.g. [Redgrave and Gurney, 2006](#)) and the memorization of novel information (e.g. [Lisman and Grace, 2005](#)), is starting to uncover the basic brain mechanisms underlying intrinsic motivations.

Recently, the concept of intrinsic motivation has been introduced in machine learning and developmental robotics ([Barto et al., 2004](#); [Oudeyer et al., 2007](#); [Schembri et al., 2007c](#)) as a means for developing artificial systems that can autonomously learn several different skills. The idea is that intelligent machines and robots could autonomously acquire skills and knowledge under the guidance of intrinsic motivations, and later exploit such knowledge and skills so to accomplish the tasks that are useful for the user in a more efficient and faster way than if they would have to acquire them from scratch. This possibility would clearly enhance the utility of intelligent artificial systems. For this reason, the investigation of intrinsically motivated learning systems is rapidly becoming a central research topic in machine learning and autonomous robotics.

Notwithstanding the importance of this emerging field, up to now there has not been any attempt to draw a unifying picture. The present book provides a general and highly interdisciplinary overview of recent research on intrinsic motivations. While the principal target of the book is the community that tries to synthesize autonomous learning artificial systems, either for applicative purposes or as a means to understand real organisms, the book contains also a number of contributions from researchers in the empirical sciences, because we think that a close dialog between the natural and the artificial sciences is the most fruitful way to address the problems related to intrinsically motivated learning. As a result, the book contains contributions from several different fields of researchers, including robotics ([Hart and Grupen, 2012](#); [Natale et al., 2012](#); [Nehmzow et al., 2012](#); [Oudeyer et al., 2012](#)), machine learning ([Barto, 2012](#); [Schmidhuber, 2012](#)), computational modeling ([Baldassarre and Mirolli, 2012](#); [Gurney et al., 2012](#); [Merrick, 2012](#); [Mirolli and Baldassarre, 2012](#)), psychology ([Schlesinger, 2012](#); [Stafford et al., 2012](#); [Taffoni et al., 2012](#)), and neuroscience ([Dayan, 2012](#); [Otmakova et al., 2012](#); [Redgrave et al., 2012](#)).

The research on intrinsic motivations is indeed very new and as such it is still not fully structured. For example, within conferences and projects concerning intrinsic motivations it is not rare to see researchers animatedly debating the very notion of intrinsic motivation, its functions, and the possible mechanisms that can be used to implement them. In this respect, much effort was spent to improve the clarity of the presentation and the use of a common language to facilitate comparison. As a result, the book is also a means to foster the emergence of common views at the broad theoretical and taxonomical level.

Overall, the book has two fundamental goals. The first aim is to present the state of the art of the research on intrinsically motivated learning. To this

purpose, it collects the contributions of the majority of the most important researchers of the field, who have achieved the most important milestones within it. The authors were asked to review their major contributions and to frame them within the broader relevant literature. These criteria assured that the book covers the important research in the field. The second aim of the book is to identify the scientific and technological open challenges related to intrinsically motivated learning and the most promising research directions to tackle these challenges. In this respect, all contributors were asked to not only presents a review of the work done within their topic of interest, but also to leverage on this to present the most important challenges raised by such investigations, and, when possible, to envisage possible solutions for them.

2 Structure of the book

Beyond this introductory chapter, the book is structured in five Parts. Part ?? contains contributions that introduce the reader to the general issues of intrinsic motivations. This allows the reader to familiarize with the notion of intrinsic motivation, with the literature on them, with the different functions they can play within organisms and intelligent artificial system, and with the different classes of mechanisms that can be used to implement them.

Parts ??, ??, and ?? constitute the “central” parts of the book. Each of these parts is dedicated to a different major class of intrinsic-motivation mechanisms. While according to us the fundamental function of intrinsic motivations is that of guiding the learning system in the acquisition of novel skills, different mechanisms can be envisaged to subserve this function (see [Mirolli and Baldassarre, 2012](#)). In particular, the first two classes of mechanisms (prediction-based and novelty-based, discussed in part ?? and ??, respectively) are based on some form of knowledge that the learning system is acquiring, while the third class (discussed in ??) is directly based on the acquisition of competence. We now briefly expand the relation of the three parts with these classes of mechanisms.

Part ?? focuses on intrinsic motivation mechanisms relying on “diachronic knowledge”, that is on the *prediction* of future states based on current states and (possibly) planned actions: in this respect, experiences that are less predictable are potentially more interesting than those more predictable, i.e. well known. For example, these types of intrinsic motivations may allow a child to focus her attention on a bottle that unexpectedly felt on the ground after she hit it by mistake. This focusing may allot her to learn a new potential effect of her actions, and the knowledge so acquired might be useful in a later stage to perform various tasks as an adult. This is a key capability at the core of human intelligence, flexibility, and creativity, maximally apparent in children at play. To date, prediction-based mechanisms are surely the most developed and studied kind of intrinsic motivation mechanisms.

Part ?? focuses on intrinsic motivation mechanisms relying on “synchronic knowledge”, that is on the sensory appearance and inner structure of world states in a given instant: in particular, these mechanisms classify new experiences on

a continuous scale from novel, hence interesting, to familiar, hence less/not interesting. These mechanisms are for example related to the fact that you see a picture for the first time (e.g., an original painting) and this attracts your attention so that your brain can store information about it; or to the fact that you see two familiar objects together for the first time (e.g., a pigeon on your desktop when you get to your office in the morning). This novelty detection is fundamental to focus learning resources on novel patterns. This kind of knowledge-based mechanisms have been less studied in robotics than those based on prediction, but might nonetheless play an important role for intrinsically motivated learning both in natural and artificial systems.

Part ?? focuses on competence-based intrinsic motivation mechanisms that measure the acquisition of competence, that is of the skills (*sensorimotor mappings* or *action policies*) that allow the agent to act on, and change, the world as desired. For example, succeeding in inserting a toy-cube in the appropriate hole of a toy-box can motivate a child to keep on playing with the toys (Ornkloo and Hofsten, 2006). Competence-based intrinsic motivations, whose existence and importance has recently been highlighted, are an emergent new important sub-topic within the field of intrinsic motivations. Their importance is due to the fact that they directly measure the acquisition of skills that allow to best act in the world. For this reason, they represent fundamental mechanisms for driving the maximization of skill acquisition.

Part ?? is dedicated to mechanisms complementary to intrinsic motivations, namely to mechanisms that, working in synergy with intrinsic motivations, may enhance their functioning. Indeed, learning processes supporting cumulative learning in organisms, in particular in children, do not work in isolation but rely upon other mechanisms that can implement other important sub-functions that in turn can support learning. Examples of this are mechanisms for bottom-up attention guidance that support intrinsic motivations in finding novel patterns and information, and developmental constraints like the use of appropriate sensorimotor primitives and space representations, maturational processes, and social guidance.

Finally, Part ?? is dedicated to research “tools” that can be used for investigating intrinsic motivations. Different kinds of tools are covered by the different contributions: the first contribution is on a behavioural experimental paradigm usable with rats, humans, and robots; the second contribution is on a mechatronic board usable for running behavioural experiments with monkeys, children and robots; the third contribution involves a robotic hardware platform (the iCub robot) that can be used to build and test developmental models on intrinsic motivations. The significance of this part of the book lies in the “infancy” of the field of intrinsic motivations: the development of new tools of research are of great value, and their spreading among the community is of the most importance for the further development of the field.

3 Contents

3.1 Part ??: General overviews on intrinsic motivations

This first part of the book contains three contributions that introduce the reader to the general issues related to intrinsic motivations.

Intrinsic motivation and reinforcement learning. The first contribution is by **Andrew Barto**, who introduced the notion of intrinsic motivations in the fields of machine learning and robotics (e.g., [Barto et al., 2004](#); [Singh et al., 2005](#)). As you will see from the various contributions of the book, a fair amount of computational work on intrinsic motivations is based on the computational reinforcement learning framework. After introducing both reinforcement learning and the psychological literature on intrinsic motivation, Barto’s contribution explains why the reinforcement learning framework is particularly well suited for incorporating intrinsic motivations. Moreover, it discusses an evolutionary perspective according to which there is a continuum between extrinsic and intrinsic motivations as the two kinds of motivations differ only in how directly they are related to the evolutionary success of the behaviours they prompt: while extrinsic motivations drive the agent to pursue activities more directly related to survival and reproduction, intrinsic motivations lead the agent to develop behaviours less directly related to basic needs, such as play (see [Baldassarre, 2011](#) for a proposal of a stronger distinction between the two types of motivations done on this basis).

Functions and Mechanisms of Intrinsic Motivations: The Knowledge vs. Competence Distinction. The second contribution of this part, by **Marco Mirolli** and **Gianluca Baldassarre**, clarifies the distinction between knowledge and competence with respect to both the functions and the mechanisms of intrinsic motivations. The contribution starts by reviewing both the psychological and the computational literature on intrinsic motivations on the basis of this distinction. Then, it proposes that the *general function* of all kinds of intrinsic motivations is allowing intelligent systems to acquire skills rather than knowledge. However, especially in hierarchical and modular learning systems, different sub-function can be identified, which can be implemented by different mechanisms. In particular, the paper suggests that knowledge-based intrinsic motivation mechanisms can be used for discovering which skills can be acquired, whereas competence-based mechanisms can be used for deciding which skills to train and when.

Exploration from Generalization Mediated by Multiple Controllers. The last contribution of this part is from **Peter Dayan**, one of the most active computational theoreticians on brain and behaviour. In his contribution, he gives a general perspective on intrinsic motivations from the point of view of recent computationally-oriented research on brain and behaviour, in particular based on reinforcement learning. In particular, the contribution relates intrinsic motivations to the reinforcement learning notions of exploration bonuses ([Kakade and](#)

Dayan, 2002) and generalization. Bonuses can arise from the generalisation of knowledge from known worlds to successively experienced environments. Different controllers endowed with different architectures and implementing different functions can realize different bonuses. This normative grounding contributes to define the structure of the field, and to establish links between various concepts in reinforcement learning, psychology, and neuroscience.

3.2 Part ??: Prediction-based intrinsic motivation mechanisms

This part is formed by three contributions focused on intrinsic-motivation mechanisms based on measures of knowledge intended as prediction capability.

Maximizing Fun by Creating Data with Easily Reducible Subjective Complexity. This contribution is by **Jürgen Schmidhuber**, who has made seminal work on intrinsic motivations in the 90s who has made seminal work on intrinsic motivations in the 90s (Schmidhuber, 1991a,b) and has later been developing a coherent theory on intrinsic motivations from the machine learning point of view (Schmidhuber, 2010). The core idea of his work is that an agent can use the improvement of prediction capability or, at a more abstract level, the improvement in the compression of the internal representation of information, as an intrinsic reinforcement. In this way, the agent will learn to act so as to maximize the agent’s information intake. In his chapter he presents an extensive review of all the works that generated such theory during the last 20 years, gives a clear and direct presentation of the key principles of such theory, reviews the major experimental results, and proposes interesting theoretical implications of the theory (e.g., to explain a number of different phenomena, from scientific discovery to creativity in art).

The Role of the Basal Ganglia in Discovering Novel Actions. The second contribution of this part is from **Peter Redgrave** (a neuroscientist) and colleagues, and presents an important neuroscientific theory on action discovery and learning (Redgrave and Gurney, 2006; Redgrave et al., 2011) that is clearly related to the notion of intrinsic motivation mechanisms related to prediction. The theory challenges the major tenets of the standard theory related to the neuromodulator dopamine and reinforcement learning processes in the brain. By taking into account empirical data that are difficult to explain with the received view, the proposed theory provides a coherent picture that explains how action discovery might be realized in real brains on the basis of intrinsic motivation signals. These signals are produced by events (e.g., the sudden onset of a light detected by the superior colliculus, a brain area closely connected with the retina) that are unpredicted by the agent, and are capable of guiding the learning processes that lead the agent to acquire the actions that cause the unexpected events to happen.

Action Discovery and Intrinsic Motivation: A Biologically Constrained Formalisation. The last contribution of this part is by **Kevin Gurney** and colleagues,

and develops and formalize the theoretical framework proposed in the previous contribution. The main idea of this contribution is that intrinsic motivation learning signals, in particular those caused by unexpected salient events, drive the acquisition of different types of information in the brain, in particular the acquisition of “forward models” (that is, the capacity to predict future states on the basis of current states and planned actions), and “inverse models” (that is, the capacity to associate the correct actions to desired outcomes and perceived world states). While forward models contribute to establish the unpredictability of salient events, and hence to generate the related (intrinsic) learning signals, inverse models implement the capacity to act so to realize desired states, which can be acquired on the basis of intrinsic motivations. In general, this contribution represents a biologically motivated formal framework for modeling intrinsically motivated action discovery and acquisition.

3.3 Part ??: Novelty-based intrinsic motivation mechanisms

This part is formed by three contributions that are focused on intrinsic-motivation mechanisms based on measures of knowledge intended as novelty/familiarity of percepts.

Novelty Detection as an Intrinsic Motivation for Cumulative Learning Robots. The first contribution of this part is from **Ulrich Nehmzow**’s research team. Nehmzow, who recently passed away, has done pioneering work with autonomous robots and *habituable neural networks* (Neto and Nehmzow, 2007; Vieiraneto and Nehmzow, 2007). These are systems that allow a robot to learn new perceptual patterns, get progressively familiar with them, and then exploit this knowledge for accomplishing various tasks (e.g., a patrolling robot might understand if a certain infrastructure has been damaged from the last check). This paper first clarifies the relationship existing between intrinsic motivations and novelty detection, for example as studied in machine learning and statistical research. Then it presents a critical review of the methods of novelty detection previously developed by the authors. Finally it illustrates the key open challenges that need to be considered in the design of novelty detection systems, and the types of solutions that might be investigated to face them.

Novelty and Beyond: Towards Combined Motivation Models and Integrated Learning Architectures. The second contribution of this part is by **Kathryn Merrick**, who has pioneered the application of novelty-based intrinsic motivations based on habituable neural networks in both robotics and interactive software and video-games (Merrick and Maher, 2009). This contribution presents a critical review of the author’s work in the field, including attempts to model competence seeking motivations. In particular, it discusses four architectures that combine components related to intrinsic motivation mechanisms in different ways. Moreover, the contribution frames the reviewed applications within the broader perspective of “motivated agents”, proposing research roadmaps to build agents that integrate various types of intrinsic and extrinsic motivations.

The Hippocampal-VTA Loop: The Role of Novelty and Motivation in Controlling the Entry of Information into Long-Term Memory. The last contribution of the chapter is by **John Lisman** (a theoretical neuroscientist) and co-workers, who have carried out fundamental work on the brain circuits and mechanisms underlying the acquisition of novel information (Lisman and Grace, 2005). The main idea proposed here is that the complex cerebral loop that involves the hippocampus (fundamental for the acquisition and recognition of novel patterns) and the dopaminergic neurons (providing motivational and learning signals) allows the brain to understand whether a certain perceptual pattern is new so to enhance its memorisation. The contribution also discusses the most advanced open problems in the biological investigation of these issues, and possible ways to face them. Interestingly, the mechanisms presented in this work may represent the brain basis of novelty-based intrinsic motivations, which complement the mechanisms discussed in the contributions by Redgrave and coworkers related to prediction-based intrinsic motivations.

3.4 Part ??: Competence-based intrinsic motivations mechanisms

This part is formed by two contributions that focus on intrinsic-motivation mechanisms based on measures of the systems ability to reach a certain goal.

Deciding Which Skill to Learn When: Temporal-Difference Competence-Based Intrinsic Motivation (TD-CB-IM). The first contribution of this part, by **Gianluca Baldassarre** and **Marco Mirolli**, reviews one of the first models on competence-based intrinsic motivations (Schembri et al., 2007a,b,c). This is also the first model casting intrinsic motivations in an evolutionary perspective, where intrinsic motivations play the key adaptive function of guiding learning of skills in the absence or paucity of extrinsic motivations (e.g., as it happens in childhood of humans); the skills so acquired can be readily exploited in later phases of life (e.g., in adulthood). This contribution first discusses the importance of competence-based intrinsic motivations for autonomous skill acquisition by focusing, in particular, on the problem of deciding which skill to train in each moment. Then the contribution reviews the work of the authors on a competence-based intrinsic motivations mechanism that is based on the idea of using the TD-error signal of standard reinforcement learning models as an intrinsic reinforcement for a selector that has to decide which skill to train when. Finally, the contribution discusses a number of open issues related to competence-based intrinsic motivations.

Intrinsically Motivated Affordance Discovery and Modeling. The last contribution of this part is by **Stephen Hart** and **Roderic Grupen**, two roboticists who have proposed some of the most advanced robotic applications based on intrinsic motivations (Hart and Grupen, 2011). Their system is based on a two-level hierarchical system where the lower-level is formed by dynamic motor primitives and the higher level is a reinforcement learning system that learns to suitably assemble the primitives of the lower level represented in an abstract format. The

system includes a model that learns to predict in which states the motor primitives can be executed with success. The enlargement of the set of states where this happens is used to generate learning signals that so mark the improvement of the primitives and the recognition of their enhanced applicability (called here “affordance”). The authors review various experiments showing that the system can guide a robot to autonomously acquire a rich repertoire of primitives (e.g., for object reaching, grasping and displacement), and discuss possible improvement of the systems in future work.

3.5 Part ??: Mechanisms complementary to intrinsic motivations

This part is formed by two contributions that present some mechanisms that can work in synergy with intrinsic-motivation systems to foster the learning processes guided by them.

Intrinsically Motivated Learning of Real World Sensorimotor Skills with Developmental Constraints. The first contribution of this part is from **Pierre-Yves Oudeyer, Adrien Baranes, Frédéric Kaplan**, who have had a fundamental role in promoting the study of intrinsic motivations in the robotic community (e.g., [Oudeyer and Kaplan, 2007](#); [Oudeyer et al., 2007](#)). The contribution reviews the principles and main outcomes of recent research by the authors on how intrinsic motivations can be complemented by other factors so to improve the capacity of robots to undergo an truly open-ended development. In particular, the authors present several developmental constraints that can be useful for overcoming the difficulties given by the high dimensionality and unboundedness of the sensorimotor space of robots, including sensorimotor primitives, task space representations, maturational processes, and social guidance. The nature and operation of these constraints is illustrated on the basis of robot experiments, which are also used to highlight the open challenges of this research.

Investigating the Origins of Intrinsic Motivation in Human Infants. The second contribution of this part is by **Matthew Schlesinger** (a developmental psychologist), who proposes a new perspective on novelty-based intrinsic motivations by investigating the mechanisms that guide attention itself. Consider for example the eyes of a child who does not have a particular goal in mind and is engaged in exploring the environment. With high chances the child will direct the gaze to areas where the surface of objects has a high contrast, a high luminosity, or motion: all these areas have a high saliency as they are informative with respect to the surrounding portions of the environment (e.g., a high contrast on an homogeneous surface, or movement in a static background, etc.). Goal directed, top-down attentional skills are gradually acquired by building on this basic capabilities and under the guidance of more sophisticated forms of novelty detection as those reviewed in previous Parts. Three theoretical developmental perspectives that try to explain the development of visual exploration are compared. Furthermore, the contribution reviews a model on the development of perceptual completion which offers a case study on the development of

visual exploration and the role of oculomotor skills. Finally, a number of open challenges that are suggested by this work are discussed.

3.6 Part ??: Tools for research on intrinsic motivations

The final part of the book presents three contributions that illustrate new important research tools for investigating intrinsic motivations.

A Novel Behavioural Task for Researching Intrinsic Motivations. The first contribution is from **Tom Stafford** and co-workers (a team of experimental psychologists) and illustrates a new experimental paradigm, called the “joystick task”. In contrast to standard reinforcement learning paradigms, which have been developed to study how the deployment of already acquired actions is affected by extrinsic rewards, the joystick task has been explicitly designed to investigate intrinsically-motivated action discovery and acquisition. Flexibility is one of the important features of the novel experimental paradigm: the paradigm allows generating a multitude of experimental conditions with the desired level of difficulty. Moreover, different versions of the paradigm are proposed for studying action learning with different types of participants and species including rats, monkeys, humans, human patients, and robots. Beyond discussing the rationale of the task, its potential for future research on intrinsically motivated action learning, its main features, and its range of applicability, the contribution reviews also some preliminary results that have been obtained with different kinds of subjects (rats and humans), thus showing the potential of the task for comparative research.

The “Mechatronic Board”: A Tool to Study Intrinsic Motivations in Humans, Monkeys, and Humanoid Robots. The second contribution of this part is by **Fabrizio Taffoni** and colleagues, which include researcher from four different groups: a group of bio-engineers led by **Eugenio Guglielmelli**, a group of developmental neuroscientist and psychologists led by **Flavio Keller**, a group of primatologists led by **Elisabetta Visalberghi**, and a group of computational modelists led by **Gianluca Baldassarre** and **Marco Mirolli**. The contribution illustrates a new experimental tool (called the mechatronic board) that has been explicitly designed to investigate the acquisition of actions on the basis of free exploration and intrinsic motivations. The tool has been designed (in terms of materials, size, modularity, complexity, etc.) to be used with different types of participants and species: monkeys, children, human adults, and robots. In the mechatronic board, different types of manipulators can be inserted (the number and kind of objects that might be used to elicit interest and exploration is open), and a feedback panel provides different kinds of effects (sounds, lights, opening of boxes) depending on the actions of the participants. All aspects of the board are programmable and all events can be automatically recorded. For these reasons, the mechatronic board represents a important research tool for the investigation of intrinsic motivations.

The iCub Platform: A Tool for Studying Intrinsically Motivated Learning. The last contribution is from **Lorenzo Natale** and co-workers, a group of roboticists leaded by **Giorgio Metta** and **Giulio Sandini**. The contribution presents the humanoid robot iCub, a recently developed robotic platform designed to model developmental phenomena in children. The robot’s size, physical structure, level of sophistication of the sensors, and the motor plant have all been designed to this purpose. The design of the robot hardware, its drivers, and the library of control programs that have been developed for it are all open-source, as the robot has been developed with the aim of becoming an open research tool for developmental robotic research. Indeed, the robot is already in use in several robotic labs and is becoming the standard tool for research on bio-inspired autonomous robotics in Europe. For all these reasons the iCub represents a valuable tool for robotic research on intrinsic motivations.

4 Conclusion

In summary, this book aims at becoming a reference point in the field of intrinsically motivated learning by: (a) establishing the theoretical foundations of the field, (b) describing the most significant achievements accomplished so far and the most important open challenges, (c) illustrating the scientific and technological potentials of the field, and (d) presenting new methods and tools for carrying out investigations in the field.

Given the importance of intrinsic motivations for humans and for autonomous intelligent robots and machines, we hope the book will recruit new researchers and foster further investigations in this emergent and fascinating research field.

Acknowledgments

This paper and a large part of the research reported in this book have been supported by the Project “IM-CLeVeR - Intrinsically Motivated Cumulative Learning Versatile 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 “Acknowledgment” 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

- Baldassarre, G. (2011). What are intrinsic motivations? a biological perspective. In Cangelosi, A., Triesch, J., Fasel, I., Rohlfing, K., Nori, F., Oudeyer, P.-Y., Schlesinger, M., and Nagai, Y., editors, *Proceedings of the International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob-2011)*, pages E1–8. IEEE, Piscataway, NJ.
- Baldassarre, G. and Mirolli, M. (2012). Deciding which skill to learn when: Temporal-difference competence-based intrinsic motivation (td-cb-im). In Baldassarre, G. and Mirolli, M., editors, *Intrinsically Motivated Learning in Natural and Artificial Systems*. Springer-Verlag, Berlin. (this volume).
- Barto, A., Singh, S., and Chentanez, N. (2004). Intrinsically motivated learning of hierarchical collections of skills. In *International Conference on Developmental Learning (ICDL)*, LaJolla, CA.
- Barto, A. G. (2012). Intrinsic motivation and reinforcement learning. In Baldassarre, G. and Mirolli, M., editors, *Intrinsically Motivated Learning in Natural and Artificial Systems*. Springer-Verlag, Berlin. (this volume).
- Dayan, P. (2012). Exploration from generalisation mediated by multiple controllers. In Baldassarre, G. and Mirolli, M., editors, *Intrinsically Motivated Learning in Natural and Artificial Systems*. Springer-Verlag, Berlin. (this volume).
- Deci, E. (1975). *Intrinsic motivation*. Plenum, New York.
- Deci, E. L. and Ryan, R. M. (1985). *Intrinsic motivation and self-determination in human behavior*. Plenum, New York.
- Gurney, K., Lepora, N., Shah, A., Koene, A., and Redgrave, P. (2012). Action discovery and intrinsic motivation: a biologically constrained formalisation. In Baldassarre, G. and Mirolli, M., editors, *Intrinsically Motivated Learning in Natural and Artificial Systems*. Springer-Verlag, Berlin. (this volume).
- Harlow, H. F. (1950). Learning and satiation of response in intrinsically motivated complex puzzle performance by monkeys. *Journal of Comparative and Physiological Psychology*, 43:289–294.
- Hart, S. and Grupen, R. (2011). Learning generalizable control programs. *IEEE Transactions on Autonomous Mental Development*, 3(1).
- Hart, S. and Grupen, R. (2012). Intrinsically motivated affordance discovery and modeling. In Baldassarre, G. and Mirolli, M., editors, *Intrinsically Motivated Learning in Natural and Artificial Systems*. Springer-Verlag, Berlin. (this volume).

- Harter, S. (1981). A new self-report scale of intrinsic versus extrinsic orientation in the classroom: Motivational and informational components. *Developmental Psychology*, 17:100–112.
- Houkes, I., Janssen, P., de Jong, J., and Nijhuis, F. (2001). Specific relationships between work characteristics and intrinsic work motivation, burnout and turnover intention: A multi-sample analysis. *European Journal of Work and Organizational Psychology*, 10:1–23.
- Kakade, S. and Dayan, P. (2002). Dopamine: generalization and bonuses. *Neural Netw*, 15(4-6):549–559.
- Kohn, A. (1993). *Punished by Rewards*. Houghton Mifflin.
- Lisman, J. E. and Grace, A. A. (2005). The hippocampal-vta loop: controlling the entry of information into long-term memory. *Neuron*, 46(5):703–713.
- Merrick, K. and Maher, M. (2009). *Motivated reinforcement learning: curious characters for multiuser games*. Springer-Verlag, Berlin.
- Merrick, K. E. (2012). Novelty and beyond: towards combined motivation models and integrated learning architectures. In Baldassarre, G. and Mirolli, M., editors, *Intrinsically Motivated Learning in Natural and Artificial Systems*. Springer-Verlag, Berlin. (this volume).
- Mirolli, M. and Baldassarre, G. (2012). Functions and mechanisms of intrinsic motivations: The knowledge versus competence distinction. In Baldassarre, G. and Mirolli, M., editors, *Intrinsically Motivated Learning in Natural and Artificial Systems*. Springer-Verlag. (this volume).
- Natale, L., Nori, F., Metta, G., Fumagalli, M., Ivaldi, S., Pattacini, U., Randazzo, M., Schmitz, A., and Sandini, G. (2012). The icub platform: a tool for studying intrinsically motivated learning. In Baldassarre, G. and Mirolli, M., editors, *Intrinsically Motivated Learning in Natural and Artificial Systems*. Springer-Verlag, Berlin. (this volume).
- Nehmzow, U., Gatsoulis, Y., Kerr, E., Condell, J., Siddique, N. H., and McGinnity, M. T. (2012). Novelty detection as an intrinsic motivation for cumulative learning robots. In Baldassarre, G. and Mirolli, M., editors, *Intrinsically Motivated Learning in Natural and Artificial Systems*. Springer-Verlag, Berlin. (this volume).
- Neto, H. V. and Nehmzow, U. (2007). Visual novelty detection with automatic scale selection. *Robotics and Autonomous Systems*, 55(9):693–701.
- Ornkloo, H. and Hofsten, C. v. (2006). Fitting objects into holes: on the development of spatial cognition skills. *Developmental Psychology*, 43(2):404–416.
- Otmakova, N., Duzel, E., Deutch, A. Y., and Lisman, J. E. (2012). The hippocampal-VTA loop: the role of novelty and motivation in controlling the entry of information into long-term memory. In Baldassarre, G. and Mirolli, M., editors, *Intrinsically Motivated Learning in Natural and Artificial Systems*.

- Springer-Verlag, Berlin. (this volume).
- Oudeyer, P.-Y., Banares, A., and Frédéric, K. (2012). Intrinsically motivated learning of real world sensorimotor skills with developmental constraints. In Baldassarre, G. and Mirolli, M., editors, *Intrinsically Motivated Learning in Natural and Artificial Systems*. Springer-Verlag, Berlin. (this volume).
- Oudeyer, P.-Y. and Kaplan, F. (2007). What is intrinsic motivation? a typology of computational approaches. *Frontiers in Neurorobotics*, 1:6.
- Oudeyer, P.-Y., Kaplan, F., and Hafner, V. V. (2007). Intrinsic motivation systems for autonomous mental development. *IEEE Transactions in Evolutionary Computation*, 11(2):265–286.
- Redgrave, P. and Gurney, K. (2006). The short-latency dopamine signal: a role in discovering novel actions? *Nature Review Neuroscience*, 7(12):967–975.
- Redgrave, P., Gurney, K., Stafford, T., Thirkettle, M., and Lewis, J. (2012). The role of the basal ganglia in discovering novel actions. In Baldassarre, G. and Mirolli, M., editors, *Intrinsically Motivated Learning in Natural and Artificial Systems*. Springer-Verlag, Berlin. (this volume).
- Redgrave, P., Vautrelle, N., and Reynolds, J. N. J. (2011). Functional properties of the basal ganglia’s re-entrant loop architecture: selection and reinforcement. *Neuroscience*.
- Ryan, R. M. and Deci, E. L. (2000). Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. *Am Psychol*, 55(1):68–78.
- Schembri, M., Mirolli, M., and Baldassarre, G. (2007a). Evolution and learning in an intrinsically motivated reinforcement learning robot. In Almeida e Costa Fernando, Rocha, L. M., Costa, E., Harvey, I., and Coutinho, A., editors, *Advances in Artificial Life. Proceedings of the 9th European Conference on Artificial Life (ECAL2007)*, volume 4648 of *Lecture Notes in Artificial Intelligence*, pages 294–333. Springer Verlag, Berlin. Lisbon, Portugal, September 2007.
- Schembri, M., Mirolli, M., and Baldassarre, G. (2007b). Evolving childhood’s length and learning parameters in an intrinsically motivated reinforcement learning robot. In Berthouze, L., Dhristiopher, P. G., Littman, M., Kozima, H., and Balkenius, C., editors, *Proceedings of the Seventh International Conference on Epigenetic Robotics*, volume 134, pages 141–148. Lund University, Lund.
- Schembri, M., Mirolli, M., and Baldassarre, G. (2007c). Evolving internal reinforcers for an intrinsically motivated reinforcement-learning robot. In Demiris, Y., Mareschal, D., Scassellati, B., and Weng, J., editors, *Proceedings of the 6th International Conference on Development and Learning*, pages E1–6, London. Imperial College.

- Schlesinger, M. (2012). Investigating the origins of intrinsic motivations in human infants. In Baldassarre, G. and Mirolli, M., editors, *Intrinsically Motivated Learning in Natural and Artificial Systems*. Springer-Verlag, Berlin. (this volume).
- Schmidhuber, J. (1991a). Curious model-building control systems. In *Proceedings of the International Joint Conference on Neural Networks*, volume 2, pages 1458–1463.
- Schmidhuber, J. (1991b). A possibility for implementing curiosity and boredom in model-building neural controllers. In Meyer, J.-A. and Wilson, S., editors, *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*, pages 222–227, Cambridge, MA. MIT Press.
- Schmidhuber, J. (2010). Formal theory of creativity, fun, and intrinsic motivation (1990–2010). *IEEE Transactions on Autonomous Mental Development*, 2(3):230–247.
- Schmidhuber, J. (2012). Maximizing fun by creating data with easily reducible subjective complexity. In Baldassarre, G. and Mirolli, M., editors, *Intrinsically motivated learning in natural and artificial systems*. Springer-Verlag, Berlin. (this volume).
- Singh, S., Barto, A., and Chentanez, N. (2005). Intrinsically motivated reinforcement learning. In Saul, L. K., Weiss, Y., and Bottou, L., editors, *Advances in Neural Information Processing Systems 17: Proceedings of the 2004 Conference*, Cambridge, MA. The MIT Press. Vancouver, BC, Canada.
- Singh, S., Lewis, R., Barto, A., and Sorg, J. (2010). Intrinsically motivated reinforcement learning: An evolutionary perspective. *IEEE Transactions on Autonomous Mental Development*, 2(2):70–82.
- Stafford, T., Walton, T., Hetherington, L., Thirkettle, M., Gurney, K., and Redgrave, P. (2012). A novel behavioural task for researching intrinsic motivation. In Baldassarre, G. and Mirolli, M., editors, *Intrinsically Motivated Learning in Natural and Artificial Systems*. Springer-Verlag, Berlin. (this volume).
- Taffoni, F., Formica, D., Schiavone, G., Scordia, M., Tomassetti, A., Polizzi di Sorrentino, E., Sabbatini, G., Truppa, V., Mirolli, M., Baldassarre, G., Visalberghi, E., Keller, F., and Guglielmelli, E. (2012). The “mechatronic board”: a tool to study intrinsic motivation in humans, animals and robots. In Baldassarre, G. and Mirolli, M., editors, *Intrinsically Motivated Learning in Natural and Artificial Systems*. Springer-Verlag, Berlin. (this volume).
- Vieiraneto, H. and Nehmzow, U. (2007). Visual novelty detection with automatic scale selection. *Robotics and Autonomous Systems*, 55:693–701.
- von Hofsten, C. (2007). Action in development. *Dev Sci*, 10(1):54–60.