

Preface

Traditionally, artificial evolution, complex systems, and artificial life were separate fields, each with its own research community, but we are now seeing increased intertwinement and hybridization among them. It is now hard to imagine any work in one of these areas that does not refer to techniques or theoretical results normally considered to belong to one of the other two.

Evolution and complexity characterize both biological life and artificial life, whenever direct modeling of biological processes is pursued or populations of interacting artificial biologically inspired entities are created, from which complex behaviors can emerge and evolve.

The latter consideration, besides offering a proof of the tight connections existing between these disciplines, also gives an idea of the breadth of the related topics of interest, and of the different study viewpoints, ranging from purely scientific and exploratory approaches aimed at verifying biological theories to technology-focused applied research aimed at solving difficult real-world problems raised by practical and industrial tasks.

As a result of the hybridization between these disciplines, the same is happening to the corresponding research communities, and common conference tracks and workshops are being organized worldwide.

We conceived the idea of editing a book to collect contributions offering a wide panoramic view of the opportunities that cooperation among the three disciplines can produce when we started organizing the fourth edition of WIVACE, the Italian Workshop on Evolutionary Computing and Artificial Life, which was held in Parma in February 2012.

This edited book includes invited chapters from leading scientists in the fields of artificial life, complex systems, and evolutionary computing, aimed at authoritatively introducing readers to some of the main research topics that are not only shared by the three research fields, but that, in some cases, can only exist, thanks to the contribution of all three disciplines.

The book also contains a selection of the best papers presented at WIVACE 2012, thoroughly revised and extended by the authors.

The contributions, ranging from fundamental theoretical issues to state-of-the-art real-world applications, have been organized into five parts, based either on their kind (research-oriented or application-oriented) or topics (biological modeling, mind and society, evolution).

This subdivision denotes how the modeling of biological systems, both anatomical-functional and social, constitutes the wide general field in which evolution, complexity, and artificial life coexist most closely and which is studied most frequently.

This is not surprising, as artificial life and evolution are biologically inspired disciplines per se, while it is hard to imagine something which is more complex than the functioning of the human body and, possibly even more so, human society.

In the following we provide a brief overview of each chapter. On one hand, to allow readers to figure out a general picture of the composite research field induced by the interactions of the three disciplines, and, on the other hand, to let readers quickly locate the chapters which look most interesting to them.

Part I: Research Issues

The first chapter, by Domenico Parisi, discusses the challenges posed by constructing embodied artificial agents (i.e., robots) not as practical applications but as a means to understand human behavior. Since there are many ways to construct an artifact that reproduces a single (human) phenomenon, the author argues that robots constructed as theories of behavior should attempt to reproduce many different phenomena at the same time, as the more phenomena an artifact reproduces, the more likely it is that the artifact actually explains reality rather than being a “toy.” Hence, Parisi discusses a number of different phenomena that robots should try to reproduce in order to explain human behavior and that are currently underinvestigated in robotics and artificial life.

The second chapter, by Riccardo Poli and Christopher R. Stephens, deals with the difficult problem of building a theory of evolutionary systems with which we can understand and predict natural evolutionary dynamics. In particular, as a means to develop such a theory, their chapter proposes a technique originally used in statistical physics named coarse graining. The method consists of finding a set of collective variables that may offer a computationally feasible description of a system that has too many degrees of freedom to be analyzed. After describing the general technique, the authors show how to apply it to evolutionary systems in order to describe and understand their complex dynamics.

Part II: Biological Modeling

This part of the book examines the important issue of the organization principles of living systems. In particular, the main focus is on the organization of the relationships among the parts inside living beings: complementary to the reductionist and physical approaches, this kind of strategy searches for the structure of feedbacks that constitute the organizational processes of living beings, and that drive and channel their future changes, i.e., their evolution. In doing so, signal exchange and noise play a fundamental role: what these chapters interestingly show is that these drivers do not counteract each other but rather they integrate in order to allow the emergence of recurrent patterns of activation that, in turn, are an important part of the cells' regulatory processes. Notably these regulating entities are dynamical objects whose monitoring requires us to introduce new concepts and ideas.

The chapter "Models of Gene Regulation: Integrating Modern Knowledge into the Random Boolean Network Framework" introduces the theme of regulatory genetic networks, modeled by means of random Boolean networks—RBN for short—an abstract framework introduced four decades ago and now becoming one of the major models for complex systems due to their interesting dynamical behavior. In recent years interest in this approach has been renewed through important theoretical advances and also, as far as the application to genetics is concerned, by the availability of genome-wide expression data which can be properly described by RBNs. Moreover, the new versions of this framework can now describe complex phenomena like cell differentiation and whole organisms or tissues. This framework is a common feature of the first three chapters of this part of the book. Thus, Christian Darabos, Mario Giacobini, Jason H. Moore, and Marco Tomassini introduce into RBNs abstractions inspired by recent advances in genetics and biology. In particular, they discuss the topological structure of gene relations and the effects of the adjournment strategy on the model results.

Marco Villani and Roberto Serra discuss the stability properties of RBNs, introducing a new measure (attractor sensitivity) that seems particularly relevant for their application to the dynamics of gene regulatory networks. They also review results that show that RBNs can properly account for data on perturbations induced by gene knock-out in real organisms, thus revealing that living beings tend to live in, or close to, critical states. Last but not least, the authors show that adding noise to RBN framework can lead to a nice model of cell differentiation.

Stefano Benedettini, Andrea Roli, Roberto Serra, and Marco Villani show that it is possible to generate (evolve, train) an ensemble of Boolean networks that can accomplish particular requirements, while keeping the other main relevant statistical features of classical RBNs. This ensemble can be designed by means of optimization processes in which metaheuristics can optimize suitable objective functions.

The chapter by Pasquale Stano, Giordano Rampioni, Luisa Damiano, Francesca D'Angelo, Paolo Carrara, Livia Leoni, and Pier Luigi Luisi shares the spirit of these studies, focusing on the chemical communications among cells, seen as autopoietic objects, stressing again the importance of the structure of feedbacks

that constitute the current organizational processes of living beings. The authors moreover introduce the theme biological/chemical information and communication technology (bio/chem-ICT), which aims at extending the well-known field of ICT, up to now classically based on the transmission of electrical or electromagnetic signals, to the biological/chemical world of molecules.

Part III: Mind and Society

This part of the book presents models that deal with organism behavior and sociocultural phenomena. The chapters that make up this part reflect different approaches to artificial life: from biomimetic models (Santucci et al.), to models that try to reveal general principles (Pugliese), to models that are directly related to natural phenomena (Acerbi et al.), to the use of artificial life techniques for applicative purposes (Gigliotta et al.).

Building artificial agents able to autonomously learn new skills and to easily adapt to different and complex environments is an important goal for robotics and machine learning. In their chapter, Vieri G. Santucci, Gianluca Baldassarre, and Marco Mirolli propose that artificial agents with a learning signal that resembles some characteristic of dopaminergic neurons would be an advancement in the development of more autonomous and versatile systems, thanks to the enabling of cumulative learning abilities. To validate this hypothesis they perform experiments with a simulated robotic system that has to learn different skills to obtain rewards, showing that the proposed learning signal is able to drive the cumulative acquisition of different skills in a way that would not be possible otherwise.

The contribution by Francesco Pugliese presents evolutionary robotics models of the development of categorization abilities. Two different experiments are described, one involving mobile robots that perceive the color of the floor through ground sensors, and the other involving a robotic camera that can move on images. In both cases, the environments contain noisy images that must be categorized, and in both cases evolved robots are able to perform the task. More importantly, in both experiments, the robots that during evolution are facilitated by receiving linguistic signals that tell the robot the category of the perceived image achieve better performance than the robots that did not receive any help, even when, after evolution, the facilitating signals are not provided. Hence, the reported simulations suggest that social linguistic input may exert a facilitating role for the individual development of categorization abilities.

It is not only individual behavior that can be explained through computational models, but also social and cultural phenomena. The next contribution, by Alberto Acerbi, Stefano Ghirlanda, and Magnus Enquist, presents simple models of cultural evolution that try to explain cultural dynamics. In particular, the contribution deals with regulatory traits, that is traits that are culturally transmitted but that, in turn, regulate cultural transmission, such as the propensity to copy others or the ability to influence others. The authors study how the evolution of these traits influences

cultural phenomena like the emergence of open or conservative societies or the ups and downs of cultural traits, i.e., fashions, and they conclude that the presence of regulatory traits renders cultural evolution more flexible than genetic evolution, requiring substantially different models.

In the last chapter of this part, Onofrio Gigliotta, Orazio Miglino, Massimiliano Schembri, and Andrea Di Ferdinando show two ways to build serious game systems exploiting the power of artificial life (AL) techniques for educational purposes; in particular, the authors apply agent-based models, neural networks, genetic algorithms, and robotics. In the first approach, neural networks and genetic algorithms are utilized as open tools to guide artificial organisms design, in order to make it possible for users to learn the fundamental principles of autonomous robotics. In the second case AL techniques are used to model game mechanics—e.g., artificial team dynamics and avatar behavior—whereas the user learns psychological leadership theories by governing a team of artificial agents, the followers. The two cases show how AL techniques can boost serious game systems toward a new level of usability in the context of bioinspired evolutionary design processes and in the field of management training.

Part IV: Applications

This part of the book contains five chapters that give readers an idea of the wide range of applications that can be tackled using techniques derived from the disciplines under consideration and from their hybridization. As shown, apart from the immediately perceptible diversity of the application fields, it is interesting to note how the very nature of results belong to different domains, from very concrete and direct industrial applications, among which signal and image processing and analysis techniques are rather popular, to the creation of models which can be used to forecast the actions of drugs and to assess patient life expectations, from software agents developed to live in and patrol virtual environments such as communication networks to physical agents that can interact with biological tissues.

An example of this latter, futuristic kind of application is offered in the chapter by Oleg Semenov, Darko Stefanovic, and Milan N. Stojanovic, where they study the behavior of synthetic nanoscale walkers made with catalytic DNA legs attached to a rigid body, called molecular spiders, which are able to move across a surface propelled by the multivalent chemical interactions of their multiple legs with the surface itself. Molecular spiders may find important use in biomedical applications, such as searching for clinically relevant targets on the surface of a cell. The authors present simulation-based results on the efficiency of concurrent search for multiple targets by multiple molecular spiders, which influence each other's motion through stigmergic processes.

The two following chapters provide examples of the use of evolutionary computation techniques to solve pattern recognition problems of biomedical interest.

Mario Giacobini, Paolo Provero, Leonardo Vanneschi, and Giancarlo Mauri use genetic programming (GP) to analyze the genetic profile of cancer patients to forecast the outcome of the pathology and to tailor therapy individually. The so-called 70-gene signature was analyzed using a number of pattern-recognition techniques, against which the authors' GP-based approach was compared and shown to outperform the others.

Antonia Azzini, Mauro Dragoni, and Andrea G.B. Tettamanzi apply the results of their years-long research into the hybridization between EC and neural nets to a dataset of ECG recordings, acquired using cheap devices and transmitted over low-band connections with a configuration designed to meet the needs of African countries where cardiologists are not available on-site and the quality of communications is often very low. The dataset on which the evolutionary networks have been tested was previously used in a contest, so the results of the hybrid evolutionary neural system could be compared with those obtained by a large number of other approaches, performing comparably to the best of these.

The chapter by Yvonne Bernard, Lukas Klejnowski, David Bluhm, Jörg Hähner, and Christian Müller-Schloer describes a system which also relies on evolutionary computation techniques. Their approach is based on the ideas of a discipline, called organic computing, which studies the development of agents that are able to cooperate and continuously self-adapt to cope with changing environmental conditions. In their chapter the authors develop evolutionary agents that act in the Trusted Desktop Grid, a distributed computing environment, with no central control, which optimizes the sharing of computing resources.

Massively multiplayer online games (MMOGs) are increasingly successful, since they allow players to explore huge virtual worlds and to interact in many different ways, either cooperating or competing. Given the huge and ever-growing number of users, game designers have to apply strong scalable real-time strategies in order to maintain control of the system. Stefano Sebastio, Michele Amoretti, Jose Raul Murga, Marco Picone, and Stefano Cagnoni present a middleware called PATROL, based on a structured peer-to-peer overlay scheme. Among other features, PATROL provides AI-based modules to detect cheating attempts that the decentralized communication infrastructure may favor: in particular, the authors show how honest bots can detect cheating bots in real time, using strategies based on neural networks. The evolutionary agents' community is evaluated in three different situations: coexistence and competition with the best-performing adaptive agents developed in previous studies; behavior in the presence only of evolutionary agents; and, finally, in an environment in which the presence of egoistic agents introduces disturbance into the system.

Part V: Evolution

The last part of the book is dedicated to research issues in evolutionary computation and deals with two hot topics in the field, of which the former is more theory-oriented, while the latter is more technology-related.

Multiobjective optimization consists of the search for a set of so-called non-dominated solutions which provide a sampling of the Pareto front for the problem, i.e., the set of optimal values which can be obtained for each criterion, once the values of the other, independent and usually counteracting, criteria have been set.

Parallel implementation of evolutionary algorithms has always been a popular issue, thanks to their intrinsically parallel structure, but it is presently booming since multicore CPUs and handy development environments have become available, making multicore or general-purpose Graphics Processor Unit (GP-GPU) computing a higher-level task; this contrasts with the effort that was required to develop similar programs for exploiting the low-level features of processors and, especially, of GPUs, using traditional programming languages that did not have specific support for parallel computation.

Hernán Aguirre, Akira Oyama, and Kiyoshi Tanaka describe an evolutionary multiobjective algorithm which tries to bias the convergence of the population of trade-off solutions onto the Pareto front such that they assume a desired statistical distribution. To this end, they propose Adaptive ϵ -Ranking, which iterates a sampling procedure that applies ϵ -dominance with a suitable mapping function. An analysis of the experimental results, made on the functions of the DTLZ family with six objectives, shows that recombination plays a crucial role in finding a set of solutions with the desired distribution.

Finally, Kiyoharu Tagawa proposes a parallel Java implementation of differential evolution for multicore processors, of which two versions are compared: one is demonstrated to be computationally preferable, as it is able to exploit the multiprocessor's parallel capabilities more efficiently, while the other appears to provide more consistent results over different functions and thread configurations. Both versions benefit from a significant speed-up with respect to a sequential implementation, up to 60 \times when high-dimensional problems are tackled.

Parma, Italy
Rome, Italy
Reggio Emilia, Italy

Stefano Cagnoni
Marco Mirolli
Marco Villani

Evolution, Complexity and Artificial Life

Cagnoni, S.; Mirolli, M.; Villani, M. (Eds.)

2014, XVII, 280 p. 97 illus., 38 illus. in color., Hardcover

ISBN: 978-3-642-37576-7