

UNIVERSITÀ DI BOLOGNA
DIPARTIMENTO DI FISICA

Dottorato di ricerca in FISICA
CICLO XIX

Microscopic Dynamics of Artificial Life Systems

Francesco Zanlungo

Ph.D. Thesis

Bologna, 2007

Relatore:
Prof. Giorgio Turchetti

Coordinatore del dottorato:
Prof. Fabio Ortolani

Settore disciplinare FIS/02 MAT/07

Ai miei genitori

「じゃあ俺の人生がうらやましいか？」

「うらやましくないですね」と僕は言った。

「僕はあまりに僕自身に馴れすぎてますからね。それに正直なところ、東大にも外務省にも興味がない。ただひとつうらやましいのはハツミさんみたいに素敵な恋人を持っていることですね。」

「私たちがまともな点は」とレイコさんは言った。

「自分たちがまともじゃないってわかっていることよね」

「こいのが革命なら、私革命なんていらないわ。私きとおにぎりに梅干ししか入れなかったって理由で銃殺されちゃうもの。あなただっけきっと銃殺されちゃうわよ。仮定法をきちんと理解しているというような理由で。」

「ありうる」と僕は言った。

村上春木「ノルウェイの森」

Contents

Introduction	11
Purpose of the thesis	11
Description of the thesis	13
1 Basic concepts in artificial life	17
1.1 Introduction	17
1.2 Life and Alife	17
1.2.1 Definitions of life	17
1.2.2 History and purposes of Alife	19
1.3 Evolution	20
1.4 Self organisation	22
1.4.1 How self organisation works	24
1.4.2 Characteristic of self organising systems	25
1.4.3 Self organisation and evolution	26
1.5 Emergence	26
1.6 Complex (adaptive) systems	27
1.7 Distributed autonomous systems	28
1.8 Agent models	28
1.8.1 Space time properties of our models	30
2 Computational tools of artificial life	33
2.1 Introduction	33
2.2 Genetic algorithms	35
2.2.1 Biological terminology	35
2.2.2 Original formulation of GA	36
2.2.3 Schemata and building block hypothesis	38
2.2.4 Alternative formulations	40
2.3 Neural Networks	41

CONTENTS

2.3.1	Real nervous system	41
2.3.2	Artificial neural networks and supervised learning . . .	42
2.3.3	Ecological Networks	48
2.3.4	Evolution of neural networks	49
2.4	Application to the models in this thesis	51
2.4.1	Nature and purposes of the models	51
2.4.2	Justification of the use of GA and NN	51
2.4.3	Encoding	52
2.4.4	Mutation	52
2.4.5	Fitness function	53
2.4.6	Selection	53
2.4.7	The problem of defining a fitness landscape	54
2.4.8	A remark about our terminology	57
3	A gas of automata	61
3.1	Introduction	61
3.2	The model	62
3.2.1	The sight cone	62
3.2.2	The two automata problem	63
3.2.3	The geometry of switching	65
3.3	The automaton in a central force field	68
3.3.1	The stop condition for $\alpha = \pi/2$	70
3.4	The 1D model	74
3.4.1	Harmonic potential	75
3.4.2	The infinite potential well	76
3.4.3	Constant force	77
3.5	Numerical analysis of the two body problem	79
3.5.1	General behaviour for different α values	79
3.5.2	Stable orbits for $\alpha = \pi/2$	79
3.6	The N automata problem	84
3.6.1	Memory	86
3.6.2	Numerical analysis of the N body problem	88
3.7	Heterogeneous population of automata with different “social spheres”	98
3.8	Conclusions	99

4	An agent model inspired by clonal expansion	105
4.1	Introduction	105
4.2	The immune system	106
4.2.1	Structure of the immune system	106
4.2.2	The innate immune system	106
4.2.3	The adaptive immune system	107
4.2.4	Recognition of a specific pathogen	107
4.2.5	Tolerance of self	109
4.2.6	Antigen presentation	110
4.2.7	Detection of intracellular pathogens	110
4.3	Description of the model	111
4.3.1	Differential equations model	111
4.3.2	Microscopic model	113
4.3.3	Mean field equations	116
4.4	Results of the simulations	120
4.4.1	Acute antigenic impulse	120
4.4.2	The clonal repertoire model	121
4.5	Conclusions	122
4.6	Appendix	124
5	An evolutionary crowd dynamics and ToM model	127
5.1	Introduction	127
5.2	Crowd dynamics	128
5.3	Theory of Mind	130
5.4	A previous model on ToM evolution	132
5.5	Description of our model	134
5.5.1	Description of the collision avoiding “cognitive” dynamics	135
5.5.2	Description of the evolutionary and experimental settings	141
5.5.3	Comparison with Takano’s Model	144
5.6	Description of an evolutionary experiment	145
5.6.1	Self organising properties	150
5.7	The role of pain	152
5.7.1	Evolution in a high density environment	153
5.7.2	Evolution with low density and generalising abilities	155
5.8	Comparison between different l values	157
5.8.1	Behavioural difference between even and odd levels	157
5.8.2	Level dependence of the continuous parameters	158

CONTENTS

5.8.3	Evolution of agents with high l in high ρ environment	160
5.9	Evolution of ToM recursion level	161
5.9.1	Co-evolution of l and continuous parameters	161
5.9.2	Evolution of l keeping fixed the other parameters	166
5.10	Conclusions	169
6	An evolutionary urban scale traffic model	175
6.1	Introduction	175
6.2	Social insects	176
6.2.1	Some tasks performed by social insects	176
6.2.2	Self organisation	176
6.2.3	Stigmergy	178
6.2.4	Pheromone based trail following in ants	178
6.2.5	Ant colony optimisation	179
6.3	Description of the model	180
6.3.1	Dynamics of motion on the grid	181
6.3.2	Pheromone	183
6.3.3	Decision mechanism and its evolution	184
6.3.4	An interpretation of dynamics and communication	186
6.3.5	Parameters	187
6.4	Evolutionary experiments	188
6.4.1	Description of the experiments	188
6.4.2	Simulations without traffic jams, $\gamma = 0$	189
6.4.3	Traffic jams in a homogeneous network, $\gamma > 0$	195
6.4.4	Heterogeneous grid, $\gamma > 0$	197
6.4.5	Resume of the results	200
6.4.6	Notes on the evolution of the structure of the network	204
6.5	Emergence of traffic flow rules	204
6.5.1	Traffic flow rules	204
6.5.2	Corrections to the model	206
6.5.3	Clock wise and counter-clock wise behaviours	207
6.5.4	Experimental results	209
6.6	Conclusions	217
	Conclusions	221
	Acknowledgements	225

Introduction

Purpose of the thesis

As the title clearly states, the purpose of this thesis is to study artificial life (ALife) systems starting from the dynamics of their microscopic or “fundamental” constituents. Nevertheless these terms can be quite misleading if their are not specified more precisely.

From a physicist’s point of view, to study the microscopic dynamics of a system composed of many “fundamental” parts (the definition of fundamental constituents of a system is not absolute but depends on the contest, as we discuss below) means to rely not on equations regarding the dynamics of macroscopic quantities (which can be deduced from the dynamics of the constituents using, for example, statistical mechanics methods), but on the direct integration of the equation of motion of the constituent “particles”. Macroscopic observables are then obtained by averages over the microscopic ones, and compared with the results given by the macroscopic (mean field or statistical) theory.

From an historical point of view, the success of thermodynamics and statistical mechanics in physics is due to the impossibility, from both a mathematical and an experimental point of view, of obtaining information about the initial conditions and evolution of the state of a system composed of an extremely high number of constituents (of order 10^{23} atoms or molecules for macroscopic systems). Thermodynamics describes the rules governing the macroscopic observables (i.e. the only actual available observables for these systems, as pressure, temperature or volume) and statistical mechanics explains how these macroscopic observables are connected to the microscopic ones through the Hamiltonian function describing the microscopic dynamics of the system.

Nevertheless, the advent of computers has made possible the numerical in-

Purpose of the thesis

tegration of the equations of systems with many constituents (by many, we mean a number that can vary with the development of computers and with the complexity of the problem, but, although being extremely lower than the number of particles in a actual macroscopic system, is a few orders of magnitude higher than the number suitable to an analytical study, which is in many cases just two), and thus physicists, starting from the pioneering work by Fermi, Pasta and Ulam [1], have tried the numerical study of many particles systems from a “microscopic” point of view, in order to compare the results with those obtained in the realm of statistical mechanics (this operation can be viewed, at least in the case of Pasta Fermi Ulam problem and the like, as a numerical test of the ergodic hypothesis).

In the extremely vast realm of sciences that study the multiple expressions of life (from biochemistry to biology, from neurosciences to psychology, from animal behaviour studies to sociology and economics), the advent of computers has actually created a new science, artificial life, or at least has contributed in a fundamental way to its development.

According to Parisi [2], the purpose of artificial life is to build man made “artifacts” that posses at least some of the properties of life (information exchange, reproduction, evolution, learning, etc.). These “artifacts” can be robots, chemicals or, as in the case of the systems studied in this thesis, “virtual worlds” that exist in computers (computer simulations).

Given these premises, the title “microscopic dynamics of artificial life systems” could suggest that the purpose of this thesis is to study how the “property of life” (whatever this means) can emerge from the interactions of basic units that do not posses it, in the same way as life in biological system emerges from atoms and molecules that “are not alive”. But this surely fascinating study is not the subject of this thesis.

As we said before, the term “fundamental constituent” is not an absolute one, but it depends on the context and on the scale of the problem. For example, in physics, when studying the statistical properties of a gas, we can assume molecules to be its fundamental constituents, but from the point of view of high energy particle physics, a molecule is a very complex entity composed of many elementary particles (leptons, quarks). On the contrary, in fluid dynamics, the most appropriate approach consists in forgetting the discrete nature of matter, and study it as a continuum, due to the scale of the problem. Depending on the scale of the problem, in studying the motion of celestial objects, an entire planet can be considered as a point like body.

The same problem is present in all the sciences that study the possible “expressions” of life. For example, studying the working of human or animal body (or of parts of it), the “fundamental constituents” could be considered the cells (if not aggregates of cells), even if cells are by themselves complex entities composed of an high number of molecules, and already posses the aggregate property of “life”. Furthermore, when studying animal or human societies, the fundamental entities are the single individuals (if not groups of them), i.e. extremely complex entities able of involved information processing and behaviours. The advent of computer simulations has made possible the study of the dynamics of these systems on the base of a direct simulation of the interactions of their (complex) constituents, through the so called multi agent systems. This thesis is focused on the study of these systems. As we will see, this field of research is not completely unrelated to the study of how an aggregate property, as life, emerges from the interactions of constituents that do not posses it. Even if cells and animals are extremely complex, often their interactions can be expressed through simple (and local) rules. Nevertheless, these interactions result in the emergence of complex global properties at the level of the whole system.

Description of the thesis

In chapter 1 we provide a short overview of the basic concepts regarding the subject of this thesis, as artificial life, complex systems, agents, adaptation, emergence, self organisation. Even if for some of these subjects a formal and mathematical treatment is possible, it is not exposed in this presentation (if not under the form of references), since these methods are not used in the thesis.

Chapter 2 gives a description of some of the “tools” that are used in the models of the thesis, as genetic algorithms and neural networks. Even in this case, the depth and complexity at which these subjects are treated does not go beyond what is needed to understand the material exposed in the thesis. Chapters from 3 to 6 are actually devoted to the research projects developed during the Ph.D. activity. In chapter 3 and 4 we try to establish a connection between the microscopic (agent to agent) dynamics and the macroscopic one. In chapter 3 we modify a physical system (Coulomb interaction) in order to make it resemble a perception based interaction (introducing a non Newtonian effect), and we perform an analytical study of the two body problem,

Description of the thesis

and a numerical study of the statistical properties of the many body system. In chapter 4 we study a model inspired by the immune system, in order to show how is possible to establish a connection between the agent based microscopic dynamics and the macroscopic population dynamics, governed by differential equations.

Chapters 5 and 6 represent an actual excursion in the realm of “traditional” artificial life, since the attention is not, as in the first two chapters, to study the connection between extremely simple microscopic behaviours and macroscopic observables, but to evolve actual complex (even if idealised) adaptive systems. The purpose of the two models (which are focused on the concept of mobility in traffic systems, one at a crowd level, and the other on a larger scale) is both to evolve a complex behaviour at the individual and local level, and to observe the results of this behaviour as emergent properties at the global level. This result has been obtained using genetic algorithms on large heterogeneous populations.

Chapter 5 presents an evolvable crowd dynamics model in which pedestrians avoid collisions in order to safely reach their goal. The model is based on the ability of agents to predict the actions of the others using a “Theory of Mind”. The model studies the emergence of global self-organisation patterns, and the evolution of the Theory of Mind.

Chapter 6 studies a mobility system on a larger scale, using a discrete space-time representation. Agents are controlled by evolving neural networks, and the emergence of global and local strategies (based or not on communication) that optimise traffic flow is studied.

Each chapter regarding the research projects developed during the Ph.D. activity is structured as an as independent as possible unit, containing all the information needed to understand the subject given a basic knowledge of physics, artificial life and computer science. The basics concepts and computational tools of artificial life, which may not be known to all the readers with a background in physics, are introduced, as we said before, in chapters 1 and 2.

Bibliography

- [1] E. Fermi, J. Pasta and S. Ulam, Los Alamos Sci. Lab. Rep., LA-1940 (1955).
- [2] D. Parisi, communication to the participants to the third Italian Workshop of Artificial Intelligence (2006).

BIBLIOGRAPHY

Chapter 1

Basic concepts in artificial life

1.1 Introduction

In this thesis we will often refer to some basic concepts related to our research field, as *artificial life*, *evolution*, *complex adapting systems*, *emergence*, *self organisation*, *distributed autonomous systems* and *agents*. In this chapter will try to provide a short introduction to these subjects. Actually this chapter is more similar to a glossary than to a sound introduction, that would have resulted in a quite hard and long work. Some of these concepts are still missing a universally accepted formal definition, which means that many definitions have been tried, and that many other objections to these definitions have arisen. In other cases, as for self organisation, which has been initially studied in physical systems, a formal and mathematical theory is present but its description would go beyond the purpose and the possibilities of this thesis. Furthermore, a “linear” exposition of these subjects is quite difficult, since many of them are closely correlated, and a short description of one of them is usually very hard without recalling one of the others.

1.2 Life and Alife

1.2.1 Definitions of life

According to the definition by Parisi [1], the purpose of artificial life is to build artifacts that posses some of the characteristics of living systems. But which are these characteristics? What is life?

Probably no one would doubt that a cat is alive, and that a stone is not.

1.2 Life and Alife

It is also probable that many people would agree that a robot able to navigate safely in a room in an autonomous way (not guided by a human using remote control) *presents some of the characteristic of living systems*. Even more people would agree with the last sentence if told that no one has programmed the robot to do that, but that the robot has learnt to (or evolved to be able to) do that (probably most people would be surprised that such a thing is possible).

But how are we going to answer to the question “is the robot alive?”?. And what if we were questioning about the life not of cats or stones, but of more “border line” entities like viruses?

Boden [2] says that “*There is no universally agreed definition of life. The concept covers a cluster of properties, most of which are themselves philosophically problematic: self organization, emergence, autonomy, growth, development, reproduction, evolution, adaptation, responsiveness and metabolism*”. Part of this list is due to Aristotle, which, according to Matthews [3], listed as properties characterising life “*self-nutrition, growth, decay, reproduction, appetite, sensation or perception, self-motion and thinking*” (interestingly all these properties tend towards the preservation of the species, as noted by Matthews).

Many other lists of properties have been proposed. Bedau [4], more than presenting a long list of properties, associates the concept of life to that of adaptation to a changing environment, i.e. to evolution. This definition includes also systems that are not usually considered as alive, such as human cultures and economic markets. Furthermore, the focus changes from the individual to the species. Life, according to Bedau’s definition, is a property of cats, not of a single cat. We could say that a cat is alive since it is part of a species, cats, that being the result of evolution and being able to evolve, holds the “property of life”.

Also Ray [5] defines a living system as “*self-replicating, and capable of open-ended evolution*”. According to him, Alife “creatures”, if meeting these requests, are alive.

Adami [6] classifies definitions of life in five groups: *physiological*, which centre on functions performed by organisms (breathing, moving, etc.); *metabolic* which centre on the exchange of materials between the organism and its surroundings; *biochemical*, focusing on the capability to store hereditary information in nucleic acid molecules; *genetic*, focusing on the process of evolution; *thermodynamic*, which defines living systems in terms of their ability to

maintain low levels of entropy. While the definitions in the first two classes resulted to be too restricting to describe all known forms of life, the third one seems to be apt to describe terrestrial life, i.e. life as it is present on earth. Even if it could be that terrestrial life is the only existing form of life, it does not mean that other forms of life are not possible. The fourth and fifth classes of definitions try to encompass *any possible form of life*, the fifth one trying to define life on the base of thermodynamical quantities, i.e. creating a connection with physical laws in order to investigate how the transition to the “living state” can happen.

One of the main problems in understanding which are the general properties of life is that we only know a single form of life, terrestrial life, whose manifestations share the same basic biochemistry and a common genetic descent (the DNA code is biologically universal).

1.2.2 History and purposes of Alife

According to Langton [7], artificial life can help us in understanding *what life is* and *how life is possible*. Since we just know a single manifestation of life, we can try to create new forms of life to gain a better understanding of its nature.

Of all the properties of life according to the proposed lists, metabolism is probably the only one that cannot be defined in informational terms. Is thus not surprising that the advent of modern computers has led many researchers to try (and partially to succeed) to recreate these properties through informational concepts and computer modelling.

From an historical point of view, we could say that modern Alife (but see [7] for a brief review of previous attempts to build machines that could “behave like living things”) was born with the pioneering works of Turing [8] and von Neumann [9]. Using the concept of universal computer (or *Turing machine*) developed by Turing, von Neumann investigated and demonstrated the logical possibility of self reproduction. To von Neumann is also due the introduction of cellular automata, while Turing [10] published a mathematical paper on the development of biological forms that had a great influence on modern embryology and that was at the base (many years later) of many works in the field of artificial life.

Another fundamental contribution of Turing and von Neumann to artificial life was their role in the development of computer. Nevertheless, the com-

1.3 Evolution

computational power of computers in the 1950s was too low to develop their ideas, and thus artificial life did not emerge as an unitary branch of science until the 1980s, when the term “Alife” was coined and the first international conference (1987) was held.

As we said before, the purpose of artificial life is to reproduce in artifacts the characteristics of living system, and each work in that direction should be considered as part of Alife. Nevertheless, the true “spirit” of Alife is a little bit more specific, since central concepts in Alife research are self organisation and emergence. We will return again on these concepts, but as a first simplification we could say that an Alife researcher expects her model to show more properties than those that she put in it. This philosophy is well expressed in the usual distinction between artificial life and artificial intelligence (AI).

We could define the purpose of AI as to build artifacts (robots or computer programs) that present some form of intelligence, a definition that resembles the one we gave for Alife. If we admit intelligence to be one of the characteristics of (some advanced) living systems, we could say that any work in AI is also a work in Alife (despite that, artificial intelligence emerged as an independent discipline many years before artificial life). Nevertheless, the approaches to problems (at least between “classical AI” and Alife) are completely different. AI researchers proceed top-down, developing an high-level representation of the problem and trying to pass it as a computer program to the machine. We could say that they “impose order” on the system. Alife researchers proceed bottom-up, they just fix a few low level rules, and expect these rules to lead to an adaptation of the system to the problem, and we could thus say that they expect order to “emerge by itself” in the system. (This is obviously just a very rough simplification, that reflects some of the main original differences between the two disciplines, but does not capture all the complexity of modern day research in the two fields.)

1.3 Evolution

We have seen that some of the attempts to give a concise definition of life rely on the concept of evolution. Adami [6] gives a thermodynamical definition of life, according to which living system are those that are able to maintain low disorder (entropy) levels. Even if this definition does not make any reference to the mechanisms that allow living systems to do that, we know that for terrestrial life this property relies strongly on evolution, meant as the com-

bination of *self replication*, *survival of the fittest* and *random search of new solutions* through mutation.

Any individual being seems to be doomed to lose its battle against entropy, nevertheless if a self replicating ordered structure emerges for some reason, it would be able to persist in the environment through time under the form of its copies. Even if the ordered structure has a finite “life”, i.e. it is ordered only during a finite time, in case its life is long enough to include a reproduction cycle, the structure will persist. Furthermore, if the life is long enough to include many reproduction cycles, the structure will “invade” the environment.

Obviously both the ability of the structure to “stay ordered for long enough”, and to self reproduce are not absolute and depend on the interaction with the environment. If the self reproduction process includes some mistakes or *mutations*, sometimes slightly different forms of the original structure will be created. If these forms are not able of self reproduction, they would disappear as soon as they “die” (lose their ordered structure), while if they are able to reproduce, they will persist in the environment. As we already said, life and reproduction depend on the interaction with the environment. They are not “for free”, and we can imagine that when a large number of ordered structures is present in the environment, a competition for resources arises. In case that in the environment a single species of ordered self reproducing structures is present, each one capable of many reproducing cycles during its life, competition will cause the reproducing ability of the structure to get lower, until each structure is able of just a single reproducing cycle, and the number of structures saturates to the maximum *capacity* of the environment. But if in a saturated environment different structures are present, the structure with the higher reproducing ability in that environment (the one that is more fit or has the higher *fitness* to that environment) will invade the environment, while the others (if relying just on the same resources) will disappear. Self reproduction and random mutations thus lead to the appearance of structures that are “very fit” to the environment, even if no tendency to “get fitter” is present in the structures. Most mutations are less fit, but they just disappear from the environment, while the most fit known solutions persist, waiting for a fitter solution to emerge and take their place. (Adami [6] introduces a measure for the complexity of a string with respect to the environment, which is based on the definition of complexity for a string by Kolmogorov [11], taking in account the information which is provided by the

1.4 Self organisation

environment, represented as a tape of a Turing machine).

As we have already seen, the proliferation of ordered structures, as the appearance of new structures, can change the fitness to the environment, since it introduces a change in the environment itself. This is just one of the possible causes of environmental changes. Self reproduction by itself would not allow the structures to resist to major changes, and thus a “perfectly” self reproducing structure would not be stable and would soon disappear. Nevertheless, self reproduction with mutations allows the emergence of structures apt to self reproduce and thus persist in the new environment, and it is stable to changing environments. (Obviously the mutation rate should be low enough to allow persistence in a stable environment over short time scales, and high enough to allow adaptation. Since also the reproduction system is fixed by the nature of the self reproducing structure, we can think that the mutation rate too adapts in order to be fit to the environment.).

In our search for a mechanism that allows a persistence of ordered structures, we came out with a description of evolution quite similar to that given by Darwin [12]. (Our description does not refer to the particular mechanism used by terrestrial organisms for reproduction, which is briefly exposed in chapter 2 as an introduction to genetic algorithms. We recall that Darwin had no knowledge -obviously- of modern genetics nor of Mendelian inheritance). Evolution [13] accounts for the astonishingly diverse and well apt to the environment forms of life on earth (which originate probably from common ancestor, given the universality of the DNA genetic code). Nevertheless, evolution does not provide an answer to the question “how did life arise?”. According to some researchers, the spontaneous emergence of life is an extremely improbable event. Kauffman [14] says that “*Creditable arguments by respected scientists have led to the unfortunate conclusion that we cannot exist*” and “*plausible calculations demonstrate that insufficient time has elapsed for life to have originated by chance*”. According to him, “*these arguments fail by failing to utilize the self organized collective properties of simple and complex systems*”. Self organisation would thus be at the origin of life. Furthermore, it would have a strong influence on evolution.

1.4 Self organisation

Discussing the role of self organisation in the emergence of life and its influence on evolution would be beyond the purpose of this thesis (and the

capabilities of its writer). Nevertheless, self organisation has a fundamental role in Alife research and applications, and thus deserves a brief discussion. Theories of self organisation originally developed in physics and chemistry to describe the emergence of macroscopic patterns out of processes and interactions described at the microscopic level [15], [16]. For the purpose of this thesis a qualitative description, based on the exposition of self organisation in biological system and on their relevance for computer modelling, as provided by Camazine et al. [17], will be sufficient.

Many physical and biological systems show the formation of global or “high level” organised patterns. By global patterns, we mean that the scale of these patterns is typically a few order of magnitude larger than that of the interaction range between the components of the system. Two different explanations are possible: the patterns have been created by some kind of external agent (or even internal to the system, but whose interaction properties go beyond the range and capabilities that we had supposed), or they are the result of some “unexpected” non-linear summation of the basic interactions.

According to the definition of Camazine et al., “*Self-organization is a process in which pattern at the global level of a system emerges solely from numerous interactions among the lower level components of the system. Moreover, the rules specifying interactions among the system’s components are executed using only local information, without reference to the global pattern*”.

When observing a global pattern in a physical system, if there is no evidence of the intervention of an external agent or force, is quite straightforward, assuming a certain degree of “simplicity” in the constituents, to assume that the pattern is due to self organisation. Nevertheless, when observing a biological system, whose components have a certain degree of complexity, at a first glance other solutions could seem more plausible. In chapter 6 we will describe briefly some of the self organising properties of social insects. It was once believed that some of the complex tasks performed by social insects (for example nest building) could be attained only under the supervision of a leader, which was identified with the queen. Nevertheless, more recent researches show that the queen, at least in large colonies and for what concerns a large number of activities, has no supervision role.

Four alternatives to self organisation in biological systems are proposed in [17]: leaders, blueprints, recipes and templates. By a *leader*, they mean an individual who has a knowledge of the desired pattern, has access to information regarding the state of the whole system and directs the work of each

1.4 Self organisation

individual. By a *blueprint* they mean a *compact* representation of the desired pattern, i.e. some kind of map or project. A *recipe* is a list of sequential instructions that specify the actions to be performed in order to build the desired pattern. Finally a *template* is a full size guide that specifies the final pattern.

Is not difficult to think about examples in which human agents use one of these methods to establish a global pattern. The authors provide also a few examples of biological systems in which these methods are used. Nevertheless they also provide a long list of examples in which experiments show that these methods are not used and thus the global pattern emerges on the base of self organisation properties of the system (examples regarding social insects are provided, as we said before, in chapter 6, while in chapter 5 we treat self organisation in human societies and in computer models of human behaviour).

1.4.1 How self organisation works

Self organising systems usually rely on a combination of *positive* and *negative feedback*. Negative feedback is a mechanism well known in biology, used to stabilise physiological processes. When a parameter of the system grows, it originates by negative feedback a process that leads to its diminution and vice versa. Positive feedback has exactly the opposite result: as soon as a given parameter of the system subjected to positive feedback starts to grow, it causes a process that leads to its further growth (“snowball effect”).

Usually, in self organising biological systems, positive and negative feedback act on different ranges of the parameters. An example is given by the nesting activities of many species of birds. These birds follow the individual and local rule: “nest where others nest”. If by chance a small aggregate of neighbouring nests arises, positive feedback causes the formation of a large colony of nests. Let us now suppose that this colony is located on a small island. If a too large number of birds nests on the island, they will probably exhaust its resources. It is thus probable that the actual rule would be “nest where others nest, if not overcrowded”. When the density of nests gets too high, negative feedback starts to regulate it. (Notice that also the overall number of birds provides a form of negative feedback. This could seem a very trivial example, but it is quite typical of self organisation properties in biological systems: while positive feedback is due to behavioural rules, negative feed-

back is just caused by the exhaustion of some kind of resource.)

How can positive and negative feedback generate global patterns? The process can be described as “*amplification of fluctuations under control*”. The global pattern has to emerge over a uniform medium (clusters of nests in an area initially without nests). At the beginning the process is random (bird nests are randomly distributed), but then a fluctuation breaks the symmetry (by random fluctuation, one or a few small clusters are formed) and positive feedback causes the emergence of the pattern (small clusters turn big due to the snowball effect). When the final structure of the pattern has been reached, it has to be kept fixed by negative feedback (overcrowding and exhaustion of the number of birds fix the size of clusters).

Another characteristic of self organisation, which is included in its definition, is the local nature of the behavioural rules and information gathering. A rule suggesting to “form large clusters” of nests would not be admissible for a self organising system, since it includes information on the global pattern, while a rule as “nest near other nests” is admissible, since it only requires local information.

The behavioural local rules causing the emergence of global patterns are based on multiple interactions between the individuals. Nevertheless, these interactions need not to be direct, but could also be mediated through the environment. This phenomenon is called *stigmergy*: each individual acts on the environment and modifies it, but bases its actions on the state of the environment: thus the “work” performed by an individual influences (on a local base) that of the others (stigmergy, or “information gathered from work in progress” is explained in detail in chapter 6.2.2 where a remarkable example of self organising pattern formation, termite nest building, is described).

1.4.2 Characteristic of self organising systems

In general self organising systems are *dynamic*: the pattern is reached and *maintained* by the continual interactions between its constituents. They possess *emergent properties*, i.e. they have some properties which are qualitatively different from those of their components and cannot be understood by a simple addition of individual contributions. They are *multistable*: depending on initial conditions and random processes they can reach different final states (patterns). They often exhibit *bifurcation* as the consequence of the change of some parameters.

1.5 Emergence

The dynamic process led by positive and negative feedback usually leads to the formation of *stable patterns*: the system is robust to perturbations.

1.4.3 Self organisation and evolution

The authors of [17] do not enter in the discussion on the role of self organisation in the emergence of life nor on its role “in competition” with selection as discussed by Kauffman [14]. They just suggest that self organisation can be a very powerful tool through which evolution can find a simple solution to complex problems: in self organising systems, evolution just needs to encode the (simple, and thus probably more easy to find) local rules, and the solution to the complex problem emerges as a consequence of those rules.

This is the same spirit that leads us to investigate self organising properties in the models studied in this thesis (and in particular in chapters 5 and 6 where these properties are the result of an evolutionary process).

1.5 Emergence

The word *emergence* has been used many times during this chapter. Its basic meaning has been probably already captured: an emergent property is some kind of complex characteristic of a system resulting from a combination of apparently simple rules. Nevertheless a formal definition is more difficult to obtain.

From a *mathematical* point of view emergence is just a result of non linearity: if the rules that govern the interaction between the components are not linear, their contributions to the behaviour of the system cannot simply be summed up. Furthermore the dynamics of very few non linear systems, and in general only for given initial conditions, can be solved using analytical methods. For these reasons the behaviour of a non linear system cannot be understood analysing simply the functioning of its components, and the global behaviour “emerges” when these components are put together.

Nevertheless it is clear that not all the effects of non linearity deserve to be called emergent. From the point of view of a researcher studying biological or Alife systems an emergent property has to be “*interesting*”. Emergent properties are often defined in Alife as “*unexpected properties*”. This definition, despite being unsatisfactory (what is unexpected for a researcher could be anticipated by another one on the base of greater understanding of the sub-

ject or mathematical knowledge) reveals the spirit underlying the notion of emergence in artificial life: an emergent property has to be surprising. (Clark [18] says that between the central themes of Alife there is “*the emergence of **interesting** or adaptive features from **clever** couplings between agents and the environment (including other agents)*”. We added the boldface.)

A philosophically more sound definition is provided by Steels [19], that suggests that the word emergent can be used when the results of multiple interactions require description in a vocabulary different from that used to describe the properties of the inner components. According to Clark [18], emergence involves behaviours that are not (directly) due to any controllable part or environmental interaction, but that are side-effects of directly controllable actions or interactions. He cites an example, due to Steels and Hofstadter [20]: if an operative system begins to “trash around” as soon as 35 users are on line, it would be a mistake to go to the programmer and ask him to raise the trashing number to 60, since the trashing number, though being a feature of the model, is an emergent or uncontrolled variable. All that the programmer can do is to change the controllable variables in order to prevent the system to trash.

1.6 Complex (adaptive) systems

A system composed of many interacting units, showing self organisation and emergent properties, is a *complex adaptive system* (CAS). Holland [21] gives a few examples of CAS: metropolises, the immune system, the nervous system, ecosystems, Internet, economies etc.

One of the properties of these systems is *adaptation*, the process that leads the system to fit itself to the environment. We have already seen how evolution can lead to adaptation to the environment, but other mechanisms are possible. For example, if we are considering a system whose fundamental units are complex enough, as human societies, we can expect these units to be able to learn, i.e. to adapt to the environment during their life on the base of their experience. It is obviously not necessary to hypothesise very complex units with “black box” learning properties in order to have adaptation on the base of experience: the nervous system is composed of relatively simple units that are able to adapt their connection weights on the base of their interaction with the environment (more on this in chapter 2). Also the immune system (see chapter 4) is able to adapt to the menaces of the en-

1.7 Distributed autonomous systems

vironment to the body, recognising and “remembering” external pathogens, while being able to distinguish the “self”, i.e. the cells of the body. These learning abilities are based on the interaction rules between the constituents of the systems, rules that were obtained as the result of an evolution process. Deneuborg et al. [22] provide us an interesting example of how *fixed* behavioural rules can result in the adaptation of the system to different environments, showing that the different foraging patterns of some army ant species can be obtained using the same local rule given a different prey distribution. The rules are always the same but they are “clever enough” (though very simple) to provide different and efficient behaviours in different environments.

1.7 Distributed autonomous systems

When we have introduced the concept of self organisation, we have stressed that a real self organising process emerges in a system whose units interact on the base of local and individual rules, without supervisors (leaders) or external influences. A system with these interaction rules is said to be *distributed* (since its dynamics is due to the individual actions of the basic units) and *autonomous* (since there is no leader, internal or external, influencing these dynamics) [23].

In this thesis we are interested in studying distributed autonomous systems, systems composed of many fundamental and interacting units which are not supervised by any external mechanism. As we said in the introduction, we will assume this fundamentals units to represent “complex entities” provided with perception and data processing capabilities. These capabilities are introduced in our model as “black box” properties, i.e. we do not investigate how they can emerge from more fundamental interactions. Nevertheless these interactions will be introduced in the models as simple rules, and we will expect emerging and self organising properties to be present at the global level, in some cases as the result of an evolution process.

1.8 Agent models

As we said in section 1.5, when the interactions between the constituents of the model are non linear, analytical methods cannot provide much infor-

mation about the evolution of the system and computer models are needed. Different approaches can be used to study complex systems with computers. A traditional approach, based on *differential equations* (DE), consists in writing equations for the dynamics of a few average quantities in the system. These are usually ordinary differential equations (ODE), but more complex models, taking in account the spatial structure of the system, can be based on partial differential equations (PDE). These equations are in general non linear, and thus can only be solved using a numerical approach, but analytical methods can provide useful information about the evolution of the system (attractors, stable orbits, etc.). Furthermore, the use of ordinary differential equations strongly reduces the number of degrees of freedom in the system, and the computational cost of numerical integration. The major shortcoming of this approach (for what concerns the systems studied in this thesis) is evident: how can we study distributed systems whose emerging properties are obtained by individual contribution that cannot be summed up on the base of equations for the dynamics of average (i.e. summed up) quantities? (This does not mean that average, or *mean field*, equations cannot provide any insight on the nature of the problem, at least in some cases. A comparison between an approach based on differential equations with an approach based on microscopic dynamics is provided in chapter 4.)

An approach based on *cellular automata* (CA) is more adequate to describe a distributed system. Cellular automata [24] are grids composed of many cells, whose state changes taking in account the state of neighbouring cells on the base of simple rules. CA are shown to be in principle equivalent to Turing machines [25] and have been used in many Alife projects. Also cellular automata have a shortcoming for what concerns our approach to the problem, since they allow only a discrete description of space. Furthermore, the basics constituents of CA are the units of the grid, and not moving individuals (suitable CA evolution rules can describe moving individuals, but this approach is less intuitive than an approach based directly on the description of the individual).

An approach that allows a description based on individual and independent units, without restrictions on the space-time representation, is based on the use of *agents*. Agents are units capable of perception (i.e. information from the environment is passed to the agent according to its particular sensory properties, which are decided as part of the model), of data processing (even this process can be arbitrarily complex, and its nature is fixed by the model),

1.8 Agent models

and on the base of these can perform actions that modify the environment (for example, they can move).

1.8.1 Space time properties of our models

In the four models presented in this thesis, we use four different space time representations of agents and their physical interactions (by physical interactions we mean the interactions that modify the motion of the agents and are not determined by its decision mechanism: for example, the decision mechanism could suggest to the agent to go in a certain direction, but its actual motion could be different due to a collision with a wall or another agent).

In chapter 3 we describe interactions between agents in similarity with the interactions between charged particles. These repulsive interactions are considered to be due to the decision mechanism of the agent. Since these agents are *point like* particles they do not have a inner physical structure and thus physical interactions (i.e. collisions) are not introduced in the model. The agents move in *continuous two dimensional space and in continuous time* (obviously numerical integration requires the use of a discrete time step, but this is not introduced as part of the nature of the model).

In chapter 4 we use a discrete space time to describe the immune system. Agents occupy cells of two overlapping *discrete two dimensional space grids*, and they are thought to have *finite dimension*. When an agent occupies a cell no other agent on the same grid can be present on that cell. Interactions happen between agents on overlapping cells of the two grids, or at the sides of two occupied neighbouring cells (“the surface of the agents”).

In chapter 5 each agent moves in a *continuous two dimensional space and time* (in this case we can say that also time is realised in a continuous way, or at least using floating point numbers, since the algorithm that regulates physical interactions does not need a discrete time step), and has a *finite body size* (it is represented as a disc). Elastic collisions (physical interactions) of agents with other agents and with the walls are exactly resolved.

Finally in chapter 6 space is represented as a *discrete space time grid*, (in two dimensional space), and agents are *point like*. The physical interactions between agents and of agents with the environment, which have a fundamental role in the evolution of the system, are introduced as modifying the transition probabilities of agents from a site to another of the grid.

Bibliography

- [1] D. Parisi, communication to the participants to the third Italian Workshop of Artificial Intelligence (2006).
- [2] M. A. Boden, introduction of *“The Philosophy of Artificial Life”*, Oxford University Press (1996).
- [3] G. B. Matthews, *“Aristotle on Life”*, in M. A. Boden (ed.), *“The Philosophy of Artificial Life”*, Oxford University Press (1996).
- [4] M. A. Bedau, *“The Nature of Life”*, in M. A. Boden (ed.), *“The Philosophy of Artificial Life”*, Oxford University Press (1996).
- [5] T. S. Ray, *“An Approach to the Synthesis of Life”*, in M. A. Boden (ed.), *“The Philosophy of Artificial Life”*, Oxford University Press (1996).
- [6] C. Adami, *“Introduction to Artificial Life”*, Springer-Verlag (1998).
- [7] C. G. Langton, *“Artificial Life”*, in M. A. Boden (ed.), *“The Philosophy of Artificial Life”*, Oxford University Press (1996).
- [8] A. Turing, *“On computable numbers, with an application to the Entscheidungsproblem”*, Philosophical Transaction of the Royal Society (London), Series B, 237: 37-72 (1936).
- [9] J. von Neumann, *“The general and logical theory of automata”*, in *“Cerebral Mechanism in Behavior-The Hixon Symposium”*, John Wiley (1951).
- [10] A. Turing, *“The Chemical Basis of Morphogenesis”*, Proc. Lond. Math. Soc. Ser. 2 43, 544; 42, 230 (1952).
- [11] A. N. Kolmogorov, *“Three approaches to the definition of the concept of quantity of information”*, Probl. Inform. Transmission 1, 1 (1965).

BIBLIOGRAPHY

- [12] C. R. Darwin, *“On the Origin of Species”*, John Murray (1859).
- [13] M. Ridley, *“Evolution”*, Blackwell (2004).
- [14] S. A. Kuaffman, *“The Origins of Order”*, Oxford University Press (1993).
- [15] H. Haken, *“Synergetics”*, Springer-Verlag (1983).
- [16] G. Nicolis and I. Prigogine, *“Self-Organization in Non-Equilibrium Systems”*, Wiley and Sons (1997).
- [17] S. Camazine, J. L. Deneubourg, N. R. Franks, J. Sneyd, G. Therlauz and E. Bonabeau, *“Self-Organization in Biological Systems”*, Princeton University Press (2001).
- [18] A. Clark, *“Happy Couplings: Emergence and Explanatory Interlock”*, in M. A. Boden (ed.), *“The Philosophy of Artificial Life”*, Oxford University Press (1996).
- [19] L. Steels, *“The Artificial Life Roots of Artificial Intelligence”*, Artificial Life, 1: 75-110 (1994).
- [20] D. Hofstadter, *“Metamagical Themas: Questing for Essence of Mind and Pattern”*, Penguin (1985).
- [21] J. H. Holland, *“Hidden Order”*, Basics Books (1995).
- [22] J. L. Deneunburg, S. Goss N.R. Franks and J. M. Pasteels, *“The Blind Leading the Blind: Modelling Chemically Mediated Army Ant Raid Patterns”*, J. Insect Behaviour 2: 295-311 (1989).
- [23] L. A. Segel and I. R. Cohen, preface of *“Design Principles for the Immune System and Other Distributed Autonomous Systems”*, Oxford University Press (2001)
- [24] S. Wolfram, *“Cellular Automata and Complexity”*, Perseus (1994).
- [25] C. G. Langton, *“Studying Artificial Life with Cellular Automata”*, Physica D, 10: 120-149 (1986).

Chapter 2

Computational tools of artificial life

2.1 Introduction

In studying complex adaptive systems with computers we face the problem of designing systems composed of many autonomous parts from whose interaction emerges some kind of global self organisation pattern. A similar problem is faced by the designers of distributed autonomous intelligent systems, i.e. systems used to solve some kind of complex problem and that are not governed by a central unit, but are based on the autonomous actions of many constituents. As we have seen, an emergent property is a feature of the overall system which is caused by local agent behaviour, but that cannot be deduced directly from the latter since it is the result of many non-linear interactions. If the deduction of the global pattern from the local rule is not easy, it is clear that the reverse problem of discovering which rule causes the global pattern is usually even harder. We can obviously rely on intuition: we can guess a local rule and use computer simulations to verify if the wanted behaviour emerges. (In case we are studying a complex system actually present in nature, for example a biological one as a social insect colony, this process of modelling the system can be helped by direct and experimental observations of the single agent's behaviour. It is thus not surprising that these biological systems have lately been an inspiration for the design of intelligent systems, see for example [1].). Nevertheless finding the local behaviour which can account for a given global pattern is usually almost impossible especially for

2.1 Introduction

complex tasks, and even more if we want the system to be actually adaptive, i.e. resisting to changing conditions.

A similar problem is faced by the designers of intelligent systems, in particular when these systems are in relation with an actual and complex environment as the real world is (for example in robotics). The traditional top-down methodology used in AI, i.e. to create large centralised complex programs, has showed to be very effective in theorem-proving and game playing, but not very satisfactory when facing interaction with a real physical environment. The “animat path to AI” or bottom-up approach [2] consists in constructing simple agents, which are able to adapt to the environment they are located in. Here arises a paradox (similar, in some way, to the problem of obtaining a complex global behaviour as the sum of simple local rules): how can we construct simple agents able to deal with a complex environment if we are hardly able to construct complex robots that can perform well in that environment?

The solution to these problems resides in the imitation of nature. In nature we can observe many complex adaptive systems, which arose on the base of the strength of evolution, i.e. of self-replication, survivor of the fittest, and emergence of new solutions due to random genetic mutations. The application of these principles in the realm of design of intelligent systems and optimisation problems is usually referred as genetic algorithms (GA).

One of the most striking results of evolution is the animal nervous system, through which animals are able to respond with appropriate actions to the stimuli provided by the environment. Control systems inspired by the natural neural system and used in order to provide effective agent to environment interaction, and also to introduce the ability to learn about the environment (and thus to quickly adapt to its eventual changes), are usually called Neural Networks (NN). The use and study of neural network based computer programs goes often under the name of *connectionism*.

These methods are not only useful tools to solve complex problems, but, since they are inspired by and mimic actual biological processes, they provide us useful insight about how these mechanisms work. If in the realm of intelligent system design GA and NN are stimulating options whose actual usefulness needs always to be analysed in relationship with the particular problem to be solved, in artificial life, i.e. in the attempt to reproduce biological complex systems or other systems that present some of their features, they are natural if not obliged choices.

As we will see, GA and NN are not unrelated subjects, since a useful way to obtain neural networks that can perform efficiently in a given environment is to make them evolve through GA.

2.2 Genetic algorithms

Genetic algorithms in their traditional form have been developed by John Holland [3], and were aimed also to solve optimisation problems but in particular to study the phenomenon of adaptation as it occurs in nature. An introduction to the subject, which served as a base for this short presentation, can be found in [4].

2.2.1 Biological terminology

The original formulation of GA, and the usual explanation of their working, is so strongly related to the process of biological evolution that is probably useful to introduce some biological terminology and to revise briefly how evolution in actual biological organisms works.

Each cell of a living organism contains the same set of *chromosomes* (strings of DNA), that serve as a blueprint for the organisms. A chromosome is divided in *genes*, each gene being located in a given *locus* or position and encoding for a particular protein. Each gene can assume different forms or *alleles*. One could think (even if it is just a simplification) that each gene encodes a *trait* or characteristic of the organisms, and that to different alleles correspond different expressions of the trait (for example, a given gene could encode the colour of the eyes, the trait, and to different alleles could correspond different colours). The collection of an organism's chromosomes forms its *genome*, while the particular set of genes in a given organism's genome is the *genotype*. The genotype gives rise to the *phenotype*, the collection of the organism's physical and mental characteristics. The distinction between genotype and phenotype is a fundamental one: even if the genotype encodes the phenotype, there is not a one to one correspondence between the two, since the phenotype is the result of a development process, during which the organisms undergoes constant interactions with the environment. (Just to make it more clear, a human being has always the same genotype during her life, but her phenotype changes during development from the foetal and child forms to the adult one. Furthermore, this phenotype depends on the

2.2 Genetic algorithms

conditions under which development happened, as for example nutrition). Most sexually reproducing organisms are *diploid*, i.e. chromosomes are arrayed in pairs. For example, each human cell has 23 pairs of chromosomes. In sexual reproduction *crossover* occurs: genes are exchanged between each pair of chromosomes of each parent in order to obtain a single chromosome (*gamete*), and gametes from the two parents form a full set of diploid chromosomes. Offspring are subject to *mutations*, i.e. due to copying errors elementary bits of DNA are changed from parent to offspring.

The ability of the organisms to perform in the environment (its *fitness*) is basically determined by its ability to reproduce (its number of offspring or its *fertility*) and thus pass and diffuse its characteristics in future generations, an ability which is due to its phenotype.

Holland [3] gives an estimate of the complexity of finding a genotype that encodes a good phenotype: assuming 2 alleles for 10000 genes for a vertebrate genotype, we have roughly 10^{3000} different solutions. Furthermore, as we said before, thinking that to each gene corresponds a trait is just a simplification, and actually to each allele corresponds a different protein. The form of cells is determined by the highly non linear interactions between these proteins. As always happens in complex systems, the final product is not just the sum of the contribution of different alleles, a phenomenon known in biology as *epistasis*.

How does nature manage to solve this optimisation problem? As we have already seen, the “strength” of life resides in the combination of self reproduction, survival of the fitness and search of new solutions through casual mutations. The use of these strategies in computation problems goes under the name of evolutionary computation.

2.2.2 Original formulation of GA

Holland’s genetic algorithms (it should be noticed that the term genetic algorithm is used nowadays, and also in this thesis, to refer to a more broader class of evolutionary algorithms, usually quite different from Holland’s original formulation) stresses the importance of the particular encoding and *genetic operators* used by nature, and thus tries to reproduce them in a simplified but faithful way.

The characteristics of the “candidate solution” to the problem (for example the characteristics of an agent) should be encoded as a bit string, a chro-

mosome. The genes are the single bits (or groups of them), and the alleles are 0 or 1 (larger alphabets, i.e. a larger number of symbols to represent an higher number of alleles, are also allowed, even if in the following we will always refer to the binary representation). The performance of a given solution is measured using a *fitness function*, and in each *generation* of the genetic algorithm the fitness of a *population* (a fixed number) of solutions is tested. While the first generation is randomly created, the following ones are obtained from the previous one using three genetic operators: *selection*, *crossover* and *mutation*.

Selection consists in the choice of the chromosomes to be passed to the following generation: since in biology fitness is the number of offspring, also in GA the number of offspring of a given solution should be proportional to the fitness. If we define f_i as the fitness of solution i , and N as the number of solutions in each generation, the probability of solution i in generation n to be the parent of a solution in generation $n + 1$ is

$$p_s^i = \frac{f_i}{\sum_{j=1, \dots, N} f_j} \quad (2.1)$$

and thus its expected number of offspring is Np_s^i . (This process is called *roulette wheel selection*, since it is equivalent to assign to each individual a slice of a “roulette wheel” proportional to its fitness, and choose N parents spinning the wheel N times). The selection is obviously made with replacement, i.e. a solution can be the parent of many offspring.

Crossover consists in pairing the chromosomes of two selected parents and with probability p_c (crossover probability) cross over the pair at a randomly selected point, to create two offspring (since crossover is performed in a single point of the chromosome, this is usually referred as “single point crossover”, figure 2.1.). In case crossover does not happen, the chromosomes of the parents are simply copied in the offspring.

Mutation consists in randomly flipping each gene in each offspring with probability p_m . Using these operators a new generation of N “solutions” is obtained and tested. Since the described process generates an even number of offspring, in case of odd N a chromosome is randomly chosen and discarded. A whole evolutionary process, or a *run* of the GA, consists of a fixed number of generations (the actual number of generations depends on the nature of the problem, including computational cost, but in general at least a few hundreds of generations are necessary). Since GA strongly depend on random

2.2 Genetic algorithms

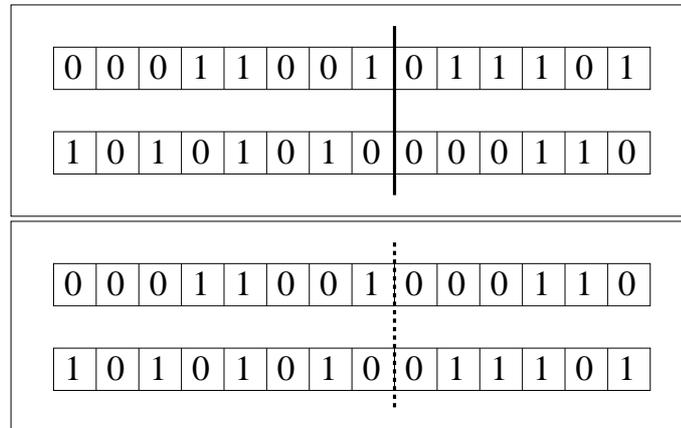


Figure 2.1: Chromosomes before (top) and after (bottom) single point crossover. The crossing point is identified by the thick line.

processes, different runs give often different results, and thus an evolutionary experiment in general consists of many runs, from which averaged results can be obtained.

A (very trivial) application of genetic algorithms could be the optimisation of a real positive function with a single variable. The “candidate maximum” would be a real number, coded as a string of bits, and the fitness function would just be the value that the function assumes at the candidate solution. To each possible solution (string) corresponds a value of the fitness, and the set of all these values forms “the fitness landscape” of the problem. We expect the genetic algorithm to explore effectively this “landscape” and reach its maximum.

When facing an real problem, many questions about how to use GA arise, and actually different evolutionary strategies (corresponding to different encodings, or different forms of the operators) could be necessary. Before proceeding in the analysis of these problems, let us first study how the original formulation of GA works.

2.2.3 Schemata and building block hypothesis

According to Holland the success of evolution depends on its ability to find *coadapted* sets of alleles, i.e. sets of alleles from different genes which to-

gether significantly augment the performance of solutions. These sets should be the “building blocks” of good solutions, and should be evaluated and recombined (without being destroyed) in an effective way. Holland formalises these concepts introducing the notion of “schemata”. A schema is “family” of bit strings, i.e. a string of three symbols, 0, 1 and *. The additional * stands for “any possible symbol”, i.e. the schema $1*0$ stands for a family composed of two strings, 110 and 100, while $1*****1$ stands for all the (2^5) possible 7 bit strings that begin and end with a 1.

The introduction of the concept of schemata poses the attention, more than on the particular strings, on classes of strings that share some features. If these features are exactly those that make a solution to be a good one, i.e. if the schema corresponds to an useful “building block”, all the strings that represent that particular schema will receive a high fitness, independent of the actual values assumed by the bits marked with an *. We could thus say that GA, while explicitly testing strings, are implicitly testing schemata. Since each string of length l is an instance of 2^l schemata (to each gene or bit can be substituted either the actual allele of the string or an * in order to obtain a valid schema; for example string 11 is in schemata **, $1*$, $*1$ and **), a population of N strings contains a number of schemata between 2^l and $2^l N$. Holland shows that while testing N strings, GA process at least N^3 different schemata at time (“implicit parallelism”). The schema theorem (see [3, 4] for mathematical details) shows how selective reproduction gives a larger number of offspring to schemata with higher fitness. It also shows that schemata whose components are closer (for example $****11*$) have a larger probability to be preserved by single point crossover than schemata whose components are far apart ($1*****1$), since the latter are more easily broken. The “*building block hypothesis*” suggests thus that GA are able to find the building blocks of the solutions, i.e. useful groups of close alleles (evaluating them as schemata through implicit parallelism), to reproduce them with selective reproduction, and to recombine while preserving them through single point crossover.

According to this interpretation, crossover has the main role. The only role of mutation, whose probability should be kept at very low values, is to prevent loss of diversity at a given position. Let us assume, for example, that the best possible solution to our problem is a string that begins with a 1, and is thus in schema $1*****\dots$. If at a given point during evolution all the strings in the population start with a zero and are in schema $0*****\dots$,

2.2 Genetic algorithms

crossover will not be able to generate the best possible solution. Mutation solves this problem randomly creating instances of 1*****....

2.2.4 Alternative formulations

Nevertheless, schema theorem and its implications for the behaviour of GA have been the subject of critical discussion in the GA community, as described by [4]. Furthermore, in their book on evolutionary robotics (a subject closely related to the research projects described in this thesis, as we will see), Nolfi and Floreano [5] state that *“Several practitioners reported that crossover had no influence in their experiments, or even that it lowered the performance of the genetic algorithm. In those experiments, crossover operates as a sort of macro mutation operator and may generate instability. Over the last years, much more attention is being paid to the mutation operator. In most of the experiments described in this book the crossover probability is zero or very small”*.

The details of the particular “evolutionary strategy” to be used depend on the kind of problem that is faced. Of course a first serious question, when trying to solve an optimisation problem, is “should I use GA or another search strategy?”. Nevertheless, when trying to study the evolution of an artificial life complex adaptive system, like those studied in this thesis, GA are more than an option, and in some way impose themselves, and thus the question of the performance of GA with respect to other optimisation strategies is not relevant in this thesis (we’ll return on this point in section 2.4, and in particular we will stress on the impossibility to define a traditional fitness landscape for the systems studied in this thesis).

Different kinds of encoding (some of them introducing a difference between genotype and phenotype and thus the concept of development in interaction with the environment), selection, crossover and mutation operators have been proposed in different optimisation and Alife problems. Without giving an overview of all these different options (we refer once again to [4, 5] and to a large GA and Alife literature) we will expose in section 2.4 the choices that we have made in our models and the reasons that led us to those choices.

2.3 Neural Networks

2.3.1 Real nervous system

The human brain [6] consists of roughly 10^{11} nerve cells or *neurons*. Neurons communicate using short impulsive electric signals, that propagate as variations of potential difference between the interior and exterior parts of the neural membrane. The connections between neurons are mediated by electrochemical junctions called *synapses*, which are located on branches of neurons named *dendrites*.

Each neuron has thousands of synapses and thus receives incoming signals from thousands of other neurons. These signals are “summed up” and if the total exceeds a threshold, the neuron “fires” i.e. it produces a new impulse that propagates through a branch of the neuron called *axon* and reaches other neurons.

Synapses are the neuron’s input: at synapses arrive the impulses that have propagated through the axons of other neurons. The influence of these impulses on the activity of the receiving neuron depends on the strength of synaptic connections (there is a weighted sum of inputs); furthermore some impulses have an excitatory effect on the firing of neurons, while others have an inhibitory effect (the weights have a “sign”).

The axon is the the neuron’s output, and it is unique: each neuron has a single axon (even if it branches in order to reach synapses of different neurons), and all the pulses have the same characteristic profile, height and width. Neurons can thus be viewed as communicating with yes/no (1 or 0) binary signals.

(This short introduction reflects obviously only a minimal part of the actual complexity of real neurons, and focuses on the features that have been considered fundamental in the computing abilities of neurons and have been introduced in neural networks models. In particular no mention has been done about the role of time, even if propagation delays and pulses frequencies influence the neuron activity).

2.3 Neural Networks

2.3.2 Artificial neural networks and supervised learning

The term *connectionism* refers to the approach of using *neural networks*, both for computational purposes and for the study of the actual human and animal brain and nervous system (in a broader sense, connectionism, in neuroscience and psychology, is an approach that models mental or behavioural phenomena as emergent properties of networks of simple units).

A neural network (NN, see [7] for a throughout exposition of the subject, and [8] for a simple introduction to the main ideas), is “*an interconnected assembly of simple processing elements, **units** or **nodes**, whose functionality is loosely based on the animal neuron. The processing ability of the network is stored in the inter-unit connection strengths, or **weights**, obtained by a process of adaptation to, or **learning** from, a set of training patterns*” (We added the boldface to the definition by [8]).

Basically neural networks try to reproduce the fundamental computational features of actual neurons, and thus their nodes are computational units which take a few inputs, weight them and produce an output based on the inputs (in general the output is a non linear function of the inputs). As for actual neurons, inputs are provided by other neurons, and thus various nodes are connected between them (obviously also external inputs, that can be considered as the inputs of the whole network, are allowed). Equally, there are also outputs which are not, or not only, passed to other nodes, and have the role of external outputs or outputs of the network (figure 2.2).

Perceptrons

A first, and the most simple, example of neural network is a *threshold logic unit* (TLU) or *perceptron* (figure 2.3). As the name suggests, it is a network that consists of a single unit, takes n external inputs, weights them, and produces a “logic” or binary output (if the weighted sum is higher than a threshold θ the output is 1, otherwise it is 0). The choice of a binary output reflects the yes or no response of biological neurons, while the external inputs can be in principle binary or real numbers. The behaviour of a TLU is obviously determined by its weights. Denoting these weights with a n component vector \mathbf{w} and the inputs as the vector \mathbf{x} , the output of the network is

$$y = \vartheta(\mathbf{w} \cdot \mathbf{x} - \theta) \tag{2.2}$$

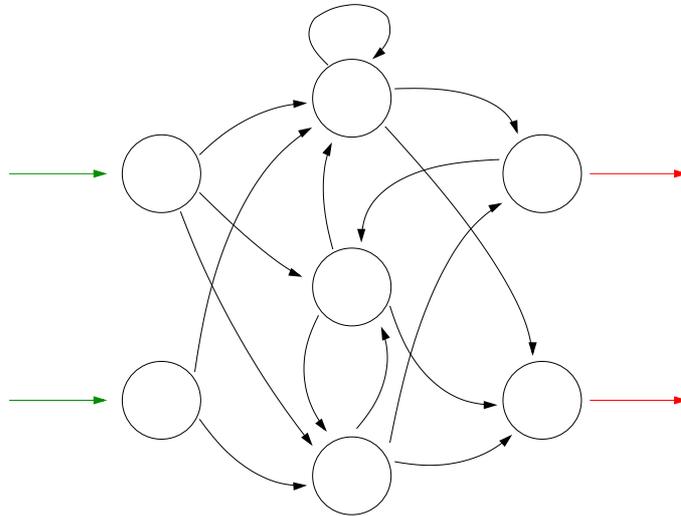


Figure 2.2: A neural network. Circles represent units, lines represent connections, arrows synapses. All the lines exiting from a unit have the same value. Green lines are external inputs, red ones external outputs (for these lines arrows do not stand for synapses, and just show the “direction” of the line).

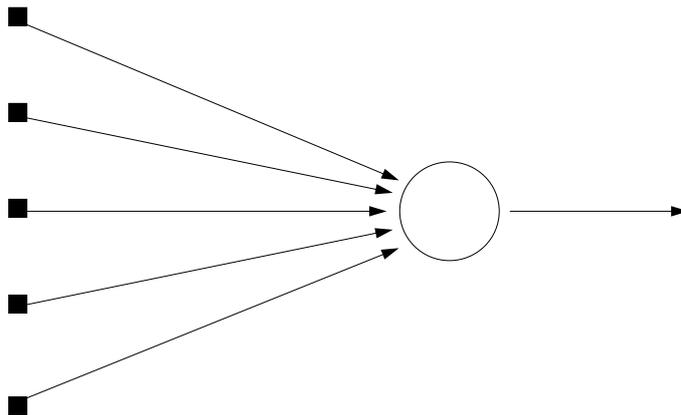


Figure 2.3: A perceptron or TLU with 5 inputs.

2.3 Neural Networks

where by ϑ we have denote the step function

$$\vartheta(x) = \begin{cases} 0 & \text{if } x \leq 0 \\ 1 & \text{if } x > 0 \end{cases} \quad (2.3)$$

The negative of the threshold $-\theta$ is usually treated as an “input independent weight” or a *bias*, i.e. as a weight that is multiplied for a constant (equal to 1) input. From now on we will use the latter notation, and assume that a TLU with n input has a $n + 1$ dimensional weight vector $\mathbf{w} = (w_1, \dots, w_n, -\theta)$, and that equation (2.2) becomes

$$y = \vartheta(\mathbf{w} \cdot \mathbf{x}) \quad (2.4)$$

where $\mathbf{x} = (x_1, \dots, x_n, 1)$. It is clear that the hyperplane $\mathbf{w} \cdot \mathbf{x} = 0$ separates the n dimensional space of inputs in two sets, one that is classified by the TLU with 1, and another which is classified with 0.

Perceptron and delta rules

The main feature of neural networks is their ability to learn. An application of a TLU is thus that of learning how to classify the inputs. Let us suppose that we have two sets of objects, A and B , whose characteristics can be expressed as n dimensional real vectors, that the two sets are linearly (i.e., by an hyperplane) separable, and that we want to teach a neural network to recognise which element is in set A (which will correspond, for example, to output 1) and which one is in B (0). The procedure (called *supervised learning* since there is a teacher that provides a few examples and tells to the network which answers are right), consists in giving to the network a *training set*, a group of vectors \mathbf{x}_i together with the “right answers” or targets t_i ($i = 1, \dots, M$ where M is the size of the training set.) The weights (and biases) of the network are thus initialised to values next to zero, and the network is tested on the training set. After each test on vector \mathbf{x}_i the weights are adjusted ($\mathbf{w}' = \mathbf{w} + \Delta\mathbf{w}$) following the *perceptron rule*

$$\Delta\mathbf{w} = \alpha(t_i - y_i)\mathbf{x}_i \quad (2.5)$$

The idea is to correct the weights, in a direction that compensates the error, every time that the network gives a wrong answer, the amount of the correction being determined by a learning parameter α . It can be shown that, following this procedure, the network will end having a vector of weights that

classifies exactly the two sets.

Until now we have always considered networks with a binary output, but in some situations it is useful to have networks with a continuous output. For example, instead of using a neural network to classify two (or more) sets, we would like to use it to approximate or fit a function. In these cases, instead of applying a step function on the *activation* (the weighted sum of inputs and bias) of the node $a \equiv \mathbf{w} \cdot \mathbf{x}$, a sigmoid function is used

$$y = \sigma(a) \equiv \frac{1}{1 + e^{-ka}} \quad (2.6)$$

This function can be viewed as a smooth version of the step function, which is obtained in the $k \rightarrow \infty$ limit. (The output can also be considered as the probability to have output $y = 1$ in case of a yes or no unit. The use of continuous outputs can be justified, from a biological point of view, observing that the activity of real neurons is actually determined by the frequency of impulses, which is a real number).

In general it can be showed that, using a smooth $y = \Phi(a)$ function for the output, the *delta rule* (where $\dot{\Phi}(a) \equiv \frac{d}{da}\Phi(a)$)

$$\Delta \mathbf{w} = \alpha (t_i - y_i) \mathbf{x}_i \dot{\Phi}(a) \quad (2.7)$$

corresponds to a procedure of minimisation of the *error function*

$$E = \frac{1}{2} \sum_i (t_i - y_i)^2 \quad (2.8)$$

of the network on the training set. In case we use a sigmoid output function, the delta rule can be expressed using the values of the outputs, since

$$\dot{\sigma}(a) = k y (1 - y) \quad (2.9)$$

If we want classify more than two sets, or fit vector functions, we just need more than a single output unit, and all the previous formulae remain valid under a substitution of the scalar values y, t with vectors \mathbf{y} and \mathbf{t} (see figure 2.4).

Feedforward networks with hidden layers and backpropagation algorithm

Obviously not all the problems are linearly separable, and thus we should expect perceptrons to fail in many situations. Nevertheless, it can be showed

2.3 Neural Networks

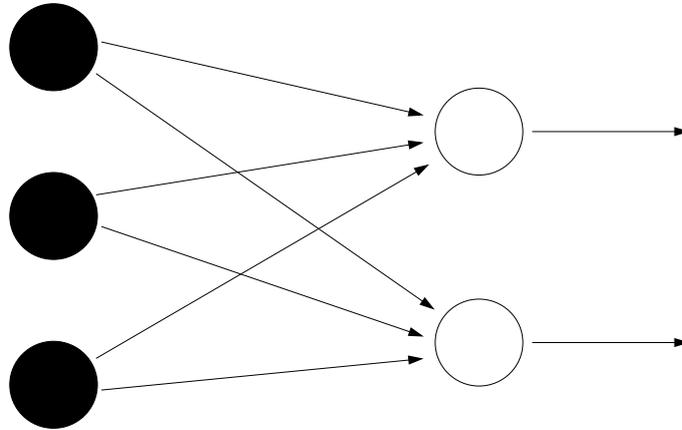


Figure 2.4: Two perceptrons used in a neural network with 3 inputs (black circles) and two outputs.

that a network that uses an hidden layer of nodes that elaborate the external inputs, and pass their output as inputs for the output units, can identify any region (including non convex and non connected ones) of \mathbb{R}^n . Furthermore these *feedforward neural networks with an hidden layer* (figure 2.5) can approximate any continuous function as closely as wanted.

The delta rule generalises to the *backpropagation algorithm* in case of the presence of an hidden layer. Following [5], denoting \mathbf{x} with n components x_k , $k = 1, \dots, n$ as a vector in the training set (external inputs); \mathbf{h} as the vector of outputs of the l nodes in the hidden layer with components h_j , $j = 1, \dots, l$, \mathbf{y} as the vector of m external outputs with components y_i , $i = 1, \dots, m$; the matrices v_{jk} and w_{ij} as the connections (weights), respectively, between the inputs and the hidden layer units, and of hidden layer units with the output nodes; we have

$$h_j = \Phi \left(\sum_k v_{jk} x_k \right) \quad (2.10)$$

and

$$y_i = \Phi \left(\sum_j w_{ij} h_j \right) \quad (2.11)$$

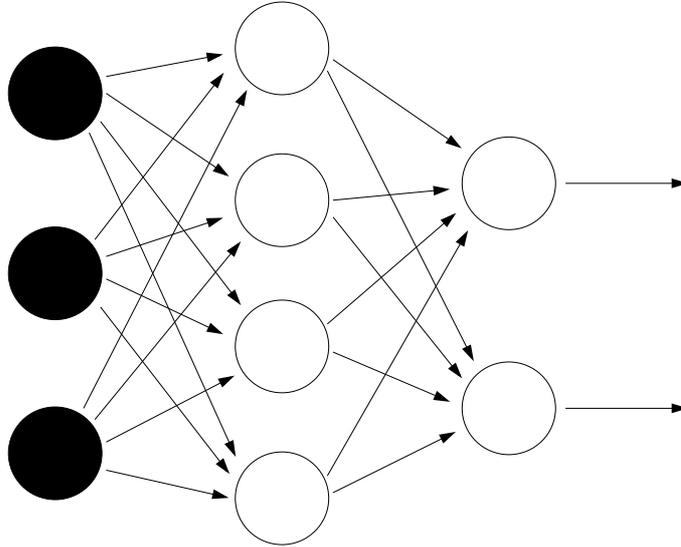


Figure 2.5: A (fully connected) feedforward neural network with 3 inputs (black circles), 4 units in the hidden layer and 2 outputs.

We define the *delta error* of the output units as

$$\delta_i = \dot{\Phi} \left(\sum_j w_{ij} h_j \right) (t_i - y_i) \quad (2.12)$$

where t_i are the m components of the target (expected) output \mathbf{t} .

Assuming that the error “propagates back” through the connection weights we define the delta error at the hidden units as

$$\delta_j = \dot{\Phi} \left(\sum_k v_{jk} h_k \right) \sum_i w_{ij} \delta_i \quad (2.13)$$

Finally the weights are corrected according to

$$\Delta w_{ij} = \alpha \delta_i h_j \quad (2.14)$$

$$\Delta v_{jk} = \alpha \delta_j h_k \quad (2.15)$$

and the procedure is repeated for any \mathbf{x} in the training set.

Networks with a suitable architecture and trained with backpropagation can learn, in principle, any arbitrary mapping between input units and output

2.3 Neural Networks

units (even if problems of overfitting of the data and convergence to local error function minima are present).

Also architectures different from feedforward networks are currently used in many applications (for example, to introduce memory effects *recurrent connections* are necessary, see for example [5]), but are not used in this thesis.

2.3.3 Ecological Networks

Until now we have exposed *supervised learning* algorithms, i.e. algorithms in which neural networks are trained on sets of examples which are provided, together with the right answers, by a teacher (the programmer). Two major shortcomings are present in this approach from the point of view of studying complex adaptive systems, i.e. if we want to use neural networks as the control system of agents in a distributed autonomous system.

The first problem is due to the fact that in this situation we are usually interested in agents moving freely in their environment, exploring it and “learning” about it (here learning is meant in a broader sense, since, as we will see, in many situations we are interested in agents that can evolve but not learn, or, even if the subject is not treated in this thesis, that can evolve and learn), and thus learning is not limited to a set of training problems given by a teacher. The “problems” that the agent controlled by the network experiences are determined by the direct interaction of the environment with the agent. Furthermore, the agent is not just passive in this process, since its interaction with the environment modifies the environment around the agent and thus the kinds of “experiences” that it can have.

Parisi [9] talks about *ecological* neural networks: these are the networks which are situated in their environment and thus are not simple passive receivers of inputs but can control their inputs through their actions. A typical strategy that ecological networks (i.e. agents which are controlled by networks) can follow to influence their input is movement: let us suppose that agents are located in an environment in which they have to perform a large number of tasks (they have to solve some problems), that each task affects in the same way their possibility to “survive” in the environment (the notions of “ecology” and “environment” are leading us back to the concept of fitness...not surprisingly GA will be back soon), and that these problems present themselves in spatially distinct regions of the environment. A trainer will probably

try to give to the network-agent a general training set, in which all kinds of problems are present, in order to obtain an individual able to perform in any possible situation. Typically an ecological neural network follows a quite different strategy: it specialises on a single task, and tries to move in the environment in order to have to face only that task. The strategies that ecological neural networks use to solve problems are at the same time simple and surprising, (they are usually emergent, using in this case the meaning of “unexpected”, since the researcher is not usually able to predict how the agent will solve the problem). In many cases these networks are able to solve problems that the researcher who set up the experiment thought that they could not solve, for too simple. (Many cases are described in [9], [5] and also in chapter 6. In particular, chapter 5 in the book by Nolfi and Floreano is extremely interesting, with the description of how, by *sensory motor coordination*, i.e. moving in the right way, ecological networks -in their case the control system of real robots- are able to classify objects that trained networks are not able to classify).

2.3.4 Evolution of neural networks

In describing ecological neural networks we deliberately used the expression “possibility to survive” to describe the ability of networks to perform in the environment. This could seem a little surprising after the previous discussion about training sets, targets and delta rules, in which no notion of survival was present. This leads us to the second problem that “traditional” networks present if used in a distributed (adapting) autonomous system: not only we cannot (nor want to) provide to the agents a training set, since they find the problems to solve directly on the environment, but in general we neither know which is the “right” output that the network should correspond to a certain input. In fact, if we are interested in studying a system in which some kind of self organised global property emerges from the local behaviour of agents, in general we do not know which local rule will correspond to an high level global performance, and thus we do not have any criterion to judge the single actions of agents, and thus if the network gave a good output. This does not mean that we cannot “judge” in any way the performance of a single agent (ecological network): it is possible to establish a criterion to judge how well an agent is “fit” to a given environment, but often it will be impossible to establish a correspondence between this “fitness” and a single evaluation of

2.3 Neural Networks

the neural network based control system, since the performance of the agent depends on a long succession of actions.

We are now back to the usual language of genetic algorithms: when using ecological neural networks, i.e. NN as the control system of agents interacting with the environment, the focus is usually shifted from learning to evolution: the weights of networks are evolved using GA. Instead of using a single network, whose weights are updated at each evaluation in order to obtain a better response (learning), a population of networks is created. The weights of the network are kept fixed during the whole life of an agent (length of a generation), and are evaluated using a fitness function related to its overall performance. At the end of the generation a new population is created using the genetic operators.

In this formulation adaptation does not happen on an individual level, during a single life, but at a population level, during many generations.

Many applications of neural networks in artificial life rely strongly on the use of evolutionary methods. Very interesting results have been obtained in the field of evolutionary robotics, i.e. using evolving neural networks as control system of real robots. The fact that evolution and GA have a predominant role in these applications does not mean that is not possible to use also learning methods, even if they are not used in this thesis. Chapter 7 of the book on evolutionary robotics by Nolfi and Floreano [5] treats the subject of the interaction between evolution and learning, which is interesting also from a biological point of view (Baldwin effect, [10]). Furthermore, evolution is not the only way to treat NN when supervised learning is not possible: *reinforcement learning* techniques are possible (see for example [11]). Finally, we have to notice that, even if in this thesis by evolution of neural networks we will always mean only evolution of the weight of the networks, actually GA have been used also to evolve learning rules and the structure of the networks. This last subject, in particular, is very interesting because it is strictly connected to the problem of open ended evolution (in order to evolve the structure of the network to arbitrarily complex structures chromosomes of unfixed length are necessary), and of the introduction of an actual distinction between genotype and phenotype (when the structure of the network gets more complex, it cannot be completely described by the genotype). Both [4] and [5] (chapter 9) give an introduction to these subjects, which have been object of many research works in the last years.

2.4 Application to the models in this thesis

2.4.1 Nature and purposes of the models

Two of the models studied in this thesis, and in particular the models on crowd dynamics and the evolution of the Theory of Mind (chapter 5) and that on the emergence of a communication system and of traffic flow rules in a mobility system (chapter 6), concern adaptive distributed autonomous systems with a large number of individuals.

Both models regard a similar problem, motion in a medium (a mobility system) with finite carrying capacities, i.e. optimisation of a traffic flow problem. Even if we use the word “optimisation” the problem is not faced from an engineering point of view, i.e. our purpose is not to find an optimal solution to an existing traffic flow problem. Neither the purpose of these models is to reproduce the behaviour of agents (people or even animals) in real mobility systems, even if we hope that the proposed models can give some insight about the nature of these systems (and at least in the crowd dynamics model we found some resemblance with the behaviour of actual pedestrians).

Our models are thought as experiments on emergent properties in the field of artificial life, in which a number of agents are located in an environment that changes according to their actions, and agents are expected to find a way to adapt to this environment. According to the interest of our research group to the study of problems related to the motion of pedestrians and vehicles in urban areas, we decided to use environments that, even if in a extremely simplified way, reflect the nature of these systems.

2.4.2 Justification of the use of GA and NN

Given these premises (a theoretical interest in general emergent properties and thus the search for a simple formulation of the problem in which those properties could arise with the smaller possible external -i.e. by the researcher- intervention) an approach based on genetic algorithms has been considered as the most natural. The suppression of learning could seem a major shortcoming for models describing the behaviour of agents in mobility systems (after all pedestrians and car drivers are complex individuals with learning ability) but this simplification seemed compatible with the request of simplicity and the pure theoretical nature of the work (the focus is more on the emergence of effective strategies and global properties than on the

2.4 Application to the models in this thesis

reproduction of the actual mechanisms that can lead to these properties). While in the crowd dynamics model, in order to introduce the concept of Theory of Mind, the qualitative nature of the control mechanism has been designed by us, letting to the GA the role of evolving a few parameters that fixed its quantitative behaviour, in the model on the evolution of traffic flow rules the request for a general and not biased control mechanism has led us to the use of feedforward neural networks with an hidden layer. Since the number of the neurons (units) in the hidden layer can affect the performance of the network, we took advantage of the high number of agents in our model to analyse in a parallel way which structure was more suitable to the solution of the problem (we could say that we used an extremely simplified version of a GA in order to evolve the structure of the network, a GA in which we used only selection and neither mutation nor crossover, i.e. we did not create any new structure but just choose between the existing ones).

2.4.3 Encoding

In both models the control mechanism of the agents was determined by a string of *real* numbers, the parameters of the collision avoiding mechanism in the crowd dynamics model and the weight of the neural networks in the model on emergence of communication and traffic flow rules (actually in both models also a single integer value was present, the “level of Theory of Mind” in the first model and the number of hidden units in the second one).

In order to use the original GA formulation by Holland, it should have been necessary to encode these numbers as binary strings. Nevertheless, the remarks that we have already expressed about the actual effectiveness of the crossover operator, at least in what regards these applications of genetic algorithms, have led us to the decision to encode these number directly as floating point numbers.

2.4.4 Mutation

We thus decided not to use at all a crossover operator, and the search in the space of solutions was performed only by the mutation operator. Mutation was performed adding to the real value (the “gene”) a number chosen using a Gaussian probability distribution with mean zero and width chosen in relation to the possible range of values assumed by the mutated gene. This

formulation seemed to us the most appropriate one under the assumption that the fitness of the agents is in some way a smooth function of the real parameters. The use of a Gaussian probability distribution assures that a large number of small mutations will lead to solutions which are similar to the previous ones (these mutations allow the system to climb up local maxima in the fitness landscape) while a few major mutations prevent the system to be stuck in local minima (for simplicity's sake we used the expressions "fitness of the agents" and "fitness landscape", even if these concepts are not well defined for the problems that we have studied, as we discuss below).

2.4.5 Fitness function

The choice of the fitness function is one of the most important steps in the use of a genetic algorithms, not only from a computational point of view (the use of an non appropriate one could prevent the evolution of the system) but also from a "philosophical" point of view, i.e. the choice of the fitness function has to reflect the "spirit" of the research project (see also the discussion about *fitness space* in [5]). To state it more clearly, since in our research projects we are interested in studying general emerging self organisation properties of the system, and want these properties to emerge in a "spontaneous way" (i.e. with a minimum possible intervention of the researcher), the fitness function has to be very simple, i.e. it has to include a very low number of terms. In particular, since in our mobility models the purpose of the agents is to reach a goal point in the shortest possible time, we just used as a fitness function the ratio between the distance from the starting to the goal point over the the time elapsed to reach the goal. Since in one of the two models (the crowd dynamics model) the physical interactions between the agents (collisions) have been described in detail, in the fitness function of that model we introduced also a negative term due to these collisions (since collisions were represented as elastic bounces between rigid discs, we considered the amount of exchanged momentum as the most straightforward way to represent these collisions in the fitness function).

2.4.6 Selection

An appropriate choice of the selection operator is very important in the development of a GA. Roulette wheel selection assigns to each agent a probability

2.4 Application to the models in this thesis

to reproduce which is proportional to its fitness. Even if this choice is consistent with the biological definition of fitness, it depends very strongly on the details of the fitness function. Furthermore, this choice presents some problems in the first stages of evolution (when usually few agents, or even a single one, have a fitness considerably higher than that of the others, and thus are the only ones to be selected, leading to a genetically very poor population) and in the last stages, when the differences in fitness between the agents are present but very small, in which case all the agents have almost the same probability to be selected, even if their performance is actually different.

For these reasons, *rank based selection operators* have been introduced, i.e. selection operators in which the probability to be selected was not directly proportional to the agent's fitness, but just to its position in the fitness rank. In our opinion, the most "natural" selection operator (i.e. the one that has the smaller dependence on particular choices made by the researcher) is *tournament selection*. In our implementation of tournament selection, a number s_t (the *size of tournament*) of agents from the previous generation is picked up and their fitness is compared. The winner, i.e. the agent with the highest fitness, "takes all" and his genes are passed to the following generation. This procedure is iterated N times in order to obtain a complete new generation. The size of tournament s_t has been fixed to 2 in case of small populations ($N \approx 100$) but could reach higher values for larger N (up to 40 for $N = 4000$).

2.4.7 The problem of defining a fitness landscape

Heterogeneous and homogeneous populations

When an agent is tested alone on the environment, i.e. not in presence of other agents, it is possible to talk about the agent's fitness according to the environment, and thus we can introduce the concept of fitness landscape. Obviously, if both the environment and the actions of the agents are probabilistic, we cannot expect to obtain always the same results under different tests of the same agents, nevertheless if our fitness function and our test conditions are general enough to reflect the nature of our problem, we can consider the fitness of the agent as a good estimate of its "absolute" ability to perform in the environment.

The situation is quite different when other agents, *with a different genome*, are present in the environment, competing and/or collaborating with them. In this situation the fitness of an agent cannot be considered anymore as

absolute (i.e. depending only on the structure of the environment) but it is relative to the other agents. (This situation is present, for example, in the experiments in which a *coevolution* of prey and predators is performed, see for example chapter 8 in the book by Nolfi and Floreano [5]).

When we are studying the behaviour of a crowd, or of a large scale mobility system, we obviously have to test agents in an environment in which other individuals are present. In this situation two contrasting choices are possible, the first one consisting in using *homogeneous* populations, i.e. populations of clones with an identical genome, and the second one in using *heterogeneous* populations in which each agent has a different genome.

In the first case we can talk about a fitness landscape: to each genome corresponds a fitness value (that can slightly vary due to random oscillation between different tests), which is the average fitness of the clones in the population. Also in the second case we can talk about the overall (and average) fitness of the population: to each *collection* of genomes in the population corresponds a fitness value. Nevertheless this information is of no use to the GA: the selection operator can rely only on individual fitness values, which are always relative to the composition of the population.

Two reasons led us to the choice of using heterogeneous populations. The first one is purely computational and due to the large number of agents used in our models. A genetic algorithm has always to operate on a population of genomes, even in the case in which population of clones with identical genomes are used, and thus in order to test all the genomes a larger number of simulation tests have to be performed. For example, in the case of the model studying the emergence of communication and traffic flow rules, we used an heterogeneous population with 4000 agents. Using 20 homogeneous populations with different genomes would have increased the computational time by a factor 20, while lowering the diversity of the population by a factor 200. In the case of the crowd dynamics model, in which $N = 100$, to an increase of a factor 20 of the computational cost would have corresponded a decrease in diversity of a factor 5. (Nevertheless in some situations the use of homogeneous populations has some benefits due to a higher reliability of the fitness test. For example, in the model on the emergence of communication and traffic rules, given the high dependence of the performance of agents on initial conditions, each agent had to be tested 20 times to obtain a reliable fitness. Using homogeneous populations we would have obtained, with a single simulation, a very reliable test, and thus the computational cost

2.4 Application to the models in this thesis

of the two experiment would have been the same. Furthermore, the loss of diversity would have been partially compensated by the higher reliability of the fitness function. In the crowd dynamics model this effect was not present because, even if also in that model agents have to be tested under different conditions to obtain a reliable fitness function, these conditions regard the nature of the environment, and thus no benefit comes from the use of homogeneous populations.).

The second reason is related to the nature and “spirit” of our work. Real biological environments (including crowds and urban areas) are not populated by clones, and thus diversity can be the source of some of the most interesting phenomena in these systems. For this reason we think that the use of heterogeneous populations is at the same time more sound (at least for what concerns the kind of problems that we want to solve) and more challenging (for example, in the case of the evolution of communication, at least in our model, one of the most interesting points was to study how and if a common communication system could be developed between heterogeneous agents).

Criteria for evaluating the results

Obviously the use of heterogeneous populations for which a purely individual fitness function cannot be defined presents some problems to a genetic algorithm that performs the search of a better solution on the base of individual fitness, and for this reason an efficient solution is more difficult to be found than in the case in which homogeneous populations are used. Furthermore, the results are more difficult to analyse. Since fitness can be properly defined only in relation to populations, we always used three criteria to test the agents evolved by our GA. The first criterion consisted in *analysing the results obtained by the best heterogeneous population during evolution* (i.e. the highest average fitness value obtained during the experiment). The second criterion consisted in *testing a population of clones, whose genes where obtained performing an average over the corresponding gene positions in the generation with the highest fitness*. The third criterion consisted in *testing a population of clones of the individual agent with the higher fitness during the experiment* (in general the fitness of the best individual was lower, once tested in a population of clones, than that of the average individual in the best population, and thus sometimes these results, as in the case of the model on emerging communication, are not shown).

2.4.8 A remark about our terminology

A last remark regarding the terminology has to be done. As we said before, to obtain a reliable fitness function agents in our models had to be tested under different conditions. Since in each one of these “sub-test” our agents had to “run” from their starting point to their goal, we call these tests effected under different conditions *runs*. In each run all the agents are tested together, and thus we can say that a generation (a complete test of all the agents) is composed of “many runs”. This notation is quite different from the standard one in GA, that, as we said before, uses the term run to refer the succession of all the generations in a single repetition of the experiment. (We will use in this case the term *evolutionary process* or just *repetition of the experiment*). Notice also that we will often use the term generation to refer to a complete test of the fitness of agents. Even if these tests represent the most computationally demanding part of our programs, actually a generation consists of the test *and* the application of the genetic operators.

2.4 Application to the models in this thesis

Bibliography

- [1] E. Bonabeu, M. Dorigo and G. Theraulaz, “*Swarm Intelligence*”, Oxford University Press, (1999).
- [2] S.W. Wilson, “*The animat path to AI*”, in J. A. Meyer and S. W. Wilson (eds.), “*From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behaviour*”, MIT Press/Bradford Books, (1991).
- [3] J. H. Holland, “*Adaptation in Natural and Artificial Systems*”, MIT Press, (1992).
- [4] M. Mitchell, “*An introduction to Genetic Algorithms*”, MIT Press, (1996).
- [5] S. Nolfi and D. Floreano, “*Evolutionary Robotics*”, MIT Press/Bradford Books, (2000).
- [6] R. F. Thompson, “*The brain: a neuroscience primer*”, Freeman, (1993).
- [7] S. Haykin, “*Neural Networks. A Comprehensive Foundation*”, Macmillan, (1994).
- [8] K. Gurney, “*An introduction to Neural Networks*”, UCL Press, (1997).
- [9] D. Parisi, “*Artificial life and higher level cognition*”, *Brain and Cognition*, 34:160-184, (1997).
- [10] J. M. Baldwin, “*A new factor in evolution*”, *American Naturalist*, 30:441-451, (1896).
- [11] R. S. Sutton and A. G. Barto, “*Introduction to Reinforcement Learning*”, MIT Press, (1998).

BIBLIOGRAPHY

Chapter 3

A gas of automata

3.1 Introduction

We propose the “gas of (von Neumann) automata” as a basic model for complex systems formed by a large number of interacting individuals provided with a sensory system, such as a crowd or a swarm [1]. By “automata”, in this context, we mean a point-like particle in a 2 dimensional space whose interaction properties are used to model in a very simplified way the interactions of agents provided with a sensory and decision system. The purpose of the model is not to build an accurate nor even a toy model of an actual complex agent system, but just to create a bridge between physical systems, that can be studied with the analytical methods of dynamical systems and statistical mechanics, and systems constituted of agents with a more complex perception and decision mechanisms that can be hardly studied with these methods. To do that we modify a known physical system (a gas of Coulomb particles) in order to describe an interaction based on perception, in particular sight. We thus assume that particles (automata) interact with the others only when these fall in their “sight cone”, a condition that introduces a *non Newtonian effect* (the third law of dynamics does not apply) and makes the behaviour of the system quite different from that of usual physical particles. We then assume, to introduce a very simplified “decision mechanism”, that the interaction between the agents is purely repulsive or “misanthropic”, i.e. intended to maximise the distance to other automata. In other words, we will just assume that, when the automata see each other, they interact as usual charged particles with the same charge sign and repel with central forces. The structure of the model is very simple and cannot describe neither ac-

3.2 The model

tual perception (sight is simply introduced as a sharp on or off condition for interaction) nor actual social interactions (that are more complex than pure misanthropy), but these simplifications allow us to focus on the effect of non Newtonian “perception like” interactions and to use the analytical methods applied to physical particles.

3.2 The model

We consider a set of automata A_1, \dots, A_N moving in the plane. Each automaton is point-like and its dynamic state is specified by its phase space coordinates $\mathbf{x}_i = (x_i, y_i, v_{xi}, v_{yi})$. The collective behaviour of the system is determined by external forces and by the mutual interactions. We introduce as an extremely naive representation of the interaction of two people a repulsive force which decreases with the distance and eventually vanishes beyond a given range R . Two automata A_1 and A_2 behave as two charged particles if they see each other. As a consequence in the asymmetric situations when A_1 sees A_2 but A_2 does not see A_1 the third principle of dynamics is manifestly violated and the behaviour is substantially different with respect to a usual system of interacting particles.

3.2.1 The sight cone

The forces \mathbf{F}_1 and \mathbf{F}_2 acting on the automata A_1 and A_2 are defined by

$$\mathbf{F}_1 = \mathbf{F}^{\text{ext}}(\mathbf{r}_1) + \frac{\mathbf{r}_1 - \mathbf{r}_2}{r_{12}} f(r_{12}) \vartheta(C_1) \quad (3.1)$$

$$\mathbf{F}_2 = \mathbf{F}^{\text{ext}}(\mathbf{r}_2) + \frac{\mathbf{r}_2 - \mathbf{r}_1}{r_{12}} f(r_{12}) \vartheta(C_2) \quad (3.2)$$

where C_1 and C_2 are defined by

$$C_1 = \mathbf{v}_1 \cdot (\mathbf{r}_2 - \mathbf{r}_1) - v_1 r_{12} \cos \alpha \quad C_2 = \mathbf{v}_2 \cdot (\mathbf{r}_1 - \mathbf{r}_2) - v_2 r_{12} \cos \alpha \quad (3.3)$$

By ϑ we have denoted the step function

$$\vartheta(u) = \begin{cases} 0 & \text{if } u \leq 0 \\ 1 & \text{if } u > 0 \end{cases} \quad (3.4)$$

and by v_1, v_2, r_{12} the norm of the vectors $\mathbf{v}_1, \mathbf{v}_2, \mathbf{r}_1 - \mathbf{r}_2$. Since the force is repulsive we have $f(r_{12}) > 0$. Letting ϕ_1 be the angle between the vector \mathbf{v}_1

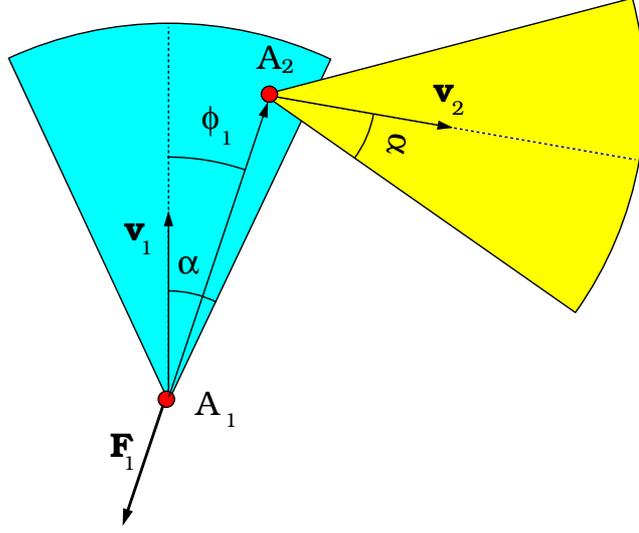


Figure 3.1: Definition of cones. The non Newtonian effect is due to the fact that A_1 “sees” (interacts with) A_2 , but A_2 does not “see” A_1 .

and the vector $\mathbf{r}_2 - \mathbf{r}_1$ pointing to particle 2 we have $\vartheta(C_1) = 1$ if $\phi_1 < \alpha$, see figure 3.1.

In case the automata have a finite sight (interaction) range, i.e. if the cone of vision is given not only by the angle α , but also by a radius r_v , the cone conditions (3.3) are multiplied by a $\vartheta(r_v - r_{12})$ factor

$$C_1^{r_v} = C_1 \vartheta(r_v - r_{12}) \quad C_2^{r_v} = C_2 \vartheta(r_v - r_{12}) \quad (3.5)$$

(we will assume $r_v = \infty$, and thus apply conditions 3.3, any time that we omit its value).

3.2.2 The two automata problem

We distinguish three cases to which correspond different dynamical properties.

I) The automata see each other: $\vartheta(C_1) = \vartheta(C_2) = 1$.

Supposing the forces are conservative the energy function is a first integral and the total angular momentum is conserved.

$$H_I = \frac{\mathbf{v}_1^2}{2} + \frac{\mathbf{v}_2^2}{2} + V^{\text{ext}}(\mathbf{r}_1) + V^{\text{ext}}(\mathbf{r}_2) + V(r_{12}) \quad (3.6)$$

3.2 The model

II) The automata do not see each other: $\vartheta(C_1) = \vartheta(C_2) = 0$.

The energy function is different because the mutual interaction potential has been switched off. The total angular momentum is still preserved.

$$H_{II} = \frac{\mathbf{v}_1^2}{2} + \frac{\mathbf{v}_2^2}{2} + V^{\text{ext}}(\mathbf{r}_1) + V^{\text{ext}}(\mathbf{r}_2) \quad (3.7)$$

III) Automaton A_1 sees A_2 but A_2 does not see A_1 : $\vartheta(C_1) = 1$ $\vartheta(C_2) = 0$. In this case there is no longer a preserved energy function nor the angular momentum is preserved. We notice that the power of the interaction force $\mathbf{F}^{\text{int}}(\mathbf{r}_1 - \mathbf{r}_2)$ acting on A_1 is always negative if $\alpha < \pi/2$, so that H_{II} decreases, since

$$\frac{dH_{II}}{dt} = \mathbf{F}^{\text{int}} \cdot \mathbf{v}_1 = \frac{f(r_{12})}{r_{12}} (\mathbf{r}_1 - \mathbf{r}_2) \cdot \mathbf{v}_1 < -\frac{f(r_{12})}{r_{12}} v_1 r_{12} \cos \alpha < 0 \quad (3.8)$$

IV) Automaton A_2 sees A_1 but A_1 does not see A_2 : $\vartheta(C_1) = 0$ $\vartheta(C_2) = 1$. This case is similar to the previous one, with non conservation laws and a dissipative effect on H_{II} .

We consider the equations of motion

$$\frac{d\mathbf{v}_1}{dt} = \mathbf{F}_1 \quad \frac{d\mathbf{v}_2}{dt} = \mathbf{F}_2 \quad (3.9)$$

where the forces \mathbf{F}_1 , \mathbf{F}_2 are given by (3.1), (3.2). As for the usual two body problem the analysis is more conveniently performed by introducing the relative and centre of mass positions \mathbf{r} , $\mathbf{R}/2$ and velocities \mathbf{v} , $\mathbf{V}/2$

$$\mathbf{r} = \mathbf{r}_1 - \mathbf{r}_2 \quad \mathbf{R} = \mathbf{r}_1 + \mathbf{r}_2 \quad \mathbf{v} = \frac{d\mathbf{r}}{dt} \quad \mathbf{V} = \frac{d\mathbf{R}}{dt} \quad (3.10)$$

We shall consider two different types of external confining forces: linear attracting or impulsive. In the first case the system moves in an harmonic potential well and the advantage is that the forces are linear. The second case corresponds to rigid boundaries where the reflection condition applies. Even though in the absence of mutual interaction we have a free motion, the potential of the unilateral forces does not split into the sum potential for the relative motion and centre of mass coordinates as in the usual case.

Choosing $\mathbf{F}^{\text{ext}}(\mathbf{r}_1) = -\omega^2 \mathbf{r}_1$ and $\mathbf{F}^{\text{ext}}(\mathbf{r}_2) = -\omega^2 \mathbf{r}_2$ the equations for the relative motion read

$$\frac{d\mathbf{r}}{dt} = \mathbf{v} \quad \frac{d\mathbf{v}}{dt} = -\omega^2 \mathbf{r} + \frac{\mathbf{r}}{r} f(r) \left(\vartheta(C_1) + \vartheta(C_2) \right) \quad (3.11)$$

The centre of mass equations read

$$\frac{d\mathbf{R}}{dt} = \mathbf{V} \quad \frac{d\mathbf{V}}{dt} = -\omega^2\mathbf{R} + \frac{\mathbf{r}}{r} f(r) \left(\vartheta(C_1) - \vartheta(C_2) \right) \quad (3.12)$$

The geometric picture of motion is simple. Phase space is split into four regions by the manifolds $C_1 = 0$ and $C_2 = 0$. Each region in phase space is labelled by $\text{sign}(C_1), \text{sign}(C_2)$ and consequently according to the previous classification we have

$$I = (+, +) \quad II = (-, -) \quad III = (+, -) \quad IV = (-, +)$$

We notice that the relative angular momentum $\mathbf{L} = \mathbf{r} \times \mathbf{v}$ is preserved and that three different energy functions H_I, H_{II}, H_{III} are preserved in the regions I, II and III, IV respectively. Taking into account the conservation of the angular momentum norm L we can write

$$H_I = \frac{\dot{r}^2}{2} + V_I^{\text{eff}}(r) \quad V_I^{\text{eff}}(r) = \frac{L^2}{2r^2} + \omega^2 \frac{r^2}{2} + 2V(r) \quad (3.13)$$

$$H_{II} = \frac{\dot{r}^2}{2} + V_{II}^{\text{eff}}(r) \quad V_{II}^{\text{eff}}(r) = \frac{L^2}{2r^2} + \omega^2 \frac{r^2}{2} \quad (3.14)$$

$$H_{III} = \frac{\dot{r}^2}{2} + V_{III}^{\text{eff}}(r) \quad V_{III}^{\text{eff}}(r) = \frac{L^2}{2r^2} + \omega^2 \frac{r^2}{2} + V(r) \quad (3.15)$$

As a consequence we are back to one dimensions problems where the standard analysis applies.

3.2.3 The geometry of switching

The non trivial aspect is the switch from one region to another since the boundaries are manifolds defined in the space \mathbb{R}^6 whose points are identified by $(\mathbf{r}, \mathbf{v}, \mathbf{V})$. The motion develops on the 5-dimensional manifold defined by $xv_y - yv_x = L$ for any given initial condition. Since the boundaries between the regions are not defined in the radial phase space \mathbb{R}^2 whose points are identified by (r, \dot{r}) , the picture is not geometrically intuitive even though the overall mechanism is clear. The problem is solved stepwise by quadratures. We determine $r(t), \dot{r}(t)$ by a quadrature from $H_I(r, \dot{r}) = E$ and then $\phi(t), \dot{\phi}(t)$ by another quadrature. The determination of $\mathbf{R}(t), \mathbf{V}(t)$ follows by two quadratures of (3.12) since $\mathbf{r}(t)$ is known. Starting from region I or II the centre of mass motion is given by harmonic oscillations so that

$$\mathbf{V}(t) = -\omega \mathbf{R}(0) \sin(\omega t) + \mathbf{V}(0) \cos(\omega t) \quad (3.16)$$

3.2 The model

$$\begin{pmatrix} V_r \\ V_\phi \end{pmatrix} = \begin{pmatrix} \cos \phi & \sin \phi \\ -\sin \phi & \cos \phi \end{pmatrix} \begin{pmatrix} V_x \\ V_y \end{pmatrix} \quad (3.17)$$

Starting from region *III* or *IV* the solution for $\mathbf{V}(t)$ is obtained by

$$\mathbf{V}(t) = -\omega \mathbf{R}(0) \sin(\omega t) + \mathbf{V}(0) \pm \int_0^t \cos(\omega(t - \tau)) f(r(\tau)) \frac{\mathbf{r}(\tau)}{r(\tau)} d\tau \quad (3.18)$$

As consequence, evaluated at the value \mathbf{V} takes at time t , the 5D manifold becomes 3D, and the intersection with the (r, \dot{r}) plane gives a line. Thus at any instant t we have in the (r, \dot{r}) plane two lines defining the four regions $(++, +-, -+, --)$. When time varies these lines move: we denote (see figure 3.2) by $\Gamma_1(t)$ and $\Gamma_2(t)$ the lines corresponding to $C_1 = 0$ and $C_2 = 0$ respectively. The point $P(t)$, representative of the relative motion, moves along its trajectory defined by the level line Γ of the corresponding energy function. Denoting by $Q_1(t)$ and $Q_2(t)$ the intersections of $\Gamma_1(t)$ and $\Gamma_2(t)$ with Γ , let us suppose that $P(t)$ and $Q_1(t)$ move towards each other on Γ . When they meet a change occurs from one region to another. The point will start moving along another trajectory defined as the level line of the corresponding energy functional until a new encounter with the intersection of one of the two moving lines $\Gamma_1(t)$ and $\Gamma_2(t)$ occurs again. The process can end when some sort of equilibrium point is reached. Since we are in the radial space this will be a relative equilibrium. This asymptotic situation, which is to be expected because in region II and III the overall energy is monotonically decreasing when α is smaller than $\pi/2$, will in general correspond to a location intermediate between the equilibria of the effective potentials $V_I^{\text{eff}}(r)$ and $V_{II}^{\text{eff}}(r)$.

To be more explicit the cone conditions are given by

$$C_1 = -\mathbf{r} \cdot (\mathbf{V}(t) + \mathbf{v}) - r \|\mathbf{V} + \mathbf{v}\| \cos \alpha \quad (3.19)$$

$$C_2 = -\mathbf{r} \cdot (\mathbf{V}(t) - \mathbf{v}) - r \|\mathbf{V} - \mathbf{v}\| \cos \alpha \quad (3.20)$$

Recalling that the transverse component of \mathbf{v} is $r\dot{\phi} = L/r$, the equations of the curves which split the radial phase plane into the four different regions are

$$\frac{L^2}{r^2} + V_\phi^2(t) \pm 2 \frac{L}{r} V_\phi(t) = \tan^2 \alpha \left(\dot{r}^2 + V_r^2(t) \pm 2V_r(t) \dot{r} \right) \quad (3.21)$$

We notice that when $\alpha \rightarrow \pi$ the cone condition is always satisfied ($C_1 > 0$, $C_2 > 0$), whereas when $\alpha \rightarrow 0$ this condition is never satisfied. In the

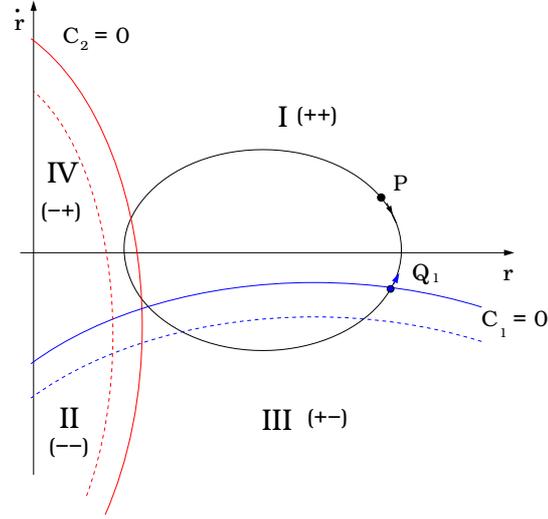


Figure 3.2: Radial phase space with a trajectory defined by the level lines of the Hamiltonian H_I supposing we start with initial conditions in the region I. In blue and red we plot the (time varying) cone conditions.

intermediate case $\alpha = \pi/2$ the cone condition takes a very simple form, namely

$$C_1 = -\dot{r} - V_r(t) \quad C_2 = -\dot{r} + V_r(t) \quad (3.22)$$

As a consequence in this specific case we find that if $V_r(t) > 0$ then

$$\text{Region I (+ +)} \quad \dot{r} < -V_r$$

$$\text{Region IV (- +)} \quad -V_r < \dot{r} < V_r$$

$$\text{Region II (--)} \quad V_r < \dot{r}$$

If $V_r < 0$ we have instead

$$\text{Region I (+ +)} \quad \dot{r} < V_r$$

$$\text{Region III (+ -)} \quad V_r < \dot{r} < -V_r$$

$$\text{Region II (--)} \quad -V_r < \dot{r}$$

3.3 The automaton in a central force field

Even though the cone condition is very simple in this case the actual analytical calculation is cumbersome because it requires integration (3.18) in order to obtain $\mathbf{V}(t)$ and the determination of the time law $\phi = \phi(t)$ in order to compute its radial component by using (3.16).

In order to understand in a simple way the dynamical implications of the cone condition we shall first consider a related but simpler problem, namely a particle moving in a central repulsive field, active only when the force centre falls within the cone of the particle.

3.3 The automaton in a central force field

We consider an automaton which moves in a central confining field and a repulsive potential which acts only when the force centre falls within the cone of the automaton. We choose the confining field to be linear corresponding to the potential $V^{\text{ext}}(r) = \frac{1}{2} \omega^2 r^2$ attractive field. It may also be replaced by a circular repulsive boundary at $r = R$ which corresponds to the limit of a non linear confining field defined by the potential

$$V^{\text{ext}}(r) = \lim_{m \rightarrow \infty} \begin{cases} 0 & \text{if } r < R \\ (r - R)^{2m} & \text{if } r > R \end{cases} = \begin{cases} 0 & \text{if } r < R \\ +\infty & \text{if } r > R \end{cases} \quad (3.23)$$

The cone condition is defined by $\eta < \alpha$ where η is the angle between the velocity, along which we choose the axis of the cone, and the vector $-\mathbf{r}$ pointing from P to the origin.

$$-\mathbf{r} \cdot \mathbf{v} = r v \cos \eta \quad \eta < \alpha \quad (3.24)$$

As a consequence the cone condition can be written as $\cos \eta > \cos \alpha$, since the angles ϕ and α are defined in the interval $[0, \pi]$, so that equation (3.24) becomes

$$-\mathbf{r} \cdot \mathbf{v} > r v \cos \alpha \quad (3.25)$$

Using polar coordinates $\mathbf{r} = r \mathbf{e}_r$ and $\mathbf{v} = \dot{r} \mathbf{e}_r + r \dot{\phi} \mathbf{e}_\phi$ we have

$$-r \dot{r} > r (\dot{r}^2 + r^2 \dot{\phi}^2)^{1/2} \cos \alpha \quad (3.26)$$

We distinguish two cases:

a) Small cone: $0 \leq \alpha < \pi/2$

$$\dot{r} < 0 \quad |v_\phi| \equiv r |\dot{\phi}| < |\dot{r}| \tan \alpha \quad (3.27)$$

This condition is well defined even when $r = 0$, a condition that is never reached if the repulsive potential diverges at the origin as it is the case on the Coulomb like potential $V = -q^2 \log r$.

b) Half plane cone: $\alpha = \pi/2$.

The cone condition is simply $\dot{r} < 0$.

c) Large cone: $\alpha > \pi/2$.

In this case $\cos \alpha < 0$ and the cone condition is always satisfied if $\dot{r} < 0$. When $\dot{r} > 0$ we write equation (3.26) as $\dot{r} < (r^2 + r^2 \dot{\phi}^2)^{1/2}$ from which we obtain

$$|v_\phi| \equiv r |\dot{\phi}| > \dot{r} \tan(\pi - \alpha) \quad (3.28)$$

Since the forces are always central the angular momentum is preserved and we define $L = \|\mathbf{r} \times \mathbf{v}\| = r^2 |\dot{\phi}|$ as its norm, observing that the sign of $\dot{\phi}$ is preserved in any orbit. As a consequence we can summarise the cone conditions in the form

$$\mathbf{a)} \quad -r \dot{r} > L \cot \alpha \qquad \mathbf{b)} \quad -\dot{r} > 0 \qquad \mathbf{c)} \quad r \dot{r} < L \cot(\pi - \alpha)$$

Finally the cone condition can be expressed in a compact form

$$\vartheta(C) \qquad C = -r \dot{r} - L \cot \alpha \quad (3.29)$$

As a consequence the line separating the cone from the no-cone region is just

$$C = 0 \qquad r \dot{r} = -L \cot \alpha \quad (3.30)$$

The equations of motion read

$$\frac{d\mathbf{r}}{dt} = \mathbf{v} \qquad \frac{d\mathbf{v}}{dt} = -\omega^2 \mathbf{r} - \frac{\mathbf{r}}{r} \frac{dV(r)}{dr} \vartheta(C) \quad (3.31)$$

Taking into account the angular momentum conservation we introduce two energy functions: $H_I(r, \dot{r})$ which is preserved in the cone region $C > 0$ and $H_{II}(r, \dot{r})$ which is preserved in the no-cone region $C \leq 0$

$$H_I(r, \dot{r}) = \frac{\dot{r}^2}{2} + V_I^{\text{eff}}(r) \qquad V_I^{\text{eff}}(r) = \frac{L^2}{2r^2} + \omega^2 \frac{r^2}{2} + V(r) \quad (3.32)$$

and

$$H_{II}(r, \dot{r}) = \frac{\dot{r}^2}{2} + V_{II}^{\text{eff}}(r) \qquad V_{II}^{\text{eff}}(r) = \frac{L^2}{2r^2} + \omega^2 \frac{r^2}{2} \quad (3.33)$$

The geometric construction is rather simple: after drawing the hyperbola separating the cone and no-cone regions and supposing we start in region

3.3 The automaton in a central force field

II we draw an arch of the trajectory defined by $H_{II}(r, \dot{r}) = E_0$ where $E_0 = H_{II}(r_0, \dot{r}_0)$. Supposing that the trajectory intersects the hyperbola at the point (r_1, \dot{r}_1) , the orbit continues on the arc defined by $H_I(r, \dot{r}) = E_1$ where $E_1 = H_I(r_1, \dot{r}_1)$. The first intersection (r_2, \dot{r}_2) of this new arc with the hyperbola defines the transition to a new arc in the region *II* defined by $H_{II}(r, \dot{r}) = E_2$ where $E_2 = H_{II}(r_2, \dot{r}_2)$. The process usually ends on one orbit of H_I or H_{II} (see figure 3.5).

In figures 3.3, 3.4 we show the plot in the radial and configuration space for a given random choice of the initial conditions and different values of α . In figure 3.5 the asymptotic orbits in configuration space are shown.

In this case we chose a Coulomb repulsive potential namely

$$V_I^{\text{ext}}(r) = \frac{L^2}{2r^2} + \frac{r^2}{2} - \log r^2 \quad (3.34)$$

where $L = x(0)\dot{y}(0) - y(0)\dot{x}(0)$.

To summarise we can say in a synthetic way that the 2D model is governed by one conservation law, the angular momentum, and one alternating conservation law for the energy function. As a consequence we may resume the motion by introducing the following energy $H = H_{II} + (H_I - H_{II})\vartheta(C)$

$$H(r, \dot{r}) = \frac{\dot{r}^2}{2} + \frac{L^2}{2r^2} + \omega^2 \frac{r^2}{2} + V(r)\vartheta(-r\dot{r} - L \cot \alpha) \quad (3.35)$$

where

$$L = r^2 \dot{\phi} \quad (3.36)$$

3.3.1 The stop condition for $\alpha = \pi/2$

The simplest case to analyse is when $\alpha = \pi/2$ since the cone condition becomes $\dot{r} < 0$ and the change between regions *I* and *II* occurs at $\dot{r} = 0$ namely at the inversion points. It is not hard to see that in this case as soon as the inversion point r_n falls between the minimum r_I^{min} of $V_I^{\text{eff}}(r)$ and the minimum r_{II}^{min} of $V_{II}^{\text{eff}}(r)$ the motion is arrested and thus this inversion point becomes a stopping point. As a consequence the sequence of inversion points is

$$r_2 < \dots < r_{2n-2} < r_I^{\text{min}} < r_{2n} < r_{II}^{\text{min}} < r_{2n-1} < \dots < r_1 \quad (3.37)$$

and the orbit is formed by the arcs $r_0 - r_1$ in region *I*, $r_1 - r_2$ in region *II*, $r_2 - r_3$ in region *I*, and so on until $r_{2n-1} - r_{2n}$ in region *II*, with the point r_{2n}

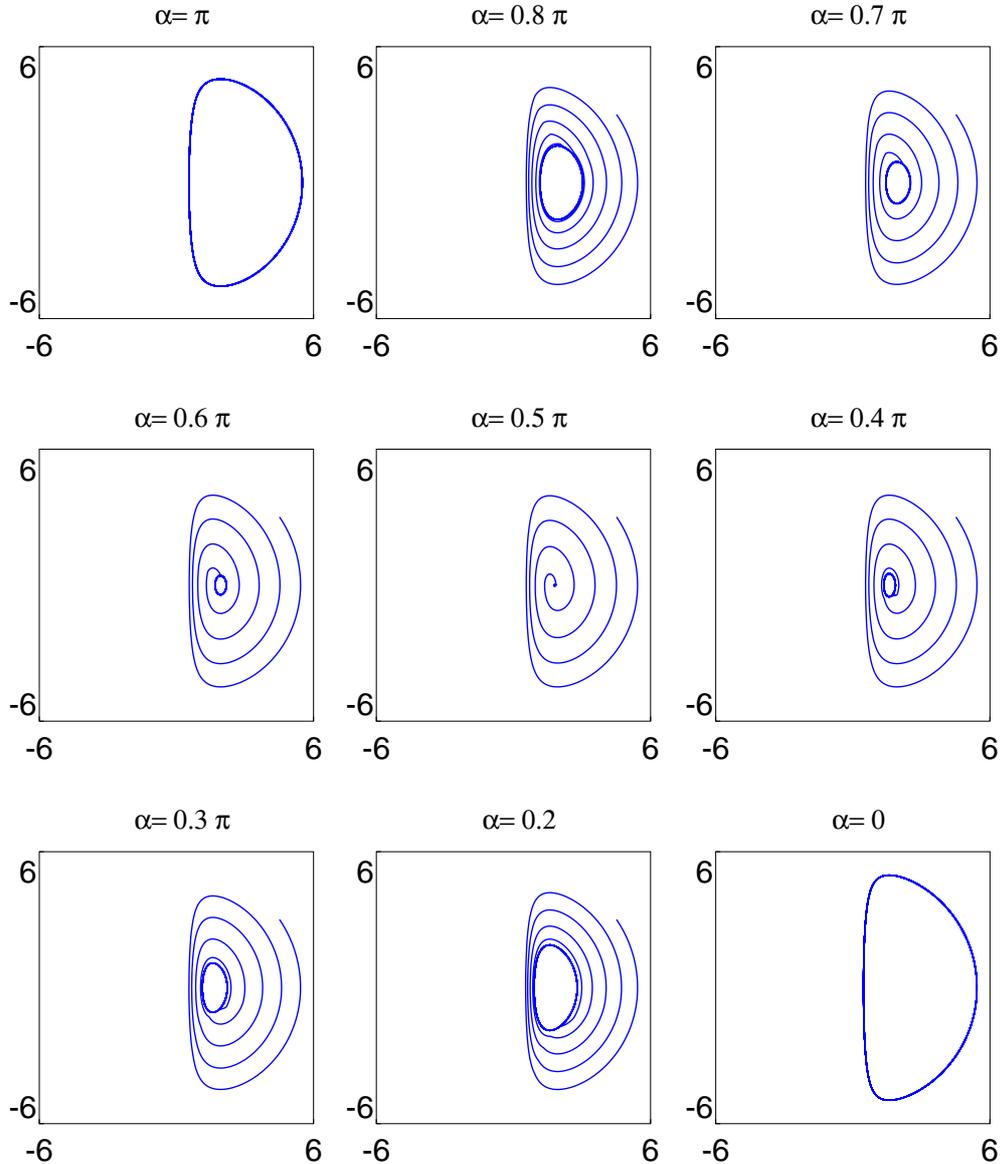


Figure 3.3: Plot of the orbits in the radial phase space (r, \dot{r}) corresponding to the initial conditions $x(0) = 3.721$, $\dot{x}(0) = 2.221$, $y(0) = 2.496$, $\dot{y}(0) = 2.171$ for different values of α . The potential is $V = \frac{1}{2}r^2 - \log r^2 \vartheta(C)$.

3.3 The automaton in a central force field

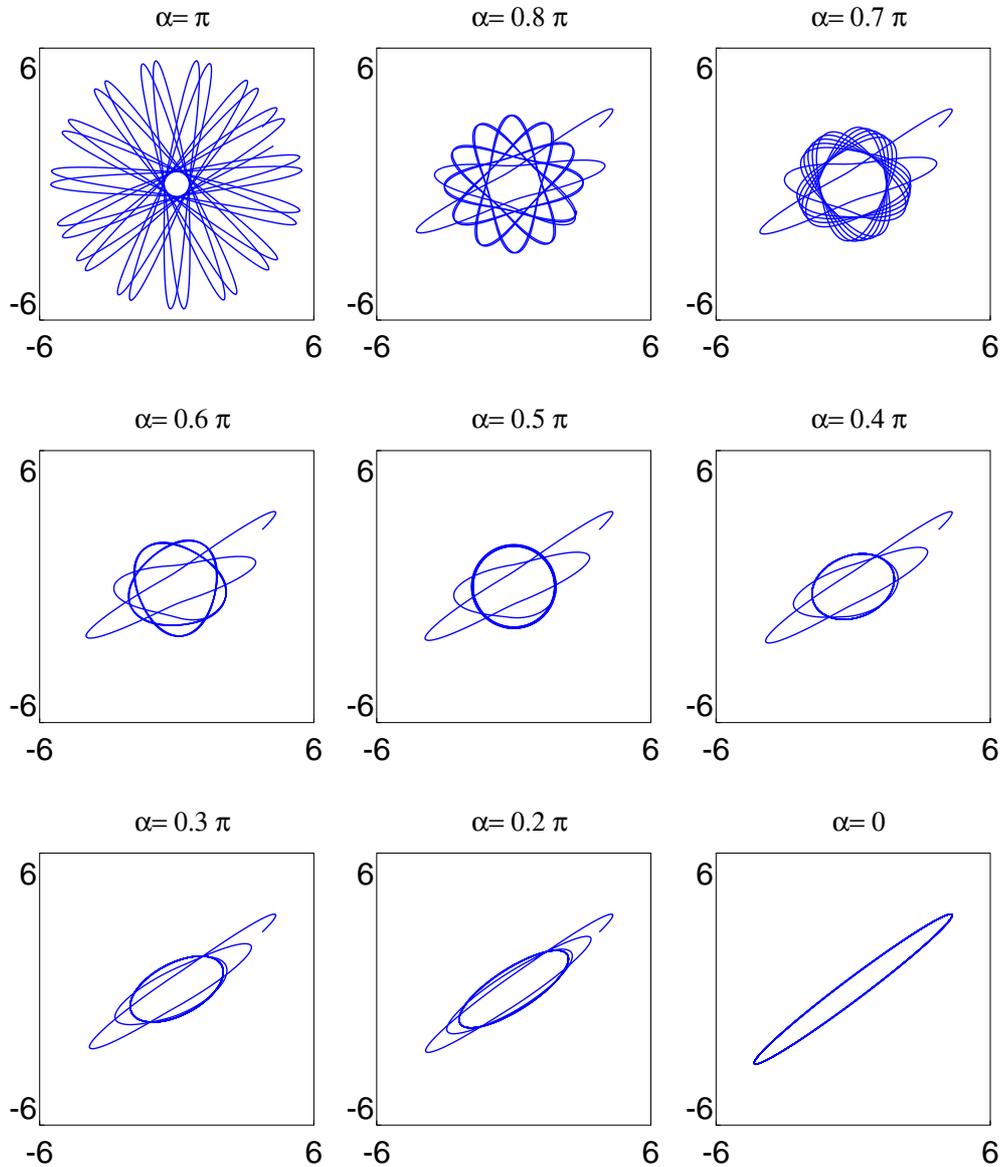


Figure 3.4: Plot of the orbits in the configuration space (x, y) corresponding to the initial conditions $x(0) = 3.721$, $\dot{x}(0) = 2.221$, $y(0) = 2.496$, $\dot{y}(0) = 2.171$ for different values of α .

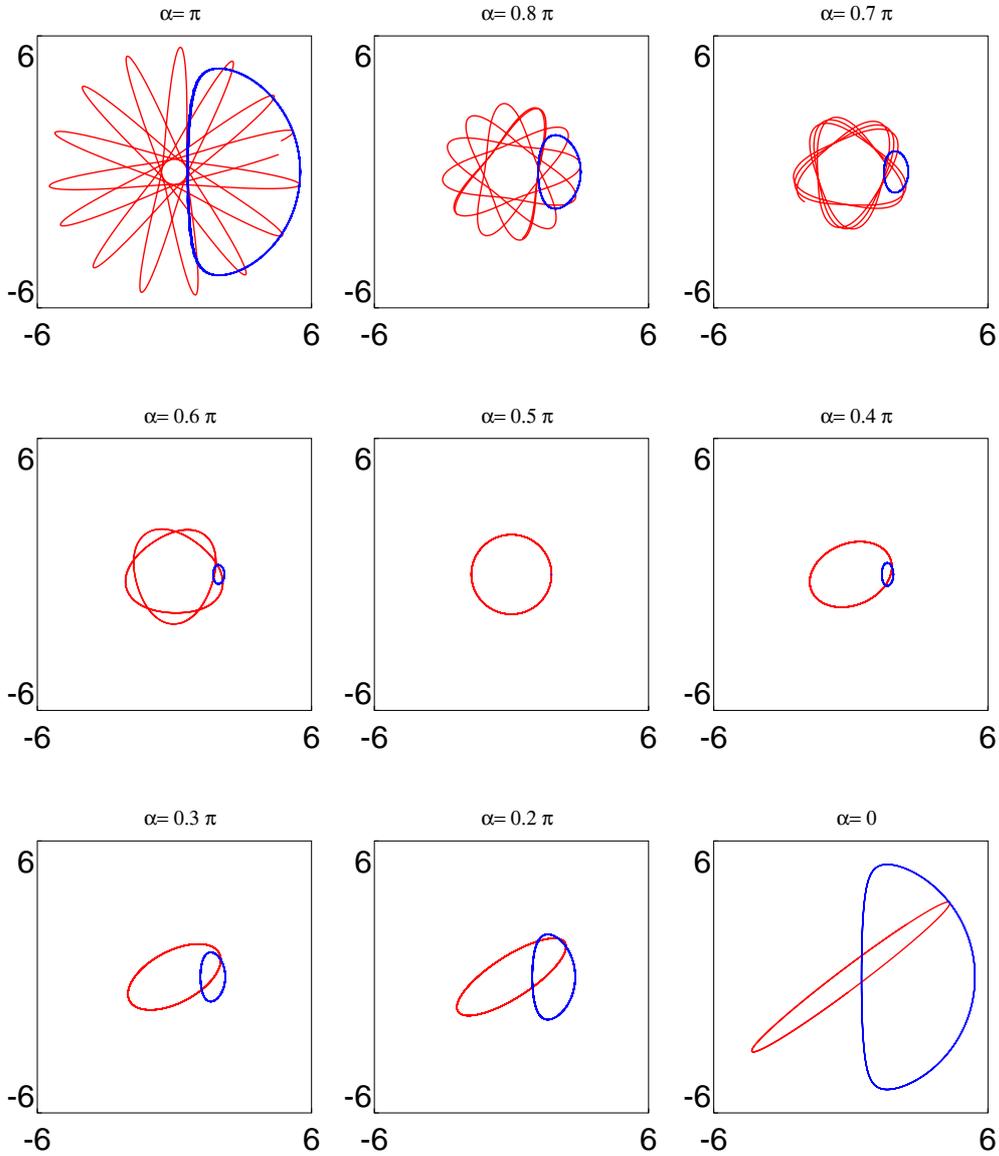


Figure 3.5: Plot of the asymptotic orbits $150 \leq t \leq 200$ in the configuration space (x, y) (red) and the radial phase space (r, \dot{r}) (blue) corresponding to the same initial conditions as for figures 3.3 and 3.4.

3.4 The 1D model

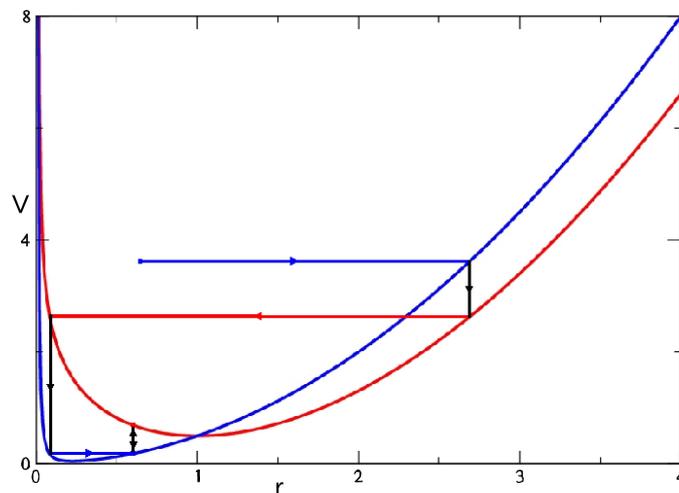


Figure 3.6: Dynamics in the r coordinate for an automaton in 2D central field, $\alpha = \pi/2$. The automaton feels the red potential ($V_I^{\text{eff}}(r)$) when $\dot{r} < 0$, the blue ($V_{II}^{\text{eff}}(r)$) one when $\dot{r} > 0$. At each inversion point there is a switch with energy loss, until it reaches a point between the two minima where it stops.

being the last point of the orbit, because this inversion point falls between the two minima and is a stopping point. This can be easily seen by drawing the corresponding orbits and taking into account that on any orbit we move clockwise. The dynamics in the r coordinate is described in figure 3.6. (See also the discussion about the stopping point in 1 dimension, section 3.4.1).

3.4 The 1D model

The 1D model is basically equivalent to the 2D model with central fields in the $\alpha = \pi/2$ case. The cone condition is given by $\vartheta(-x\dot{x})$ and the motion is governed by two alternating energy functions H_I in the cone region $x\dot{x} < 0$ and H_{II} in the no cone region. Using a compact notation we introduce the function $H = H_{II} + (H_I - H_{II})\vartheta(-x\dot{x})$. More explicitly we have

$$H(x, \dot{x}) = \frac{\dot{x}^2}{2} + V^{\text{ext}}(x) + V(x)\vartheta(-x\dot{x}) \quad (3.38)$$

If the potentials are even functions of x we may consider only the phase space defined in the half plane $\mathbb{R}_+ \times \mathbb{R}$ and we recover the correspondence with the

2D case with $\alpha = \pi/2$ by replacing $V^{\text{ext}}(x)$ with $V^{\text{ext}}(x) + L^2/(2x^2)$.

We shall consider a Coulomb-like repulsive potential $V(x) = -q^2 \log|x|$ and two types of confining forces corresponding to a harmonic potential and to repulsive barriers.

$$V^{\text{ext}}(x) = \frac{\omega^2}{2} x^2 \qquad V^{\text{ext}}(x) = \begin{cases} 0 & \text{if } |x| < R \\ +\infty & \text{if } |x| > R \end{cases} \quad (3.39)$$

3.4.1 Harmonic potential

In the left frame of figure 3.7 we show the potentials $V_I(x)$ and $V_{II}(x)$ which act when the cone condition is on and off respectively, for a confining harmonic potential, choosing $\omega = 1$, $q^2 = 2$. Since the potentials are even we refer to the x positive axis $x \geq 0$

$$V_I(x) = \frac{1}{2} x^2 - 2 \log x^2 \quad \text{if } \dot{x} < 0 \quad (3.40)$$

$$V_{II}(x) = \frac{1}{2} x^2 \quad \text{if } \dot{x} \geq 0 \quad (3.41)$$

In the right frame of figure 3.7 we show the phase space trajectories. Since the minimum of $V_{II}(x)$ is at $x = 0$ and the minimum of $V_I(x)$ is at $x = 2$ when the particle moving towards the origin reaches its inversion point x_1 it stops because this is at the left of the minimum of $V_I(x)$ and on the right of the minimum of $V_{II}(x)$. This can be proved in different ways: as soon as the particle reaches the inversion point x_1 of $V_I(x)$ at $t = \bar{t}$ then the potential is switched at $V_{II}(x)$. Since the force due to this potential is attractive the particle should start moving towards the origin with $\dot{x} < 0$. But this causes the immediate switch to the potential $V_I(x)$ and so on. The final effect is that the inversion point becomes a stopping point as previously claimed. (When a numerical integration scheme is used the automaton has small excursions to the left and right of x_1 whose amplitude decreases with the integration step. The presence of this inversion point is a source of instability for any numerical scheme, since a time step that is sufficient to resolve the smooth potential faces some problems when it reaches the stopping point. See also the relevance of this problem in the numerical study of the two automata problem, section 3.5.2).

Also the analysis of the phase space trajectories clearly explains the stopping mechanism: when x_1 is reached the motion along any arc of the curves defined

3.4 The 1D model

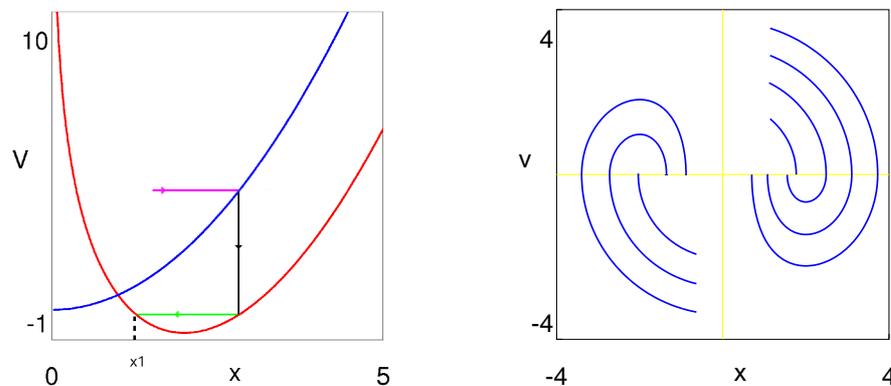


Figure 3.7: 1D model. Left: plot of the potentials $V_I(x)$ (red) and $V_{II}(x)$ (blue). The purple line corresponds to the energy level of the initial conditions. When the intersection with the blue ($V_{II}(x)$) potential is reached, the motion switches to the red ($V_I(x)$) potential and energy drops to the green level. Then when the intersection at left between the green line and the red curve is reached, the motion is arrested: this is the x_1 stopping point. Right: phase space orbits for different initial conditions. When the $v = 0$ axis is crossed in the $(-2, 2)$ interval, the motion is arrested.

by $H_I(x, \dot{x}) = V_I(x_1)$ and $H_{II}(x, \dot{x}) = V_{II}(x_1)$ terminates. One might say that the automaton continuously looks on the right and on the left but it cannot move.

3.4.2 The infinite potential well

A similar picture occurs when the oscillator potential is replaced by rigid barriers (infinite potential well). In figure 3.8 we show the mechanism which brings to the stopping point which occurs after that the particle is reflected on the barriers and reaches its inversion point. At $x = x_1$ the particle is switched from the repulsive potential, that we choose to be $V(x) = -\log|x|$, to the free state and since its velocity is zero it will stand still forever. In this case the turning back and forth while standing is avoided.

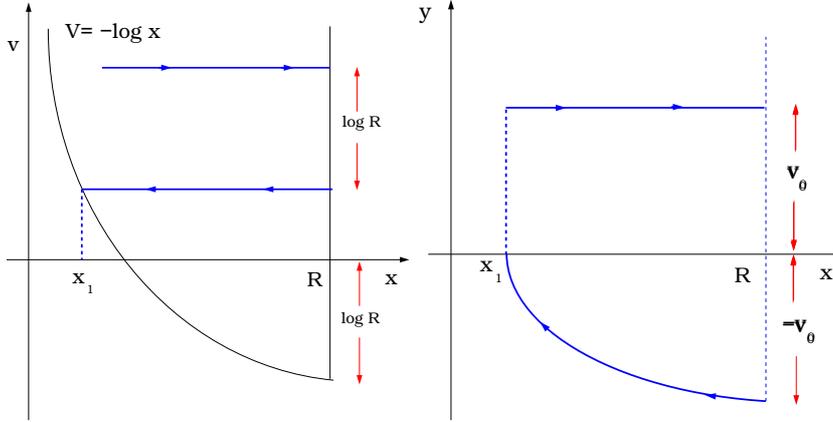


Figure 3.8: Left: plot of the potential $V(r) = -\log r$, the infinite barrier at $x = R$ and the inversion point x_1 after reflection. Right: the corresponding orbit in phase space.

It is easy to show that, denoting with (x_0, v_0) the initial position and velocity of the automaton, and with R the position of the barrier, the stopping point is given by

$$x_1 = \begin{cases} R e^{-\frac{v_0^2}{2}} & \text{if } v_0 > 0 \\ x_0 e^{-\frac{v_0^2}{2}} & \text{if } v_0 < 0 \end{cases} \quad (3.42)$$

3.4.3 Constant force

Another situation in which an analytic solution is possible is that in which we consider an attractive harmonic potential and a repulsive constant force with the cone condition

$$H = \frac{v^2}{2} + \frac{x^2}{2} - gx \vartheta(-x\dot{x}) \quad (3.43)$$

In this case the automaton is allowed to reach $x = 0$. To study the potential as symmetric, the y axis should be treated as a infinite barrier at which the velocity is reversed and the potential switched. Otherwise if we decide to use an asymmetric potential and the $\dot{x} < 0$ cone condition even for $x < 0$, we obtain that when $\dot{x} > 0$ the particle feels the potential $V_I(x)$, while it feels $V_{II}(x)$ when $\dot{x} < 0$ where

$$V_I(x) = \frac{x^2}{2} \quad V_{II}(x) = \frac{(x-g)^2}{2} - \frac{g^2}{2} \quad (3.44)$$

3.4 The 1D model

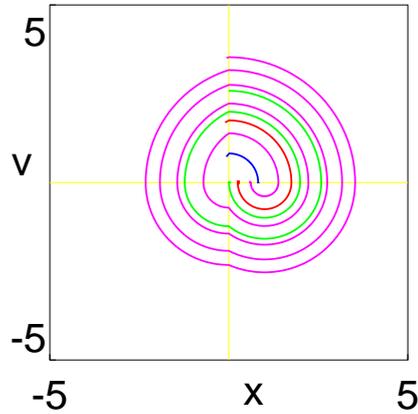


Figure 3.9: Phase space curves for different values of v_0 in case of a constant force potential with $g = 1$.

Suppose that our initial condition is $x(0) = 0$, $\dot{x}(0) = v_0 > 0$. Then the first inversion point is $x_1 = v_0$. The next one is x_2 where $V(x_2) = V(x_1)$ namely $x_2 = 2g - x_1$. The sequence is

$$\begin{array}{ll}
 V_I & V_{II} \\
 x_1 = v_0 & x_2 = 2g - v_0 \\
 x_3 = v_0 - 2g & x_4 = 4g - v_0 \\
 \dots & \dots \\
 x_{2n+1} = v_0 - 2ng & x_{2n+2} = (2n + 2)g - v_0
 \end{array}$$

The motion is arrested when the condition that x_k falls between the two minima is verified. As a consequence $0 < x_k < g$ implies the arrest and k is given by

$$k = \left[\frac{v_0}{g} \right] + 1 \quad (3.45)$$

where the square brackets denote the integer part. In figure 3.9 we show the phase space curves for different values of v_0 ($g = 1$). In case of a symmetric potential (3.45) applies only for $0 < v_0 < 2g$.

3.5 Numerical analysis of the two body problem

3.5.1 General behaviour for different α values

We have analysed numerically the two body problem for a repulsive potential $V(r_{12}) = -\log r_{12}^2$ (which corresponds, if $\alpha = \pi$, to a 2D Coulomb potential with $q^2 = 2$) and a harmonic confining potential with $\omega = 1$. We have chosen a set of random initial conditions and integrated the system by using a second order symplectic symmetric integrator and a fourth order Runge-Kutta. The time step was typically chosen equal to $\Delta t = 0.01$ but this was lowered to $\Delta t = 0.001$ in some critical situation. We recall that $\Delta t = 0.1$ is enough to integrate the system when $\alpha = \pi$ with a relative variation of the first integrals within 10^{-4} for $t \leq 10^3$.

The orbits for the relative motion of the two automata resemble in some way the behaviour of the single automaton moving in a central force, perceived only when the centre of forces falls in its cone (figures 3.3, 3.4). In Figure 3.10 we show the asymptotic curves for the centre of mass and relative motion, which result to be ellipses for any $\alpha < \pi$. The simulations show that, in a small neighbourhood of $\alpha = \pi/2$, the centre of mass comes to rest asymptotically. Nevertheless, this zone is numerically quite instable for a standard integrator and should be investigated by using exact algorithms based on the use of the first integrals (or using a very short integration time step, as we do in the next section).

3.5.2 Stable orbits for $\alpha = \pi/2$

The analytical analysis performed for the 1D model and its equivalent 2D model in the $\alpha = \pi/2$ case can help us in the analysis of stable orbits for the two body problem, at least for $\alpha = \pi/2$.

Let us put ourselves in the centre of mass reference frame and suppose that the two automata perform a circular orbit with velocities of equal norm and opposite in direction, as in figure 3.11, and that the relative distance between automata r (the diameter of the orbit) is such that $R_{II} < r < R_I$, where by R_I and R_{II} we denote the minima of potentials (3.13), (3.14).

This configuration is stable because, in a way similar to what occurred for the 1D stopping point, if the automata switch to cone zone I they feel a repulsive

3.5 Numerical analysis of the two body problem

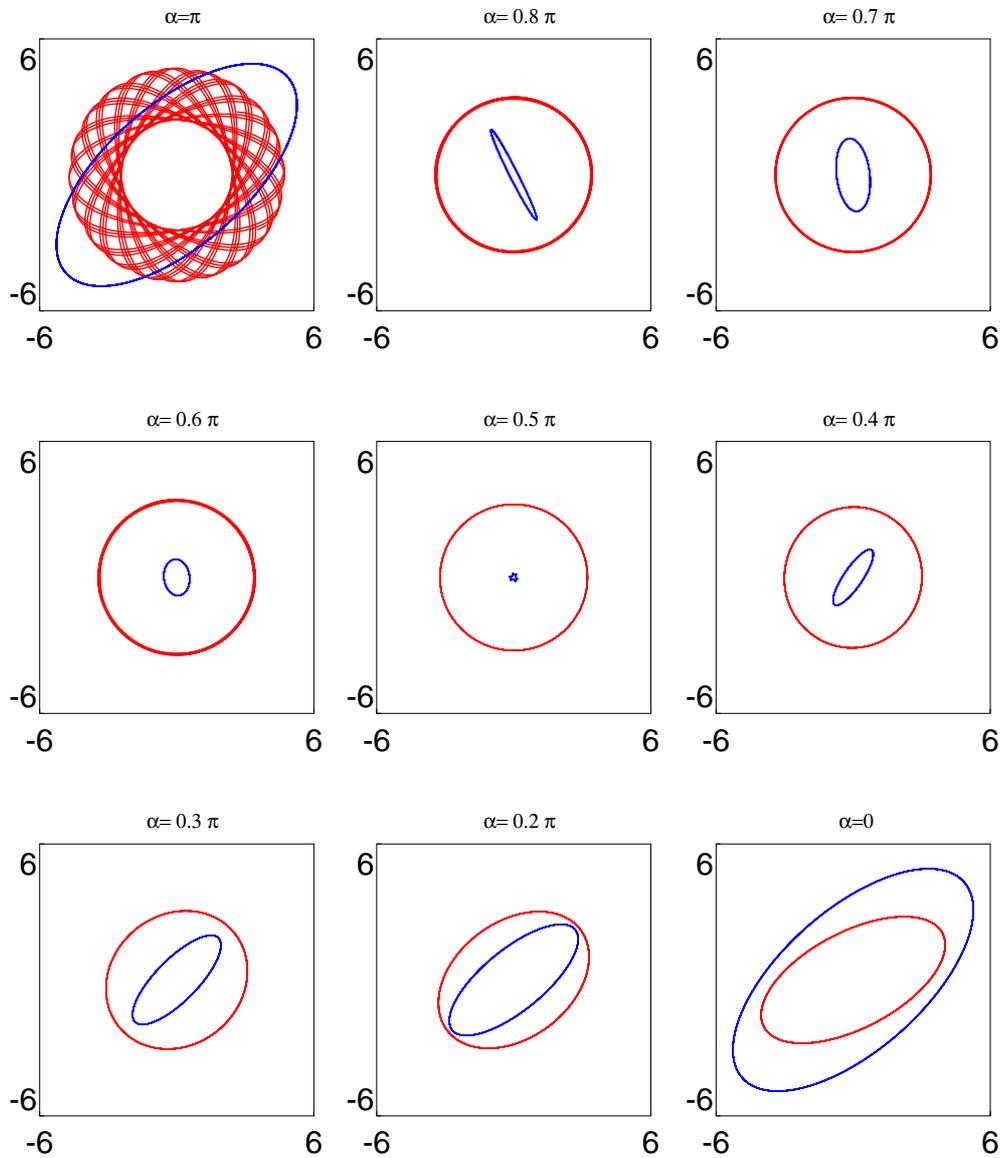


Figure 3.10: Asymptotic orbits for the centre of mass (blue) and the relative motion (red) taken in the time interval $150 \leq t \leq 200$. The confining potential is harmonic with $\omega = 1$ while the repulsive potential is $-\log r_{12}^2$. The various frames refer to different values of α . The initial conditions randomly chosen are $x_1(0) = 3.721$, $\dot{x}_1(0) = 2.221$, $y_1(0) = 2.496$, $\dot{y}_1(0) = 2.171$ for particle 1 and $x_2(0) = 1.425$, $\dot{x}_2(0) = -1.091$, $y_2(0) = -0.2983$, $\dot{y}_2(0) = 2.220$ for particle 2.

interaction and are pushed towards zone II , while the reverse happens in zone II , in which they feel an attractive force and are pushed towards zone I . The situation is somehow made more complex by the possible switches to zones III and IV and probably for this reason actual stable radii r_s are usually, according to numerical simulations, next to the minimum R_{III} of potential (3.15).

In figure 3.12 we show four (asymptotic) orbits for a two automata system with a $-\log r^2$ interaction term and an harmonic confining potential ($\omega = 1$). The two orbits with larger radius were obtained for systems with $L = 10$ (one with initial kinetic energy $T_0 = 10$, the other with $T_0 = 1$), while those with smaller radius for systems with $L = 1$ ($T_0 = 10$ and $T_0 = 1$).

The minima of the potentials are given by

$$R_I = \frac{L}{\left((4 + L^2)^{1/2} - 2\right)^{1/2}} \quad R_{II} = L^{1/2} \quad R_{III} = \frac{L}{\left((1 + L^2)^{1/2} - 1\right)^{1/2}} \quad (3.46)$$

We have performed a large number of simulations with different values of L , initial kinetic energy T_0 and different randomly chosen initial conditions for the automata positions. The resulting values for the radius of the stable orbit \bar{r} resulted to be always centred around R_{III} (the average value, once fixed L and T_0 , resulted to be in a first approximation equal to R_{III} , the difference being lower than the mean square deviation. The average value of \bar{r} resulted thus to depend only on L). Using $L = 7$ we have $R_I \approx 3.046$, $R_{II} \approx 2.645$, $R_{III} \approx 2.841$. Figure 3.13 shows the occurrence of \bar{r} values for 300 simulations with different initial conditions and $L = 7$, $T_0 = 10$. These results show that also in the two body problem, as it happened for the stopping point in the one dimensional problem with a single automaton, the value of the radius of the stable orbit \bar{r} depends in general from the initial conditions, and the only constraint is given by the conservation of L , i.e. $R_{II} \leq \bar{r} \leq R_I$. Nevertheless, for a very large number of initial conditions, we have $\bar{r} \approx R_{III}$ (i.e., $|\bar{r} - R_{III}| \ll R_I - R_{II}$).

While the radial asymptotic orbits always result in circular uniform motion, the centre of mass orbit, in whose reference frame the simple orbits in figure 3.11 are described, are in general not so trivial (figure 3.14).

These results were obtained using $\Delta t = 10^{-5}$ and the overall integration time was $T = 10^3$. If a longer ($\Delta t \geq 10^{-4}$) integration step was used, the asymptotic orbits resulted to be unstable, since dissipation led to a slow but

3.5 Numerical analysis of the two body problem

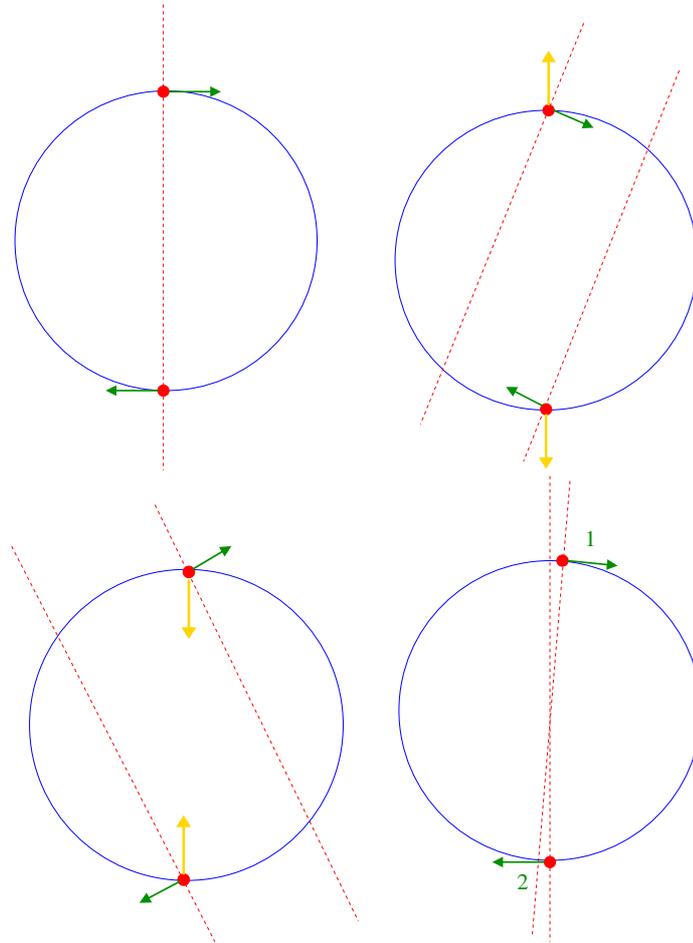


Figure 3.11: Top, left: a stable orbit for the two automata $\alpha = \pi/2$ problem. Right: if the automata get closer switching to cone zone *I* (they see each other) they feel a repulsive force, since their distance is lower than R_I . Bottom, left: when they switch to cone zone *II* the force is attractive, since their distance is greater than R_{II} . Right: this situation corresponds to cone zone *III* (1 sees 2 but 2 does not see 1).

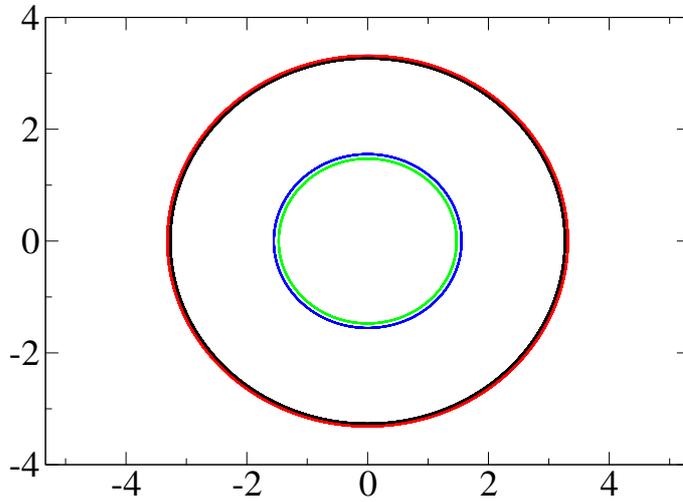


Figure 3.12: Four asymptotic orbits ($900 < t < 1000$) for the $\alpha = \pi/2$ automata system, radial motion. The red and black curves correspond to $L = 10$, the green and blue ones to $L = 1$.

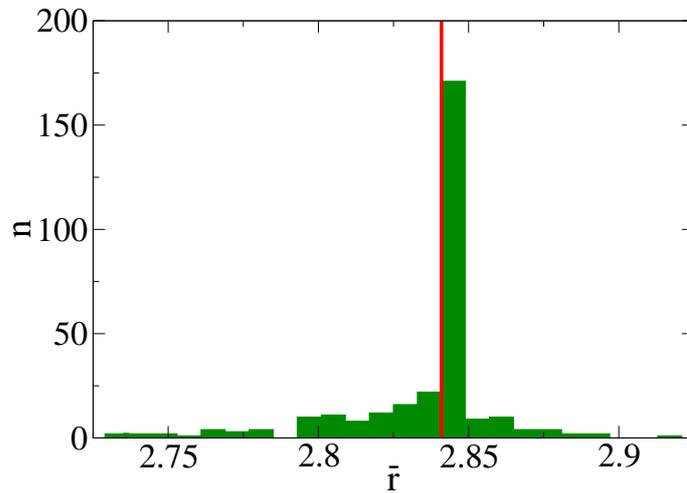


Figure 3.13: Occurrence of \bar{r} over 300 simulations with different initial conditions, $T_0 = 10$, $L = 7$. The red line shows the value of R_{III} .

3.6 The N automata problem

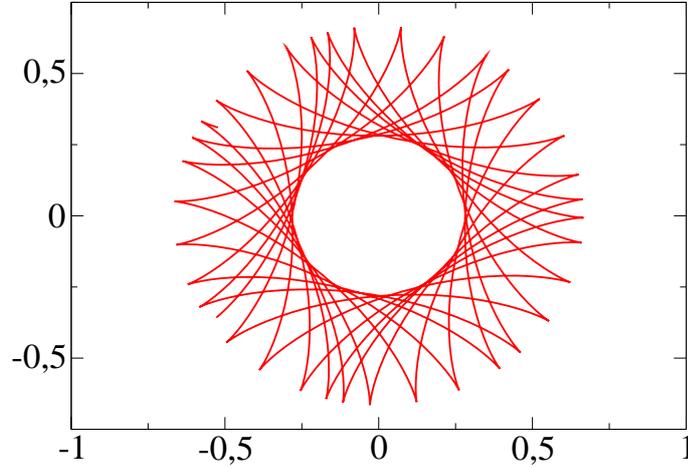


Figure 3.14: An asymptotic orbit ($900 < t < 1000$) for the $\alpha = \pi/2$ 2 automata system, centre of mass motion. Initial condition corresponded to $L = 1$ and kinetic energy $T_0 = 10$.

progressive decrease of the radius \bar{r} of the orbit, due to numerical instability. Even if the chosen integration method (second order symplectic integrator) would be apt to integrate a similar problem when using a smooth potential, the “switching” nature of the automata potential gives rise to numerical instabilities (see also the discussion about the difficulties of the numerical integrators when the one dimensional system reaches the stopping point, section 3.4.1). Nevertheless, since the integrator respects the conservation of L , the condition $R_{II} \leq \bar{r} \leq R_I$ is respected up to $\Delta t \approx 10^{-1}$.

Figure 3.14 shows that also the complete loss of kinetic energy for the centre of mass motion in the $\alpha = \pi/2$ case, as shown in figure 3.10, where due to numerical instability.

3.6 The N automata problem

The $N = 2$ problem is in some way unique since it is the only one that presents dynamically stable asymptotic orbits, like those described in section 3.5.2, while in all the $N \geq 3$ systems that we have numerically tested the kinetic energy dropped quickly to zero, i.e. the system converged to a state in which all the automata had velocity $\mathbf{v}_i = 0$. Figure 3.15 shows the time

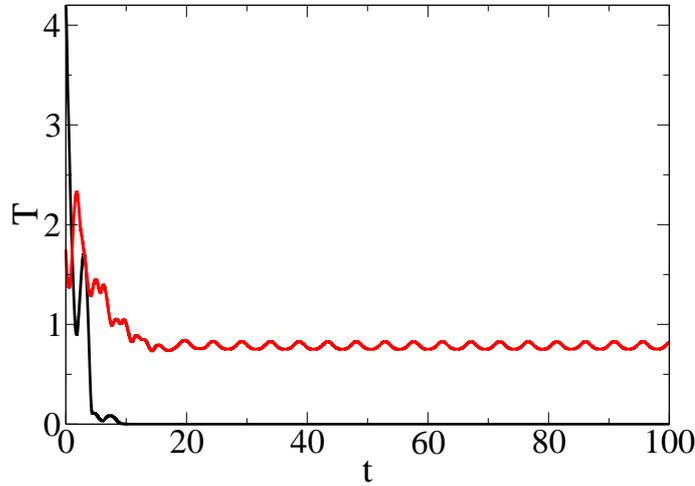


Figure 3.15: Time evolution of kinetic energy for a $N = 2$ (red line) and $N = 3$ (black line) system.

evolution of the average kinetic energy T for a system with $N = 2$ and $N = 3$ automata. While T oscillates around a stable equilibrium value T_∞ for $N = 2$ (corresponding to the relative motion and centre of mass orbits in figures 3.11, 3.12, 3.14), it quickly drops to zero for $N = 3$.

The equilibrium configuration for the $N = 3$ system corresponds, for any initial condition, to a (roughly) equilateral triangle centred around the centre of forces (figure 3.16), but this feature is unique and is not present for $N \geq 4$ (see figure 3.21) since as the number of automata grows dissipation gets stronger and the equilibrium state depends strongly on initial conditions and does not show any ordered structure (as for the one automaton 1D system the stopping condition is not unique but depends on initial conditions, as also does the radius of stable orbits in the $\alpha = \pi/2$ one and two automata problems).

Due to this strong energy dissipation the $N \gg 1$ automata system as we formulated it does not show any interesting dynamics. Dissipation can be attenuated choosing the right values of α and r_v (the radius of view, as defined in equation (3.5)) since conservative systems are obtained in the $r_v \rightarrow 0$, $\alpha \rightarrow 0$ and $\alpha \rightarrow \pi$ limits (the first two cases correspond to non interacting particles moving in the confining potential, while the latter to regular interacting particles, or Coulomb oscillators in case the interaction potential is chosen

3.6 The N automata problem

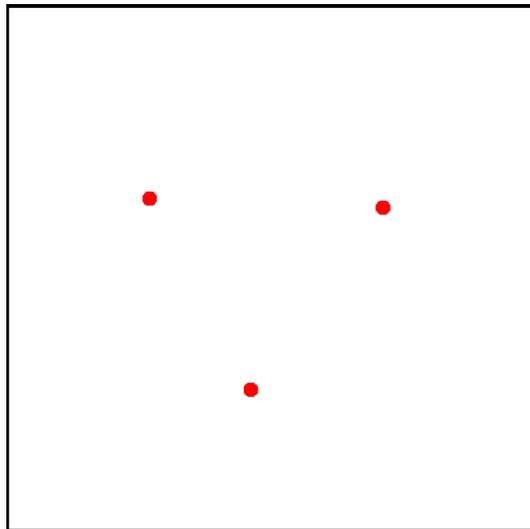


Figure 3.16: Equilibrium configuration for $N = 3$.

as $-q^2 \log r$) but this does not introduce substantial changes in the dynamics of the system (figures 3.17, 3.18).

3.6.1 Memory

In order to have an equilibrium state with positive average kinetic energy (or “temperature”) $T = \left(\sum_i \frac{v_i^2}{2}\right)/N$, we have introduced a “memory mechanism”. The idea at the base of this mechanism is that an agent with some kind of memory can retain for some time information about the position of another agent even after that this one has exited its sensory system (figure 3.19). A very simple way to introduce this mechanism would be to let the “observer” interact with the “observed” automaton for a time τ as if it were located in the point of space in which it has left the cone of vision (the observer just remembers the last point where it has seen the observed automaton). In a more complex and computationally expensive formulation, the “observer” would calculate an approximate trajectory for the “observed” automaton, in order to predict its position for a time τ after that it has left the cone of vision. In our model we use a more powerful and computationally not expensive (even if highly unrealistic) formulation: we suppose that the “observer” is able to calculate and thus to know precisely the position of the “observed” automaton for a certain amount of time after that the latter

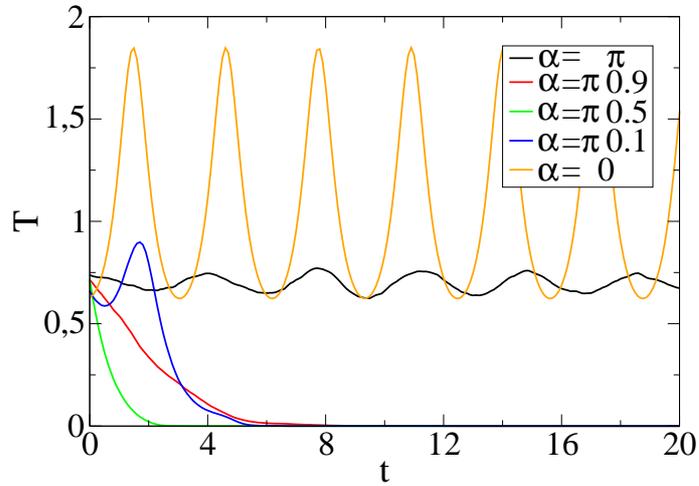


Figure 3.17: Time evolution of average kinetic energy for $N = 100$ automata, $r_v = \infty$, for different values of α . Initial condition were obtained as described in section 3.6.2.

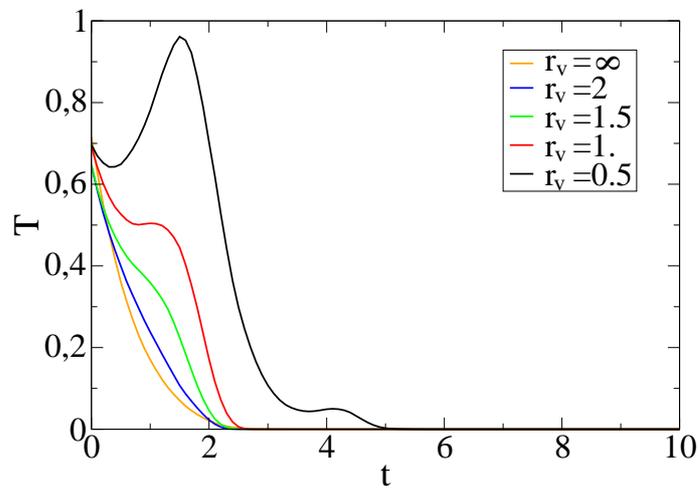


Figure 3.18: Time evolution of average kinetic energy for $N = 100$ automata, $\alpha = \pi/2$, for different values of r_v . Initial condition were obtained as described in section 3.6.2.

3.6 The N automata problem

has left the cone of vision and thus interacts with the observer for that time interval τ .

Therefore we introduce the memory mechanism in the following way: automaton A will continue to interact with automaton B for a time τ after B has left its cone of vision, at least until their distance will be inferior to the radius of vision $r < r_v$ (i.e. memory acts only on the α part of cone conditions (3.5), but not on the r_v part. In this interpretation, while α stands for a physical limit to the vision of the automaton, r_v is to be interpreted as a “zone of interest”: automaton A decides not to interact with B if their distance is greater than r_v).

In the $\tau \rightarrow \infty$ limit the conservative case is restored, at least after a transient phase (since not all the automata see each other at $t = 0$). Using finite values of τ , we found that an interesting dynamics was restored, since for high enough values of the memory time the system reached an equilibrium state with $T > 0$ (or at least a value of T which remained almost constant in the time scale that we investigated, see figure 3.20), while for very low but positive $\tau > 0$ values in the equilibrium configuration the automata had velocities $\mathbf{v}_i \approx 0$ (and thus $T \approx 0$) but were distributed in a spatially ordered structure (a “crystal”, see figure 3.21).

3.6.2 Numerical analysis of the N body problem

Initial Conditions

We have thus proceeded to explore the dependence on the control parameters α , r_v and τ of the equilibrium state of the system. In our numerical simulations we have used $N = 100$ agents located in a 2D open space, confined by an harmonic potential with $\omega = 1$, and interacting with a 2D logarithmic Coulomb potential $-q^2 \log r$.

In the $\alpha = \pi$, $r_v = \infty$ case the system corresponds to a system of 2D Coulomb oscillators [2] and is described by the Hamiltonian function (the mass of the particles has been fixed to $m = 1$)

$$H = \sum_i \frac{p_{xi}^2 + p_{yi}^2}{2} + \sum_i \omega^2 \frac{x_i^2 + y_i^2}{2} - \frac{\xi}{N} \sum_{i < j} \log r_{ij} \quad (3.47)$$

where, following the convention used in beam dynamics, we have defined $\xi = 2Nq^2$. In the $N \rightarrow \infty$ limit of this equation the one particle phase space

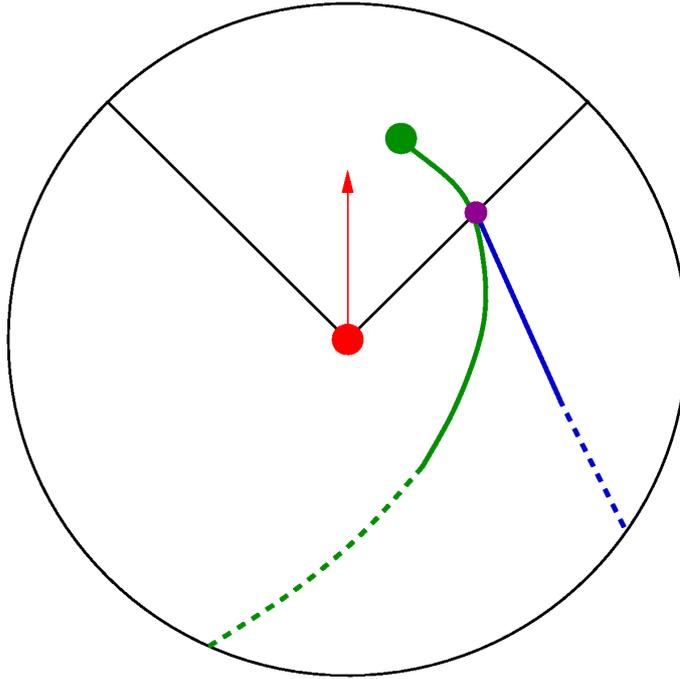


Figure 3.19: Three different mechanisms for memory. The red ball is the “observer”, while the green ball is the “observed” automaton, whose trajectory is given by the green line. A first memory mechanism consists in “remembering” the observed automaton in the point where it exits the cone, i.e. to interact for time τ with the automaton *as if it were located* on the violet ball. A second mechanism consists in remembering also its velocity, and thus *extrapolate its position as the blue (straight) line*. A third method, the one we used, to let the observer interact with the observed automaton located in its actual position. Continuous lines represent the part of orbits which is remembered (for a time interval τ). If τ is very high, the interaction continues until the distance between automata is $r < r_v$.

3.6 The N automata problem

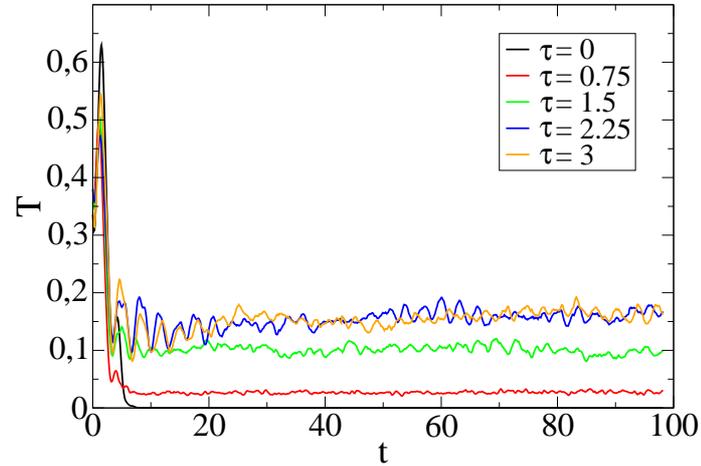


Figure 3.20: Time evolution of the “temperature” T of a system with $N = 100$ automata, $r_v = 1$, $\alpha = 0.1$, for different values of τ . Dissipation decreases as τ grows and for high enough values of τ the system reaches an equilibrium state after a transient.

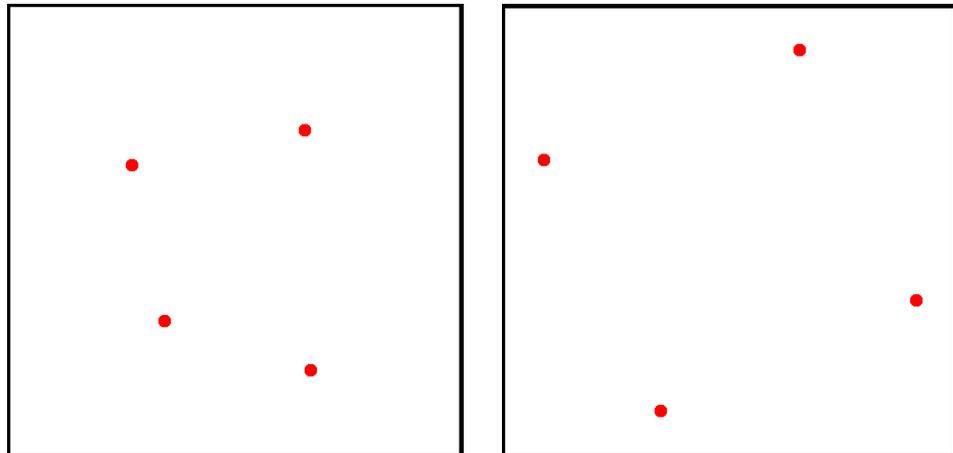


Figure 3.21: Equilibrium configurations for $N = 4$, $\alpha = \pi/2$, $r_v = \infty$. Left: equilibrium configuration for automata without memory, $\tau = 0$. Right: spatially organised configuration obtained using $\tau = 0.01$, in which all the automata have (roughly) the same distance from their first neighbours.

distribution $\rho(x, y, p_x, p_y)$ is governed by the Liouville equation

$$\frac{\partial \rho}{\partial t} + [H_1, \rho] = 0 \quad (3.48)$$

where H_1 is the one particle Hamiltonian whose interaction potential is

$$V(x, y) = - \int \rho_s(x', y') \log[(x - x')^2 + (y - y')^2]^{1/2} dx' dy' \quad (3.49)$$

that satisfies the Poisson equation

$$\Delta V = -4\pi \rho_s(x, y) \quad (3.50)$$

where

$$\rho_s(x, y) = \int \rho(x, y, p_x, p_y) dp_x dp_y \quad (3.51)$$

A self consistent solution of these equations is the KV distribution

$$\rho(x, y, p_x, p_y) = C^{-1} \delta(\bar{\omega}^2(x^2 + y^2) + v_x^2 + v_y^2 - \bar{\omega}^2 R^2) \quad (3.52)$$

where R is the radius of the self consistent charge distribution, while $\bar{\omega}$ gives the frequency of oscillations of particles in the charge distribution, and is given by

$$\bar{\omega}^2 = \omega^2 - \frac{\xi}{R^2} \quad (3.53)$$

In our simulations for the N automata problem, we used equation (3.52) as a probability distribution for our initial conditions, for any value of the control parameters, using $R = 1.84$, $\xi = 2$. The overall integration time has been chosen as 10 periods of Coulomb oscillators in the self consistent charge distribution, as given by (3.53), and the integration time step as $\Delta t = 10^{-3}$ using a second order symplectic integration method, which allowed, for Coulomb oscillators, a energy variation of order 10^{-3} in the studied time scale. Obviously the initial conditions correspond to a self consistent charge distribution for the Coulomb oscillators (at least in the $N \rightarrow \infty$ limit, i.e. the distribution is quite stable if $N \gg 1$) but neither for the non conservative automata systems nor for the $\alpha = 0$ non interacting system (see for example figure 3.18, where the kinetic energy oscillates strongly and to higher values for $\alpha = 0$ while it is almost constant for $\alpha = \pi$).

We studied the behaviour of the system in the $0 \leq \alpha \leq \pi$, $0 < r_v < 4$ and $0 < \tau < 20$ range. The maximum chosen value of r_v is larger than the diameter of the charge distribution, while the maximum τ value is more than

3.6 The N automata problem

twice the period of oscillation of particles, and thus this range of parameters should allow us to describe any feature of the behaviour of the system. The α range has been divided in 10 regular intervals, while the r_v and τ ones in 30 intervals, for an overall number of more than $8 \cdot 10^3$ simulation with different values of the parameters.

Temperature

We define the equilibrium temperature T^{eq} of our system as the average (over time) of the average (over the number of automata) kinetic energy

$$T^{\text{eq}} = \frac{\overline{\sum_i \frac{v_i^2}{2}}}{N} \quad (3.54)$$

where the time average is performed over the second half of the integration time, in order to avoid the contribution of the transient. (See figure 3.20. Actually not all the simulations reached such a clear equilibrium state within the chosen integration time and for those values the obtained “equilibrium value” has just an indicative significance).

The results are shown in figure 3.22. For any value of α , the highest T^{eq} value is reached when $r_v = 0$, i.e. for conservative non interacting particles (the red vertical line at left in each figure), while the lowest values are obtained for $\tau = 0$, i.e. for automata without memory whose kinetic energy drops to zero (deep violet bottom lines).

As a general rule we can say that T^{eq} , while keeping the other parameters fixed, grows with τ , decreases as r_v grows and has a minimum for $\alpha = \pi/2$. The reason of these rules can be easily explained. We know from the discussion in the previous sections, and in particular from equation 3.8, that non Newtonian interactions are dissipative since their direction is (always if $\alpha \leq \pi/2$ and in prevalence if $\pi/2 < \alpha < \pi$) opposite to the direction of the motion and thus have negative power. Introducing a memory effect we allow the automata to restore at least some part of the “positive power” interactions, and thus obtain an increase of energy. As we said before, when τ is high enough, the system behaves exactly as a conservative one (at least after a transient).

Also the decrease of T^{eq} for growing values of r_v is due to the negative power of interaction, since increasing the radius of view increases the range of interaction and thus dissipation. A particular case of this rule is obtained for

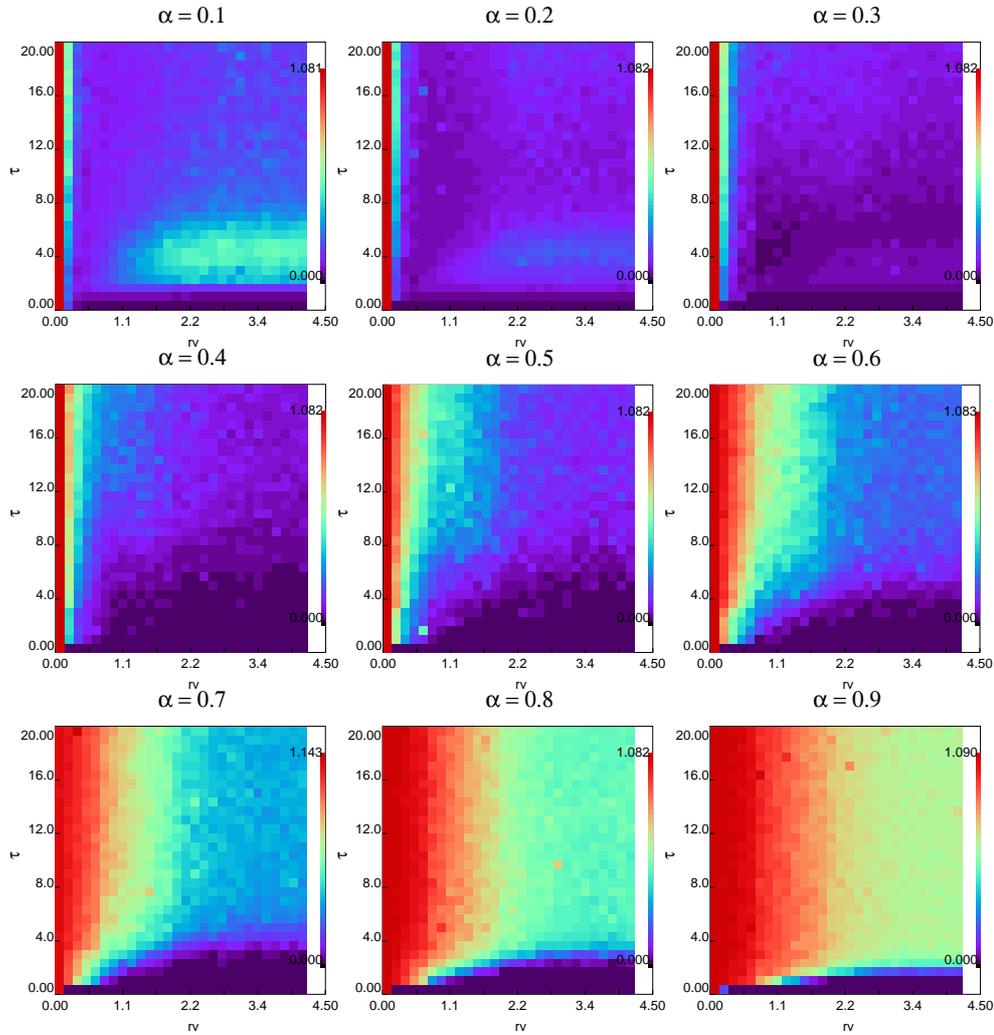


Figure 3.22: Dependence of the equilibrium value of the temperature T^{eq} on α , varying from 0.1 (top left) to 0.9 (bottom right); r_v (x axis) and τ (y axis). Red corresponds to higher values, violet to lower ones, as indicated on the colour bars.

3.6 The N automata problem

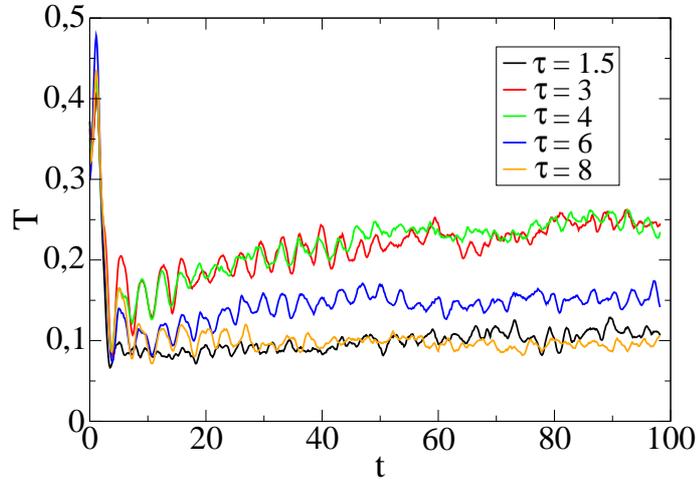


Figure 3.23: Time evolution of the “temperature” T of a system with $N = 100$ automata, $r_v = 2$, $\alpha = 0.1$, for different values of τ . Inside a given τ range the system reaches an higher equilibrium value.

high values of τ and α (top of the $\alpha = 0.9\pi$ graph in figure 3.22): for these values the system behaves as a conservative one, and as r_v grows a larger part of the energy is stored in potential energy, and thus kinetic energy (T^{eq}) decreases.

As we have already said, in the $\alpha \rightarrow \pi$ and $\alpha \rightarrow 0$ limits we obtain conservative systems, while for $\alpha = \pi/2$ all and only the interactions with negative power are present and thus dissipation is maximised, which explains the third law.

Nevertheless a few exceptions to these laws are found, in particular for low values of α . For example, when $\alpha = 0.1\pi$, we can see an “island” of high temperature in the $2 < \tau < 6$ range (figure 3.22, top-left), a temperature that decreases for higher τ values. An analysis of the dynamics of the system for these values of the parameters shows that, after a transient phase in which the behaviour is the same for all the τ values, for some values of α the system regains part of the energy it had lost (figure 3.23).

A possible explanation of this energy gain can be obtained on the base of figure 3.24. Let us consider an automaton moving in a central field (i.e., we suppose for the sake of simplicity that the other automaton it sees is not moving). The red circle in the figure represents r_v , the radius of view or

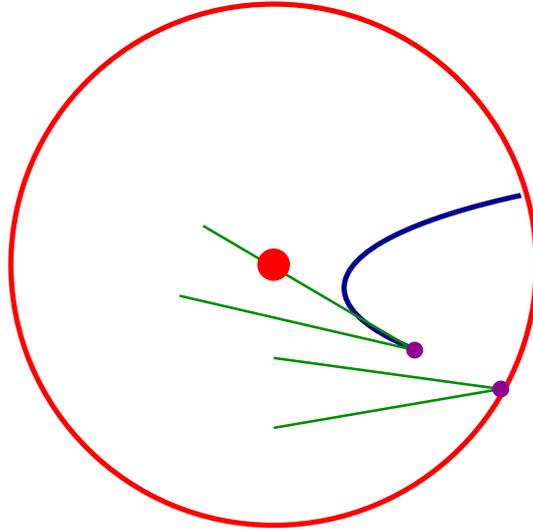


Figure 3.24: Energy gain mechanism for an automaton with low α and high τ .

range of interaction of the automaton. When the orbit of the automaton crosses the red circle it cannot see the centre of forces (red ball) because, due to the low α value, it does not fall into its cone of vision; but as soon as it sees the centre of forces, if τ is high enough, it will continue to interact until it exits the red circle. The violet curve in the figure stands for the portion of the automaton orbit in which it interacts with the centre of forces, and shows that the direction of the interaction is prevalently directed as the motion, and thus leads to an increase of kinetic energy. Nevertheless, if τ assumes an even higher value, the system behaves as a conservative one, and this process does not apply. Since this mechanism is based on memory, it occurs with some delay and is not present in the transient phase, in which all the systems behave in the same way.

Crystals

Coulomb systems are known [3] to assume crystal configurations when frozen to very low temperature $T \approx 0$. In our system, which is equivalent to Coulomb oscillators when $\alpha = \pi$, freezing occurs spontaneously as a result of dissipation for $0 < \alpha < \pi$. Moreover, we have verified that the $T \approx 0$ configuration corresponds to an uniform configuration in case of a few body

3.6 The N automata problem

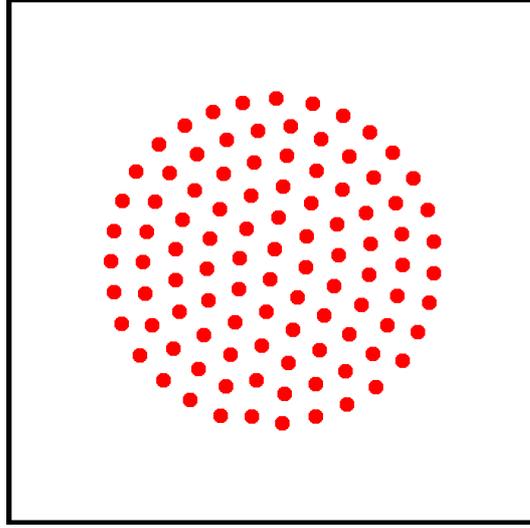


Figure 3.25: Equilibrium configuration for $N = 100$, $r_v = \infty$, $\alpha = \pi/2$ and $\tau = 0.1$.

system (figures 3.16, 3.21). We thus studied the possible occurrence of regular structures as equilibrium configuration for $N \gg 1$ automata systems. In figure 3.25 we show one of these $T^{\text{eq}} \approx 0$ “crystal like” configurations, obtained using $N = 100$, $r_v = \infty$, $\alpha = \pi/2$ and $\tau = 0.01$. In these structures automata are uniformly distributed inside a disc of given radius, and thus, in order to study the occurrence of these crystals, we studied a parameter, defined as

$$\gamma = \frac{\Delta d_f}{\langle d_f \rangle} \quad (3.55)$$

where $\langle d_f \rangle$ is the mean value of the distance to first neighbours, while Δd_f is its mean square deviation, which denotes in some way the “disorder” of the distribution, since its value is next to 1 for a random distribution, but drops to zero for uniform distributions as that of figure 3.25. The equilibrium values are shown in figure 3.26, which shows that crystals (violet lines at the bottom of graphs) arise in case of low but strictly positive memory $\tau > 0$ (as we have seen before, in the $\tau = 0$ case dissipation is too fast to form ordered structures, while for high enough τ the system does not freeze). Furthermore, in order to have an order structure, automata have to interact properly between them, and thus crystal do not arise for low values of α and r_v .

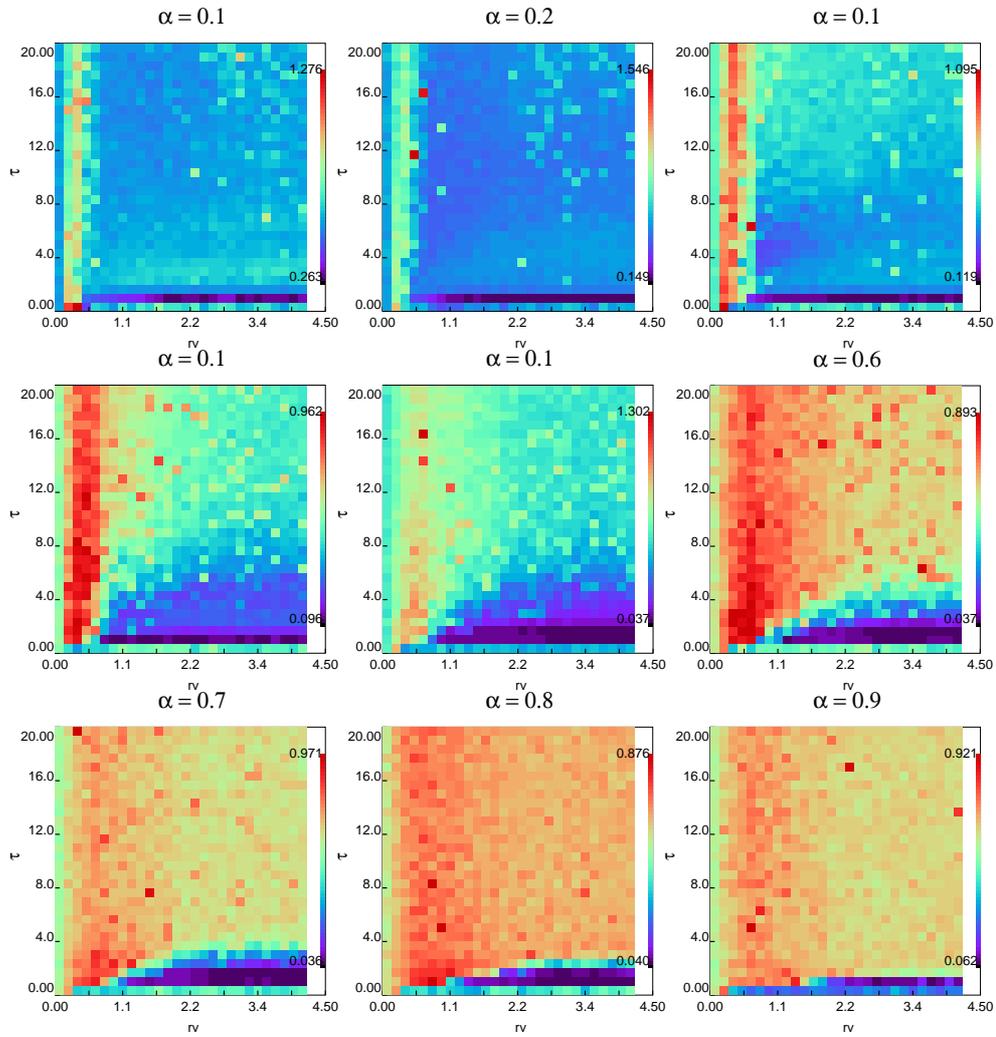


Figure 3.26: Dependence of the equilibrium value of the “disorder” parameter γ (defined in equation 3.55) on α , varying from 0.1 (top left) to 0.9 (bottom right); r_v (x axis) and τ (y axis). Red corresponds to higher values, violet to lower ones, as indicated on the colour bars.

3.7 Heterogeneous population of automata with different “social spheres”

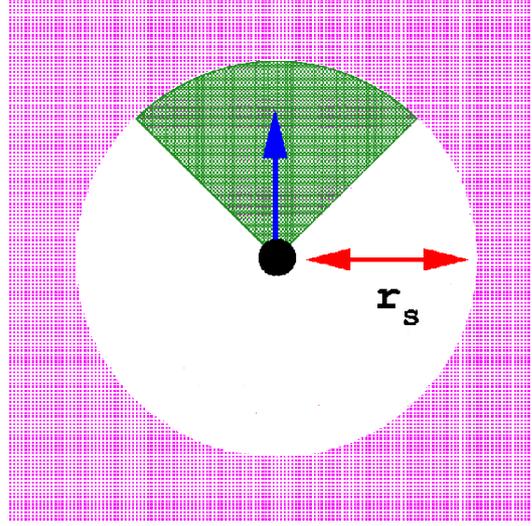


Figure 3.27: The force field around an automaton moving with the velocity denoted by the blue arrow. Outside its social radius r_s the force is attractive (violet zone), while inside the disc of radius r_s the force is null outside the cone of vision, and repulsive inside it.

3.7 Heterogeneous population of automata with different “social spheres”

We have used the results of this model in order to develop a slightly more complex system in which the interaction between automata was attractive if their distance r was higher than a given “social radius” r_s , while it was repulsive when $r < r_s$. The “force field” of the agent, as explained in figure 3.27, is given by equation

$$\mathbf{F}_{12} = q^2 \frac{\mathbf{r}_{12}}{r^2} \left(1 - \frac{r_{12}}{r_s} \right) C \quad (3.56)$$

where the cone condition C is defined by

$$C = \begin{cases} 1 & \text{if } r_{12} > r_s \\ \vartheta(\mathbf{v}_1 \cdot (\mathbf{r}_2 - \mathbf{r}_1) - v_1 r_{12} \cos \alpha) & \text{if } r_{12} \leq r_s \end{cases} \quad (3.57)$$

The force can be interpreted in this way: automaton A_1 is attracted to automaton A_2 by a “vision independent” attractive force that goes to zero as soon as $r_{12} \leq r_s^1$, where r_s^1 is the “social radius” of A_1 , which can be

different from r_s^2 . Then, if its distance from A_2 gets lower than r_s^2 , A_1 applies a repulsive force in order to restore its “social distance”, but only if it sees A_2 . Since the α dependence of the cone condition is present only for the repulsive part of the force, the system is once again dissipative and, in absence of memory, for $0 < \alpha < \pi$ its kinetic energy drops to $T^{\text{eq}} = 0$.

In our simulations we split the agents in two groups with $r_s^2 > r_s^1$. In this way we have introduced another non Newtonian effect, besides that due to the cone condition, since now the intensity and direction of \mathbf{F}_{12} and \mathbf{F}_{21} can be different even when both cone conditions are switched on.

Figure 3.28 shows the equilibrium configuration for a system of 50 automata with $r_s^1 = 2$ and 50 automata with $r_s^2 = 1.5$. The values of the other parameters were $\alpha = \pi/2$ and $\tau = 0.1$.

Due to the low $\tau > 0$ value, the kinetic energy of the system goes quickly to zero, but the presence of a memory effect allows the system to reach a minimum energy ordered configuration. Nevertheless, this configuration is not the regular lattice in figure 3.25, since the interaction between automata is asymmetric, but a configuration like that of figure 3.28 (right) in which the two groups are clearly separated, those with the higher r_s value being more farther from the centre of mass and at a larger distance one from the other (no confining force was applied for this system since the interaction potential is attractive for large distances). The stable configuration was reached in a quite large time, since the system remained for a long period in metastable states as that in the left panel of figure 3.28.

Similar configurations can be obtained also for completely heterogeneous automata, in which social radii are Gaussian distributed (figure 3.29).

3.8 Conclusions

We have introduced a simple model in which the interaction between physical particles was modified in order to introduce non Newtonian effects that could simulate perception. We have performed an analytical and numerical study of the one and two body problem, studying in particular the effect of dissipation, which leads to a stopping condition (in one dimensional models) or to particular stable orbits that can be analysed or even predicted with the help of analytical methods.

We then performed a numerical study of the N body system, after having in-

3.8 Conclusions

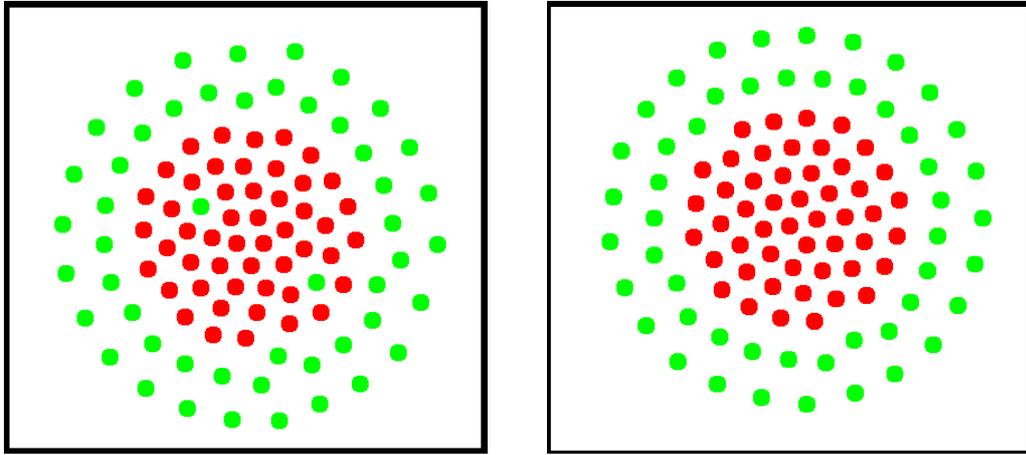


Figure 3.28: Metastable (left) and stable (right) $T^{\text{eq}} \approx 0$ configurations for automata split in two groups with different “social radii” r_s . Green automata have $r_s = 2$, red ones $r_s = 1.5$.

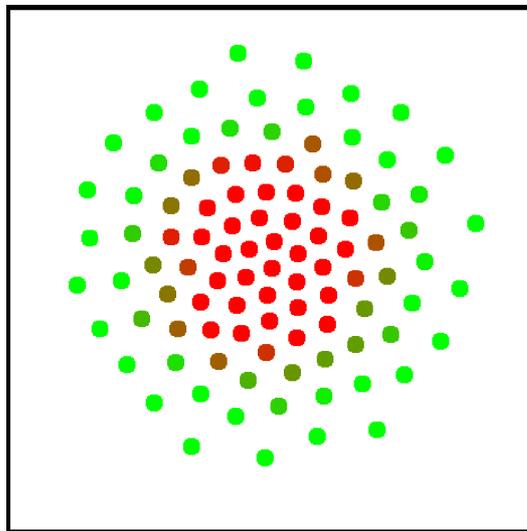


Figure 3.29: Stable $T^{\text{eq}} \approx 0$ configuration for heterogeneous automata whose social radii are given by a Gaussian distribution (pure green correspond to the highest values, pure red to the lowest ones).

roduced a memory effect which contrasts dissipation, studying in particular the kinetic energy (temperature) of the equilibrium system and the eventual formation of ordered structures at zero temperature.

We finally performed a preliminary study of a more complex model, in which heterogeneity and attractive social forces have been introduced.

3.8 Conclusions

Bibliography

- [1] C. W. Reynolds, “*Flocks, Herds, and Schools: A Distributed Behavioral Model*”, Computer Graphics, 21(4) (SIGGRAPH '87 Conference Proceedings) pages 25-34, (1987).
- [2] C. Benedetti, S. Rambaldi and G. Turchetti, “*Relaxation to Boltzmann equilibrium of 2D Coulomb oscillators*”, Physica A 364 197-212, (2006).
- [3] R. W. Hasse, “*Theoretical Confirmation of Coulomb Order of Ions in a Storage Ring*”, Phys. Rev. Lett., 83, 3440 (1999).

BIBLIOGRAPHY

Chapter 4

An agent model inspired by clonal expansion

4.1 Introduction

In this work we try to create a bridge between two different ways to study clonal expansion in the immune system (IS). One kind of approach consists in studying concentrations of different species of cells, whose behaviour and interaction is modelled through a system of differential equation (DE), the other one in studying microscopic interactions between single cells, that are usually modelled as cellular automata (CA).

It is our opinion that in immunology as in other fields of research the languages of microscopic (CA or agent) and macroscopic (DE) models could be integrated, in order both to use the analytical results to explain and partially predict the behaviour of the simulated models, and to utilise simulations to enrich with microscopic details the assumptions of macroscopic models.

In this paper we present a simplified model of clonal expansion, in which we stress our attention on spatial interaction between T cells and antigen presenting cells (APC), while omitting the details of the T cell-antigen and APC-antigen interactions, and ignoring many other important agents of the IS. The aim of this paper is thus not to present a new model for clonal expansion, but to start a project of work in which two different ways to model it could be combined.

4.2 The immune system

Hofmeyer [1] provides a short introduction to the immune system for researchers without background in immunology, but that are interested in studying the immune system as a distributed autonomous system (i.e. as the kind of system that is studied in this thesis). In our brief introduction to the main features of this system we will mainly follow his description, which we have completed with the help of [2], which is a short introduction to the immune system for undergraduate students in immunology.

4.2.1 Structure of the immune system

The immune system is one of the most deeply studied biological distributed autonomous systems, due to its cardinal biological and medical importance, and a great deal of information is available concerning its operation, from sub-molecular events to cellular population evolution [3].

We could say that the purpose of the immune system is to protect the body from threats posed by *pathogens*, a term by which we mean a large class of microorganisms (parasites, viruses, bacteria, fungi etc.) that can cause some kind of harm (diseases) to the body. This is a very complex task, which is mainly composed of two linked but different problems: *detection*, i.e. to recognise the danger (and, equally important, recognise what is not a danger), and *elimination*, i.e. the removal of this danger.

The immune system is constituted of different defence layers of growing complexity. The first and most elementary is the skin, while the second barrier is physiological, since the body environment (pH, temperature) provides inappropriate living conditions for many foreign organisms. Once pathogens fit to survive in the inner body environment enter the organism, they have to face the *innate* and *adaptive* immune systems, that are in charge of detection and elimination tasks.

4.2.2 The innate immune system

The innate immune system is called in this way both because it is present in all animals and because it remains the same through time and does not change or adapt to specific pathogens. It consists primarily of a chemical response system, the *complement system*, and of the *phagocytic system*, which involves cells as *macrophages*, whose task is to detect and engulf extracellular

molecules and materials, cleaning the system of both debris and pathogens. The molecules of the complement system can eliminate bacteria in two ways, directly through *lysis* (i.e. rupturing the membrane of the bacterium) or indirectly through *opsonisation* (i.e. binding to the bacterium surface and signalling it as a harmful pathogen to macrophages). *Self cells* (the cells that are part of the organism, and should not be attacked by the immune system) have proteins on their surface that prevent complement from binding.

As said before, macrophages are able to engulf and thus eliminate extracellular materials, have receptors for some bacteria and in particular for complement, and thus recognise as harmful and eliminate opsonised bacteria. Once activated, they secrete *cytokines*, whose effect is to create an inflammatory response and to increase the body temperature, in order to recruit a large number of immune system cells and to provide an appropriate environment for the “battle”.

4.2.3 The adaptive immune system

The adaptive immune system is able to learn to recognise specific kinds of pathogens, and also to retain a memory of them for future response. Learning occurs during the so called *primary response*, i.e. during a slow response to a kind of pathogen not encountered before. After that the adaptive system has learnt to recognise the pathogen, the *secondary response* is so quick and efficient that often there are no medical indicators of reinfection.

The most important components of the adaptive system are *lymphocytes*, or *white blood cells*, that have a role in both pathogen detection and elimination.

4.2.4 Recognition of a specific pathogen

Each lymphocyte has on the order of 10^5 identical receptors on its surface. These receptors have a complex three dimensional structure which can bind more or less easily to *epitopes*, which are locations on the surface of a pathogen or protein fragment, a *peptide* (the strength of this bond is called the affinity between the receptor and the epitope). If the number of bound receptors exceeds a threshold, it means that the lymphocyte is in an environment in which a large number of pathogens to which it has an high affinity is present, and thus the lymphocyte is *activated*. This activation

4.2 The immune system

threshold could be lower in a certain class of *memory cells* that are responsible of secondary responses. The structure of all the receptors is the same on a given cell, but these receptors can recognise a small set of similar structures, and thus there is no need to have a different lymphocyte for every possible epitope pattern. Nevertheless one of the most serious problems that the immune system has to face is to provide a repertoire of lymphocytes diverse enough to recognise any pathogen, since according to estimates the human body manufactures a number of proteins ten orders of magnitude lower than that of the foreign patterns to be recognised. The problem is partially solved through a pseudo-random DNA recombination, and partially through *dynamic protection*, i.e. a continual turnover of lymphocytes.

The effectiveness of the adaptive immune system is based not only on the large number of receptors, but also on its ability to adapt to specific kinds of epitopes, and to remember these adaptations for speeding up future responses. In these tasks a core role is performed by a class of lymphocytes called *B cells*. When a B cell is activated, it migrates to a *lymph node* (one of a few hundreds glands distributed throughout the body) where it produces many short lived clones through cell division. In this process of cloning, B cells undergo a form of mutation called *somatic hyper-mutation*, i.e. the mutation rate is nine orders of magnitude higher than common, in order to have an high chance to obtain different receptor structures. In the lymph nodes, new B cells have the opportunity to bind to pathogenic epitopes which are collected from the site of infection and presented to B cells on the surface of cells called *follicular dendritic cells*. If the receptors of these new B cells do not bind to pathogens, the cells die after a short time, while if they do they leave the node and differentiate into *plasma* or *memory* B cells.

Plasma B cells secrete a soluble form of their receptors, called *antibodies*, which have a double role: they opsonise pathogens (i.e. they signal them to macrophages, as the complement does, but in this case the response is highly specific to a given kind of pathogen), and they neutralise them preventing them from binding to self cells. Since the pathogens cause the production of antibodies by plasma B cells, they are often called *antigens* (antibody-generating, a term that actually refers to anything that causes antibody production in the organism). Memory B cells are those responsible of the secondary response: since they have an higher than average affinity to the pathogen epitopes, they provide a quick response to a second infection caused by the same pathogen, or even by a pathogen whose structure is similar to

that of the one that caused the primary response (this phenomenon is at the base of *immunisation*, a process in which a subject is exposed to a benign form of a pathogen in order to be later able to develop an immune response also to the virulent form). Actual understanding of immune memory is limited by the fact that B cells live only a few days, and thus also memory cells should be quickly removed from the organism. A possible explanation assumes that activated B cells are long lived, and survive for up to the lifetime of the organism, while another one assumes that traces of the non-self proteins that caused the infection are retained in the organism for many years, causing the permanence of the corresponding memory cells.

4.2.5 Tolerance of self

The other big issue the immune system has to deal with is to avoid *autoimmunity*, i.e. to attack the cells of its own body (or to be *tolerant of the self*). *T-helper* or *Th cells* are responsible of *tolerance*. They are called “helper” since they help B cells in recognising the self from the non-self. The name T cells derives from an organ called *thymus*, in which they mature. In this organ they are exposed to most self epitopes and, through a process called *clonal deletion* or *negative selection*, they are eliminated if bind to self epitopes. Also B cells undergo a similar process in the *bone marrow* (from which the “B” in their name), but the hyper-mutation they are subjected to in the lymph nodes could produce also auto-reactive clones.

This *peripheral* (or *distributed*, as opposed to the central one in the bone marrow and thymus) *tolerisation* of B cells is performed in the following way. To be activated, B cells need co-stimulation by two signals: one (*signal I*), as already described, when an high enough number of receptors binds to the epitopes, and a second (*signal II*) is provided by Th cells. In a process known as *antigen processing*, B cells engulf pathogenic peptides and then expose them on their surface, using molecules of the *major histocompatibility complex* (MHC). The T-cell receptors binds to the MHC-peptide complex on the B cell surface. If binding is successful, the B cell is activated, while if it is unsuccessful the B cell is caused to die.

4.2 The immune system

4.2.6 Antigen presentation

B cells are not the only cells that are able to engulf endogenous peptides and then present them on their MHC molecules. This characteristic is shared with other cells in the immune system (macrophages and *dendritic cells*) which are called for this reason *antigen presenting cells* (APC).

The role of dendritic cells is to activate *virgin* T cells. Usually, when non activated, dendritic cells spend their time filtering great amounts of extracellular fluids, but when they are located in an infected area (i.e. when receptors on their surface recognise molecules characteristic of microbial invaders, or cytokines that are released during an infection) they become activated and migrate to the nearest lymph node. Here, showing on their MHC molecules the peptides they picked up in the battle zone, they can activate virgin T cells. In the lymph node dendritic and T cells are engaged in a so called “dance”. The surfaces of these cells are kept together by very weak and non specific binding forces, that allow them to be bound for very short times. Nevertheless, if the T cell receptors find on the APC surface matching antigen-MHC compounds, the bond gets stronger and T cells can be activated. After activation is complete, the dendritic and the T cell part; while the first goes to activate other cells, the second one starts its proliferation process, doubling its number every six hours. After a few days of proliferation (during which some of the members of the clone could interact newly with the APC and thus proliferate even more) T cells are released and free to perform their two main tasks. The first one, as we have already described, is to “help” or control B cells, while the second one is to release the right kind of cytokines in the battle (infected) sites, cytokines that, as we said before, will regulate the battle against the pathogen (notice, as we said before, that dendritic cells need cytokines in order to be activated, migrate to lymph nodes and activate T cells. Then T cells release cytokines that can direct the response and thus activate macrophages, dendritic cells, B cells, which later can act as APC cells and activate helper T cells. . . this complex network of co-stimulation is one of the main features of the immune systems that allows its efficiency and correct work).

4.2.7 Detection of intracellular pathogens

We have said that cells as dendritic cells, macrophages and B cells can expose endogenous peptides on their surface, on MHC molecules. These MHC

molecules, that are present only on immune system cells, are called *class II* MHC molecules. Nevertheless, a different kind of MHC molecule, called *class I* MHC, is present on every kind of cell, and these molecules are used to expose on the surface fragments of proteins (peptides) from within the cell. The class I MHC-peptide complex are recognised by the receptors of another class of lymphocytes, the *cytotoxic* or *killer T cells* (*Tk cells*). This name is due to the fact that if the Tk cell recognises the peptide of an harmful pathogen and thus is activated, it kills the infested cell. Tk cells are the most effective weapons of the immune system against *intracellular infections*.

Limits of our model

In this short introduction we have presented only a part of the complexity of the immune system, enumerating its main actors and giving a short sketch of some of their basics interactions. Nevertheless, even this short exposition shows clearly that in the following model we will describe only to little extent and a great approximation the components and interaction of the immune system, and thus that our model has not a predictive or descriptive power, but is just a preliminary work intended to study some of the spatial effects of the APC-T cell interaction, and, as already said, to create a bridge between macroscopic (based on differential equations) and microscopic (based on agents or cellular automata) models.

4.3 Description of the model

4.3.1 Differential equations model

One of the major open question in immunology is the problem of understanding clonal expansion, i.e. how T cells, that belong to a very large repertoire, are selected in response to a specific treat (the presence of an antigen) and proliferate to form a large clone, and how this proliferation is regulated.

De Boer and Perelson presented a model that justifies the maintenance of diversity in the periphery through the concept of competitive exclusion [4]. This competition between T cells (between the different clones and inside the same clone) arises as competition for the peptides presented on the surface of APC. In fact these peptides can be freely available on the surface of an APC, or be captured in the receptor of a T cell bound to an APC; in the

4.3 Description of the model

second occurrence they are no longer available to other T cells.

De Boer and Perelson imposed a quasi-steady-state condition for the number of complexes given the number of peptides, and obtained a system of differential equations for the different clone sizes, which corresponds to the well-known principle of competitive exclusion in biology (two different species cannot co-exist in equilibrium if they use just the same resource) and introduced also a capacity (equilibrium size for a single clone).

In their model the number of peptides is considered to be proportional to antigen concentration, which is assumed as fixed. This assumption is well justified in case of self antigens, while for pathogens they assumed this fixed concentration to be the equilibrium value of a prey equation for the antigens, in which T cells had the role of predators. Using these assumptions, immune memory is attained through the persistence of antigen at a controlled concentration. (See [5] and appendix 4.6 for a treatment of prey-predator equations, and [6] for an application to the immune system).

This is one of the many models that describe clonal expansion using a system of differential equations (see for example [7] for a short review) and has been further on studied and improved by the authors [8]. Our interest in the first version of the model is due to its simplicity and to the fact that its basic assumptions concern the microscopic spatial interactions between T cells and APC, averaged in the quasi-steady-state condition.

Since there are many experimental results concerning how these interactions happen [9, 10, 11, 12], we think that this model is well apt to a microscopic formulation, in which the different individual cells are represented as agents in a computer simulation .

These are the differential equations that describe our version of the De Boer-Perelson model

$$\dot{A}_i = aA_i - bA_i^2 - \sum_j c_{ij}A_iT_j \quad (4.1)$$

$$\dot{P}_i = dA_i - rP_i \quad (4.2)$$

$$\dot{N}_{APC} = 0 \quad (4.3)$$

$$\dot{T}_i^N = 2gT_i^A - hT_i^N - \sum_j k_{ij}T_i^NT_j - lf(\sum_j m_{ij}P_j)FT_i^N + oC_i + s \quad (4.4)$$

$$\dot{T}_i^A = qC_i - gT_i^A \quad (4.5)$$

$$\dot{C}_i^A = -qC_i - oC_i + lf(\sum_j m_{ij}P_j)FT_i^N \quad (4.6)$$

$$\dot{F} + \sum_i \dot{C}_i = 0 \quad (4.7)$$

Equation (4.1) tells us that the n_A species of antigens A_i follow a logistic prey equation in which the n_T T cell clones T_i have the role of predators. Equation (4.2) gives the average number of peptides of species i , P_i , (we assume for simplicity a one to one correspondence between peptides and antigens) presented on a site of an APC cell. This number grows with the number of antigens and follows a decay rule (peptides remain on the APC's surface for a finite average time). With equation (4.3) we fix the number of APC cells. Equations (4.4) and (4.5) concern the number of non-activated T_i^N and activated T_i^A T cells ($T_i \equiv T_i^N + T_i^A$). Non activated T cells are produced by duplication of activated ones with a rate g and die by apoptosis with rate h . The probability rate s represents an external source (thymus). F is the total number of free sites on the APC's surface, to which T cells can bind with a probability rate that depends on a function f of the probability to find a given species of peptides multiplied by its affinity m_{ij} to it (l is the probability of binding in case of maximum affinity). We call C_i a complex formed by a T cell T_i and a site of an APC. These complexes can unbind with probability rate q in case of successful activation (equation (4.5)) and with probability rate o in case of unsuccessful activation (equation 4.4). The terms k_{ij} in equation (4.4) rule the fratricide competition between the T cells (see for example [13]).

The number of complexes and free sites is governed by equations (4.6), (4.7) coherently with the assumptions of equations (4.4), (4.5) and with the request that their sum has to be fixed as the total number of sites ($n_s N_{APC}$ if n_s is the number of sites on a single cell).

4.3.2 Microscopic model

In the differential equations based model we tried to write explicitly an equation for each agent of the process, and we defined a probability rate for each interaction between these agents, since we want these equations to be the mean field version of a microscopic model. Given the high number of equations and parameters we won't try an analytical treatment and we will rely on numerical integration for their solution.

Our microscopic model is realised on two superposed 2D squared grids, one on which antigens move and one for APC and T cells. The physical region

4.3 Description of the model

corresponding to each layer will be the same (creating a correspondence between sites “located in the same physical space”) while the step of the grids and thus the number of sites could be different.

All the cells move by random walk obeying an exclusion principle (no more than a single cell on a given site of a layer), and the interaction between cells can happen by superposition when they are located on different layers, or by contact (i.e. if they are located on first neighbour sites) if they are on the same layer. We call these events that allow an interaction between the cells *encounters* (see also figure 4.1). An encounter between an antigen A_i and a T cell T_j leads to the elimination of the antigen with probability p^c_{ij} , while an encounter between an antigen and an APC leads with probability p^d to the presentation of a peptide on the “surface” (i.e. on one of the four sides) of the APC (in our convention we associate to the probabilistic rate x in the continuous macroscopic model the probability p^x in the discrete microscopic one). Encounters between a T cell T_i and an APC can form a complex, with a probability p^l multiplied by the affinity to the site $f(\sum_j m_{ij}\Pi_j)$ (a function of the averaged affinity to the peptides Π_i presented on the site, where Π_i is the number of peptides in species i present on the site). Encounters between the antigens lead to an over-population due “logistic” elimination of the antigen with probability p^b , while those between T cells in clones i and j lead to fratricide apoptosis with probability p^k_{ij} . These fratricide terms are in a certain sense *ad hoc* in our model (they are not present in the original formulation by De Boer and Perelson, even if they are present in other models, as in [13]), since we need them to avoid a filling of the grid. They should be chosen in such a way that they are not relevant under normal conditions (i.e. when the number of occupied sites is low with respect to the total number of sites). All the other processes are encounter independent and can happen with given probabilities at each time step (as for example T cell duplication or antigen reproduction).

We can say that our model is devoted to the description of the biological process described in 4.2.6, while all the processes regarding antigen removal are just expressed by the rule (local, i.e. based on encounters) *the more affine T cells are present, the more the antigen is removed*. It is quite clear that this model is too simple to describe all the complex processes that concern clonal expansion in the immune system. A more complete formulation should use at least two different 2D grids to describe the site of infection and the lymph-nodes (connected in some way to allow the displacement of T and

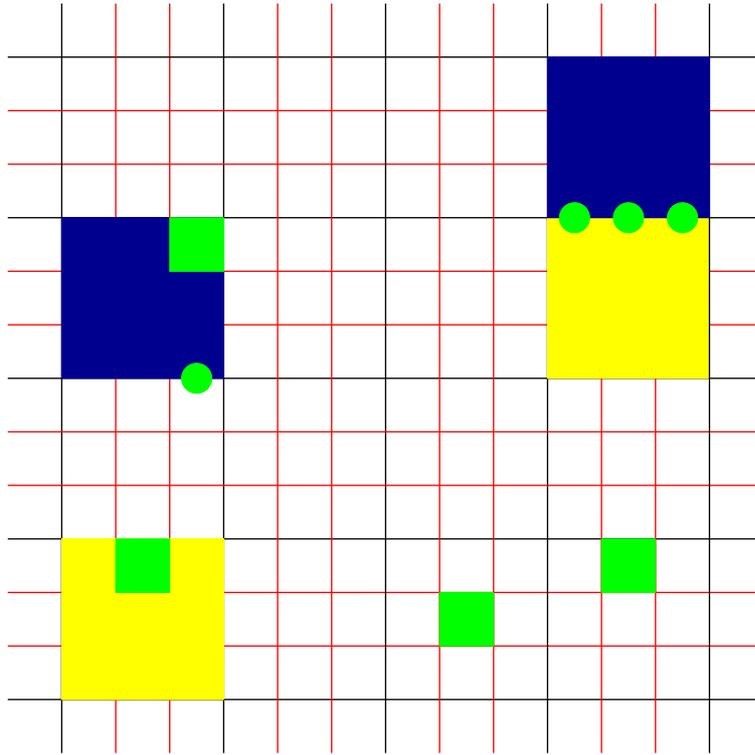


Figure 4.1: Possible interactions between different agents in the model. Antigens (small green squares) move on the red grid, while APC (blue squares) and T cells (yellow squares) on the black one. At bottom-left we see an encounter (overlapping) between a T cell and an antigen, that can lead to antigen removal. At top-left we see an encounter between an APC and an antigen, that can lead to the presentation of a peptide (green circle) on the surface of the APC. Finally, at top-right we see an encounter between an APC and a T cell, interacting on a “surface” that presents 3 peptides. This encounter can lead to the formation of a compound and to activation of the T cell, if the affinity of the T cell to the peptides is high enough.

4.3 Description of the model

dendritic cells), while for a realistic description of immunological memory a differentiation between naive and memory T cells is necessary.

4.3.3 Mean field equations

While in section 4.3.1 we have proposed a differential equations based model for clonal expansion, in section 4.3.2 we have proposed a microscopic agent model to describe the same dynamics. We now show how the equations in section 4.3.1 can be obtained as mean field equations for the model in section 4.3.2.

While all the probability rates in a macroscopic model have to be chosen on the base of macroscopic observations, in such a way that the behaviour of the solutions will correspond to the behaviour of the biological species under some given assumptions, the probabilities of the microscopic model, according to the spirit of this work, should be given on the base of microscopic observations, as reported for example in [9, 10, 11, 12]. The time step should be chosen smaller than the shorter characteristic time of the processes involved, and all these characteristic times should be expressed as probabilities. An average process would be necessary to describe 3 dimensional cells with a complex shape as 2D squared objects, and probably also minor changes on the geometry (allowing for example APC and T cells to have different size) could be necessary. Nevertheless, given the preliminary stage of this work and its general purposes, and considering also our limits in the interpretation of experimental data given our scientific formation, we just do very simple considerations that allow us to have some qualitative result, without any claim to quantitative or predictive results.

We can obtain the mean field equations for the microscopic model in the following way. Let us assume for example that the average time for antigen duplication is one day. If we choose a time step Δt of 15 minutes, the probability for antigen duplication is fixed to $p^a = 0.01$. Defining N_A as the number of sites of the antigen's grid and assuming random distribution for all the cells, the probability for an antigen to have an encounter with another antigen on one of its 4 sides is A/N_A , and thus the time evolution of the number of antigens in absence of T cells is given by

$$A(t + \Delta t) = A(t) + p^a A(t) - p^b A(t)^2 / N_A \quad (4.8)$$

The value of p^b can be fixed given the wanted maximum density of antigens (the capacity),

$$A_{max}/N_A = p^a/p^b \quad (4.9)$$

and in the continuous limit we obtain equation (4.1) through the identifications $a = p^a/\Delta t$, $b = p^b/(N_A\Delta t)$.

Proceeding in the same way the discrete version of equation (4.2) is, recalling that Π_i is the number of peptides of species i on a single side of an APC, while P_i is the average number of peptides on sides

$$\sum_{APC} \Pi_i(t + \Delta t) = \sum_{APC} \Pi_i(t) + p^d A_i(t) N_{APC}/N_T - p^r \sum_{APC} \Pi_i(t) \quad (4.10)$$

or, averaging over all the $4N_{APC}$ sides

$$P_i(t + \Delta t) = P_i(t) + p^d A_i(t)/(4N_T) - p^r P_i(t) \quad (4.11)$$

where N_T is the number of sites of the APC-T cell grid. The continuous version of (4.11) is equation (4.2), through the identification $d = p^d/(4\Delta t N_T)$, $r = p^r/\Delta t$. Equation (4.2) has solution

$$P_i(t) = e^{-rt} \left[\int A_i(t') e^{rt'} dt + \text{const} \right] \quad (4.12)$$

that reduces to

$$P_i(t) = \frac{A_i d}{r} + \left[P_i(0) - \frac{A_i d}{r} \right] e^{-rt} \quad (4.13)$$

in case of constant A_i concentration. We use $p^d = 1$ (the APC always recognises the antigen) and $p^r = 0.02$, corresponding to a permanence of the peptide on the antigen surface for an average time of 12 hours.

We compare in figure 4.2 the numerical integration of equation (4.1) (for a single species and in absence of T cells) with the corresponding results given by the microscopic model, and in figure 4.3 we present the same comparison for the analytical result of equation (4.13). (We have used for these simulations $N_T = 9 \cdot 10^4$, $N_A = 3.6 \cdot 10^5$ and $p^b = 0.05$ which corresponds, according to equation (4.9), to a capacity of an antigen every 5 sites).

While there is an almost perfect correspondence between the curves in figure 4.3, there is a slight difference between those in figure 4.2. This effect is due to the fact that while the behaviour described by equation (4.10) depends on the interaction between cells located on different layers, and thus is not actually based on microscopic spatially constrained interactions, the behaviour

4.3 Description of the model

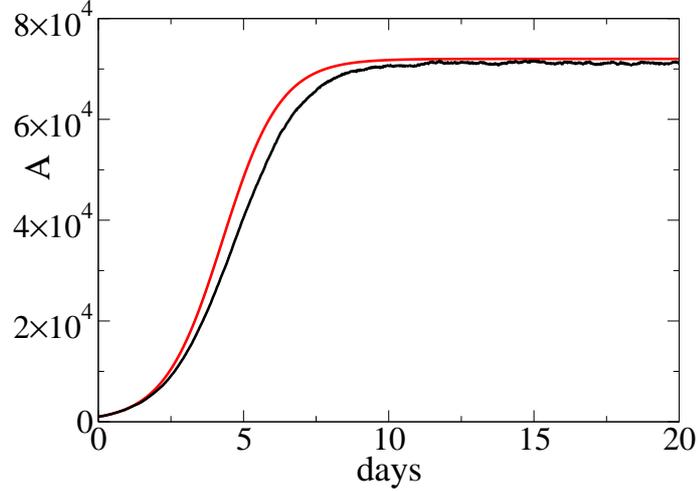


Figure 4.2: Comparison between the free growth of the antigen number $A(t)$ as obtained from the microscopic model (black line) and the mean field equations (red line). The time unit is one day, as in all the figures to follow.

described by equation (4.8) relies on and influences the spatial distribution of antigens. For this reason the mean field equation describes well the microscopic model in the initial configuration, when a uniform distribution is imposed, and at the equilibrium, while the discrepancy is stronger during the expansion.

The local effects are obviously stronger when we consider the spatial T-APC interaction. Let us fix $N_{APC} = 2 \cdot 10^3$ on the $N_T = 9 \cdot 10^4$ grid, use the sigmoid function

$$f\left(\sum_j m_{ij}\Pi_j\right) \equiv \frac{1 - e^{-\sum_j m_{ij}\Pi_j}}{1 + e^{-\sum_j m_{ij}\Pi_j}} \quad (4.14)$$

to obtain the affinity of a T cell to a site on the APC surface, $p^g = 0.05$ (an activated T cell needs 5 hours to split referring to the time step of 15 minutes), $p^h = 0.001$ (a life span of 10 days for the T cells), $p^l = 0.25$ (an hour to form a complex in case of maximum affinity), $p^a = 0.2$, $p^o = 0.04$. (These are the probabilities to unbind with and without activation in case of maximum affinity. The dependence of these microscopic probabilities on the affinity has been chosen *ad hoc* is such a way that the first one grows and the second one decreases with affinity).

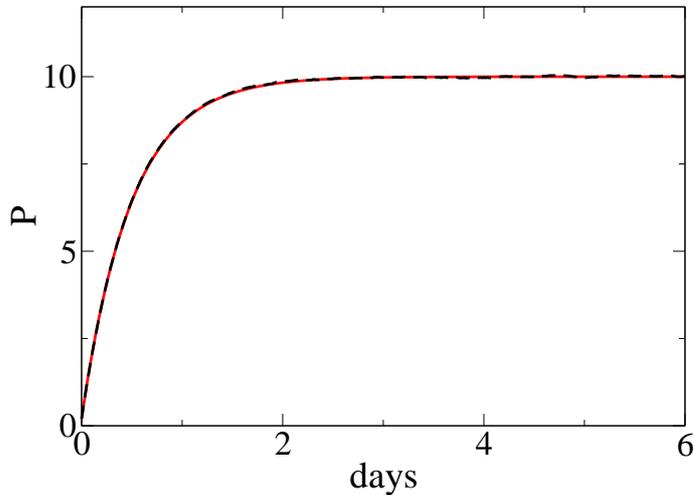


Figure 4.3: Average number of peptides in presence of a fixed number of antigens, as obtained by the microscopic model (black line) and the mean field equations (continuous red line).

We can now consider a single clone T with maximum affinity to a single species of antigen A ($m \equiv m_{11} = 1$), fix A to its maximum capacity ($c \equiv c_{11} = 0$, i.e. antigens are not removed), and obtain in the usual way the discrete mean field equations for T , F and C whose continuous limit leads to equations (4.4)-(4.6), redefining the parameters on the base of the microscopic probabilities.

Figure 4.4 refers to the growth of the clone, and compares the integration of the mean field equation with the results given by the microscopic model (the fratricide term value is fixed to $p^k = 0.1$). In this case the discrepancy is stronger, and it is also a qualitative one. The growth in the microscopic model is lower at the beginning, while the equilibrium value is higher. Two different effects are present, both due to the presence of zones around the APC in which T cells reproduce: the fratricide effect is enforced because of the higher density in these zones, but also the probability to meet an APC and thus to be activated is enhanced. Since these effects depend strongly on the density of cells, is possible to obtain the parameters of equations (4.1)-(4.7) by a process of best-fitting only on regions in which the values of A and T are almost constant (this means that those equations are able to de-

4.4 Results of the simulations

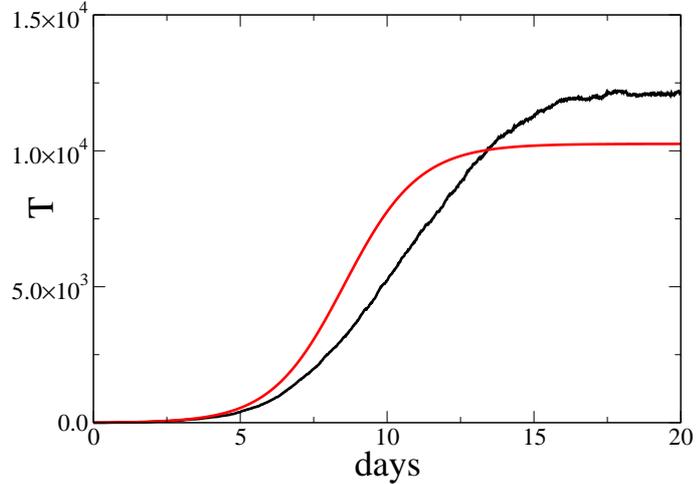


Figure 4.4: T clone expansion in response to a fixed number of antigens in the microscopic (black line) and mean field models (red line).

scribe properly the behaviour of the microscopic system only if we introduce a dependence of the parameters on A and T).

4.4 Results of the simulations

4.4.1 Acute antigenic impulse

In order to complete the model we fix the value of the parameters c_{ij} as $c_{ij} \equiv c m_{ij}$ (we are assuming that the ability of a T cell in removing an antigen is proportional to its affinity to it). We have used $c = 0.2$ in order to obtain a realistic time scale for the response of the immune system to the infection.

In figure 4.5 we plot the evolution of the clone size T and antigen A populations, comparing the results of the microscopic model with the solutions of the mean field equations. In agreement with the previous discussion the results are very similar at equilibrium values, while the agreement is only qualitative during the transient part. Damped oscillations are present in both models, and both the period and the height of peaks and valleys are of the same order of magnitude (the damping rate and the period of oscillations are higher in the microscopic model).

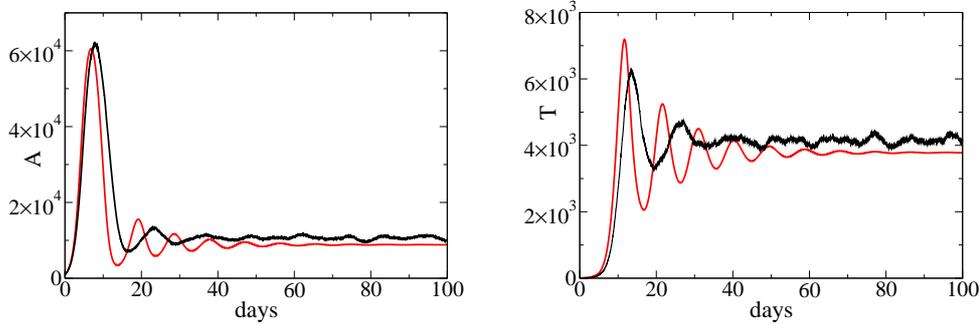


Figure 4.5: Evolution of the system under an acute antigenic stimulus. The evolution of the antigen number in the microscopic and mean field model is shown at left, while the size of the T cell clone is reported on the figure at right (black lines correspond to the microscopic model, red ones to mean field equations).

This behaviour corresponds to that of a prey-predator system (see appendix 4.6 and [5]). To an equilibrium value with $A \neq 0$, $B \neq 0$ corresponds a “memory” effect due to the permanence of the antigen. In this situation the response to a secondary stimulus is obviously quicker (figure 4.6).

4.4.2 The clonal repertoire model

We finally consider the effects of both fratricide and spatial competition terms between different clones in presence of a differentiated antigen repertoire. By using a fratricide term in which the decrease is proportional to the overall size of the clones, $\Delta_- T_i = -k T_i \sum_j T_j$, we obtain a mutual exclusion principle. In fact, if we summarise with $\Delta_+ T_i$ the growth terms, the relative variation of the clone size is

$$\frac{\dot{T}_i}{T_i} = \frac{\Delta_+ T_i}{T_i} - \frac{\Delta_- T_i}{T_i} \quad (4.15)$$

Since $\Delta_- T_i/T_i$ is the same for all the clones, supposing that there is a unique antigen with the highest affinity to the clone T_j , i.e.

$$\Delta_+ T_j/T_j > \Delta_+ T_i/T_i \quad \forall i \neq j \quad (4.16)$$

if the clone j reaches an equilibrium

$$\Delta_+ T_j = \Delta_- T_j \quad (4.17)$$

4.5 Conclusions

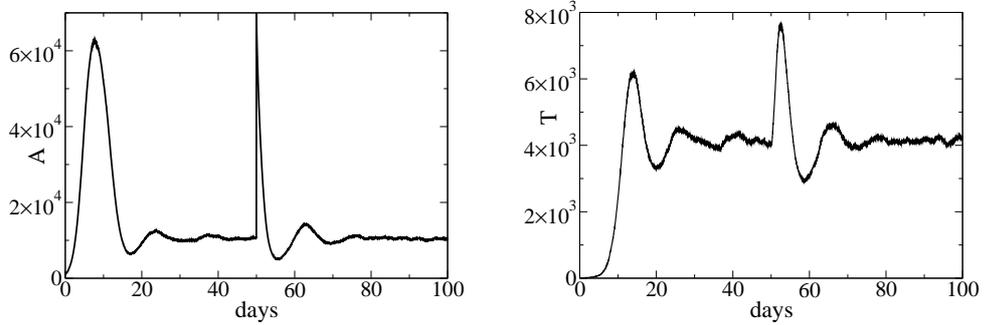


Figure 4.6: Left: evolution of the antigen population after a secondary impulse occurring 50 days after the primary, microscopic model. Right: corresponding evolution of the T cells clone.

then any other clone extinguishes since

$$\Delta_+ T_i - \Delta_- T_i < 0 \quad (4.18)$$

(These are the basics of competitive exclusion, see [5]).

To show that in our model there is competition for peptides presented at the APC surface (the mechanism investigated in [4]), we can use a “pure fratricide” term $-kT_i^2$. This is actually a “non-competitive” one since it favours the small clones. In fact, studying the expansion of 3 clones under the stimulus of a single antigen, using an affinity matrix $m_{i,1}$ such that $m_{i,1} \ll m_{1,1} = 1$ if $i \neq 1$, we have (figure 4.7) an equilibrium with $T_2 \neq 0$.

Nevertheless even in this situation the competition for the peptides on the APC surface leads to a control in the overall number of T cells, at least when the number of clones is large. To study this effect we introduce an antigen with constant concentration, to which 10 clones have maximal affinity. Once these clones have reached their equilibrium size, we introduce three different additional antigens at which three new clones are highly affine. The results of figure 4.8 show that the size of the “old” clones shrinks as a reaction to the growth of the new ones (this effect can be due only to spatial competition for resources since no fratricide competition between different clones is present).

4.5 Conclusions

Analytical models and simulations are usually treated as completely distinct fields of research, even when they face the same problem. In this work we

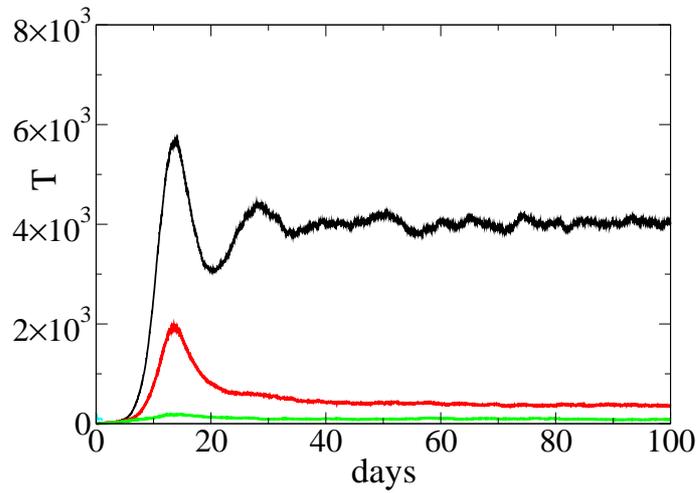


Figure 4.7: Time evolution of the size of three clones one of which (black line) has higher affinity to a given antigen.

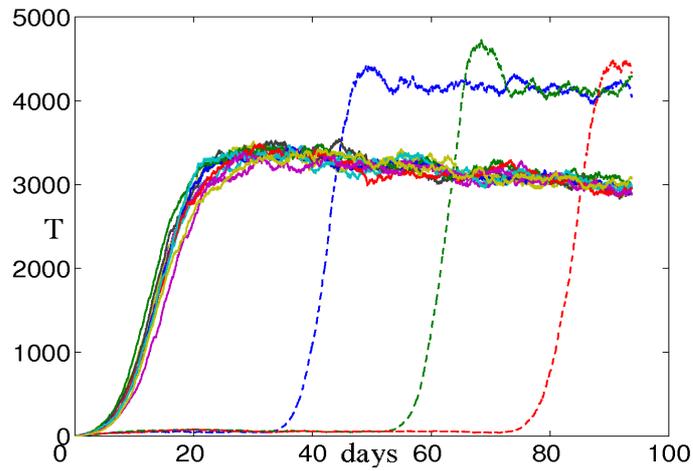


Figure 4.8: Time evolution of the size of 10 clones (continuous lines) stimulated by a single antigen and shrinkage due to the expansion of three new clones (not continuous lines).

4.6 Appendix

have presented a microscopic dynamical model inspired by clonal expansion in the immune system, together with a system of differential equations that could be interpreted as its mean field theory. We have shown how the mean field equations can be used to interpret the results of simulations, while the microscopic model can be used to add a local and spatial character to a macroscopic system based on differential equations.

We do not claim that the results of our model are biologically relevant, but we present it as a starting point for a more complex model and as a solution for a compromise between pure analytical and pure simulated models that could be used in different fields of research.

4.6 Appendix

The dynamics of the model can be described by a simplified system of differential equations for A and T . We assume that the antigen-APC-T average interaction consists of a growth term for the T clone proportional to A . The equations become

$$\dot{A} = A(a(1 - cA) - bT) \quad \dot{T} = T(-d + eA - fT) \quad (4.19)$$

These Lotka-Volterra equations with a logistic term have been extensively investigated (see [5]) and if $e > cd$ they exhibit a critical stable point

$$T_c = \frac{a(e - cd)}{eb + caf} \quad A_c = \frac{af + db}{eb + caf} \quad (4.20)$$

Every solution in the positive sector $T > 0$ $A > 0$ is attracted by this point which is topologically a focus. Convergence rate to equilibrium and the oscillations period are determined by the eigenvalues of the Jacobian matrix. From its trace and determinant

$$\text{Tr } J = -a \frac{acf + bcd + ef - cdf}{eb + caf} < 0 \quad \det J = \frac{a(e - cd)(bd + af)}{eb + caf} > 0 \quad (4.21)$$

we obtain the eigenvalues $\lambda_{\pm} = \frac{1}{2}[\text{Tr } J \pm \sqrt{\text{Tr } J^2 - 4\det J}]$ which are real negatives or complex with negative real part. We have oscillations if $\Delta = \text{tr } J^2 - 4\det J = -\omega^2 < 0$ and their period is $2\pi/\omega$.

Bibliography

- [1] S. A. Hofmeier, “*An overview of Immunology*”, in A. Segel and I. R. Cohen (eds.), “*Design Principles for the Immune System and Other Distributed Autonomous Systems*”, Oxford University Press (2001).
- [2] L. Sompayrac, “*How the Immune System Works*”, Blackwell Publishing (2003).
- [3] A. Segel and I. R. Cohen (eds.), “*Design Principles for the Immune System and Other Distributed Autonomous Systems*”, Oxford University Press (2001).
- [4] R. De Boer and A. S. Perelson, “*T cells repertoires and competitive exclusion*”, *J. Theor. Biol.* **169**, 375-390 (1994).
- [5] J. Hofbauer and K. Sigmund, “*Evolutionary games and population dynamics*”, Cambridge University Press (1998).
- [6] M. Novak, R. May and K. Sigmund, “*Immune responses against multiple epitopes*”, *J. Theor. Biol.* **175**, 325-350 (1994).
- [7] R. Antia, V. Ganusov and R. Ahmed, “*The role of models in understanding CD8⁺ T-cell memory*”, *Nature Reviews Immunology*, (2005).
- [8] R. De Boer and A. S. Perelson, “*Competitive control of the self renewing T cell repertoire*”, *International Immunology*, Vol. 9, No. 5, pp. 779, Oxford University Press (1997).
- [9] A. Lanzavecchia and F. Sallustio, “*Lead and follow: the dance of the dendritic cell and T cell*”, *Nature Immunology* Vol. 5 No.12, 1201-1202 (2004).

BIBLIOGRAPHY

- [10] S. Hugues, L. Fetler, L. Bonifaz, J. Helft, F. Amblard and S. Amigorena, “*Distinct T cell dynamics in lymph nodes during the induction of tolerance and immunity*”, *Nature Immunology* Vol. 5 No.12, 1235-1242 (2004).
- [11] R. Lindquist, G. Shakhar, D. Dudziak, H. Wardemann, T. Eisenreich, M. Dustin and M. Nussenzweig, “*Visualizing dendritic cell networks in vivo*”, *Nature Immunology* Vol. 5 No.12, 1243-1247 (2004).
- [12] R. Germain and M. Jenkins, “*In vivo antigen presentation*”, *Current opinion in Immunology* **16** 120-125 (2004).
- [13] R. Callard, J. Stark and A. Yates, “*Fratricide: a mechanism for T memory-cell homeostasis*”, *TRENDS in Immunology* Vol.24 No.7, 370-375 (2003).

Chapter 5

An evolutionary crowd dynamics and ToM model

5.1 Introduction

We develop an evolutionary agent model to study the behaviour of pedestrians in a crowd. Pedestrians are known, on the base of direct observations, to present forms of self organisation on a local base in order to optimise traffic flow. For example, if two different flows are directed in opposite directions inside a single corridor, pedestrians tend to self organise in lanes when the density is high enough (the number and dimension of lanes depend on the density of pedestrians and dimension of the corridor), while if two flows are located in corridors crossing at a given angle, self organisation emerges through the formation of stripes roughly parallel to the bisector of the angle between the two corridors.

Different computer simulation models that reproduce this kind of behaviour have been developed. In the model that we present we try to obtain the same kind of behaviour as the output of an evolutionary process, in which the self organisation behaviour emerges as a natural way to overcome a traffic flow problem due to the physical collisions between the agents. In our model agents are realised as 2 dimensional discs moving in corridors delimited by walls, and all the collisions (agent to agent and agent to wall) are exactly resolved using an event based algorithm. The agents are evolved using a simple fitness function that gives a positive value to the velocity in the direction of their goal, and a negative one to the physical momentum exchanged in

5.2 Crowd dynamics

collisions, and we expect evolution to develop the lanes and stripes formation behaviour as a way to maximise the fitness function.

The ability of the agents in avoiding collisions is based on their prediction of future positions of the others. In order to predict the movement of the others, agents need a model of their behaviour. We thus use our model also as a way to study the evolution of the Theory of Mind (ToM), meant as the ability of individuals to understand that also the others have intentions and beliefs, and in particular, on the base of a previous model, we define and study the evolution of the “level of Theory of Mind” (ToM level or recursion level) of an agent. Level 0 agents are those that ignore the behaviour of the others while level 1 agents are those that take in consideration the behaviour of other agents, assuming that they behave as level 0. Level 2 agents assume the others to be level 1 and are thus capable of “recursive thinking”, and so on.

5.2 Crowd dynamics

Pedestrian crowds have been studied on an empirical base (using direct observations, photographs and films) and with the aid of simulation models during the last four decades (a detailed description of experimental observations and simulation results can be found in [1, 2], to which we will refer during all this section). Human behaviour is supposed to be “chaotic” and very irregular but, while this can be true for extremely complex circumstances, in standard situations individuals do not usually choose between many possible alternatives, but act more or less in a automatic way relying on a previously optimised (learnt by trial and error) behavioural strategy. This seems to be the case of pedestrian behaviour. According to observations, pedestrians seem to like to walk on straight lines as long as possible, and thus their trajectory usually assumes a polygonal shape. If alternative routes are of the same length, a pedestrian prefers the one where she can go straight ahead for as long as possible and change direction as late as possible. If it is not necessary to hurry up to reach the destination in a given time, pedestrians prefer to walk at a certain individual desired speed, which minimises energy consume. The desired speeds are Gaussian distributed with a mean value of about 1.34 m/s and a standard deviation of 0.26 m/s . Pedestrians keep a certain distance from other pedestrians and obstacles, a distance that gets smaller as the pedestrian hurries and as the density gets higher. At given

densities and situations, the movement of pedestrians shows similarities with that of gas and fluids, or with the dynamics of granular fluids.

In our work we were particularly interested in phenomena of local self organisation that emerge when fluxes of pedestrians with different directions cross. It has been observed that, if the density is high enough, pedestrians spontaneously organise themselves in lanes of uniform walking direction, when two parallel but opposed fluxes are located in the same “corridor”. This behaviour can be found under the form of a (spontaneous) traffic flow rule (see the discussion in chapter 6, in particular 6.5) or as self organisation at a local level (as described in chapter 1.4). We can talk about a traffic flow rule when all the pedestrians (usually inside a given geographical area) use to choose *a priori* (or just prefer) to walk in a certain side of the “corridor” in order to minimise traffic flow problems. Studies show that this tendency actually exists, and that sort of spontaneous traffic flow rules are present, usually at the level of a national country area (these rules may or may not be related to the corresponding traffic flow rules for vehicles). Nevertheless, usually self organisation emerges on a local base and not as *a priori* rule, and the number and dimension of lanes depend on the dimension and density of fluxes. According to Helbing et al. [1], *“The mechanism of lane formation can be understood as follows : pedestrians moving against the stream or in areas of mixed directions of motion will have frequent and strong interactions. In each interaction, the encountering pedestrians move a little aside in order to pass each other. This sideways movement tends to separate oppositely moving pedestrians. Moreover, pedestrians moving in uniform lanes will have very rare and weak interactions. Hence the tendency to break up existing lanes is negligible, when the fluctuations are small. Furthermore, the most stable configuration corresponds to a state with a minimal interaction rate and is related to a maximum efficiency of motion”*.

The phenomenon of lane formation that occurs when two parallel fluxes cross can be considered as a particular case of the “stripe formation” self organisation phenomenon that as been observed ([3]) when two fluxes of pedestrians cross at an angle α . In this situation “stripes” of pedestrians moving in the same direction are observed to form in the crossing zone, the direction of the stripes being parallel to the bisector of the angle between the fluxes (see for example the behaviour of evolved agents in our model, figure 5.13).

In order to describe these local self organising phenomena it is probably necessary to use a individual based micro-simulation of crowd dynamics (i.e.

5.3 Theory of Mind

in the spirit of this thesis). Helbing ([1, 2]) has introduced a “social force model” in which the various terms of the pedestrian to pedestrian interaction were modelled as position and velocity dependent physical forces, and that can reproduce the phenomena we have just described (and other self organisation phenomena, as the behaviour at bottlenecks).

5.3 Theory of Mind

Following the paper “Does the Chimpanzee have a Theory of Mind?”, by Premack and Woodruff [4], in recent years the term *Theory of Mind* (ToM) has been used in behavioural sciences to identify the ability to understand that also the others have minds, with different beliefs, desires, mental states, and intentions, and to develop mental models about these beliefs and intentions.

This ability is surely present in human beings, and has a fundamental role in social interactions. For example, focusing on the crowd dynamics problem, while moving in a crowd trying to reach our goal, we have to realise that also the other persons have goals and are trying to reach them, and take in account and try to predict their behaviour in order to avoid collisions with them.

We can also think that the other pedestrians too are trying to predict the behaviour of people in the crowd (including ourselves), and take in consideration the influence of those predictions on their behaviour. And finally, to forecast accurately their motion, we could take in consideration the ability of the others “to read our mind” and thus to predict all the considerations that we are doing about the others’ behaviour, *et cetera*. This “recursive thinking structure” (I think that she thinks that I think. . .) is surely possible for human beings, and plays a role, if not in crowd dynamics (after all even social insects are capable of efficient traffic flow organisation, see the discussion in chapter 6), in complex social situations.

Having a Theory of Mind of the other individuals and using it to predict their behaviour means, following Dennet [5], to assume “the intentional stance” while interacting with them. According to Dennet when trying to explain and predict the behaviour of an “object” we can assume three different levels of abstraction: the *physical stance*, the *design stance* and the *intentional stance*.

In the physical stance, we analyse the properties of an object on the base

of its physical and chemical structure. This approach is the only possible when studying the basic elements of the physical world (particles, etc.) and even when analysing more complex systems which are composed of many “fundamental” components that are not part of the biological world, nor artifacts produced by the human activity. This level is the more “concrete” one, since it is related to well defined physical quantities, and allows quantitative predictions. Nevertheless, it is very expensive from a computational point of view, since it requires and processes a lot of information. In fact, just sticking to the physical stance, the analysis of a system composed of many fundamental components cannot be based simply on the laws that rule these components, but requires approximations and the study of average macroscopic components (thermodynamics, statistical mechanics).

The design stance is very useful when dealing with objects in the realm of biology and engineering. It uses concepts like purpose, function and design to predict the behaviour of a given object, assuming that it corresponds to some kind of plan the object as been built (by its designer or by evolution) for. This approach requires less information and has an higher computational power. For example, when dealing with a thermometer, we do not need to study the physical behaviour of all its components to infer the environment’s temperature, we just need to assume that someone else (the designer) has done all the work for us and built an appropriate scale to read the temperature, and read it. Obviously the design stance does not give us all the information about the object, for example it does not say anything about the behaviour of the object in situations it has not been designed for (we rely on the physical stance to know how much time it will take for the thermometer to fall on the floor when left from a given height, to know if it will break on the impact and even to study its behaviour outside the range of temperature it has been designed for, for example to predict at which temperature it melts).

The more abstract level is the intentional stance, that uses concepts like beliefs, thinking and intent, and is useful to deal with “minds”. When we are dealing with the behaviour of other human beings, but also with animals and even with complex computer software, for example when we are playing chess against a computer, we use the intentional stance. While using the intentional stance to predict the next chess move by the software, we are not necessarily assuming that the software really has beliefs or intentions, we are just using a powerful approach to predict its behaviour. Once again, the

5.4 A previous model on ToM evolution

intentional stance does not provide us all the information regarding the “object”. For example, to study the physiology of the human body, the design or physical stance, depending on the required detail or abstraction level, are obviously more appropriate.

Large part of the experimental studies in the field of Theory of Mind have been devoted to investigate if there are signs of the presence of ToM in animals (especially in non human primates) and if ToM deficits are related to autism. According to recent works some animals, as dogs [6], dolphins [7], goats [8] and crows [9], are able to follow the gaze of other individuals. Other researches seem to show the presence of a limited ToM in primates [10], but no indisputable signs have been found [11], and probably true recursive thinking distinguishes the human mind from the mind of other animals.

5.4 A previous model on ToM evolution

The idea at the base of the crowd dynamics model that we have developed was present in a model by Takano, Katō and Arita [12, 13]. These papers, assuming the perspective of evolutionary psychology, that views the human mind as a product of evolution [14], explore the dynamics and adaptivity of the mechanism of recursion in ToM, using the computer modelling methods of artificial life. A known evolutionary hypothesis regarding ToM is the social intelligence hypothesis [15] which states that intelligence has evolved not to solve physical problems, but to solve complex social problems. In their model social interactions are represented by collision avoiding between agents moving in the same environment. Each agent tries to predict the moves of the others on the base of its ToM, and uses these predictions to avoid collisions. The model does not assert to simulate realistic collision avoiding, but to be a thought experiment used to investigate under which conditions recursive thinking emerges.

Agents are represented as discs of radius r moving in a 2 dimensional continuous space. Each agent is provided with a goal and feels an attraction force towards it. Its fitness function is given by the ratio d/t of the distance from the starting point to the goal over the time it takes to reach the goal. Discs (the agents’ bodies) can overlap, and in that situation they are in collision mode; when in collision mode their velocity is diminished by a factor 400 and thus collisions severely affect the agent’s fitness.

In order to avoid collisions agents predict the future positions at the next

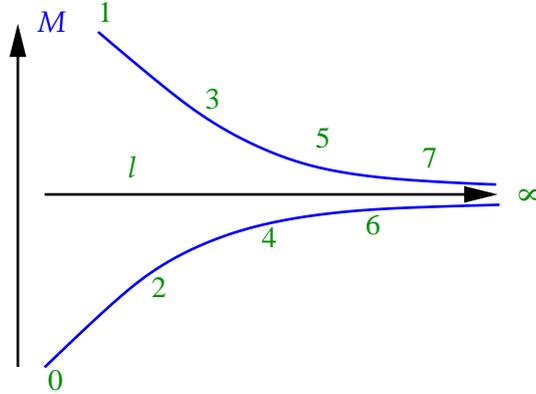


Figure 5.1: Convergence of collision avoiding magnitude M as a function of ToM level l (the figure has simply a qualitative and descriptive value).

time step Δt of all the agents that fall in their sensory system which is a fan shaped visual cone of radius r_v and angle a_v . These predictions are based on the ToM or recursion level l of the agent, which is defined as follows: level 0 agents are those that do not take in account the others' action while level n agents are those that trying to predict the others' behaviour assume them to be level $n - 1$. Following this definition a level 1 agent will try to predict the others' moves assuming that they are level 0, and thus that they do not take in consideration the actions of their neighbours, while a level 2 agent assumes also the others to take in consideration the neighbours' moves, but to have a simple model of their behaviour, since it considers them to be level 1, and so on.

The key finding of the papers was that the magnitude of collision avoidance is higher for odd levels than for even ones, the difference being greater the lower the level is. In particular, it is minimum for level 0 agents and reaches its maximum for level 1 agents, while it converges to a given value when $l \rightarrow \infty$ from above for odd levels and from below for even ones (figure 5.1). Stronger collision avoiding does not mean by itself higher fitness, since an agent that tries too strongly to avoid collision could need a very long time to reach the goal. The authors performed a throughout study of the r_v and l dependence of fitness for homogeneous populations, and verified that even level agents had an high fitness when r_v was high, while fitness was higher for odd levels in case of low values of r_v , meaning that when the visual cone

5.5 Description of our model

was reduced it was necessary to have a strong collision avoidance behaviour, which resulted harmful in case of a large visual cone. The most interesting values of r_v resulted to be at the boundary between the two zones, since they were the ones for which the highest fitness was obtained by the high l levels. Furthermore, this value represented the optimal fitness value, i.e. the highest value for any r_v .

They performed also simulations with non homogeneous populations, and found that to this “boundary” area corresponded a value of r_v for which the highest (odd) level invaded all the population.

According to the authors’ interpretation, when the degree of social interaction is very high, even low levels (0, 2) are adaptive, while in case of very low social interaction level 1 is adaptive. But when the degree of social interactions is at a given specific level, agents could evolve, from a functionalist point of view, higher and higher levels. Humans are the only species that has evolved an high level of recursion (about five according to [5, 16], but individual difference could be present since according to [17] there is a large gap in distribution between level 4 and 5 in humans).

5.5 Description of our model

In this work we perform an evolutionary simulation of pedestrians in a crowd, that move towards a given goal trying to avoid mutual collisions. A particular stress is posed on the mechanism of prediction of the movement of the other pedestrians, that is at the base of the collision avoiding mechanism and that allows us to introduce in the model the concept of Theory of Mind (ToM).

To this purpose we introduce an idealised collision avoiding mechanism in which a few free parameters can be optimised by a genetic algorithm, according to a fitness function in which a positive term is given by the velocity to reach the goal, and a negative one by collisions. Since the outcome of our model is given by evolution corresponding to a very simple fitness function, we do not claim that it can describe actual human behaviour, but we expect it to present at least at a qualitative level some of the features of the self organised motion of actual crowd dynamics, as described in section 5.2. Since our personal experience shows that the lane formation in the $\alpha = 0$ case can be obtained (at least at a very rough level) also using simply a model based on velocity independent forces, we will explicitly focus on evolving a model

which is able to show self organisation in corridors crossing at an arbitrary α value.

We decided to use, instead of a completely evolvable decision mechanism (as for example a neural network), a given decision mechanism whose basic characteristics were fixed, while a few parameters were let free to evolve, since this formulation is suitable to a more clear interpretation of the results and to the introduction of concepts like that of Theory of Mind.

5.5.1 Description of the collision avoiding “cognitive” dynamics

All the agents-pedestrians in our model are represented by hard discs in 2 dimensions that undergo elastic collision between them and with the walls. We can divide the dynamics of the agents in two parts, their physical (i.e. collisions between agents and of agents with walls) and their “decision” or “cognitive” dynamics, which is based on their sensory system and decision mechanism.

Physical interaction

While the latter one, which will be described in details below, is performed at fixed time steps of length Δt , the physical one is exactly integrated using an event-driven algorithm, which means that, regardless of the length of the time step that we use to integrate the “cognitive” dynamics, the moment of each collision is exactly calculated assuming that agents move with constant velocity when their motion is not changed by collisions or “decisions”, and the dynamics of collisions is resolved assuming elastic bounces, i.e. exchanging the components of the agents’ velocity that are perpendicular to the surface of collision.

Decision mechanism force

Each agent is provided with a “goal”, i.e. a region or point of the environment that it “wants” to reach. At each time step Δt a decision mechanism, independent of the physical dynamics, is applied simultaneously by all the agents, on the base of their goal and of their sensory perception of the other agents. The output of the decision mechanism is an impulsive force \mathbf{f} that

5.5 Description of our model

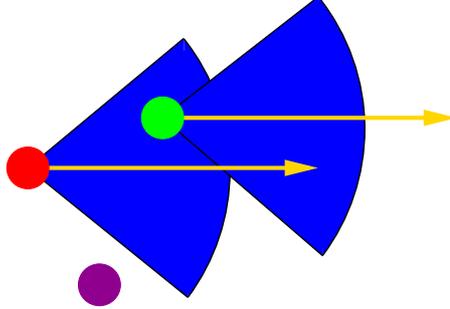


Figure 5.2: The red agent sees the green one but does not see the violet one, while the green one does not see anything. The arrows show the agents' velocities around which the visual cones are centred.

modifies the motion of the agent according to

$$\mathbf{v}(t) = \mathbf{v}(t - \Delta_t) + \mathbf{f}(t) \quad (5.1)$$

$$\mathbf{x}(t + \Delta_t) = \mathbf{x}(t) + \mathbf{v}(t) \quad (5.2)$$

(actually in our model the agents cannot exceed a maximum velocity v_{max} , which is imposed as a constraint on these equations).

\mathbf{f} is given by two terms: an “external force” \mathbf{f}_e , that does not depend on the interaction between agents (for example a driving force to the goal and eventually an obstacle avoidance force, as in the experimental setting that we describe below), plus an “interaction” one (\mathbf{f}_{int}).

Perception

The “interaction” term of the decision force is determined by the “observed” agents, i.e. those that fall in the field of view of the observer, which is a visual cone of given radius and angle centred around the agent's velocity (figure 5.2).

For each one of these observed agents, the observer is supposed to know exactly the position (of the centre of mass), velocity and direction to the goal. The third assumption could seem unrealistic, but this information can be approximately deduced, in the case of real pedestrians, by observing the gaze or the “body language” of other people (obviously an exact knowledge, not only of the goal, but also of the position and velocity, is unrealistic, and is just one of the approximations of our model).

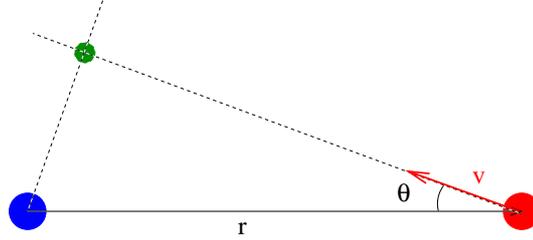


Figure 5.3: The green spot corresponds to the (future) position of minimal distance of the observed red agent with respect to the blue observer. Velocities are calculated in the frame of reference of the blue agent.

Prediction of collision

The logic at the base of the decision mechanism is to understand if there is the danger of a collision and to apply a force to avoid it. In order to do that the observer (agent i) examines all the relative positions and velocities of the observed (j) agents (\mathbf{r}_{ij} and \mathbf{v}_{ij}) and calculates the time at which the approaching agents (defined as those for which the angle between \mathbf{r}_{ji} and \mathbf{v}_{ij} is $\theta_{ij} < \frac{\pi}{4}$) will reach the minimum distance.

The minimum of these times

$$t_{pi} = \min_{j \in \text{field of view}, \theta_{ij} < \frac{\pi}{4}} \frac{r_{ij} \cos(\theta_{ij})}{v_{ij}} \quad (5.3)$$

is defined as the “time of probable impact” (figure 5.3) at which the future positions of all the observed agents (both approaching or not) are calculated (figure 5.4).

Interaction force

The interaction force \mathbf{f}_{int} is calculated as a sum of central repulsive forces depending on these future relative distances at the time of probable impact $\mathbf{d}_{ij}(t_{pi})$.

Each term will be given as (figure 5.5)

$$\mathbf{f}_{int}(\mathbf{d}) = \begin{cases} w(t_{pi}, v_{pi}) \gamma \mathbf{e}_d & \text{if } d \leq D_0 \\ w(t_{pi}, v_{pi}) \gamma \left(\frac{d}{D_0}\right)^{-\delta} \mathbf{e}_d & \text{if } d > D_0 \end{cases} \quad (5.4)$$

5.5 Description of our model

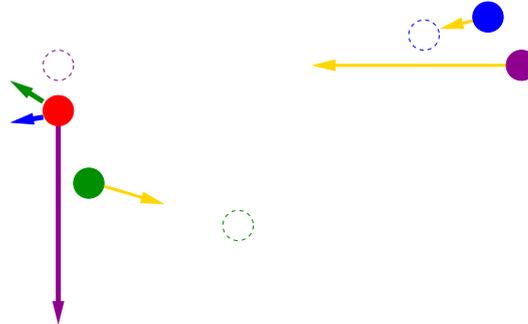


Figure 5.4: The red agent is the observer and applies the decision mechanism. t_{pi} is determined as the future position of minimal distance of the violet agent, and at that time all the (future) positions of the observed agents, violet, green and blue, are calculated. The red agent feels repulsion central forces (coloured arrows, each colour corresponds to that of the agent causing the force) determined by these future positions (dotted empty balls). Velocities in the observer's frame of reference are given by yellow arrows.

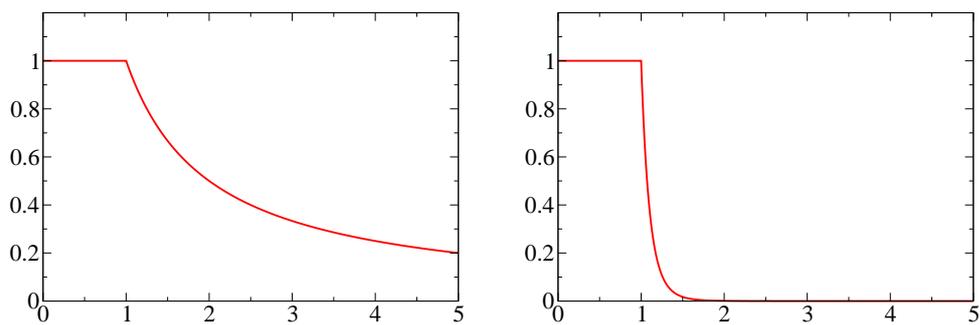


Figure 5.5: Left: absolute value of the interaction force (equation (5.4)) for $D_0 = 1$, $\delta = 1$, $\gamma = 1$ and $w = 1$ (“long range” interaction). Right: the same function with $\delta = 10$, (“hard core” interaction).

with $\delta > 0$ and $\gamma > 0$, where v_{pi} is the velocity of the agent that is going to cause the “impact” at t_{pi} and w is a term that determines the “danger” of the situation as

$$w(t_{pi}, v_{pi}) = \min \left(\frac{v_{pi}}{\gamma t_{pi}}, 1 \right) \quad (5.5)$$

where γ is the maximum force that the agent can apply. w is defined in such a way to attain a complete stop in case of a frontal impact, and in some way renormalises γ (actually these formulae have to be slightly modified in order to take in account the presence of the driving force to the goal).

In figure 5.5 we show two possible forms of the function given by equation (5.4). When δ is low ($\delta \approx 1$) we can talk about “long range” interaction, because the agents interact also when their (future predicted) distance d is higher than D_0 , while for $\delta \gg 1$ there is an “hard core” interaction, since considerable interactions happen only if $d \leq D_0$. D_0 can be interpreted as the agent’s perception of the size of its own body (its diameter, or the sum of its radius plus the other agent’s radius, assumed as equal), or as a “comfortable distance” to another agent. The first interpretation is more accurate in case of “hard core” interaction: agents interact only if they predict their bodies to overlap in the future.

Recursion level

The last point to be clarified is how the new positions of the observed agents are calculated, which is correlated to (and thus allows us to introduce) the agent’s ToM recursion level. If the agent has no ToM, the calculations just described will not be performed at all ($\mathbf{f}_{int} = 0$), and the agent changes its motion according just to its goal (the “external force”, thus a level 0 agent is a non interacting particle, or, more properly, it interacts only through elastic physical collisions, but not through a decision mechanism).

In case of a level 1 agent, the other agents are assumed to be level 0, and their future position at time t_{pi} are calculated on the base of their velocity corrected by the impulsive external force \mathbf{f}_{int} . As we said before, “following the gaze”, the observer is able to know the goal of the observed agent, i.e. the direction of the driving force. But this does not mean that it knows the intensity of this force. Actually, in this first version of the model, each agent will assume that all the agents are identical to itself, and will consider that they are driven to their goal with the same attraction that it feels for its own one, even if this is not necessarily the case. (We could say that in our model

5.5 Description of our model

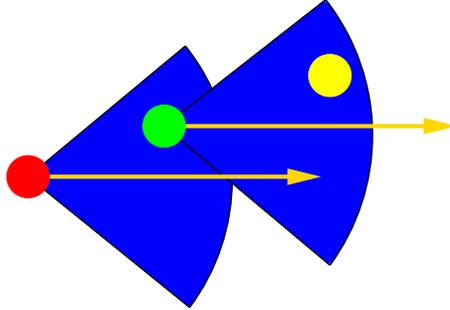


Figure 5.6: The red agent thinks that the green one does not see any agent, while it sees the yellow agent.

agents assume the others to have the same Theory of Mind and behavioural patterns that they have, while they admit the others to have different intentions and beliefs, i.e. different goals and perceptions).

Since we are talking about relative positions (following the discussion described in figure 5.4), to calculate them the observer has to consider its own future position too. This is calculated applying to itself the same Theory of Mind it applies to the others, i.e. assuming that it will act as a level 0 agent (even if it “knows” to be level 1, since obviously all the calculations about future position have to undergo some approximation).

We notice that even in this first (level 1) case a superposition principle does not apply. Due to the presence of t_{pi} the force that the observer feels when it sees simultaneously two agents is not the sum of the two forces felt when a single agent is observed.

In case of an agent with a level 2 ToM (or a level $n > 1$ one), the observer performs all the level 1 (level $n - 1$) calculations for all the observed agents (including itself). Obviously it performs these calculation on the base of its own observations, which do not necessarily imply a perfect knowledge of the other agent’s observations (see figure 5.6).

In the case of a level $n > 1$ agent, the superposition principle does not apply, not only for the presence of the “time of probable impact” t_{pi} , but also because in this case the “future predicted positions” of the observed agents are modified by the presence of other agents (since these agents are considered by the observer to be at least level 1 and thus to interact with the others).

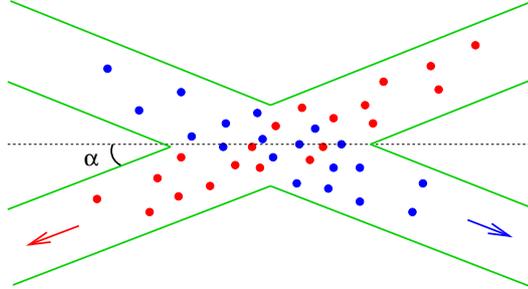


Figure 5.7: The arrows represent the goals or driving forces of the agents.

5.5.2 Description of the evolutionary and experimental settings

The environment

The model does not impose any constraint on the position of the goals, but in order to study the emergence of the self organisation properties described in section 5.2, we split agents in two groups with different goals, each group moving in a corridor. The two corridors form a crossroad with an angle α (figure 5.7).

Goals are realised with a constant driving force term \mathbf{f}_g directed along the corridors, while a tendency to avoid collisions with the walls is introduced as a term \mathbf{f}_w whose direction is normal to the walls and whose intensity is given by

$$f_w = \begin{cases} 0 & \text{if } x > d \\ c \frac{d-x}{d} & \text{if } x \leq d \end{cases} \quad (5.6)$$

where x is the distance to the wall and d the maximum distance at which the wall is “felt”, while c stands for the strength of the repulsion.

Genetic code

The “genetic code” of the agents is given by the set of continuous parameters f_g , c and d (that fix the “external” force), α_v (the angle of view), r_v (the radius of view), γ (the maximum possible interaction strength), δ (the scaling law for the interaction force in equation (5.4)) and D_0 (the “perceived

5.5 Description of our model

diameter of body”), plus a discrete one, the ToM recursion level l . This latter parameter was kept fixed in some of the experiments, devoted simply to study the evolution of the collision avoiding behaviour, while it was let free to evolve in experiments aimed to study the evolution of the ToM level.

At the beginning of the evolutionary process N agents (we used $N = 100$ in all experimental settings) are generated and the parameters are randomly chosen in a given range of values. (We fixed some maximum values for the external forces, namely 20 for c and 4 for f_g and choose the corresponding parameters between 0 and those values; let d vary between 0 and the half width of the corridors; fixed maximum values for r_v and D_0 as 10 times the actual radius of the agent r ; let α_v evolve in all the allowed range $[0, \pi]$; randomly choose δ between 0 and 10 and γ between 0 and a very high upper bound, 4000. Some of the parameters, D_0 , γ and δ , were let free to be driven by evolution over the upper bounds of these ranges).

Heterogeneous population

Throughout all the evolutionary process agents are part of a heterogeneous population, which means that the genetic code of each agent is chosen in an independent way, i.e. each agent is different from the others, but all these agents are located at the same time in the same environment. As discussed below this choice poses some problems in the evaluation of a single agent performance, and thus on the effectiveness of the genetic algorithm, because the fitness of a given agent is determined not only by its own genetic code, but also from those of the other agents it interacts with (we treated this problem, i.e. the impossibility of defining a *fitness landscape* for heterogeneous populations, also in chapter 2.4.7). Nevertheless, we chose to use this setting, instead of a more traditional setting in which the agents were evolved using homogeneous populations composed of clones, for two reasons. The first reason is merely a computational one: since we are studying crowd dynamics, we need populations with a quite large number of agents (for example we used $N = 100$). In case of a non homogeneous population each run, which is usually computationally expensive, especially if high ToM levels are involved, allows us to test N genetic codes, while it would allow us to test only a single genetic code if we used clones. The second reason is a theoretical one, since we assume that differences in the individual behaviour in the population have an important role both in crowd dynamics and in ToM evolution.

Generalisation at different angles

To test properly the ability of agents to generalise their behaviour in different situations, each generation is composed by several *runs* (see also chapter 2.4.8 for an explanation of what we mean by “run”, which is quite different from the usual GA terminology). In each run the agents are randomly located in one of the two groups with different goals, and the angle between the corridors is changed. Each time that the agents reach the end of the corridor (their goal) they are relocated at the beginning of it, in a random position (i.e. they have no memory of their transverse position).

Fitness

If we define the average velocity of the agent towards the goal $v_g = \frac{\Delta x}{\Delta t \cos(\alpha)}$ as the total displacement of the agent along the corridor over time, averaged over all the runs, and Δp as the overall physical momentum exchanged by the agent during collisions, the fitness is given by

$$f = v_g - \beta \Delta p \quad (5.7)$$

β is defined as the “pain” factor in the evolution of agents, and determines their tendency not only to reach the goal in the shortest time, but also to avoid collisions (obviously in order to reach the goal moving in a crowded environment is important to avoid collisions, but the β term makes this dependence explicit, and introduces a tendency to avoid collisions “for collision avoidance’s sake”).

Genetic operators

At the end of each generation agents are chosen for reproduction on the base of a tournament selection (two agents are randomly selected and the best one passes its genes to the next generation), and then the genetic codes are mutated with mutation probability $p_m = 0.05$ for each gene. We also used crossover, with a very low crossing probability $p_c = 0.01$. The continuous parameters are directly represented as real numbers and mutated with a Gaussian error, while the ToM level l is allowed to change up to $l \pm 2$ (the reason to allow also a two step mutation will be more clear when we will discuss the difference between even and odd ToM levels). The number of generations in a complete evolutionary process is 500.

5.5 Description of our model

Physical dimensions of parameters

In our experiments we use agents with physical radius $r = 0.5$ and maximum velocity $v_{max} = 2$. These parameters are intended as adimensional but they can also be considered as expressed in meters and meters per second. It will follow that, assuming masses to be fixed at 1, all the quantities in the model have proper dimension and can be scaled changing the space and time unities (Δt scales as time; f_g , γ and c as forces; d , D_0 and r_v as distances). We used a time integration step $\Delta t = 0.05$.

5.5.3 Comparison with Takano’s Model

Since we compare our results to the model of Takano et al. described in section 5.4, it is important to understand which are the basic differences between that model and the model that we are proposing.

In Takano’s model collisions and collision avoiding were introduced to represent and simulate social interaction. There was no intent to describe actual collisions, nor to develop a realistic or efficient collision avoiding behaviour. The interaction between agents was given and fixed *a priori*, and the behaviour of agents depended only on the extent of the sensory system and on the ToM level. As stated in their paper, keeping low the number of parameters in the model allows a deeper analysis of the parameter space and a more clear interpretation of results.

In our model, we described physical collisions in a realistic way, giving to each agent a finite, fixed and impenetrable volume, and measuring the “pain” of the agents on the base of a physically well defined function (momentum) and tried to evolve an efficient (and possibly similar to the pedestrians’ actual one) collision avoiding behaviour. As a result, in our model we have a relatively large number of parameters, whose value is fixed by the evolutionary process, which makes an exploration of the parameter space more difficult. Moreover, in our model a certain amount of “plasticity” is present, which means that given a certain value of the ToM level, the evolutionary process will fix the values of the other (continuous) parameters in order to have an appropriate collision avoiding behaviour. This means that studying the evolution of the ToM level l in our model means to study the co-evolution of l and of the continuous parameters.

We also notice that two concepts that has been fundamental in our model in order to develop an efficient collision avoiding behaviour, the “time of

probable impact” t_{pi} (equation (5.3)) and the “danger” weight w (equation (5.5)), were not present in Takano’s model, in which future positions were calculated at the next time step Δt , and thus the behaviour of the agents in the two models is quite different.

5.6 Description of an evolutionary experiment

Parameter space

As we just said the behaviour of the agents in the model is the output of an evolutionary process, and obviously depends on how the experiment is performed, i.e. on a few parameters that fix the structure of the environment in which agents are evolved. These parameters are the number of runs for each generation, the angles α between the corridors in these runs, the density ρ of the agents in the environment and the value of the pain factor β , the parameter that fixes the negative weight of collisions in the fitness function (5.7). We have observed that agents evolved using 3 runs for each generation, with values of the angles α given by $\frac{\pi}{4}$, $\frac{\pi}{2}$ and $\frac{3}{4}\pi$, were we able to generalise their behaviour to any value of α , and thus we kept these values fixed for any evolutionary experiment that we have performed, and studied the dependence of the evolutionary output on the values assumed by ρ and β .

Density

Regarding the density of the agents, since we have kept the number of agents N fixed to 100 in all experiments, ρ was actually determined by the length and width of corridors. The maximum density that we used, to which we are going to refer as $\rho = \rho_1$, was obtained locating $N = 100$ agents with radius $r = 0.5$ in corridors of length 30 and width 10 (Figure 5.8 shows the agents in an environment with density ρ_1 . The red agents start at the top-left of the figure, and their goal is at bottom-right, while the white ones move from top-right to bottom-left. Agents are initially located in random positions in their own corridor, but the evolved agents shown in the figure are self organised in large clusters.). The other values for the density ρ that we have used in the evolutionary experiments are $\rho_2 = \frac{2}{3}\rho_1$, $\rho_3 = \frac{1}{2}\rho_1$ and $\rho_4 = \frac{1}{4}\rho_1$. In order to understand in a better way the results of our experiments, before studying the β and ρ dependence, we first analyse several repetitions of single

5.6 Description of an evolutionary experiment

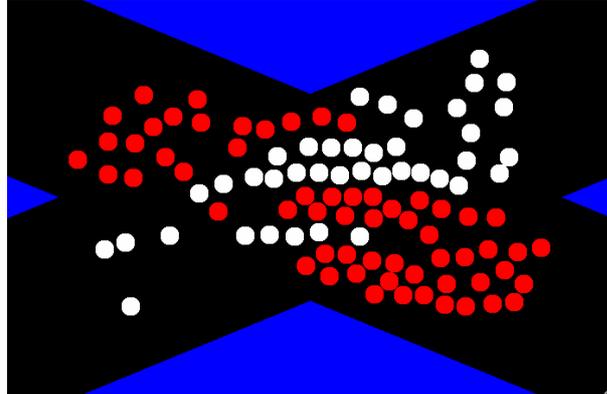


Figure 5.8: Evolved agents moving in the experimental setting using $\rho = \rho_1$.

experimental setting in which $N = 100$ agents with ToM recursion level fixed to $l = 1$ were evolved using $\beta = 1$ and $\rho = \rho_1$.

Fitness, velocity and heterogeneity

In figure 5.9 we show the fitness and average speed to the goal of the agents through generations in 5 repetitions of the experiment. Both f and v_g reach their maximum value (roughly 1 for f and 1.8 or 90% of the maximum allowed speed for v_g) after around 200 generations, but their evolution is quite unstable and shows many negative “spikes”. This instability (which is present only at high densities, and was not observed in experiments regarding agents evolved with $\rho \leq \rho_2$, see for example figure 5.16), that could be probably partially attenuated through an accurate choice of the mutation and crossing probabilities, is substantially due to the use of heterogeneous populations. As we stressed before, the fitness of a given agent could be very high when moving within agents different from itself, but the overall fitness could drop down if its characters invade the population. Moreover, mutated agents, whose presence is indispensable for the evolution of the system, can affect seriously the behaviour of the whole system, including the potentially fittest agents.

To show that these effects are actually present, we tested in different runs the performances of evolved agents, using both homogeneous and heterogeneous populations. In particular, we used three populations obtained in the

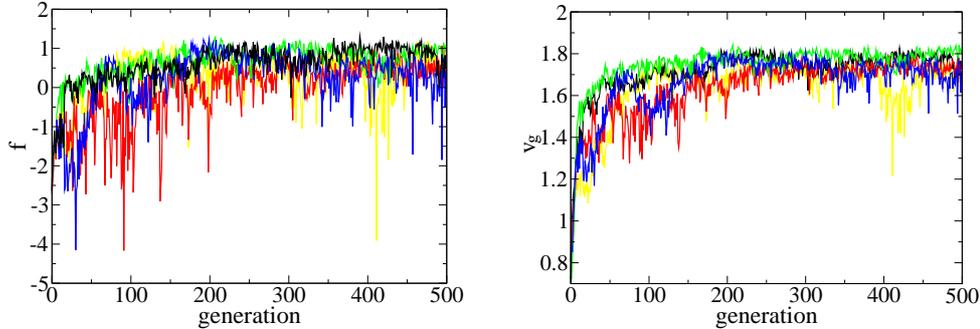


Figure 5.9: Left: f as a function of generations for 5 repetitions of the $\beta = 1$, $\rho = \rho_1$ experiment. Right: v_g as a function of generations for the same repetitions.

following way: population 1 is the non homogeneous population that had the best fitness during the whole evolution, population 2 is an homogeneous population composed by $N = 100$ clones whose genetic code was obtained averaging over all the corresponding genes in population 1, and population 3 is an homogeneous population composed of N clones of the best individual throughout evolution (see also the discussion in chapter 2.4.7). The results (figure 5.10) show that diversity affects the overall fitness and velocity of the agents (population 1 has the poorest performance) but also that the best individual (population 3) does not produce the best crowd behaviour. The fact that population 3 has a lower fitness but higher velocity suggests that the best individual had a behaviour which preferred reaching the goal than avoiding collisions, and took in some way advantage of the collision avoiding behaviour of the others.

Evolution of parameters

Studying the evolution of the (averaged over the population) values of the continuous genes, we have observed that the evolution of the interaction strength γ and of the parameters regarding the interaction with the walls is quite erratic and does not have a large impact on the behaviour of the agents, at least if their values are within a quite large range. (We recall that γ is just the maximum possible strength, while the actually applied one is modulated through the w term in equation 5.5. This means that once γ is

5.6 Description of an evolutionary experiment

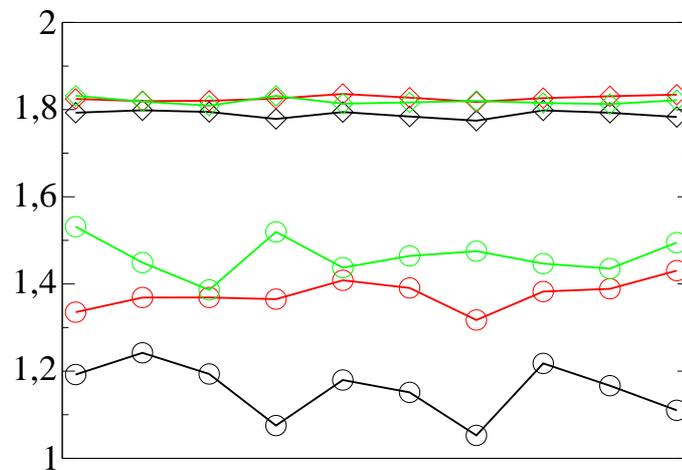


Figure 5.10: Comparison between 10 test simulations for a crowd of clones and a crowd of non homogeneous (evolved) agents. Black lines shows the results obtained by the best generation during evolution (population 1), green ones those obtained by clones of the “average member” of that generation (population 2) and red ones those obtained by the clones of the best individual throughout evolution (population 3). Circles show the values corresponding to the fitness f , while diamonds those corresponding to the velocity to the goal v_g .

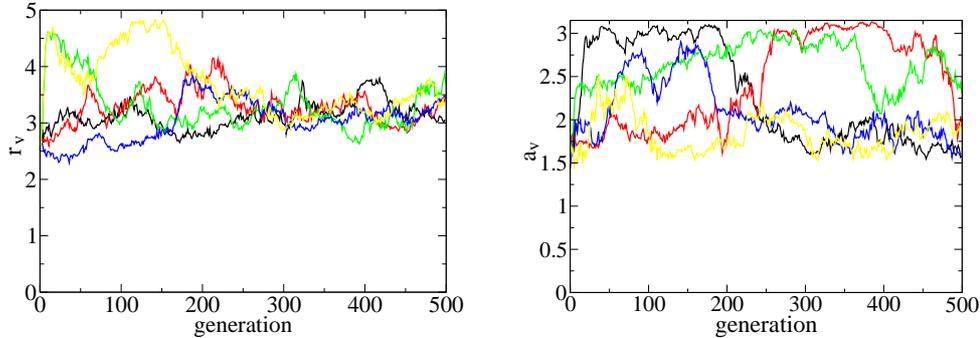


Figure 5.11: Evolution of the average value of r_v (left) and a_v (right) as a function of generations for 5 repetitions of the experiment.

high enough, its exact value is not of great importance in the model.) For this value of β the attraction to the goal f_g reaches quickly and stays next to its maximum allowed value ($f_g = 4$), while the values regarding the dimension of the cone of vision (figure 5.11) approach initially their maximum allowed values, $r_v = 5$ and $a_v = \pi$, at least in some repetitions of the experiment, but in the following stages of evolution seem to settle to lower values (in particular a_v reaches a value next to $\frac{\pi}{2}$ in most simulations, i.e. perception is limited only to the agents located in the direction of movement).

An interesting result regards the evolution of the “perceived size” D_0 , which was initially chosen as a random value between 5 (the maximum allowed value for r_v) and 0, but converges quickly to a stable value next to the actual size of the agent $2r = 1$ ($D_0 \approx 1$ for three simulations, $D_0 \approx 0.8$ for the other two). A confrontation between the graphs in figure 5.12 shows that $D_0 \approx 1$ was obtained by agents that had developed an “hard core interaction” ($\delta \gg 1$), and thus suggests that the agents had “learnt” their actual size, i.e. they interact only with agents whose future position will overlap their own. (This is quite clear in the “hard core interaction” cases, while it seems reasonable also in the other two repetitions, since D_0 is slightly lower than the actual diameter and δ assumes values around 3 and 5, which implies a quite quick decrease of interaction strength with distance. We also tested that changing the actual agents’ radius r to a different value \bar{r} while leaving the other parameters unchanged, evolution leads to $D_0 \approx 2\bar{r}$).

5.6 Description of an evolutionary experiment

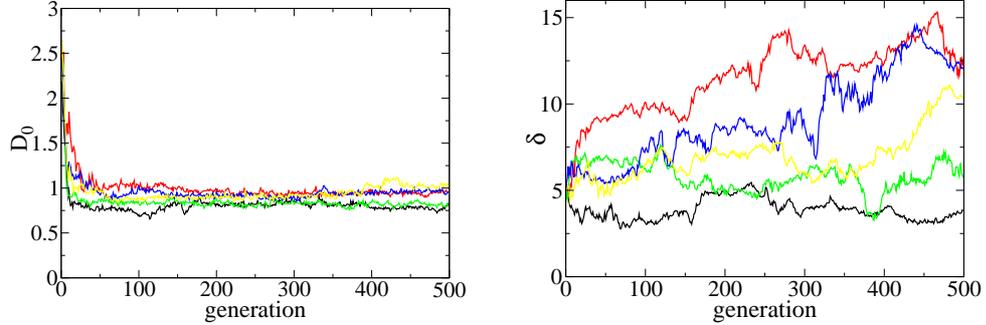


Figure 5.12: Evolution of the average value of D_0 (left) and δ (right) as a function of generations for 5 repetitions of the experiment. Notice that to “hard core interaction” $\delta \gg 1$ corresponds $D_0 \approx 2r = 1$.

5.6.1 Self organising properties

We show in figure 5.13 the behaviour of agents evolved in the $\beta = 1$, $\rho = \rho_1$ experiment, which present the patterns observed in actual pedestrian dynamics, as described in section 5.2. (The figures show the results obtained by an homogeneous population of clones of the average member of the best generation in the five repetitions of the experiment. Nevertheless, qualitatively similar results can be obtained using best agents from each repetition and also, even if the pattern formation is slightly more unstable, using heterogeneous populations). As the figures show, agents are able to organise for any value of α , i.e. to generalise their behaviour also to angles between the fluxes that have not been tested during evolution (including the extreme values $\alpha = 0$ and $\alpha = \pi$).

We have thus been able to obtain the self organising behaviour of actual pedestrians as the output of an evolutionary process in which agents with a finite size body try to maximise their velocity and minimise the “pain” felt in collisions.

We claim that these patterns are clearly the result of self organisation. It can be questionable that the interaction rules that we have introduced are simple (nevertheless, in the $l = 1$ and $\delta \gg 1$ case, they can be simplified as “assuming that everyone moves on straight lines, check which agents will be the first ones to collide, and interact with them proportionally to the ratio between their velocity and the time at which the collision will happen”) but

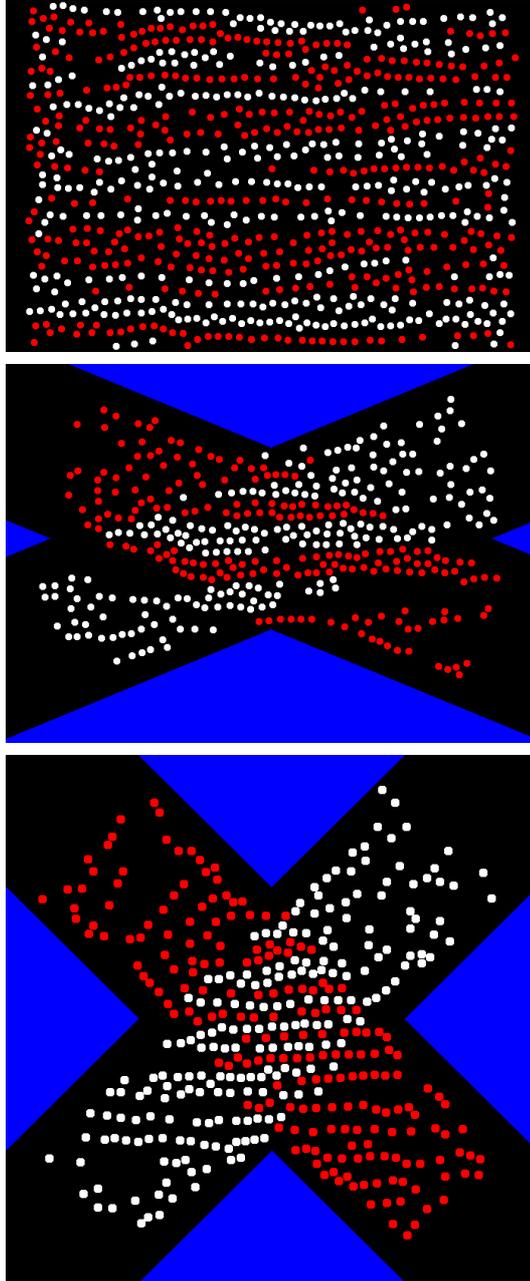


Figure 5.13: Lanes and stripes formation in evolved agents. These figures have been obtained using $N \gg 100$ agents (while keeping density fixed to ρ_1) in order to make the stripes and lanes formation more evident. The angles between corridors are $\alpha = 0$, $\alpha = \frac{\pi}{2}$, $\alpha = \pi$.

5.7 The role of pain

they are surely local, since the sensory system does not allow a complete knowledge of the state of the system (this is especially true for $\delta \gg 1$, when interaction is limited only to the agents that are predicted to collide) and individual, since the dynamics of the system is determined solely by the decision mechanisms of the agents.

Nevertheless the self organising nature of the patterns is assured by the fact that *no description of the overall pattern is included in the decision mechanism* (agents are not requested to form lanes or stripes, for example following other agents moving in their same direction). Furthermore, *the observed pattern has not been obtained as the result of a “biased” parameter fixing* (parameters are fixed by a GA which uses a *fitness function in which no description of the pattern is provided*).

The process of positive feedback (see the discussion in chapter 1.4) is given by the tendency of agents to follow other agents moving in their same direction, but we stress once again that this tendency is not included directly in the model, but arises because when a small group of agents moving in the same direction is formed as the result of a fluctuation, it is extremely profitable, from the point of view of collision avoiding, to follow it when moving in the direction of the goal of the “observer”, and to avoid it very strongly if going in the opposite direction. It is thus very likely that any efficient collision avoiding mechanism could lead to a similar pattern formation.

Negative feedback in our setting is due, besides to the finite number of agents, to the fact that clusters are destroyed as soon as agents reach the end of the corridors. Pattern formation in the model is dynamic and continuous: usually agents arriving at the crossing point enter in already formed stripes or lanes (the pattern is stable), but if some fluctuation destroys an existing pattern, a new one is formed in a short time.

5.7 The role of pain

Keeping fixed $l = 1$ we analyse now the role of the pain parameter β in evolution. To compare the performance of agents evolved with different values of β we can study the value, both reached during evolution and attained by evolved agents, of the velocity towards the goal v_g , which is the common part of the fitness function (since it does not depend on β). We could expect v_g to be a decreasing function of β , since the agents evolved with $\beta = 0$ are only requested to maximise v_g , while those evolved with high β are supposed to

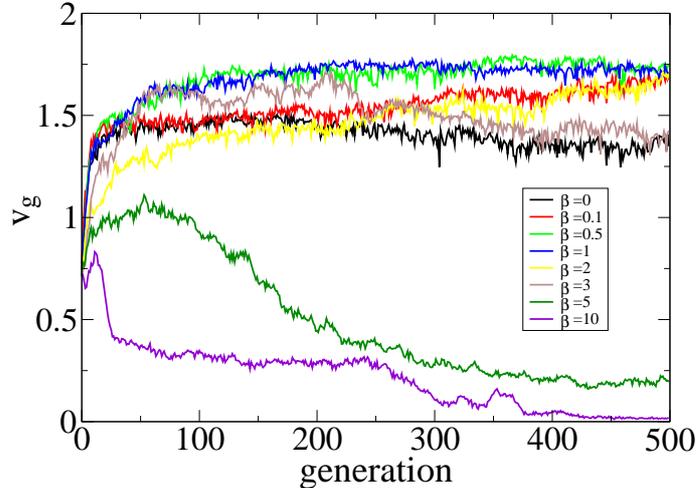


Figure 5.14: v_g as a function of generations for agents evolved with different values of β (see the legend for the correspondence between colours and values of β).

avoid strongly collisions even if this behaviour prevents them from reaching an high velocity. Quite interestingly, this is not necessarily the case.

5.7.1 Evolution in a high density environment

In figure 5.14 we show, for 100 agents with density ρ_1 , the evolution through generations of the average value of v_g (the average is performed over the $N = 100$ agents and over 5 repetitions of the experiment). This figure shows that agents evolved with $\beta = 0.5$ and $\beta = 1$ reach during evolution a v_g value higher than that reached by those evolved with other values of β . When agents are evolved using $\beta \gg 1$, v_g decreases during evolution and finally drops to a value next to 0. Nevertheless, an analysis of the corresponding fitness function shows that f is an increasing (even if always negative) function and thus that in these cases evolution leads the agents to move very slowly in order to avoid collisions that would strongly decrease their fitness (in particular, evolved $\beta = 5$ and $\beta = 10$ agents are almost still, while for $\beta = 2$ and $\beta = 3$ the behaviour can be quite different between different repetitions of the experiment and also during the same evolutionary process, and the corresponding curves are more noisy).

5.7 The role of pain

For low values of β , in most repetitions of the experiment v_g reached a quite lower value than that reached in the $\beta = 0.5$ and $\beta = 1$ experiments. Since for $\beta \approx 0$ we have $f \approx v_g$, *this means that the agents evolved with $\beta \approx 1$ can perform better than the $\beta \approx 0$ agents the task the latter have been evolved for.* Given that the initial values of the parameters and the structure of the genetic algorithm are the same, this means that *the absence of a “sense of pain” prevents the correct evolution of the system.*

It is difficult to understand why this happens, but some hints can be obtained studying the evolution of some parameters that are strongly related to the β pain factor. In figure 5.15 we can observe the average evolution of f_g (attraction to the goal) and D_0 (perceived size of the body) for different values of β . For agents evolved with $\beta \approx 0$ f_g converges very quickly to its maximum value, while for $\beta \approx 1$ this convergence is quite slower (for $\beta \gg 1$ f_g drops to 0, allowing agents to “stand still”). These data seem to suggest that, while agents evolved with $\beta \approx 1$ are “urged” by the form of their fitness function to develop a collision avoiding behaviour in the first stages of evolution, the $\beta \approx 0$ ones can have “an easy reward” just increasing their attraction to the goal, but they pay the cost of this behaviour later, being trapped in a zone of the parameters (a local maximum) that does not allow them to develop a good collision avoiding behaviour. This conjecture is confirmed also by the evolution of D_0 .

Agents evolved with $\beta \approx 1$ are able to learn quickly the size of their body ($D_0 \rightarrow 2r = 1$), while those evolved with $\beta \approx 0$ do not show an appropriate learning curve. (Also agents evolved with $\beta \gg 1$ tend to overestimate their size, but this could be due to their tendency to avoid strongly any possible collision. Furthermore, since these agents tend to have $f_g = 0$, it is quite hard to talk about any collision avoiding strategy, because evolved agents almost do not move at all. See for example the D_0 curve for $\beta = 5$: as soon as f_g goes to zero, the agents lose the perception of the size of their body that they had previously learnt).

Obviously, since evolution is a random process, we could think that some particularly large fluctuation could drag the $\beta \approx 0$ agents outside the local maximum and allow them to develop an appropriate collision avoiding behaviour, and actually this has been the case at least for few simulations. We could also think that if we let agents evolve for a larger time, an efficient behaviour would sooner or later emerge in any experiment (notice for example how at the end of the evolutionary process for $\beta = 0.1$ v_g reaches a value

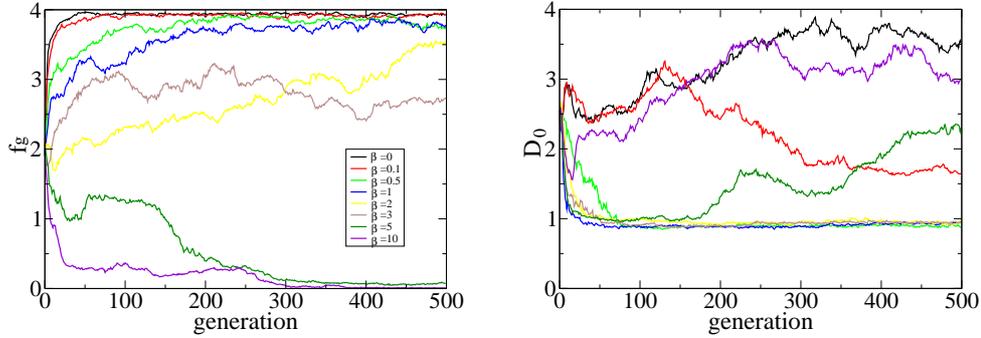


Figure 5.15: Evolution of the average (over agents and over 5 repetitions of the experiment) value of f_g (left) and D_0 (right) as a function of generations. See the legend at left (or in figure 5.14) for the correspondence between colours and values of β .

next to that of $\beta = 0.5$ and $\beta = 1$, figure 5.14). Nevertheless the results show clearly that an appropriate β pain factor leads to an easier emergence of the collision avoiding behaviour.

If we analyse the behaviour of evolved agents we see that the self organisation (lanes and stripes formation) behaviour is present only for the agents with $v_g > 1.7$, i.e for all the repetitions of the experiment in which agents had been evolved with $\beta \approx 1$, but only in some repetitions for $\beta \approx 0$.

5.7.2 Evolution with low density and generalising abilities

In the simulations that we have just examined, the density of the agents was quite high, i.e. they occupied a large portion of the overall volume of the corridors (figure 5.8). We repeated the same analysis of the effect of β on evolution (limiting ourselves to three values $\beta = 0$, $\beta = 1$ and $\beta = 10$) for an environment in which the density of agents was 2 times lower ($\rho = \rho_3$). The results of these simulations (figure 5.16) show that in this case all the agents, regardless of the value of β they were evolved with, reach an average speed around $v_g = 1.95$ in the last stages of evolution. An analysis of the evolution of the perception of the size of their body, D_0 , shows that agents evolved with $\beta = 1$ and $\beta = 10$ learnt the actual size of their body $D_0 \approx 2r = 1$,

5.7 The role of pain

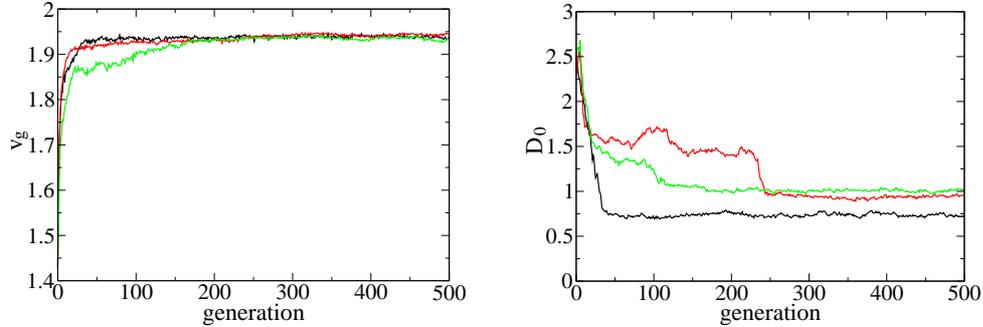


Figure 5.16: Left: evolution of v_g through generations for $N = 100$ with density ρ_3 (average over agents and repetitions of the experiment). Right: corresponding values of D_0 . The black lines stands for agents evolved with $\beta = 0$, the red ones for $\beta = 1$ and the green ones for $\beta = 10$.)

while those evolved with $\beta = 0$ tend to underestimate it by a factor $\frac{3}{4}$ (all the agents had developed a quite high value of δ and thus an hard core interaction, and we can interpret D_0 directly as the agent's perception of the size of its body).

In order to understand if this different “perception of the size of body” has any effect on the behaviour of agents, we tested them in a environment different from the one they were evolved in ($N = 400$ clones of the average agent in the most fit generation over all the repetitions of the experiment for each value of β were located in corridors of length 60 and width 20, i.e with density ρ_1 or the double of the density in the experiments they were evolved in, and the angle at which corridors crossed was chosen as $\alpha = \pi$, as in the experimental setting in the bottom of figure 5.13). Figure 5.17 shows the performances of the agents. In the environment of the evolutionary experiment, agents evolved with $\beta = 0$ exchange more momentum than those evolved with $\beta > 0$, but move with a slightly higher velocity. Nevertheless, they cannot generalise this behaviour to higher values of density and α . It is also interesting to see that while the performance of the $\beta = 1$ agents was better than that of the $\beta = 10$ ones in the first environment (they exchanged the same amount of momentum and had a slightly higher v_g) the situation is reversed in the generalised environment were the $\beta = 10$ agents exchange a sensibly lower amount of momentum. These results show that *agents evolved with an appropriate pain factor are able to generalise their behaviour to more*

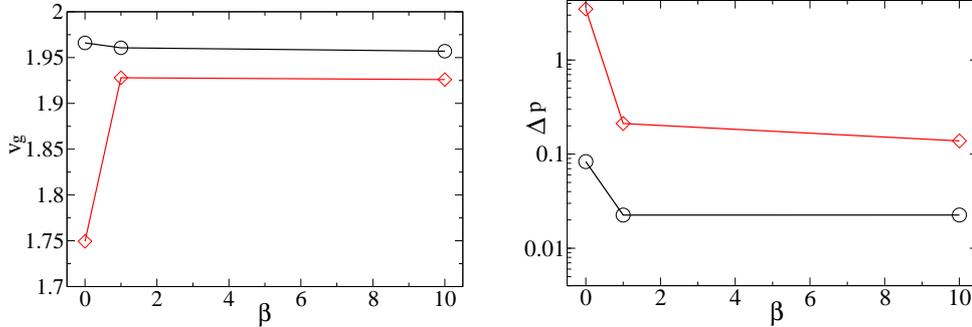


Figure 5.17: Left: average v_g of the evolved agents as a function of β . Right: exchanged momentum Δp for the same agents, in logarithmic scale. (Black and circles report the agents’ performance in the original environment, while red and diamonds in the generalised environment).

crowded environments, while those evolved with a low “sense of pain” are not able to.

5.8 Comparison between different l values

5.8.1 Behavioural difference between even and odd levels

One of the most interesting features of the model by Takano and Arita was the difference between odd and even levels in the Theory of Mind, as described in section 5.4. We find the same difference in our model. In figure 5.18 we show the average (over agents and repetitions of the experiment) momentum transferred in collisions Δp through generations for different values of l (l was equal for all the agents in a given population, β was fixed to 1 and the density was ρ_2). The results show that the amount of exchanged momentum is minimum for agents with $l = 1$, and that there is a clear distinction between odd and even levels (the results for $l = 0$, to which corresponded a noisy curve centred around $\Delta p \approx 4$, are not shown).

To explain (following the discussion in [12, 13]) this difference we can say that even levels act in a “bold” way, while odd ones are “careful”. A level 0 agent is “bold” in the sense that it doesn’t care about the others’ behaviour,

5.8 Comparison between different l values

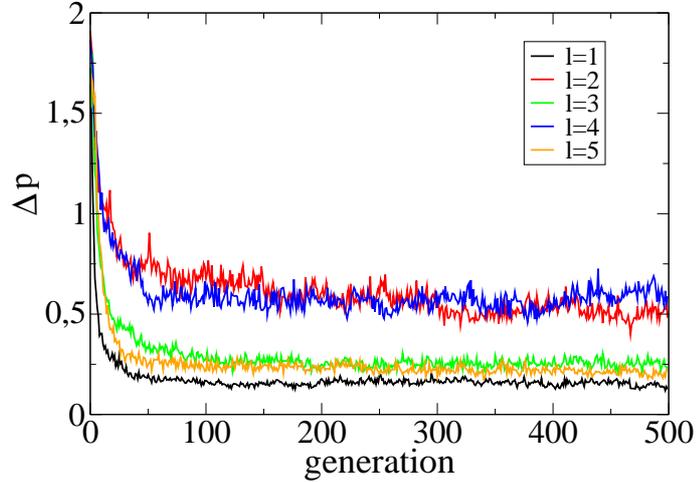


Figure 5.18: Average momentum exchanged in collisions Δp for agents evolved with a fixed value of l , for different values of l (see legend), using $\beta = 1$.

and thus a situation in which two level 0 agents have contrasting goals leads to a collision. Level 1 agents consider everyone (including themselves) to be level 0, and thus they predict the collision, and in order to avoid it their behaviour is very “careful”. But level 2 agents consider everyone to be very “careful”, and thus “try to take advantage” of this situation behaving in a “bold” way, and so on (figure 5.19).

5.8.2 Level dependence of the continuous parameters

As stated before, the main difference between our model and the model described in section 5.4 is the presence of many parameters evolved by the genetic algorithm, which means that the behaviour of an agent is not determined *a priori*, but is the consequence of the evolutionary process. For this reason, before any evaluation of the agents’ fitness, we need to evolve them, which is very expensive from a computational point of view, in particular when high ToM levels are concerned, and thus a throughout analysis of the (ρ and β) parameter space, even at fixed l , has not been possible. Nevertheless, our model has more plasticity, and at the end of the evolutionary process we can expect, for any value of $l > 0$, agents to have developed a

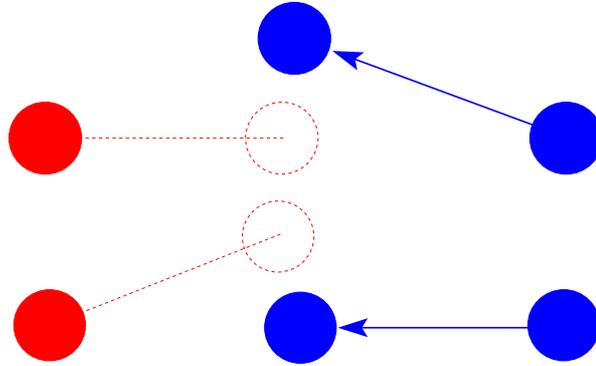


Figure 5.19: Top: the blue agent is a $l = 1$ (or in general odd) one and predicts the red one not to avoid the collision, and tries to avoid it (careful behaviour). Bottom: now the blue agent is $l = 2$ (or in general even), predicts the red agent to avoid the collision, and does not try to avoid it (bold behaviour).

certain ability to perform in the environment (for example, in figure 5.18 Δp is a decreasing function of generations for any l value, even if its equilibrium value depends on l). In fact an analysis of the behaviour of evolved agents shows the lanes and stripes formation (at least at some degree) for any value of l , including even ones. Figure 5.20 shows how this is attained. While D_0 reaches quickly a value next to the actual size $2r$ for odd levels, the value is higher for even ones. Nevertheless, the curves for even levels are different from the erratic ones that agents with too low or high β had in figure 5.15, and are similar (in particular the $l = 4$ one) to those obtained for odd levels, just with an higher asymptotic value. Remembering (see the discussion following figure 5.5) that the evolution of D_0 is tied to that of δ , since these parameters fix the form of the function in equation 5.4, a detailed analysis of the output of each evolutionary experiment shows that while odd levels developed $D_0 \approx 2r$ and $\delta \gg 1$, i.e. an “hard core” interaction within the actual size, for even levels two possible strategies are present: one with $\delta \approx 1$ and $D_0 \approx 2r$, and another one with $\delta \gg 1$ and $D_0 \approx 4r$; both resulting in an overestimation of the size of the body that attenuates the “bold” behaviour

5.8 Comparison between different l values

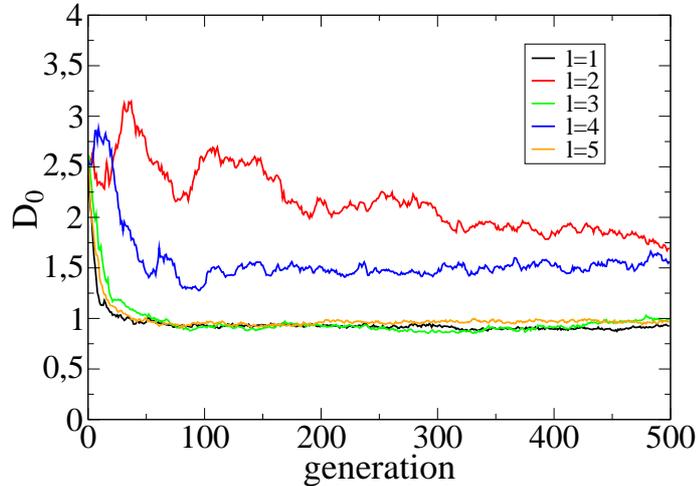


Figure 5.20: Evolution of D_0 for agents evolved with a fixed value of l , for different values of l (see legend), using $\beta = 1$.

of even levels (if in the bottom part of figure 5.19 the blue even level agent projected the future position with a radius larger than the actual one, it would consider the collision more probable to happen).

5.8.3 Evolution of agents with high l in high ρ environment

We have verified that when the density is high enough ($\rho \geq \rho_3$), the fitness equilibrium value attains always the highest value for $l = 1$, and it attains always an higher value for odd levels than for even ones. In particular (figure 5.21), when the density is quite high (for example if $\rho = \rho_1$), the fitness reached by level 1 agents is quite higher than that of the other levels, including odd ones. These results seemed to show that odd high levels could not properly self organise in case of high density. Actually, an analysis of the best agents through evolution shows that the performance of level 3 agents is comparable, even if lower, to that of level 1 agents (figure 5.22). In particular, they attain almost the same value for the velocity v_g (1.82 for level 1 agents and 1.79 for level 3 ones) and show similar self organisation properties. Moreover, the difference in performance between clones and heterogeneous population is higher for level 3 agents than for level 2 and level 1 ones (defin-

ing this difference as Δf we have $\Delta f = 0.72$ for level 3, $\Delta f = 0.16$ for level 2 and $\Delta f = 0.24$ for level 1). It is not hard to convince ourselves that high odd levels are more unstable in case of a heterogeneous population. We expect the behaviour of our agents to converge to a very accurate prediction of the others' motion for $l \rightarrow \infty$, but this assumes an exact knowledge of the values of parameters (genetic code) of the other agents, i.e., in our model in which each agent assumes for the others its same behaviour, homogeneity. The presence of different values for the behavioural parameters in the population leads to non accurate predictions. This error in the prediction of future positions probably propagates at each evaluation, and thus is higher for high l (furthermore level 1 agents are even more stable in case of heterogeneous population, since they have to take in account only the parameters regarding the “external” -attraction to the goal, interaction with the walls- interaction of the others, and not those regarding the “internal” -agent to agent- interaction). Since evolution is performed using non homogeneous population, it is possible that the evolutionary process is more difficult for high levels, and that their performance would be equal to or even better than that of level 1 agents if we used clones for evolution.

We also tried to test the performance of the evolved $l = 1$ agents as $l = 3$ agents (i.e. keeping the continuous parameters fixed and changing l) but the results were very poor (the fitness dropped to a value next to 0), showing the importance of the adaptation of the continuous parameters to the l value (co-evolution of l and of the continuous parameters).

5.9 Evolution of ToM recursion level

5.9.1 Co-evolution of l and continuous parameters

Finally we have performed a few simulation to study how the ToM level l evolved in a population that was heterogeneous both in the continuous parameters and in the ToM level. At the first stage of evolution the continuous parameters were randomly chosen in the same ranges used for the previous simulations, while l was randomly chosen between 6 integers going from 0 to 5. The genetic algorithm was performed in the same way as before, just adding l to the parameters that were inherited, crossed and mutated. Random mutation could change l to $l \pm 1$ and $l \pm 2$. We allowed also a jump of two ToM levels because we have verified that, due to the difference between

5.9 Evolution of ToM recursion level

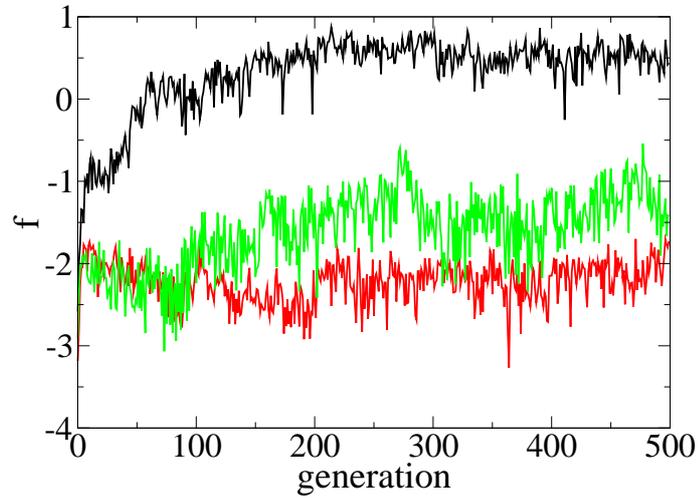


Figure 5.21: Left: Fitness through evolution (average over 5 repetitions) for $l = 1$ (black), $l = 2$ (red) and $l = 3$ (green) agents.

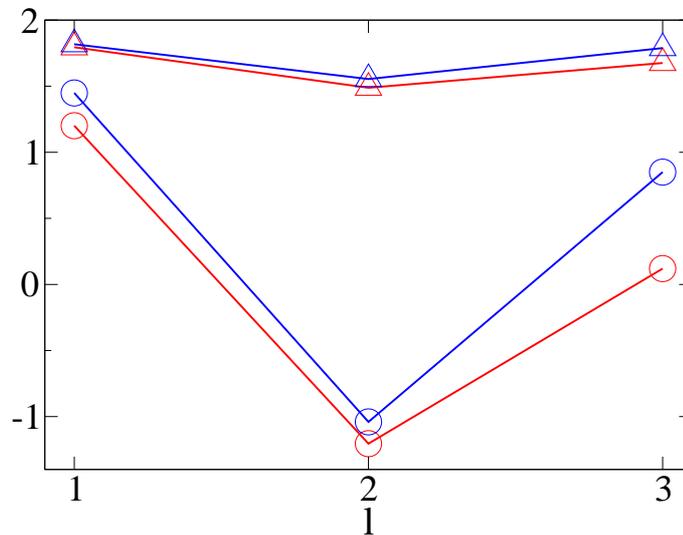


Figure 5.22: Performance of the best agents as a function of ToM level l ($\rho = \rho_1$, averages over 10 repetitions). Blue lines show the results of a family of clones of the best (average) individual, while the red ones show the results of the best heterogeneous family. Triangles show the values of the velocity towards the goal v_g , while circles those of the fitness f .

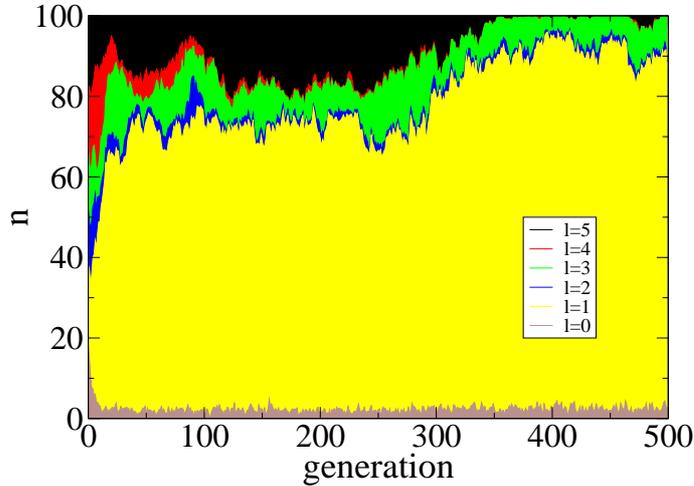


Figure 5.23: Evolution of ToM level for $\rho = \rho_1$, $\beta = 1$. The coloured surfaces stand for the number (percentage) of agents with a given value of l . The l values are shown from bottom to top in increasing order: brown stands for 0, yellow for 1, blue for 2, green for 3, red for 4 and black for 5. These results are the average over 5 different repetitions of the experiment.

even and odd levels, one step mutations resulted always to be bad mutations. We can summarise our results in this way: high enough density and $\beta \approx 1$ led to an invasion of all the population by level 1; high density and $\beta \approx 0$ or low density and $\beta \approx 1$ showed no clear prevalence of any value of $l \neq 0$, but each particular evolutionary history showed the prevalence of even or odd levels; low density and $\beta \approx 0$ led to the invasion of the population by even $l \neq 0$ levels.

In particular, figure 5.23 shows the typical result for high density ($\rho = \rho_1, \rho_2, \rho_3$) and $\beta \approx 1$. Even levels (brown, $l = 0$; blue, $l = 2$; and yellow, $l = 4$ in the figure), quickly disappear from the population. $l = 5$ has an important role in the first stages of evolution, but it also disappears as $l = 1$ invades the population. The prevalence of $l = 3$ over $l = 5$ in the last stages of evolution is probably due to mutations from 1 to 3 (mutations from 1 to 5 are not possible).

Figure 5.24 shows the evolution of the ToM level for $\rho = \rho_3$, $\beta = 0$, with no clear prevalence of any value of l (even if $l = 1$ occupies a predominant part of the population, at least in the last stages of evolution).

5.9 Evolution of ToM recursion level

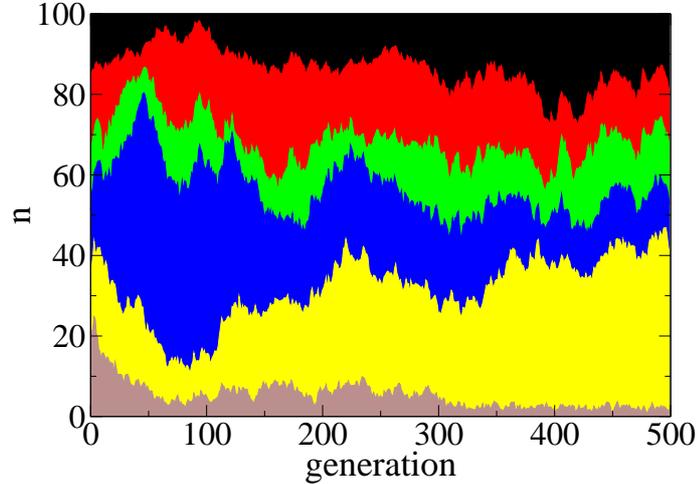


Figure 5.24: Evolution of ToM level for $\rho = \rho_3$, $\beta = 0$, average over 10 repetitions. See figure 5.23 for legend and explanation.

These results, that show the coexistence of all the levels (with the exception of $l = 0$), were obtained as the average of 10 different repetitions of the experiment, but actually single repetitions show the prevalence of a few levels, and never show the coexistence of even and odd levels. For example, in figure 5.25 we show the results of 2 repetitions of the experiment, the first one dominated by odd levels, the second one by even levels. The emergence of even or odd levels had no effect on the overall fitness of the agents, which resulted to be almost the same in the two cases.

The reason of the impossibility of coexistence between even and odd levels can be found in an analysis of the values of the evolved continuous parameters. As stressed before, these values are not l independent, and in particular have a strong dependence on the parity of l . For example, in figure 5.26 we show a comparison between the evolution of the perceived size of the body D_0 in a “odd level dominated” and in a “even level dominated” population (average over the same repetitions shown in figure 5.25). The difference in the values of the continuous parameters makes transitions between even and odd levels very difficult to happen.

The same (qualitative) results in figures 5.24, 5.25, 5.26 have been obtained using $\rho = \rho_4$, $\beta \approx 1$, while for high values of ρ and $\beta \approx 0$ we observed coexistence between even and odd levels in the same repetition of the exper-

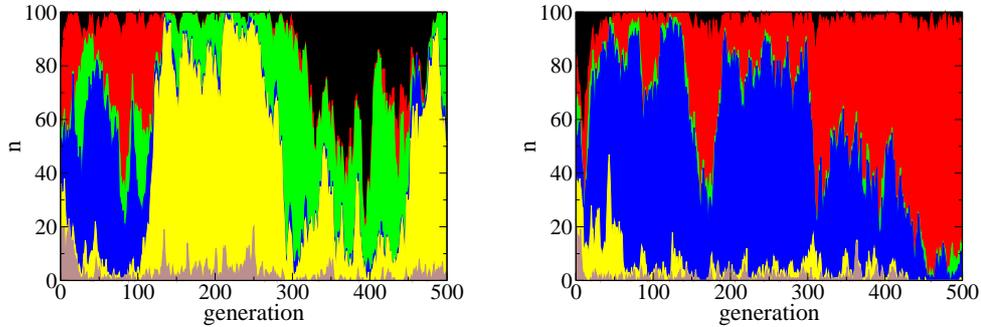


Figure 5.25: Evolution of ToM level for $\rho = \rho_3$, $\beta = 0$, in two different repetitions. Left: after around 100 generations even levels disappear, and odd ones (yellow, $l = 1$; green, $l = 3$; black, $l = 5$) invade the population. Right: even levels (with the exception of 0) invade the population, with a gradual shift from $l = 2$ (blue) to $l = 4$ (red).

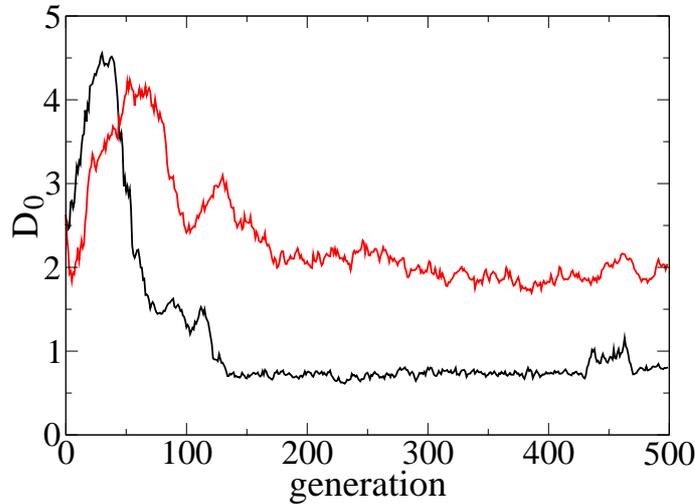


Figure 5.26: Evolution of D_0 for a population dominated by odd levels (black line, corresponding to the simulations at left in figure 5.25), and for a population dominated by even ones (red line, corresponding to the simulations at right in figure 5.25). While for odd levels D_0 converges to a value slightly inferior to $2r$, it converges to $\approx 4r$ for even levels, making the transition between even and odd levels hard to happen.

5.9 Evolution of ToM recursion level

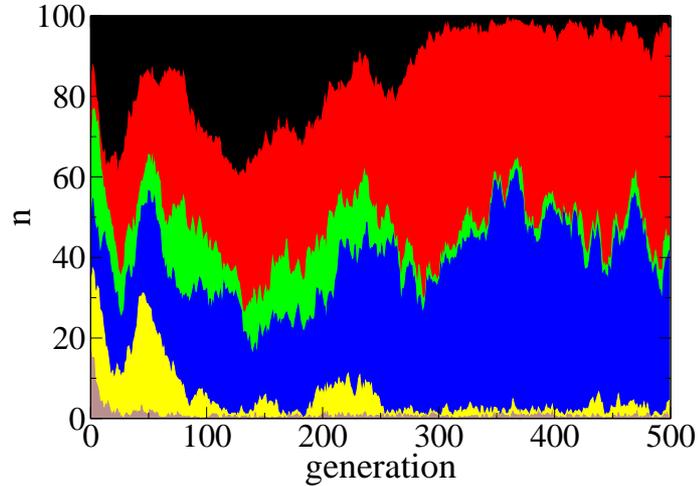


Figure 5.27: Evolution of the ToM level for $\rho = \rho_4$ and $\beta = 0$, averaged over 5 repetitions of the experiment. $l = 2$ (blue) and $l = 4$ (red) invade the population.

iment, but this result was due to the incapability of the agents evolved with these values of the parameters to develop an appropriate behaviour (see the discussion about the role of β in evolution, section 5.7.1: agents evolved with $\beta \approx 0$ and $\rho = \rho_1$ are not able to learn the size of their body, and thus the discussion following figure 5.26 does not hold).

A different situation was present for very low density $\rho = \rho_4$ and $\beta = 0$. The results shown in figure 5.27, averaged over 5 repetitions of the experiment, show that even levels invade the population in each repetition. Probably the “bold” behaviour of even levels is more apt to the solution of the problem (the disappearance of $l = 0$ shows that the problem could not be solved just “going straight to the goal”). Nevertheless, an analysis of the evolution of the perceived size of the body shows that D_0 converges to a value quite higher than $2r$, which means that evolution corrected the behaviour of even level agents in order to make it more “careful”.

5.9.2 Evolution of l keeping fixed the other parameters

As we said before, one of the problems that we had to face in reproducing the results described in section 5.4 was due to the coevolution of l (recursion

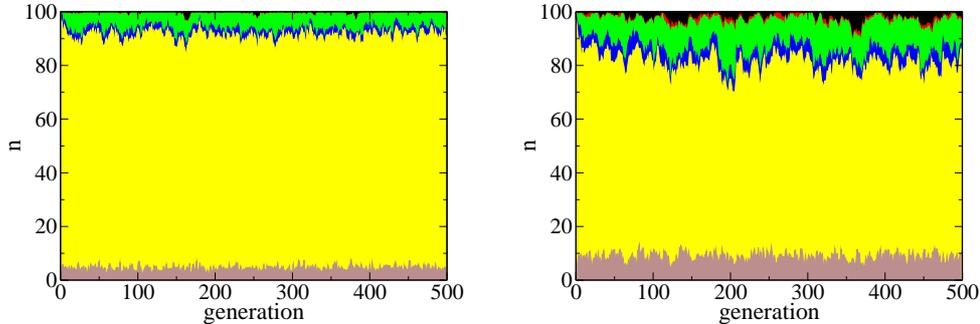


Figure 5.28: Evolution of ToM recursion level keeping fixed the continuous parameters. Left: using $\rho = \rho_1$, $\beta = 1$; right: using $\rho = \rho_2$, $\beta = 0$. The results were obtained as averages over five repetitions of the experiment. See figure 5.23 for the legend and explanation of the correspondence between colours and l values.

ToM level) and of the continuous parameters. For this reason we tried to split the evolutionary process in two steps, a first one in which we evolved the collision avoiding mechanism (i.e. the continuous parameters) keeping $l = 1$ fixed, and a second one in which we evolved l .

We thus used as a starting population the family of $l = 1$ clones that had developed the most efficient collision avoiding mechanism (more precisely, we use clones of the average individual of a family with the highest fitness through evolution, obtained using $\beta = 1$ and $\rho = \rho_1$), and performed an evolutionary simulation in which the continuous parameters were kept fixed while l could mutate (with probability $p_m = 0.1$), to see for which values of β and ρ high levels could invade the population.

When we evolve the recursion level in the same $\beta = 1$ and $\rho = \rho_1$ environment in which the continuous parameters of the collision avoiding model had been evolved, levels $l \neq 1$ are produced just as direct mutations from $l = 1$, but their fitness is always very low and they cannot occupy a considerable part of the overall population (figure 5.28 left). As density assumes lower values ($\beta = 0$ and $\rho = \rho_2$, figure 5.28 right), the fraction of $l \neq 1$ agents becomes higher, and also the $l = 4$ and $l = 5$ agents, that cannot be obtained by direct mutation from the $l = 1$ ones, appear, but the population is still almost entirely composed of $l = 1$ agents.

Figure 5.29 shows the results obtained for $\rho = \rho_3$. In the $\beta = 0$ case (right)

5.9 Evolution of ToM recursion level

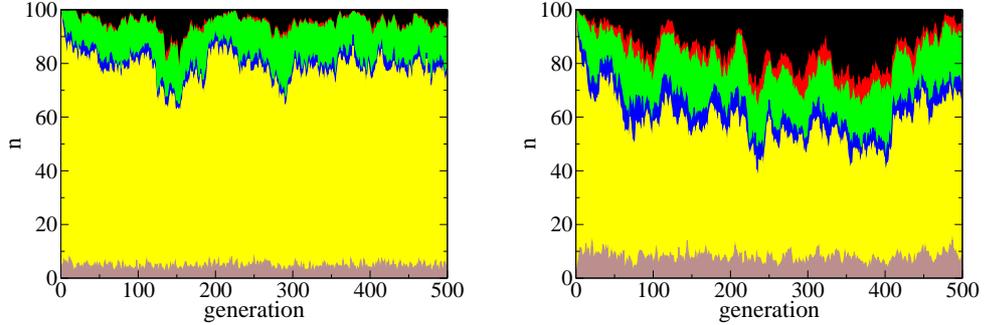


Figure 5.29: Evolution of ToM recursion level keeping fixed the continuous parameters. Left: using $\rho = \rho_3$, $\beta = 1$; right: using $\rho = \rho_3$, $\beta = 0$. The results were obtained as averages over five repetitions of the experiment.

$l = 1$, although still in a dominant position, occupies only half population. Using $\rho = \rho_4$ (figure 5.30) the high l levels invade the population, and in particular the highest allowed level ($l = 5$) has a predominant role, while $l = 1$ is reduced to a low percentage. When $\beta = 1$ (left) only odd levels are present, but if we further diminish the role of collision using $\beta = 0$ also even levels (and in particular $l = 4$) start to occupy a considerable portion of the overall population (the difference between figure 5.27, in which even levels invaded the population, and figure 5.30 at right is due to the use, in the second case, of agents whose continuous parameters had been evolved for an odd $l = 1$ recursion level).

These results confirm substantially those described in section 5.4. When agents have a strong need of avoiding collisions, $l = 1$ occupies almost the entire population, but as the role of collisions on the fitness gets smaller, gradually high (odd) recursion levels start to have a predominant role. We also found that further diminishing the role of collisions, even levels, starting from the high ones, begin to occupy a considerable portion of the population. (We used as control parameters for these simulation ρ and β as we did for all the previous simulations, following the rule that increasing these parameters increased the role of collisions in evolution. Using the right value of β resulted crucial in order to evolve the collision avoiding mechanism continuous parameters, but probably is not very important for the evolution of the recursion level alone, and thus a future throughout investigation of this phenomenon could be simplified using only ρ as a control parameter).

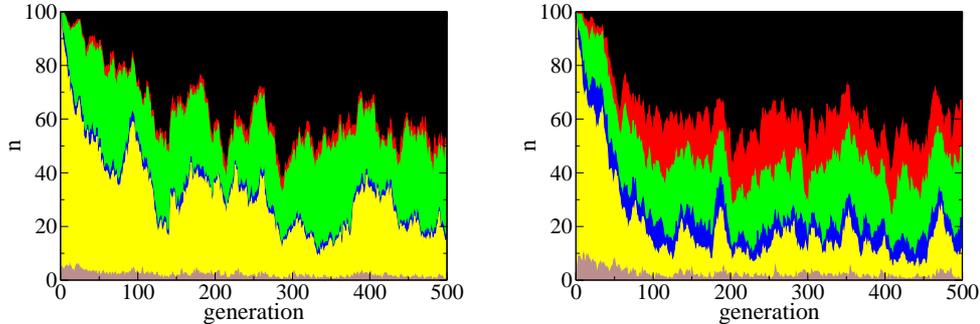


Figure 5.30: Evolution of ToM recursion level keeping fixed the continuous parameters. Left: using $\rho = \rho_4$, $\beta = 1$; right: using $\rho = \rho_4$, $\beta = 0$. The results were obtained as averages over five repetitions of the experiment.

5.10 Conclusions

We have proposed an evolutionary model to simulate crowd dynamics, using agents with a finite size impenetrable body. We have shown that using a simple fitness function, which assigned a positive value to the velocity towards the goal, and a negative one to momentum exchanged in collisions, we could evolve the same self organising behaviour observed for actual pedestrians. Agents were able to self organise even in quite crowded environments, but in order to do that it was necessary to give them a high enough sense of “pain”, i.e. to give an high enough negative weight to collision in the fitness function, to prevent agents to be kept in a local maximum that did not allow them to attain proper self organisation.

We think that a further study of the dynamics of our model is necessary. In particular, we would like to investigate which is the range of the agents’ density that gives rise to self organising patterns. We have verified that self organisation is not present at very low densities, while cannot occur at too high densities. We have also verified that the structure of patterns depends on the density, length and width of corridors (as figure 5.13 clearly shows), but we are still missing a throughout quantitative investigation. We would like also to compare our results to those obtained by other models and to perform a quantitative comparison to actual pedestrian dynamics. It would be extremely interesting to understand which are the minimal requests to have a model that presents these self organisation patterns.

5.10 Conclusions

The best results, regarding the optimisation of traffic flow, have been obtained using $l = 1$ and $\delta \gg 1$. Under these conditions the model can be easily optimised, and we intend to use it as part of a software that simulates crowd dynamics in complex and realistic topologies.

We also performed, using the same model, an evolutionary study of the emergence of recursive thinking for agents moving in a crowd and trying to avoid collisions. We verified that there is a difference between even levels, which have a “bold” behaviour, and odd ones, that show a “careful” behaviour, and due to this difference some of the parameters of the model (for example the “learnt” or “perceived” size of the agent’s own body) evolve to quite different values for odd or even levels.

Our results seem to show that actual recursive thinking cannot emerge from evolution of heterogeneous agents in a crowded environment, since $l = 1$ invades the population. Even if an accurate analysis of the parameter space has not been possible, our results seem to show also that for some values of these parameters (in particular when the agents’ density is quite low) higher recursion levels are adaptive. This confirms the results of a previous paper in which the authors used an idealised collision avoiding mechanism to simulate social interactions, and verified that choosing an appropriate range of interaction high recursion levels were adaptive. This is probably true also in our model, even if that parameter area is of no great interest for crowd dynamics, since it corresponds to very low densities.

Even if recursive thinking is not adaptive for crowd dynamics, actual pedestrians are capable of recursive thinking, and this ability could have some effect on the way they move in a crowd. We have supposed that high (odd) recursion levels could perform efficiently also in crowded environments, but that is very hard to evolve them in heterogeneous populations. This hypothesis has not still been tested for computational reasons, but could be used to obtain efficient high level agents. A future development could consist in comparing the performance of actual pedestrians with those of evolved ones, to see if their behaviour shows features which are typical of high recursion levels, or of non homogeneous populations.

We also believe that in order to perform an evolutionary study of the emergence of recursive thinking in a realistic setting, it is necessary to introduce more complex social interactions and tasks. We suppose that introducing some level of cooperation (for example, requiring agents to move in little groups inside the crowd) could require the emergence of high recursion level,

An evolutionary crowd dynamics and ToM model

and probably also break the asymmetry between even and odd levels which represented a problem in the co-evolution of the continuous and recursion level parameters in our model.

5.10 Conclusions

Bibliography

- [1] D. Helbing, P. Molnár, I. J. Farkas and K. Bolay, “*Self-organizing pedestrian movement*”, Environment and planning B: Planning and design, v. 28 361-383 (2001).
- [2] D. Helbing, L. Buzna, A. Johansson and T. Werner, “*Self-Organized Pedestrian Crowd Dynamics: Experiments, Simulations, and Design Solutions*”, Transportation Science, Vol. 39, No 1, 1-24 (2005).
- [3] K. Ando, H. Oto and T. Aoki, “*Forecasting the flow of people*”, Railway Res. Rev. 45(8) 8-13, (1988) (in Japanese).
- [4] D. Premack and G. Woodruff, “*Does the Chimpanzee have a Theory of Mind?*”, The Behavioral and Brain Sciences, 1:515-523 (1978).
- [5] D. Dennet, “*The Intentional Stance*”, MIT Press (1987).
- [6] A. Miklósi, R. Polgádi, J. Topál and V. Csányi, “*Use of experimenter-given cues in dogs*”, Animal Cognition, 1, 113-121 (1998).
- [7] A. Tschudin, J. Call, R. Dunbar, G. Harris and C. Van der Elst, “*Comprehension of signs by dolphins (Tursiops truncatus)*”, Journal of Comparative Psychology, 115, 100-105 (2001).
- [8] J. Kaminski, J. Riedl, J. Call and M. Tomasello, “*Domestic goats (Capra hircus) follow gaze direction and use social cues in an object choice task*”, Animal Behaviour, 69, 11-18 (2001).
- [9] T. Bugnyar, M. Stöwe and B. Heinrich, “*Ravens, Corvus corax, follow gaze direction of humans around obstacles*”, The Royal Society, 271, 1331-1336, (2004).

BIBLIOGRAPHY

- [10] M. C. Corballis, “*Recursion as the key to the human mind*”, in K. Sterelny, J. Fitness (eds.), “*From mating to mentality: evaluating evolutionary psychology*”, 155-171, (2003). Psychology Press, New York.
- [11] C. Zimmer, “*How the mind reads other minds*”. *Science*, 300(16) 1079-1080, (2003).
- [12] M. Takano, M. Kato and T. Arita, “*A Constructive Approach to the Evolution of the Recursion Level in the Theory of Mind*”, *Cognitive Studies*, 12(3), 221-233, (2005) (in Japanese).
- [13] M. Takano and T. Arita, “*Asymmetry between Even and Odd Levels of Recursion in a Theory of Mind*”, *Proc. of ALIFE X (10th International Conference on the Simulation and Synthesis of Living Systems)*, 405-411, (2006).
- [14] Barkow, L. Cosmides and J. Tooby, “*The adapted mind: evolutionary psychology and the generation of culture*”, Oxford University Press, (1992).
- [15] R. Byrne and A. Whiten (eds.), “*Machiavelian Intelligence: Social Expertise and the Evolution of Intelligence in Monkeys, Apes and Humans*”, Oxford University Press, (1988).
- [16] R. Dunbar, “*On the origin of the human mind*”, in P. Carruthers and A. Chamberlain (eds.), “*The evolution of mind*”, 238-253. Cambridge University Press, (2000).
- [17] P. Kinderman, R. Dunbar and R. P. Bentall, “*Theory-of-mind deficits and causal attributions*”, *British Journal of Psychology*, 89, 191-204, (1998).

Chapter 6

An evolutionary urban scale traffic model

6.1 Introduction

In this work we develop a multi-agent traffic model on a urban scale. The model is realised on a discrete grid, and agents are considered as point like, moving at discrete time steps from a site to another of the grid. This means that the interactions between agents and of agents with the environment are not described in detail, but just expressed using simple rules representing an adequate average of microscopic behaviours. Nevertheless, the model is an actual distributed autonomous system, since all the actions of the agents are performed on an individual base.

The model is aimed at understanding how self organisation patterns and communication between agents can emerge in a agent system in order to avoid a traffic flow problem. We study in particular two phenomena: the emergence of a communication system through which agents are able to choose the (time dependent) best path to reach their goals (a communication system which has been inspired by pheromone based trail laying and trail following in social insects, as ants), and the emergence of a global traffic flow rule, independent of communication and related to the geometry of the mobility system (the urban area) that leads to a minimisation of traffic flow problems.

6.2 Social insects

According to the common opinion between the researchers in the field, social insect colonies are a remarkable example of distributed autonomous system. Insects living in colonies (ants, termites, and many species of wasps and bees) are capable of accomplishing tasks whose complexity has often surprised and fascinated naturalists. Bonabeu et al. ([1, 2]) briefly expose and report a large literature about these activities, in particular those related to foraging, nest building and maintenance, and also the ability of task allocation, i.e. the (dynamic) division of the population in groups performing different duties.

6.2.1 Some tasks performed by social insects

To cite a few examples, army ants are able to organise hunting raids that involve up to 200,000 workers, while leaf cutter ants are able to forage at hundreds of meters from the nest, organising “highways” between the nest and the foraging site. Weaver ants can form chains with their own bodies in order to cross wide gaps or to accomplish results that could not be obtained by a single worker, while honey bees form chains to induce a local temperature increase and shape more easily wax combs. Nest building reaches very high levels of complexity in tropical wasps and termites.

In a social insect colony, usually a worker does not perform all the tasks, but is specialised. Often this specialisation corresponds to morphological differences between individuals, as, for example, the coexistence of minor and major workers, the first ones performing brood feeding and nest maintenance duties, the second ones defending the nest or cutting large preys. Nevertheless, social insects are capable of dynamic task allocation, modifying the size and composition of groups of workers involved in different tasks to respond to environmental changes and changing colony needs.

6.2.2 Self organisation

It was once believed that such a complex colony behaviour needs some kind of supervision or hierarchical structure, for example the queen was supposed to collect and centralise information and give direct orders to control the activities of the colony, but recent observations show that the activities are almost totally unsupervised. Even if an insect is surely, from a microscopic point of view, a quite complex creature (for example, if compared with the

agents of virtual models described in this thesis), since it can process a large number of sensory inputs and stimuli and perform decisions on the base of a large amount of information, it is today widely believed that the complexity at the level of the single agent is not sufficient to account for the complexity of the whole system, and thus that the complex behaviours of social insect colonies are due to self organisation phenomena, i.e. emerge as the result of simple interactions between the agents.

Following [2] we can say that self organisation, i.e. the appearance of global level structures on the base of local interactions between elemental units, relies on four basic ingredients: *positive feedback*, *negative feedback*, *amplification of fluctuations* and *multiple interactions*; and is characterised by a few key properties, as the creation of a *spatiotemporal structure* over a previously homogeneous medium, *multistability*, and the existence of *bifurcations* under the variation of parameters.

These properties can be illustrated with the help of an example, the construction of pillars in termites. Termites use soil pellets impregnated with pheromone (a chemical used by animals to communicate) to build pillars [3]. Two phases take place: a non-coordinate one during which pellets are randomly deposited, and a coordinated one, which starts when one of the deposits is large enough and stimulates workers to accumulate more material. This is a phenomenon of *positive feedback* and *amplification of fluctuations*: initially the medium is homogeneous, but a fluctuation makes one of the deposits get larger than the others, and then this deposit attracts more workers due to the presence of a larger amount of pheromone. The bigger the deposit becomes, the more it attracts workers, and finally a pillar emerges. To have the formation of a pillar, the number of workers has to be higher than a threshold, because if too little workers are present, the pheromone will evaporate between two passages of the workers. This is an example of *negative feedback*, since pheromone evaporation limits “the snowball effect” due to positive feedback, and also of *bifurcation*, since spatial density of termites acts as a parameter for the emergence of pillars. This macroscopic *structure* emerges over a medium that was initially uniform, and it is obviously *not unique* (it could have emerged in a different location). This process is based on *multiple interactions*, since the emergence of the pillar is not due to the “individual initiative” (the single worker does not seem to remember where pellets have been located, nor to follow a given plan, but just to react to the actions of other workers).

6.2 Social insects

6.2.3 Stigmergy

In particular, in this pillar construction process, we do not see direct termite to termite interaction, but just the response by the individual workers to particular features of the environment (pellet deposits size and corresponding pheromone densities), features which are determined by the actions of other termites. This form of indirect communication through modification of the environment is an example of stigmergy, a term which was introduced by Grassé [3] to describe how, in termite nest building, the actions of the individual termites are not determined by direct communication between workers, but by the nest structure (which is determined by the actions of the workers). Stigmergy thus means communication through actions (on the environment).

6.2.4 Pheromone based trail following in ants

Another example, more closely related to the model that we are going to expose, is that of trail tracing and following in foraging ants, which is also based on pheromone communication. Deneunburg et al. [4] showed that path selection to a food source in the Argentine ant *Linepithema humile* is based on self organisation. In their experiment, a food source is separated from the nest by a bridge with two equally long branches. Initially the two branches are equally probable to be taken, since there is no pheromone on them. Nevertheless, random fluctuations will cause one of the branches to be chosen by a few more ants, and thus to have a little more pheromone. The positive feedback due to this pheromone difference causes more ants to choose that branch, which finally turns out to be the only exploited one. This is a nice example of *multistability*: the two paths were initially equally probable to be chosen, but the system always finishes in one of the two stable states.

Deneunburg et al. [4] developed a simple mathematical model that could describe the results of the experiment. If we denote with A and B the two branches, with $A_i + B_i = i$ the number of ants that have chosen branches A and B , the probability P_A (P_B) that ant $i + 1$ chooses branch A (B) is

$$P_A = \frac{(k + A_i)^n}{(k + A_i)^n + (k + B_i)^n} = 1 - P_B \quad (6.1)$$

The probability to choose a given branch grows with the number of ants that have already chosen it, i.e. with the amount of laid pheromone (assuming no evaporation, which means that, as it actually was the case, the time scale of pheromone evaporation was considerably higher than that of the experiment). k quantifies the degree of attraction of an unmarked branch, while n determines the non-linearity of the choice function (the values of the parameters that gave the best fit to the experiment were $k \approx 20$ and $n \approx 2$).

If the two branches were of different length, the first ants to come back from the food source would be the ones that took two times the shortest path, causing at that time a larger amount of pheromone to be on that path, and thus influencing the system to converge to a state in which only the shortest path is exploited. Experiments show indeed that the chance of the shortest path to be selected grows with the ratio r between the branches.

One of the most interesting features of the behaviour of social insects colonies is their ability to adapt to environmental changes (see for example the discussion about task allocation), but in this particular case the ants were not able to switch to the short path in case this one was presented to the colony after that the longest one had already been chosen (nor they were able to switch to the shortest path in case that initial fluctuation led to the choice of the longest one). A different species of ants, *Lasius niger* [5], resulted to be able to choose systematically the shortest path, but this ability was due to the combination of pheromone based *mass recruitment* (a term that refers to a process whereby an ant chooses a path to a source food on the base of a chemical) and individual memory.

6.2.5 Ant colony optimisation

The ability of social insect colonies of solving complex and time changing problems, in a flexible and robust way, using simple and decentralised units, has attracted the interest of researchers in the field of intelligent system design, since many of these tasks have a counterpart in engineering and computer science (see [2] for a throughout exposition of the subject).

A classical example of a “social insect inspired” algorithm is the application of *ant colony optimisation* (i.e. the use of pheromone trail laying and following techniques) to the travelling salesman problem (TSP).

TSP is a path optimisation problem whose goal is to find a closed tour of minimum length connecting n given cities, each one being visited once and

6.3 Description of the model

just once. We do not enter in the details of how this problem can be solved with the aid of ant colony optimisation techniques (see for example [6, 7]) but we just stress that in order to solve this complex task, it is necessary to avoid to be trapped on suboptimal solutions, as could happen for the simpler task of choosing the shortest path from the nest to the food source for *Linepithema humile*. In ant colony optimisation this is usually accomplished by introducing negative feedback as pheromone evaporation over short time scales. Since a well marked pheromone trail is harder to maintain if the path is longer, pheromone evaporation allows the colony to exploit systematically the shortest path [2].

It is interesting to notice that in actual ant systems pheromone is persistent, i.e. evaporation is insignificant on relevant time scales. Biological systems can be an useful inspiration in computer science and engineering problem solving, but the latter do not suffer from biological limitations and do not need to be biologically plausible (in contrast with modelling, i.e. with the use of computer models in order to understand and explain biological systems). In general biological evolution has to solve very complex problems, in which optimisation undergoes many constraints, including ecological ones: for example, in the case of *Linepithema humile* it is probable that sticking to a known even if suboptimal solution is more efficient from the point of view of trail protection (from competition with other colonies or predation) than further seek for a better solution (see again [2]).

6.3 Description of the model

The model that we are going to expose is to be intended basically as an abstract model that studies how agents located in a mobility system with a traffic problem can develop strategies (in particular based on communication) aimed to optimise traffic flow. Since we want the model to be simple and general (i.e. we want it to capture the basic features of any mobility system more than describe an actual mobility system, nor a particular class of mobility systems), we decided to use point like agents moving at discrete time steps on a discrete spatial grid. This setting allows us to avoid a detailed description of the agent to agent dynamics, which is expressed by simple rules based on the number of agents on each site, without renouncing to the individual agent (distributed autonomous) approach. For the same need of simplicity, communication is not introduced in the model as a direct agent

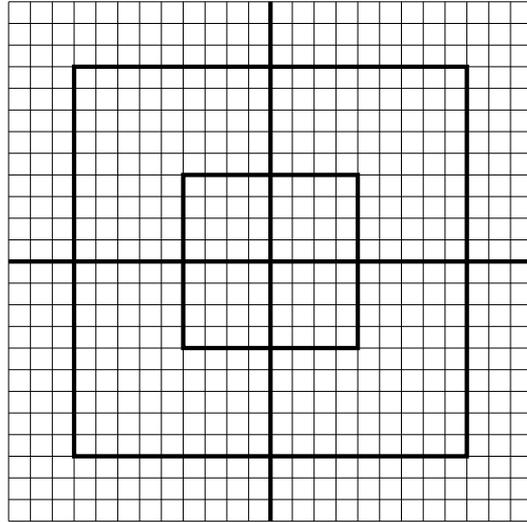


Figure 6.1: Basic structure of the mobility system. Thick lines correspond to “wide roads”, i.e. to sites on which agents are allowed to have an higher “speed” (probability to move).

to agent interaction, but as indirect communication through a modification of the environment (stigmergy).

Nevertheless we provide also an interpretation according to which our model corresponds to a given class of actual mobility systems, studied on a space-time scale some order of magnitude higher than that of actual agent to agent interactions, which are thus substituted by suitable averages.

6.3.1 Dynamics of motion on the grid

Structure of the grid

Our mobility system is constituted by a Manhattan square grid (figure 6.1). The grid can in principle be non homogeneous, i.e. the agents are allowed to have a different maximum speed on different sites of the grid (we define below what is meant by an agent speed and maximum speed). If that is the case, sites on which an higher speed is allowed are disposed in rows (thick lines in figure 6.1) in order to form “highways”.

6.3 Description of the model

Definition of an agent's velocity

At each time step agents are located on a site of the grid. In principle, any number of agents can be located on the same site, whose state is thus identified by the number of agents on it $n(i, j)$ (where (i, j) denotes the coordinates of the site). At each time step agents can perform one of five possible actions, whose result is to move on a first neighbour site, or stand still on the same site. Since actions are always performed at discrete time steps, and their result is a space step displacement or no displacement, the speed of the agent, if defined as the ratio between displacement and time, would always be equal to 1 or 0 (in adimensional space time grid units). To allow agents to have different velocities, speed is defined in our model on a probabilistic base. In particular we will define the maximum allowed speed of an agent on a given site of the grid as its maximum probability to perform an action, or $0 \leq w_{ij} \leq 1$ (since it is a site dependent quantity, we will refer to it also as the “width” of site (i, j)).

We refer to w_{ij} as the *maximum* speed or probability for two reasons: first of all, w is neither the probability to move nor the probability to move in a given direction, just the probability to perform one of the five possible actions. Which action will be performed is decided by the agent decision mechanism, which is described below. The second reason is that w gives the probability for the agent to move just in the case it is located alone on the site, i.e. if $n(i, j) = 1$.

Probability to act and the effect of crowding

In general we define, for any occupied site, a traffic function $0 \leq t_{ij}(n(i, j)) \leq 1$ as

$$t_{ij}(n(i, j)) = \left(\frac{1}{n(i, j)} \right)^{\gamma_{ij}} \quad (6.2)$$

where $\gamma_{ij} \geq 0$ is a (site dependent) traffic factor. $t(n)$ determines the dependence of the velocity of the agent on the overall number of agents on the site. It is defined in such a way that if the agent is alone on the site we have $t(1) = 1$, for any γ value. We have $t(n) = 1$ also when $\gamma = 0$, regardless of the n value. When $\gamma > 0$, t is a decreasing function of n and it is a decreasing function of γ at fixed n (we remind that $n \in \mathbb{N}$, but in the $n = 0$ case it is meaningless to define the probability to perform an action since no agent is present).

To summarise, we define the probability to perform an action or speed of an agent located in site (i, j) as the product of the width or maximum speed of that site with the traffic function

$$p_{ij}(n(i, j)) = w_{ij}t_{ij}(n(i, j)) \quad (6.3)$$

6.3.2 Pheromone

At each time step an agent “tries” to perform an action, and succeeds with probability p (equation 6.3). Depending on the success of this attempt, the agent drops, on the site on which it tried to perform the action, a given amount of pheromone, namely D_0 in case of unsuccess and D_1 in case of success (regardless of the performed action). Following [8], we define the dynamics of the dropped and the evaporated pheromone fields. The dynamics of dropped pheromone is given by

$$D(t + 1, i, j) = (1 - \beta)D(t, i, j) + \Delta D(t, i, j) \quad (6.4)$$

where by $D(t, i, j)$ we mean the amount of dropped pheromone on site of coordinates (i, j) at time t , by β the evaporation rate coefficient of dropped pheromone and by $\Delta D(t, i, j) = \sum_{k \in n} D_l^k$ the sum of pheromone dropped by the n agents located on site (i, j) at time t as the result of their success ($D_l = D_1$) or unsuccess ($D_l = D_0$) of their attempt to perform an action.

The dynamics of evaporated pheromone is given by a discrete diffusion equation

$$\begin{aligned} P(t + 1, i, j) = & \beta D(t, i, j) - 5\delta P(t, i, j) \\ & + \delta(P(t, i - 1, j) + P(t, i + 1, j) + P(t, i, j - 1) + P(t, i, j + 1)) \end{aligned} \quad (6.5)$$

where $P(t, i, j)$ is the amount of evaporated pheromone on site (i, j) at time t , and δ is the diffusion rate coefficient (the factor 5 accounts for diffusion in the vertical direction, i.e. for pheromone which leaves the plane). We allow in principle the D and P fields to assume also negative values, i.e. we let $D_0, D_1 \in \mathbb{R}$. This fields could be interpreted as the difference between two different pheromone fields, or, more properly (since this is a model inspired by a biological system but not a model of a biological system) simply as a diffusing fields with sign that our agents use to communicate.

6.3 Description of the model

6.3.3 Decision mechanism and its evolution

At the initial time, agents are located in a starting point $\mathbf{s} = (s_x, s_y)$, and have to reach a goal point $\mathbf{g} = (g_x, g_y)$.

The decision mechanism, applied, as we said before, with probability $p_{ij}(n_{ij})$ at each time step, is based on a fully connected neural network with an hidden layer (see figure 6.2 and discussion below) that takes as 6 inputs the distances in the x and y coordinates of the agent to its goal, and the “gradients” of the (evaporated) pheromone field in the direction of the 4 neighbouring sites, that is, defining $\mathbf{x} = (x, y)$ as the site in which the agent is located at the moment of decision, the six inputs I_i are given by

$$\begin{aligned} I_1 &= g_x - x \equiv \Delta_x \\ I_2 &= g_y - y \equiv \Delta_y \\ I_3 &= P(i, j) - P(i + 1, j) \\ I_4 &= P(i, j) - P(i - 1, j) \\ I_5 &= P(i, j) - P(i, j + 1) \\ I_6 &= P(i, j) - P(i, j - 1) \end{aligned} \tag{6.6}$$

The five outputs of the network correspond to the five possible actions, i.e. to move to a first neighbouring site or stand still, and the action with the highest output value is deterministically performed.

The decision mechanism of agents is evolved using a genetic algorithm. Each generation of the genetic algorithm is composed of various runs. In each run, to any agent are assigned randomly chosen start and goal points (which are requested to be at least at a given distance one from the other), in order to generalise the ability of the agent to reach any point of the grid starting from any other point. The fitness of the agent is simply obtained as the ratio

$$\frac{|s_x - g_x| + |s_y - g_y|}{\tau} \tag{6.7}$$

of the Manhattan distance from the start to the goal over the time τ needed to reach the goal (fitness is zero if the goal is not reached), averaged over runs. Defining h as the number of nodes in the hidden layer, the genetic code of agents is given by the $7 \times (h + 1) \times 5$ connections between the nodes and biases of the network, connections that are directly expressed as real numbers. As shown in figure 6.2, h is not fixed but is determined by evolution. The evolution of the neural network structure has been performed using the

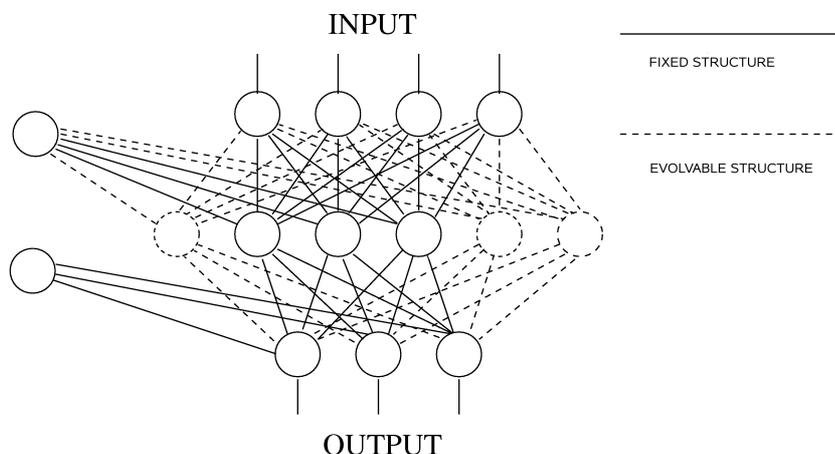


Figure 6.2: Evolvable structure of the neural network. The algorithm develops in parallel different structures and chooses the most apt to solve the problem.

following strategy, which has been inspired by [9, 10] (that actually describe more complex strategies than the simple one we used). In the first generation agents were divided in n_s families or species, each family with a fixed number h of hidden neurons. For g_s generations, at the end of each generation, tournament selection is performed inside the boundaries of a single family, and thus the size of species remains fixed (in order to “protect” families whose evolution is slower) but any g_s generations there is a “revolutionary” tournament selection which is performed without respect to the family boundaries, and thus changes the size of species (the GA does not use a mutation operator on the structure of the networks, but just a selection operator. Furthermore, the selection operator is not applied at each generation but just any g_s generations). The size of tournament selection is then adapted to the size of the family.

We decided to use this method because in the model we use a very large number of individuals and thus of genetic codes and for this reason we wanted to develop in a parallel way different structures, without bothering about which structure was *a priori* more apt to a solution of the problem, nor need to repeat many simulations with a fixed number of hidden neurons to find the optimal solution.

Mutation is performed adding a Gaussian error to the values of the connections, while the number of hidden neurons h can be inherited but not

6.3 Description of the model

mutated, and thus there is not the problem of creating new connections. The use of feedforward neural networks with an hidden layer, and the parallel test of different numbers of hidden units, should allow us to explore a very large space of solutions regarding the choice of a path from the starting point to the goal (see chapter 2.3.2 for the general properties of feedforward networks).

6.3.4 An interpretation of dynamics and communication

As we said before the aim of the model is to introduce in a simple way a traffic flow problem and a communication system aimed to solve it. Our model introduces both a static (time constant, due to the site width w) and a dynamic (time changing, due to the traffic function $t(n)$) traffic flow problem. If agents are located on a “narrow” site, or are in a large number over the same site, their probability p to act is very low. Agents can signal the local state of the grid through a “mark” left on the the environment, a mark that depends on their individual experience (their success in trying to perform an action), and this mark is an information that can reach also distant agents through pheromone evaporation and diffusion. We do not provide agents with an ability to use this information (there is not in our model an equivalent of equation (6.1) in which pheromone concentrations are directly connected to path choices), and let evolution make such an ability emerge (together with, as we will see, other communication unrelated traffic flow optimising behaviours).

Nevertheless we can provide an interpretation in which our model can be seen as a large scale model for a more realistic mobility system. Let us consider a group of agents with a finite body size moving on a network with finite carrying capabilities (two examples could be the road system and the underground system of an urban area). If a large number of agents is located in the same portion of the system, traffic turns slower (too much cars in the same road or crossroad cause a traffic jam; if too much passengers are in an underground station part of them cannot get on the train). If our purpose is to describe the overall behaviour of the system, usually we do not want to or cannot describe in detail the physical interactions between agents and of agents with the medium, and prefer to use just point like agents moving on a discrete grid. Each site of the grid represents a portion of the

actual system whose dimension are a few orders of magnitude higher than the physical dimension of agents (for example, a crossroad and a portion of the roads around it; or an underground station and the portion of railway to the neighbouring stations), and the physical interactions are replaced by simple rules that express the fact that with growing densities the motion slows down, and also eventual spatial differences in the medium (the $t(n)$ and w in our model).

“Pheromone” communication can be interpreted in the following way: agents in our model are not provided with memory (for the sake of simplicity; they could be viewed as strangers, who have never been in town before) and are not provided with any centralised information system, and thus in order to move in an efficient way they have to communicate with other agents. But since the model is a large scale model that does not describe actual agent to agent interaction, but only averaged interactions through simple rules, it also does not introduce direct communication. We can think that every time the agent experiences something about the state of the system, it “drops a comment” (dropped pheromone) about it, a comment that turns into a rumour and spreads around (evaporated pheromone) until it is picked up by another agent (as a neural network input). Which is the medium over which rumours-pheromone diffuses? We can think that in the actual mobility system two kinds of agents are present, moving ones and stationary ones. While the moving ones are represented in our model as the pheromone dropping, action taking ones, the stationary agents are just represented by the medium over which rumours diffuse.

6.3.5 Parameters

In all simulations we have used a 50×50 grid, and an overall number of $N = 4000$ agents, divided in $n_s = 10$ families with values of the number of hidden neurons h going from 5 to 14. The agent population was completely heterogeneous, i.e. in the first generation N different genetic codes were created, and all these agents were tested at the same time on the same grid. Tournament selection was initially performed over 4 randomly chosen agents, and then the size of tournament was scaled as $[N_s/100]$ (where N_s is the species size, and square brackets denote the integer part). Each generation was composed of 20 different runs in which agents had different start and goal points, constrained only by the request to be at a Manhattan distance

6.4 Evolutionary experiments

$d = |s_x - g_x| + |s_y - g_y| > 35$. Each run of each generation lasted 2500 time steps, and agents were removed from the grid when they reached their goal. The size of species was allowed to change every $g_s = 100$ generations. Each evolutionary experiment consisted of at least 500 generations (but we performed up to 2000 generations in many experiments). The pheromone diffusion and evaporation rates were kept fixed to $\beta = 0.1$ and $\delta = 0.1$.

6.4 Evolutionary experiments

6.4.1 Description of the experiments

To understand which kinds of strategies agents could develop in order to optimise their motion on the mobility system, we performed many simulations under different conditions, which means that we both used *homogeneous grids* (i.e. grids on which all the sites had the same width, $w_{ij} = w \forall i, j$) and *heterogeneous grids* (i.e. grids with a structure corresponding to figure 6.1), and with *different values of the traffic factor* γ . For each experimental setting, we evolved both *agents that could drop pheromone*, and *agents that could not drop it* ($D_0 = D_1 = 0$) and thus could not develop any communication system. Once the evolution process was over, we extracted “*evolved agents*” as the “average” agent over the best population over the whole evolutionary history (i.e., we obtained a family of 4000 clones, the genetic code of each clone being obtained as the average code in the family with the higher fitness throughout evolution). These best agents were tested in all the possible environments, the pheromone dropping agents being tested both with and without pheromone. A test corresponded to an analysis of the fitness of the 4000 clones over 20 runs of the experiment, i.e. to $8 \cdot 10^4$ individual tests. Agents selected in this way perform better than the individual agents that had the highest fitness during evolution, once the latter were tested over a large number of trials and inside a population of clones.

This process has been performed to investigate the ability of evolution to optimise traffic flow under different situations, to study the role of communication, and also to analyse if the found solutions were stable under changes of the environment. We first describe the experiments one by one, and then summarise the results in tables 6.1, 6.2 and 6.3.

6.4.2 Simulations without traffic jams, $\gamma = 0$

We first studied the behaviour of the system in case $\gamma_{ij} = 0$, i.e. when the velocity of agents on a given site did not depend on the number of agents located on that site. In this situation agents have to learn to reach their goal (on the base of the “distance to the goal” inputs) and to use properly the sites with an high w value to reach quickly their goal. Since agents are not provided with a map of the network nor know their actual position (they just know their relative position with respect to their goal), their ability to attain an high fitness strongly depends on the development of a communication system.

Different settings

The experiments were performed in the following way. First we evolved agents on a grid in which all the sites had “width” $w = 0.05$. Since the grid is homogeneous, there is no need for communication, and thus we fixed the amount of dropped pheromone to zero (both for successful and unsuccessful agents, i.e. $D_0 = D_1 = 0$). Then we modified the structure of the grid according to figure 6.1, in which “highways” had width $w = 1$. In this environment we evolved both agents that could and could not drop pheromone.

Fixed pheromone dropping

For pheromone dropping agents, after a few attempts with different D_0 and D_1 values, we arrived at the conclusion that the best performances could be attained using a positive value of D_0 and a null value of D_1 . (In particular we used $D_0 = 1$, $D_1 = 0$). This choice allows a clear distinction between zones with high and low w values, and works much better than the $D_0 = 0$, $D_1 > 0$ choice because it allows a better “description” of the grid (agents spend a longer time on sites with low w and thus is better if they can communicate “their experience” when located on these sites, i.e. when they cannot move, at least if quite high values of the δ and β evaporation and diffusion rates are used).

Evolving pheromone dropping

Finally we tried to evolve the pheromone dropping mechanism directly with a genetic algorithm. More explicitly we evolved the values of two parameters

6.4 Evolutionary experiments

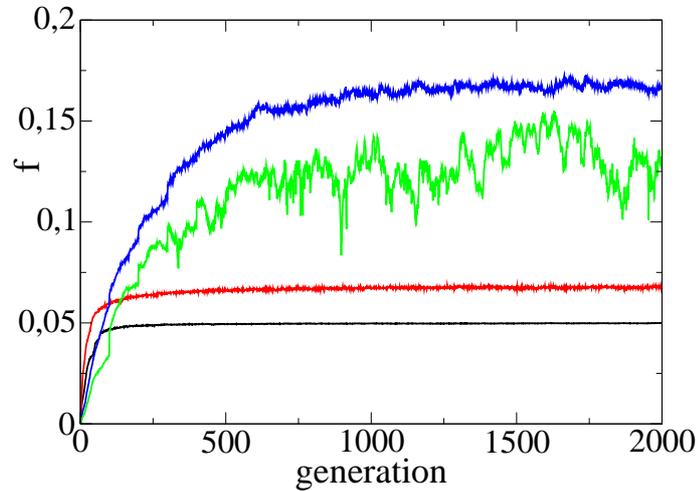


Figure 6.3: Evolution of fitness through generations for agents evolved in an homogeneous grid (black line), in a heterogeneous one without pheromone (red line), in an heterogeneous one with a fixed pheromone dropping system (blue) and with an evolving pheromone dropping system (green).

\bar{D}_C and \bar{D}_B defined in such a way that the amount of dropped pheromone D was given by

$$D = I\bar{D}_C + \bar{D}_B \quad (6.8)$$

where I is the input of an extremely simple “neural network” (with a single connection \bar{D}_C , bias \bar{D}_B and output D) that assumes value 1 if the agent is allowed to perform an action and 0 if it is not allowed to (these parameters are related to D_0 and D_1 by formulae $D_0 = \bar{D}_B$, $D_1 = \bar{D}_C + \bar{D}_B$ and thus $\bar{D}_C = -\bar{D}_B$ is qualitatively equivalent to our previous $D_0 \neq 0$, $D_1 = 0$ choice).

Results

Figure 6.3 shows the evolution of fitness through generations in the four cases. Agents evolved in the homogeneous grid reach and remain at the maximum allowed fitness $f \approx 0.050$, which means that they are able to go straightly to the goal. When the evolved best agents (according to the definition previously given) were tested in the heterogeneous environment, they had $f \approx 0.062$.

Agents evolved in the heterogeneous environment reached a maximum fitness $f \approx 0.068$, while the best evolved agents performed with the maximum available fitness $f \approx 0.050$ in the homogeneous grid and with $f \approx 0.070$ in the heterogeneous one. A surprising point is that *agents evolved in the heterogeneous grid can have, in this environment, an higher fitness than those evolved in the homogeneous one*, even without any information about neither the coordinates nor the width of the point they are located in. This is due to the fact that *even if they do not have any direct knowledge neither of the overall structure of the grid, neither of the point they are located on, they have an “historical” (i.e. due to evolution) knowledge (or better, memory) of the structure of the grid*. In fact, evolution leads agents to develop, between all the possible paths of minimum Manhattan length to the goal, those that make higher the probability to use the “highways” of maximum width (a similar phenomenon will be deeply analysed in the $\gamma > 0$ case in sections 6.4.3 and 6.5). As we will see this choice of a particular path corresponds to a *“symmetry breaking traffic flow rule”*. In this case, as shown in figure 6.4, this rule corresponds to choosing paths with a single change of direction. Since sites with an high value of w are located along straight lines (figure 6.1) if (by chance) the agent is located on one of these “highways” it will remain on it for a longer time.

Agents evolved with a fixed pheromone dropping mechanism reached a maximum fitness $f \approx 0.173$ during evolution and $f \approx 0.188$ when tested as clones, while if tested without pheromone they had the maximum possible fitness in the homogeneous environment, i.e. $f \approx 0.050$, and $f \approx 0.063$ in the heterogeneous one, exactly as the agents evolved in the homogeneous environment without pheromone, i.e. they do not develop any “traffic flow rule”. Nevertheless, *since in presence of pheromone the fitness is three times higher than in absence of pheromone, we can say that an efficient communication system has emerged between the agents*.

An analysis of the trajectories of these agents in presence and absence of pheromone (figure 6.5) shows that when the communication system is switched off the agents evolved with pheromone behave just as those evolved without pheromone (compare figure 6.4) while when they can communicate they are attracted by the “highways”, a behaviour which results into an increase of their fitness. Nevertheless, even the trajectories that agents follow when communicating are always trajectories that minimise the Manhattan distance, and not the time to reach the goal. (More precisely, in a first approximation,

6.4 Evolutionary experiments

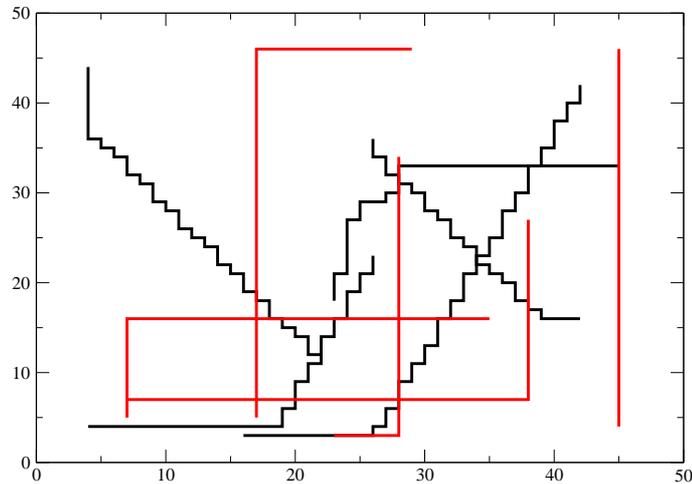


Figure 6.4: Trajectories of agents evolved in a homogeneous grid (black lines) and in a heterogeneous one (red lines). The second class of curves, with a single change of direction, makes more probable the use of “highways” of “wide” sites in figure 6.1.

agents choose, between the trajectories of minimal Manhattan distance, those that minimise the time to reach the goal). This is due to the fact that agents have only local and not global information about the structure of the grid, and thus cannot choose a path that corresponds to a global time minimum, but just a minimum between the paths of minimal Manhattan distance to the goal (their distance to the goal is the only global information provided). Deviations from these paths of minimal distance and time are present, but are probably due to the fact that this local information provided to agents is not absolute but depends on the other agents’ experience, and is thus partially incomplete (and time varying).

As figure 6.3 clearly shows, *the results obtained when evolving also the pheromone dropping system are worst and more instable than those obtained with a fixed pheromone dropping system.* Nevertheless *also in this case an effective communication system emerges* (the maximum fitness during evolution is $f \approx 0.155$, more than twice that obtained by agents evolved without pheromone). Furthermore, the behaviour of evolved agents is very similar to that obtained by agents using a fixed pheromone dropping system ($f \approx 0.190$).

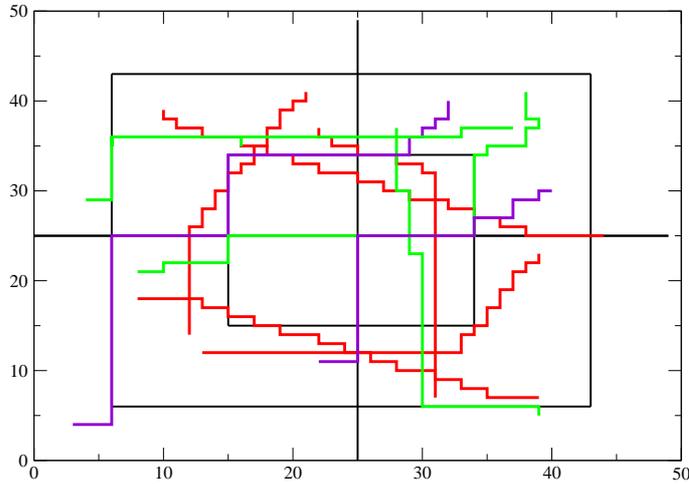


Figure 6.5: Trajectories of agents evolved with a fixed pheromone dropping mechanism. Red lines show the behaviour in absence of pheromone, while green and violet lines those of agents using pheromone. The “highways” of wide sites are shown by black lines.

Evolution of the pheromone dropping system (figure 6.6) clearly shows that the $D_1 = 0$ choice is the most effective one, since there is a clear tendency to have $\bar{D}_C = -\bar{D}_B$ (see equation (6.8)).

Comments

We can thus say that *an effective* (even if not necessarily optimal) *pheromone based communication can emerge between agents evolved in environments in which “highways” on which a higher speed is possible are present*. This communication system allows agents to choose trajectories to their goal that use effectively these “highways”. In order to have a pheromone field that describes in an appropriate way the structure of the grid, and thus effectively leads agents, an adequate exploration of the grid is necessary. It is thus not surprising that, *as the number of agents gets higher, the communication mechanism gets better* (figure 6.7). Even *when a communication system is not available, evolution leads agents to exploit the geometry of the system in the best possible way*.

6.4 Evolutionary experiments

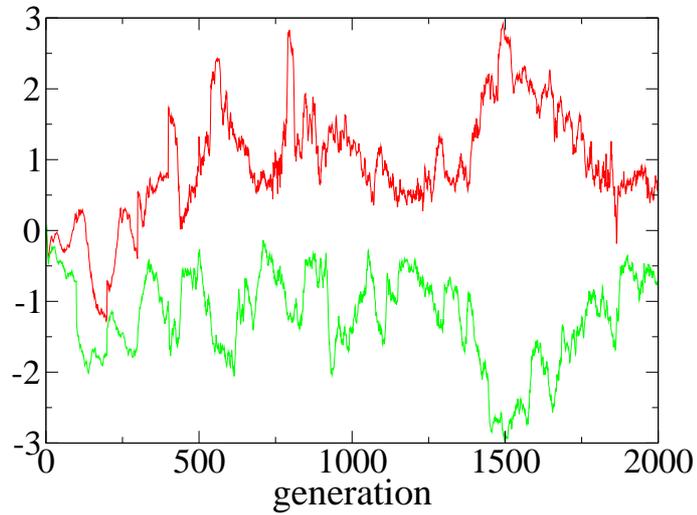


Figure 6.6: Evolution of \bar{D}_C (red line) and \bar{D}_B (green line).

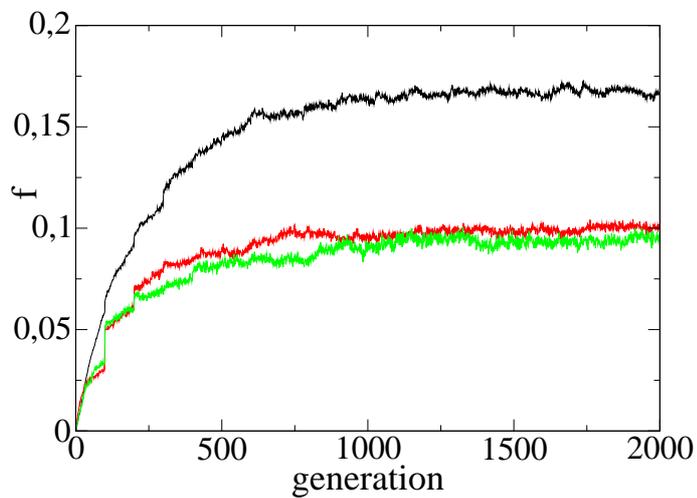


Figure 6.7: Evolution of fitness through generations for agents with a fixed dropping pheromone system. The black line shows the results obtained with $N = 4000$ agents, the red one with $N = 400$, and the green one with $N = 200$. (Discontinuous jumps in fitness are due to the particular form of the genetic algorithm that evolves the neural network structure).

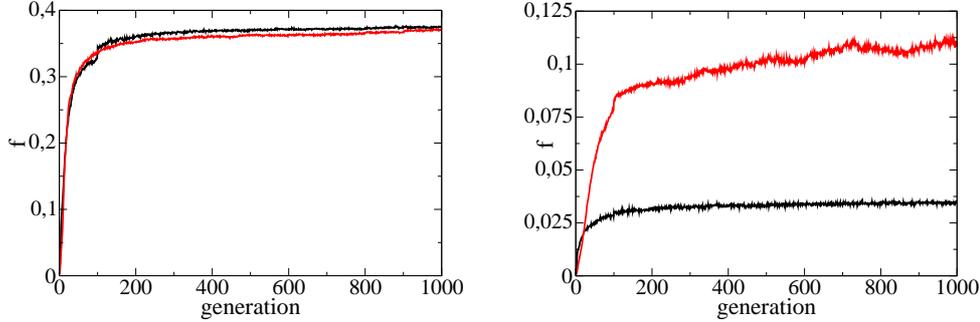


Figure 6.8: Evolution of fitness through generations for agents in an homogeneous grid with $\gamma > 0$. Left: $\gamma = 1$, right: $\gamma = 2$. Red lines show agents using pheromone, black ones those not using it.

6.4.3 Traffic jams in a homogeneous network, $\gamma > 0$

Using $w_{ij} = 1$, $\gamma_{ij} = \gamma > 0 \forall i, j$ we switch to the problem of avoiding traffic jams, i.e. overcrowded sites. This problem is qualitatively different from the previous one, since the velocity p_{ij} (equation (6.3)) is time dependent and it is not trivial (if not impossible) to find *a priori* a time minimising path to the goal.

Figure 6.8 shows the evolution of fitness in the $\gamma = 1$ and $\gamma = 2$ cases, for agents using ($D_0 = 1$, $D_1 = 0$) or not ($D_0 = 0$, $D_1 = 0$) pheromone.

Traffic flow rules

The results show clearly that effective communication emerged only in the $\gamma = 2$ case. In order to understand why agents evolved with pheromone could not develop communication in the $\gamma = 1$ case, we have to notice that *the performance of evolved clones* ($f \approx 0.380$ both for agents using or not pheromone) *was higher than that of agents going straightly to the goal* (i.e. agents evolved without pheromone in the homogeneous $\gamma = 0$ grid), which resulted to be, in the $\gamma = 1$ homogeneous setting, $f \approx 0.330$. This means that agents evolved in the $\gamma = 1$ case *had developed a “traffic flow rule” i.e. the choice of a particular path to the goal that minimises the traffic problem*. In the $\gamma = 2$ case the fitness of evolved clones was $f \approx 0.037$ for agents not using pheromone, $f \approx 0.114$ for those using it, while it was $f \approx 0.023$ for agents going straight to the goal, showing also in this case the presence of

6.4 Evolutionary experiments

a traffic flow rule (since agents going straight to the goal have lower fitness than those evolved in the $\gamma = 2$ environment but not using pheromone).

An analysis of the trajectories of agents evolved without pheromone in the homogeneous grid for different values of γ helps us to understand the nature of this traffic flow rule (figure 6.9). Agents evolved with $\gamma = 1$ follow paths with few changes of directions, similar to those of pheromone dropping agents in the heterogeneous $\gamma = 0$ case (see figure 6.5). But a deeper analysis (figure 6.10) shows two differences: first, agents evolved with $\gamma = 1$ do not follow paths of minimal Manhattan distance, since little detours are present, and furthermore, all the trajectories are followed clockwise (or all counter clockwise, depending on the output of the evolutionary process). When $\gamma = 2$, the detour from paths of minimal Manhattan distance is stronger (actually in the $\gamma = 2$ case the dependence of the trajectory on the position of the start and goal points gets stronger and an analysis of the traffic flow rule is harder to perform, as shown in figure 6.11. Probably this dependence on initial conditions optimises the overall motion of the system. It is quite clear, from figure 6.11, that the motion of agents is in some way guided by the edge of the grid. This edge had been realised as a constraint on the possible moves of the agents, but evolution led agents to exploit this constraint.). A deeper analysis of how and why these rules emerge is performed in section 6.5.

Traffic flow rules and communication

Finally an analysis of the trajectories of pheromone dropping agents evolved with $\gamma = 2$ shows clearly that in this case an effective communication system overlaps a traffic flow rule (figure 6.12). The traffic flow rule corresponds to the (evolutionary) memory of the geometry of the mobility system, and suggests to the agent the way to behave in absence of communication based inputs, while pheromone communication gives to the agent time varying information about the current state of the area of the grid it is located on. If the evolved agents are tested suppressing pheromone, their fitness drops from $f \approx 0.114$ to $f \approx 0.031$. This value is higher than that of agents without traffic flow rules, but lower than that of those who developed a pheromone independent traffic flow rule.

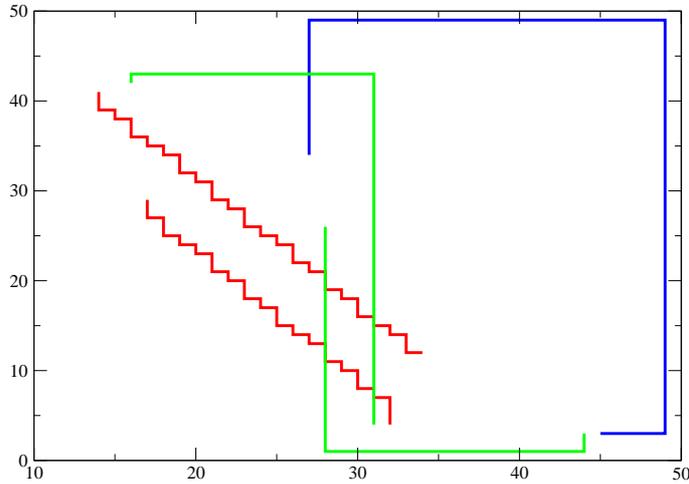


Figure 6.9: Trajectories of agents evolved in a homogeneous grid without pheromone. Red lines show trajectories of agents evolved with $\gamma = 0$, green ones those of agents evolved using $\gamma = 1$, the blue line a trajectory of the $\gamma = 2$ case.

Comments

The results show that *an effective communication system emerges also for agents moving on a homogeneous $\gamma > 0$ grid, i.e. on a mobility system whose structure is changed along time by the movement of the agents themselves.* In this case, *the communication system is developed along with a communication independent traffic flow rule, that stands for the evolutionary memory of the geometric structure of the grid.* When γ is low enough, communication does not emerge, and traffic flow is optimised only through the traffic flow rule.

6.4.4 Heterogeneous grid, $\gamma > 0$

We finally tried to evolve agents that could solve the problem of moving on a *a priori* heterogeneous and time varying mobility system. In particular we used the usual structure in figure 6.1, with $\gamma = 0$ on “narrow” ($w = 0.05$) sites, and $\gamma = 2$ on “wide” ($w = 1$) ones. This corresponds to a common situation in car mobility systems, in which the roads on which an higher speed is possible are also more prone to traffic jams. To show that in this situation

6.4 Evolutionary experiments

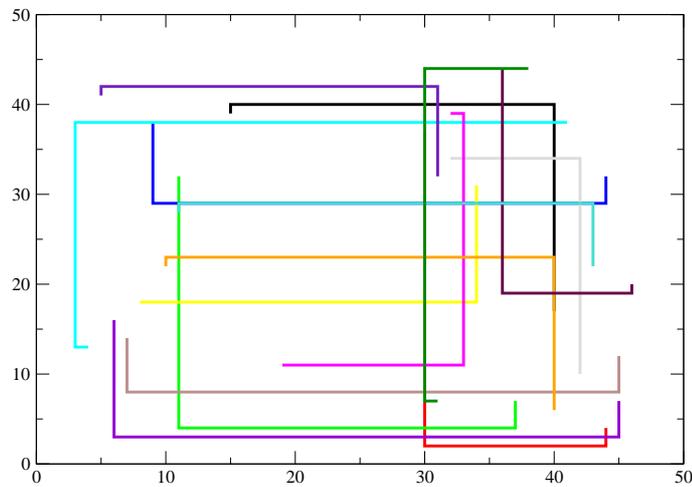


Figure 6.10: 15 randomly chosen trajectories of agents evolved in a homogeneous grid without pheromone, $\gamma = 1$. All the trajectories are followed counter clock wise, and most of them are not of minimal Manhattan distance. In the latter case, the agent always performs a change of direction few steps before arriving to the goal. The position of these changes of direction helps in understanding the direction in which trajectories are followed. See also figure 6.19 for a better understanding. To each colour corresponds a single trajectory.

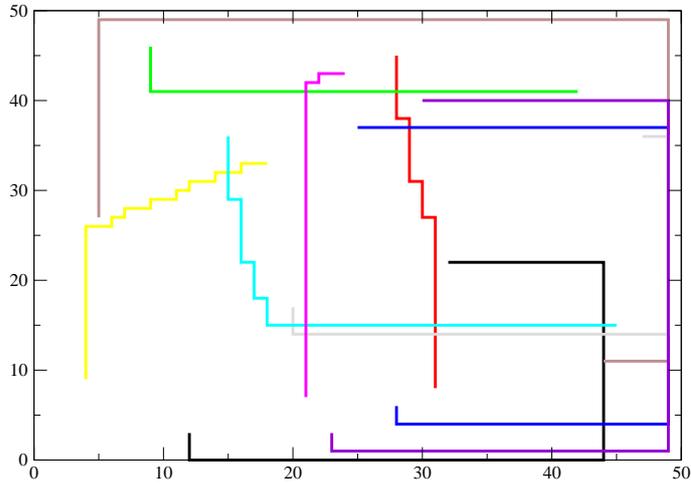


Figure 6.11: 10 randomly chosen trajectories of agents evolved in a homogeneous grid without pheromone, $\gamma = 2$. All the trajectories are followed clock wise. To each colour corresponds a single trajectory.

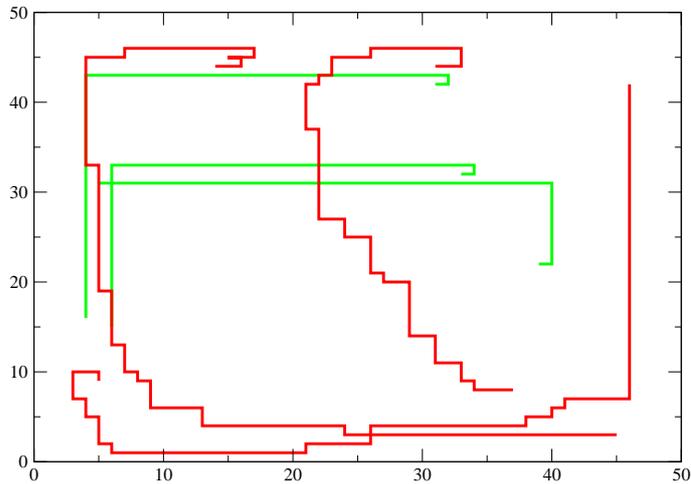


Figure 6.12: Trajectories of agents evolved in a homogeneous grid with pheromone, $\gamma = 2$. Green lines show (three) trajectories followed when communication was suppressed, while red ones (three) trajectories followed when agents could communicate. All the trajectories are followed clock wise.

6.4 Evolutionary experiments

a dilemma was actually present, we tested on these networks agents going straightly to the goal, that performed with $f \approx 0.036$. Since this value is lower than the maximum velocity allowed on “narrow” sites, this means that “highways” are harmful if not used *cum grano salis*. Agents evolved directly on the heterogeneous $\gamma > 0$ grid without using pheromone performed, once tested as clones, with $f \approx 0.046$ showing that an appropriate pheromone independent traffic flow rule could increase fitness, even if it remained under the maximum velocity allowed on narrow sites. Finally, agents dropping pheromone ($D_0 = 1$, $D_1 = 0$, but equivalent results were obtained evolving also the pheromone dropping mechanism) could perform with $f \approx 0.060$, and thus *had evolved an effective communication system*.

Figure 6.13 shows the evolution of agents in this environment, while figure 6.14 shows the trajectories of pheromone dropping evolved agents. A comparison of figure 6.14 with figure 6.5 shows that while in the $\gamma = 0$ case agents remained as long as possible on wide sites, in the $\gamma = 2$ case they have a tendency to use “highways”, but only for short stretches. These results suggest thus that evolved agents are able to use “highways” when they are free from traffic jams, but are able to leave them if they incur in a traffic jam.

6.4.5 Resume of the results

We have seen that a pheromone based (indirect) communication system can effectively emerge between agents moving on an idealised mobility system. The communication system allows agents to have a local knowledge of the structure of the mobility system and thus to choose properly a path to the goal. This communication system usually emerges alternatively (or even together) to a communication independent traffic flow rule which accounts for the evolutionary memory of the geometric structure of the mobility system. *The communication system showed to be quite general*, since when we tested evolved pheromone dropping agents in environments different from those in which they had been evolved in, they in general performed with a fitness higher than that of agents simply going straight to the goal (the only exceptions being the homogeneous environments with low γ , in which, as we have seen, is difficult to develop a communication system). For example, pheromone dropping agents evolved in the heterogeneous $\gamma = 0$ grid had $f \approx 0.041$ in the homogeneous $\gamma = 2$ grid, a fitness higher than that of

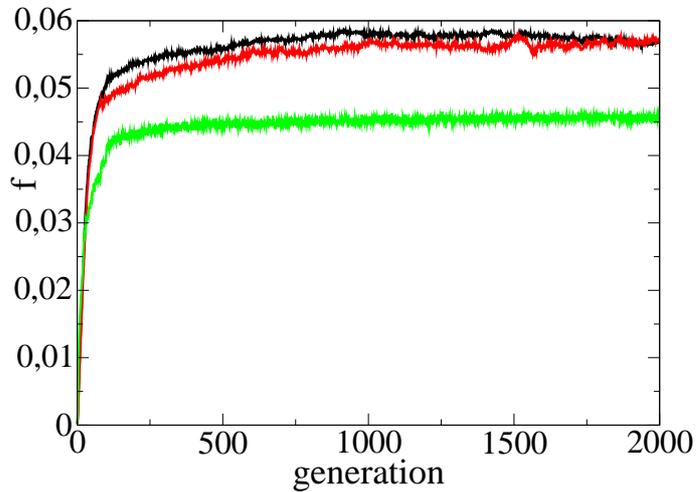


Figure 6.13: Evolution of fitness through generations for agents in a heterogeneous grid with $\gamma = 2$. The black line shows the results given by agents with a given pheromone dropping system, the red line the results for agents with an evolving pheromone dropping system, while the green line those of agents that could not drop pheromone.

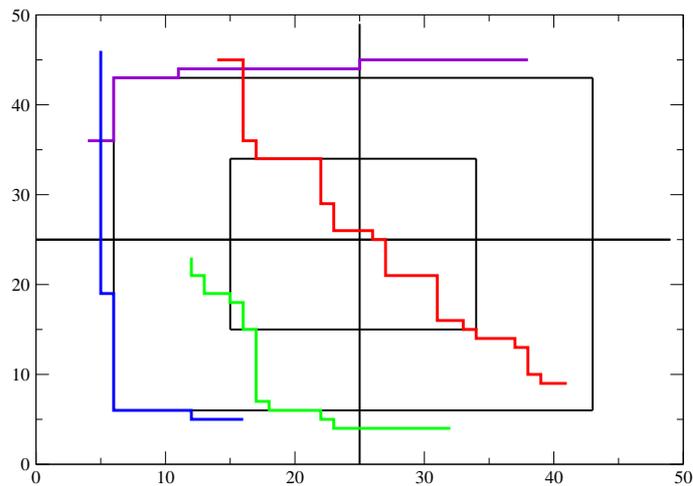


Figure 6.14: Trajectories of pheromone dropping agents evolved in a heterogeneous grid with $\gamma = 2$. The black lines show the “wide” sites (“highways”).

6.4 Evolutionary experiments

Table 6.1: Agents evolved without pheromone in homogeneous environments, tested in different environments. Columns correspond to different evolved agents, lines to different test environments. np means no pheromone, pd means pheromone dropping, ho means homogeneous, ht means heterogeneous.

	np, ho, $\gamma = 0$	np, ho, $\gamma = 1$	np, ho, $\gamma = 2$
np, ho, $\gamma = 0$	0.050	0.047	0.039
np, ho, $\gamma = 1$	0.329	0.381	0.293
np, ho, $\gamma = 2$	0.023	0.033	0.037
np, ht, $\gamma = 0$	0.062	0.065	0.048
np, ht, $\gamma = 2$	0.037	0.045	0.032

agents going straight to the goal ($f \approx 0.023$), but also of agents that had developed a pheromone independent traffic flow rule for that grid ($f \approx 0.037$). In that environment, agents evolved with pheromone in the homogeneous $\gamma = 1$ grid reached $f \approx 0.058$ and those evolved in the heterogeneous $\gamma = 2$ one an even higher $f \approx 0.075$.

We have also seen that *agents were able to evolve not only their ability to use the pheromone fields, but also the pheromone dropping mechanism*. In the latter case evolution was more unstable, but the performance of evolved agents was as high as in the case of agents evolved with a fixed pheromone dropping system. The evolution of the pheromone dropping mechanism is not trivial due to heterogeneity: an agent could be able to develop an efficient pheromone dropping mechanism, but not to exploit it, and thus the character would disappear even if increasing the overall fitness of the population. All the performances of evolved agents in the different environments are shown in tables 6.1, 6.2 and 6.3 (the results regarding agents with an evolved pheromone dropping system are not shown, since they are almost identical to those of the corresponding evolved agents with a fixed pheromone dropping system).

Nevertheless in order to develop an optimal communication system, it would be probably necessary to test (or maybe evolve) different values of the pheromone evaporation and diffusion rates, since in order to have an optimal information about the state of the grid, it is important also to optimise the

Table 6.2: Agents evolved without pheromone in heterogeneous environments, tested in different environments. See the legend of table 6.1.

	np, ht, $\gamma = 0$	np, ht, $\gamma = 2$
np, ho, $\gamma = 0$	0.050	0.050
np, ho, $\gamma = 1$	0.380	0.373
np, ho, $\gamma = 2$	0.028	0.029
np, ht, $\gamma = 0$	0.070	0.065
np, ht, $\gamma = 2$	0.047	0.046

Table 6.3: Agents evolved with pheromone, tested in different environments. See the legend of table 6.1.

	pd, ht, $\gamma = 0$	pd, ho, $\gamma = 1$	pd, ho, $\gamma = 2$	pd, ht, $\gamma = 2$
np, ho, $\gamma = 0$	0.050	0.047	0.048	0.050
pd, ho, $\gamma = 0$	0.049	0.047	0.036	0.049
np, ho, $\gamma = 1$	0.327	0.374	0.370	0.328
pd, ho, $\gamma = 1$	0.278	0.381	0.308	0.332
np, ho, $\gamma = 2$	0.021	0.034	0.031	0.023
pd, ho, $\gamma = 2$	0.041	0.058	0.114	0.075
np, ht, $\gamma = 0$	0.063	0.062	0.063	0.063
pd, ht, $\gamma = 0$	0.188	0.070	0.076	0.147
np, ht, $\gamma = 2$	0.037	0.044	0.044	0.039
pd, ht, $\gamma = 2$	0.046	0.048	0.045	0.060

6.5 Emergence of traffic flow rules

way in which information diffuses over the grid.

6.4.6 Notes on the evolution of the structure of the network

In order to obtain an optimal system, it would be probably also necessary a more accurate study of the structure of the network. It is also probable that using more complex networks different traffic flow optimising strategies would emerge in the system. Nevertheless we think that the study we have performed nicely accounts about how different strategies can emerge using quite simple settings.

As we already said, since we used a large number of agents and thus had a large number of available genetic codes, we thought that we could avoid to bother about the particular structure of the neural network, and try to develop in a parallel way different values of the number of hidden neurons h . Nevertheless we have not performed a large enough number of simulations (at least not in the same environment) to obtain a valid statistic about the evolution of the number of nodes. In general a few families invaded the population. In figure 6.15 we show a case in which the $h = 13$ family invaded all the population.

6.5 Emergence of traffic flow rules

6.5.1 Traffic flow rules

We have seen that traffic flow can be optimised in our mobility system in two ways: through the emergence of communication and through the emergence of a traffic flow rule. There are two main differences between these two phenomena: the first one relies on multiple interactions between agents (even if mediated through the environment) and on local information, while the second one relies on individual behaviour and global information about the structure of the mobility system.

Thus, to a first glance, the definition of self organisation given in section 6.2.2 and chapter 1.4 does not apply to traffic flow rules, but the distinction is more subtle, since actually no global information is provided to agents, but it is acquired through local interaction during evolution. We could say that the traffic flow rule by itself is not a self organisation phenomenon, but

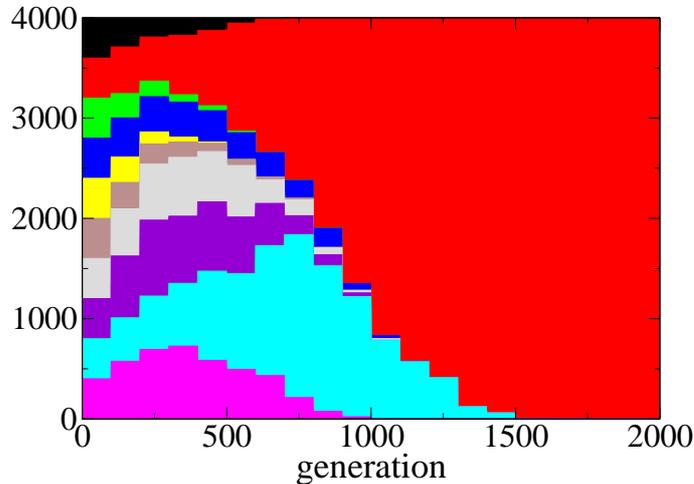


Figure 6.15: Evolution of the number of neurons in the hidden layer, in the heterogeneous $\gamma = 2$ case. $h = 13$ invades the population.

the evolutionary process that leads to the rule is in some way a form of self organisation.

We could define a traffic flow rule as a restriction on the movement of an individual, which results in an higher collective benefit, i.e. an optimisation of traffic flow. Many examples of the emergence of global patterns that optimise traffic flow are present both in human and animal societies, for example in car circulation, as described in [11] and [12], or in pedestrian flow, see [13, 14, 15] and chapter 5, and between social insects [16].

Focusing on human behaviour, we can say that while the phenomenon of “stripe formation” when pedestrians flows cross ([15] and chapter 5) is clearly due to local interaction and thus purely a phenomenon of self organisation, the situation is slightly different for pedestrians lane formation, and completely different for car circulation. As described in [13, 14], even if lane formation is basically due to self organisation, an individual tendency in all pedestrians (in a given country or geographical area) to walk always on the same side is surely present. This tendency is often related to the car circulation traffic rule of that country. At the present day, car circulation rules are present in each country as a centralised form of organisation (imposed by an authority), which has no dependence on individual interactions (cars travel always -or almost!- on the side prescribed by the rule, even if no other car is

6.5 Emergence of traffic flow rules

present). Nevertheless, they probably emerged through history as a form of self organisation (see [12]), and thus resemble the emergence of traffic flow rules in our system.

In some of these cases (car circulation, lane formation) self organisation assumes clearly the form of a symmetry breaking phenomenon, in which one of two a priori equivalent choices is performed (we refer also to the phenomenon of multistability, described in section 6.2.2 and chapter 1.4).

6.5.2 Corrections to the model

Introduction of direction

In order to study the emergence of traffic flow rules in our mobility system, we modify equation (6.3) in order to take in consideration the different directions in which agents move (since in mobility systems traffic flow rules usually emerge to avoid collisions between agents moving in different directions), and we thus define as $n^k(i, j)$ the number of agents on the site with coordinates (i, j) coming from direction k ($k = 0, 1, 2, 3, 4$ stands for the last action performed by the agents, for example 0 could stand for going in the direction of growing xs , 1 for growing ys , etc.) The overall number of agents on the site is obviously given by $n(i, j) = \sum_k n^k(i, j)$, while with $\mathbf{n}(i, j)$ we mean the set of values $n^k(i, j)$. For simplicity's sake we use only homogeneous grids, and thus we fix $\gamma_{ij} = \gamma$, $w_{ij} = 1 \forall i, j$ and remove w from our equations. We thus distinguish between two cases, *direction independent* and *collision*.

Direction independent setting

The *direction independent* setting is completely equivalent to the previous model in the homogeneous case, and thus the probability to act (or velocity) of an agent coming from direction l is given by

$$p^l(\mathbf{n}(i, j)) = \left(\frac{1}{n(i, j)} \right)^\gamma \quad (6.9)$$

and for $\gamma = 0$ we have $p = 1$ for any value of n^k , i.e. no traffic congestion can occur. When $\gamma > 0$, $p^l = 1$ occurs only if $n^l = 1$, $n^k = 0 \forall k \neq l$, i.e. if the agent is alone on the site.

Collision setting

In the *collision* setting to calculate $p^l(\mathbf{n}(i, j))$ we distinguish between the agents moving in directions orthogonal to l (or standing still), n^o , and those opposite to it, n^c , and define

$$p^l(n^k(i, j)) = \left(\frac{1}{1 + \frac{1}{2}(\sum_o n^o(i, j)) + n^c(i, j)} \right)^\gamma \quad (6.10)$$

Also in this case the probability to move in direction l is always 1 if $\gamma = 0$, while for $\gamma > 0$ it does not depend on n^l and is equal to 1 if all the agents have the same direction, i.e. agents moving in the same direction are not an obstacle for the motion of the others. The agents moving in other directions slow the motion, and in particular in this setting agents moving in the opposite direction cause a stronger traffic effect than those moving in orthogonal ones.

Neural network and GA parameters

Since we are interested in the emergence of a communication independent traffic flow rule, we remove pheromone from the model, and thus neural networks use as inputs only the distances $\Delta_x \equiv g_x - x$ and $\Delta_y \equiv g_y - y$ to the goal from the point $\mathbf{x} = (x, y)$ on which the agent is located.

We used in simulations $N = 4000$ agents with initially randomly chosen genetic codes (connections) and a fixed network structure, i.e. we did not try to evolve the number of nodes in the hidden layers. Nevertheless, in some simulations, we split the agents in $n_s = 4$ different families or species (with $h = 5, 6, 7, 8$ nodes in the hidden layer), each composed of 1000 agents, that could not exchange genetic information between them, in order to study how the evolution processes of different species mutually influenced each other. In case only one species was present, we used $h = 5$ nodes in the hidden layer. All the other parameters (and the fitness function described in equation (6.7)) remained equal to those described in section 6.3.5.

6.5.3 Clock wise and counter-clock wise behaviours

Studying the trajectories of agents evolved in a $\gamma > 0$ homogeneous grid (see for example figure 6.10) we observed that all the trajectories were followed (counter) clock wise. To study how this behaviour emerges, we classify the

6.5 Emergence of traffic flow rules

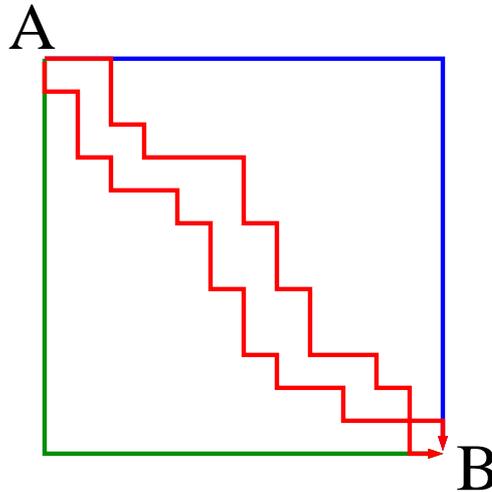


Figure 6.16: Four “good” paths ($g = 1$) from point A to B. The red ones have almost an equal number of r and l moves, while the blue one is clock wise ($l = 1$ $r = 0.5$) and the green one counter clock wise ($r = 0.5$ $l = 1$).

possible moves of agents, and define as “good move” g a move that lowers the Manhattan distance to the goal, and we split these moves as *counter clock wise* (or *right-handed* r) and *clock wise* (or *left-handed* l), corresponding to the different behaviours in figure 6.16 and table 6.4.

We thus characterise the system by the fraction of g , r and l moves over the overall number of moves, averaged over all the agents (we will usually refer to these averaged fractions simply as the g , r and l values of the systems, assuming values between 0 and 1). Nevertheless agents use more information than the one shown in table 6.4, since they also know the magnitude of the distances to the goal, and thus by line of principle any path from the start to the goal is available (in the sense that we can find a neural network that, given a starting and a goal point, will perform a given path). Furthermore, the definition of a “good move” is exact only in the $\gamma = 0$ case, since when a traffic problem is present, as we have seen for example in figure 6.11, it can be more efficient to choose longer (in Manhattan metric, but not in time) trajectories. Nevertheless, the given definitions will show to be sufficient to study the emergence of traffic flow rules in the system, at least for low values of γ .

Table 6.4: Definition of “good” (g), “left-handed” (l), “right-handed” (r) and “wrong” (w) moves on the base of the sign of the distances to the goal. $x++$ means direction of growing x s, following the C programming language convention. Notice that the l and r sets overlap and thus $g \leq l + r$.

	x++	y++	x- -	y- -
$\Delta_x > 0 \Delta_y > 0$	$g r$	$g l$	w	w
$\Delta_x > 0 \Delta_y < 0$	$g l$	w	w	$g r$
$\Delta_x > 0 \Delta_y = 0$	$g r l$	w	w	w
$\Delta_x < 0 \Delta_y > 0$	w	$g r$	$g l$	w
$\Delta_x < 0 \Delta_y < 0$	w	w	$g r$	$g l$
$\Delta_x < 0 \Delta_y = 0$	w	w	$g r l$	w
$\Delta_x = 0 \Delta_y < 0$	w	w	w	$g r l$
$\Delta_x = 0 \Delta_y > 0$	w	$g r l$	w	w

6.5.4 Experimental results

Evolution in absence of a traffic problem

We first tested that for $\gamma = 0$ no traffic rule emerged. In this situation the optimal solutions are known, and correspond to perform a g move at each time step. Following figure 6.16 we expected solutions with an almost equal number of “left handed” and “right handed” moves to be more probable than those with “*symmetry breaking*”, and actually evolution always led to an almost overall equal number of l and r percentages, both for a single species with 4000 agents and for 4 different species composed each of 1000 agents. Every path with $g = 1$, regardless of the values of r and l , is equally good, but the ones with a value of r and l around 0.5 are more numerous and thus represent a kind of “*thermodynamic equilibrium*” of the evolutionary process in case $\gamma = 0$. The fitness (and thus the percentage of g moves) quickly reached a value next to 1, showing that the evolutionary process easily solves the problem.

Evolution of one species in presence of a traffic problem

We then studied the case with $\gamma > 0$, using a single species of 4000 agents, both in the *direction independent* and in the *collision* case.

6.5 Emergence of traffic flow rules

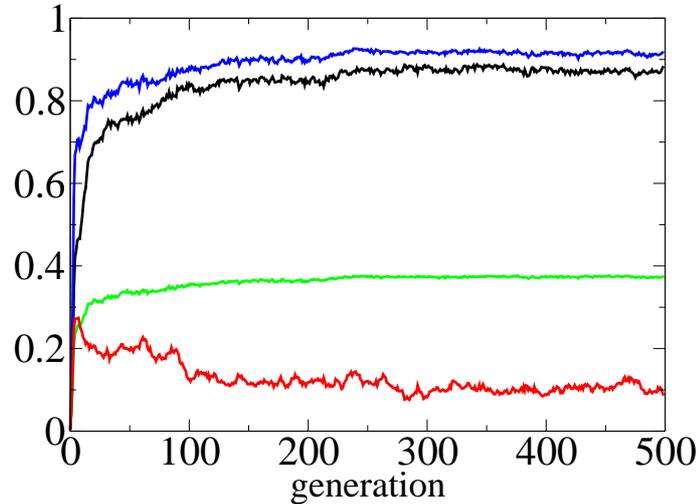


Figure 6.17: Fitness (green line), g (blue line), l (red line) and r (black line) moves as a function of the generations for $\gamma = 1$ in the case of one species. *Direction independent* setting.

In figure 6.17 we show the results of the evolutionary process in the $\gamma = 1$ *direction independent* case. We can see that the two curves for the average values of l and r moves split during the evolutionary process, reaching in one case a value next to 1 (≈ 0.9), while in the other one a value next to 0 (≈ 0.1). Since this is a symmetry breaking process, if we perform different evolutionary experiments, we observe with equal probability the emergence of a clock wise or counter clock wise motion.

In figure 6.18 we show the average values of l , r and g in the last 100 generations of the evolutionary process for different values of γ in the *direction independent* case, showing that both clock wise and counter clock wise behaviours emerge.

As we said before, when $\gamma = 0$, we have $r \approx l \approx 0.5$. When $\gamma > 0$, we can distinguish two zones. For low values of $\gamma > 0$, while the “winning” (most performed) move reaches a value next to 1, the “losing” one has a value around 0.5. This corresponds to a “perfect clock wise” or “perfect counter clock wise” behaviour according to table 1, as we have verified observing the trajectories of the agents. For higher values of γ the “losing” move assumes very low values, while the number of g moves is slightly lower. An analysis

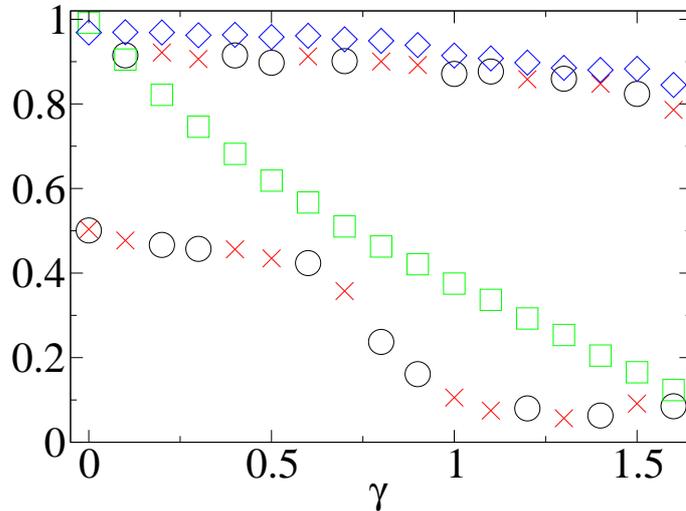


Figure 6.18: Averages over the last 100 generations as a function of γ : fitness (green squares), g (blue diamonds), l (red crosses) and r (black circles) moves. *Direction independent setting.*

of the trajectories of the agents shows that they are always clock or counter clock wise (figure 6.19 or figure 6.10), but not in the sense of table 1, since also “*wrong*” moves are performed.

According to our interpretation of the results, since the starting and goal points of the agents are chosen to be far and thus probably located in opposite parts of the mobility system, *the centre of the “town” (grid) is a critical traffic point* for agents that perform an almost equal number of l and r moves, and thus *evolution breaks the symmetry making the agents choose a path that avoids passing in the centre* (figure 6.20). While for low values of γ this choice is performed between all the paths of minimum Manhattan length, for higher values of γ the agent chooses a longer (but quicker) path that allows it to stay distant from the centre.

It is interesting to see what happens for very low values of γ (figure 6.21). In some experiments we have verified the emergence of a traffic flow rule, with a winning (l or r) choice with a value around 0.9, and a loosing choice with value around 0.5. Other experiments showed no symmetry breaking, since l and r had the same value, but this value was not 0.5 but around 0.7. In these cases agents developed only a partial circulation rule, with few

6.5 Emergence of traffic flow rules

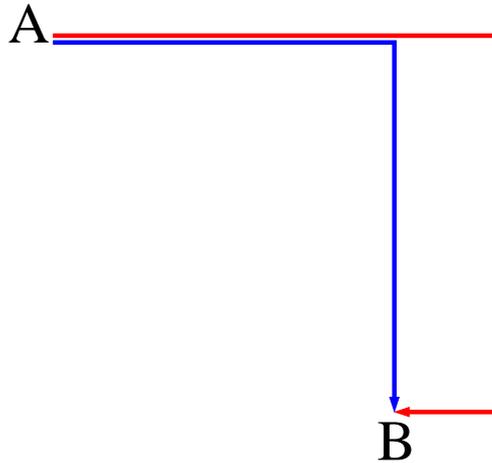


Figure 6.19: The blue path from point A to B corresponds to a perfect clock wise behaviour of minimum Manhattan length with $l = 1$, $r = 0.5$; while the red one has a l next to 1 and r next to 0, is not of minimum length and is the path chosen by evolved agents for high values of γ .

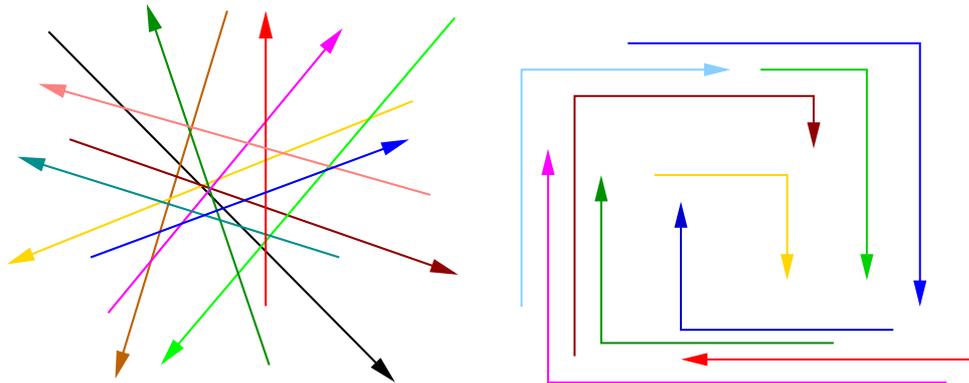


Figure 6.20: The straight (i.e., with $r \approx l$) paths at left have a high probability to pass through the centre and can cause a traffic jam, while the clock wise paths at right avoid it.

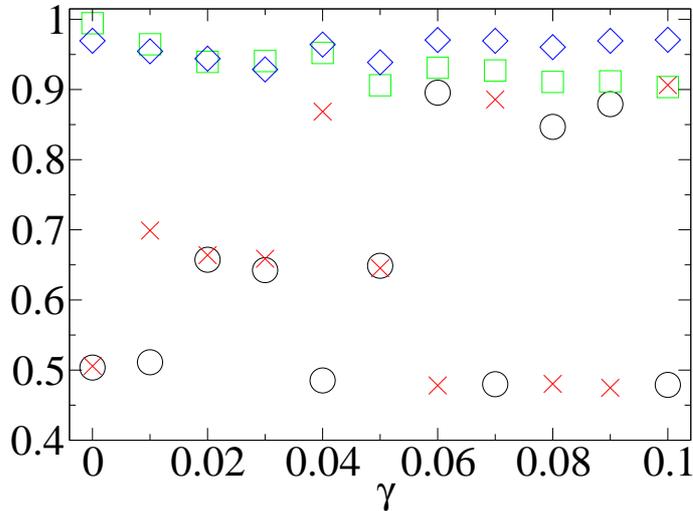


Figure 6.21: Averages over the last 100 generations as a function of γ (low values): fitness (green squares), g (blue diamonds), l (red crosses) and r (black circles) moves. *Direction independent* setting.

changes of direction but no clear predilection for clock wise or counter clock wise motion. The transition to an actual (counter) clock wise rule seems to happen around $\gamma \approx 0.05$ but the results for low values of γ seem to depend strongly on the output of the single repetition of the experiment and thus a larger number of repetitions should be necessary to obtain average values.

The *collision* case for a single species leads to almost the same results of the *direction dependent* case (figure 6.22). In this case the separation between the two zones (a first one in which paths are of minimum Manhattan length, and a second one in which detours are possible) occurs for an higher value of γ , since the traffic effect is reduced because the agents do not interact with those moving in their same direction and interact more weakly with those moving in a direction which is not opposite to their one.

Evolution of different species in presence of a traffic problem

In the experiments with a single family it was not completely clear how symmetry breaking arose. The process is surely due to the appearance of an agent that uses a clock wise or counter clock wise path to reach the goal and,

6.5 Emergence of traffic flow rules

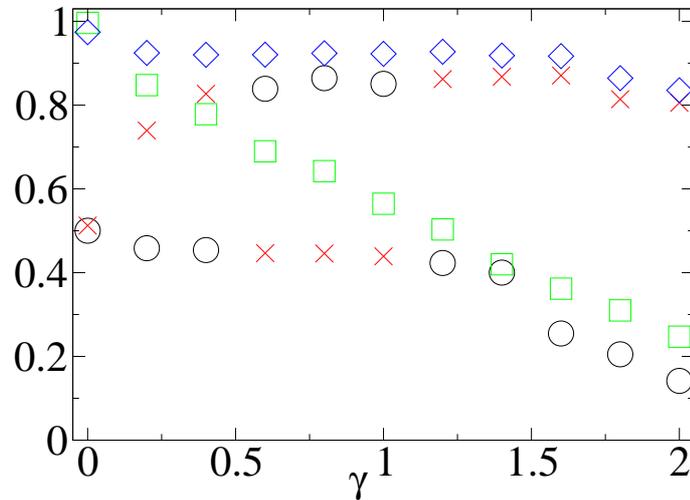


Figure 6.22: Averages over the last 100 generation as a function of γ : fitness (green squares), g (blue diamonds), l (red crosses) and r (black circles) moves, in the *collision* case.

avoiding the centre, performs with a high fitness and passes its genes to the future generations. What is not clear is if this character simply invades all the population, or if the motion of the fitter agents influences the evolution of the others.

In order to understand which hypothesis is right, we have split the population in $n_s = 4$ different species, that could not exchange genetic information between them. The results obtained in the *direction independent* case could not show clearly which hypothesis was right, since, even if in a few cases we found hints that the evolution of a species influenced those of the others (see for example figure 6.23), in general symmetry breaking arose only inside a single species, but not in the whole population.

The results for the “collision” case showed clearly the emergence of a same global traffic flow rule in all the different species. In figure 6.24 we show the evolutionary process for all the species for $\gamma = 2$, while in figure 6.25 we show the γ dependence of the average values of all the moves and of the fitness for all the species in the last 100 generations. Symmetry breaking, at least for large values of γ , emerges again as a global property in all the population, since the values of the l and r averages are always clearly separated. We can

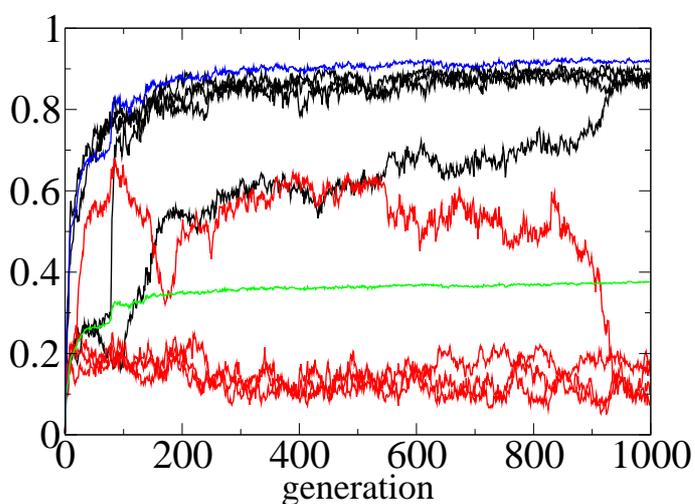


Figure 6.23: Average fitness over species (green line), average number of g (blue line), l (red lines) and r (black lines) moves for all the species as a function of generations for $\gamma = 1$ in the *direction independent* case of 4 species with 5, 6, 7, 8 hidden neurons. Three families developed a counter clock wise rule, while the fourth one initially used a clock wise rule, then switched around generation 200 to an equal distribution between l and r choices, to reach the same $r \approx 0.9$, $l \approx 0.1$ strategy of the other species around generation 900.

6.5 Emergence of traffic flow rules

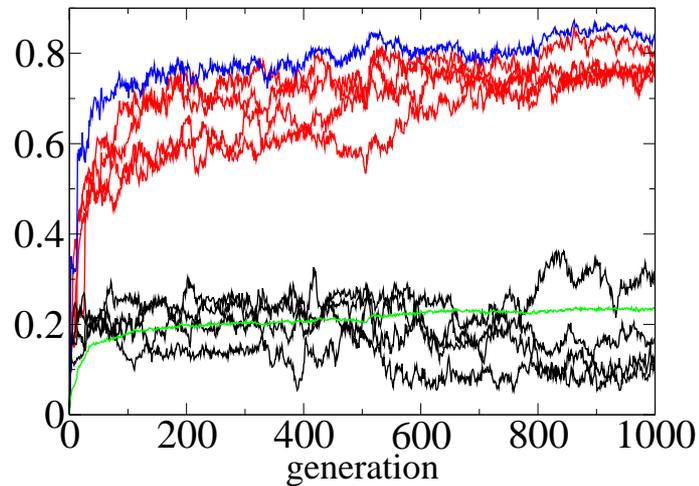


Figure 6.24: Average fitness over species (green line), average values of g (blue line), l (red lines) and r (black lines) moves for all the species as a function of generations for $\gamma = 2$ in the “collision” case with 4 species.

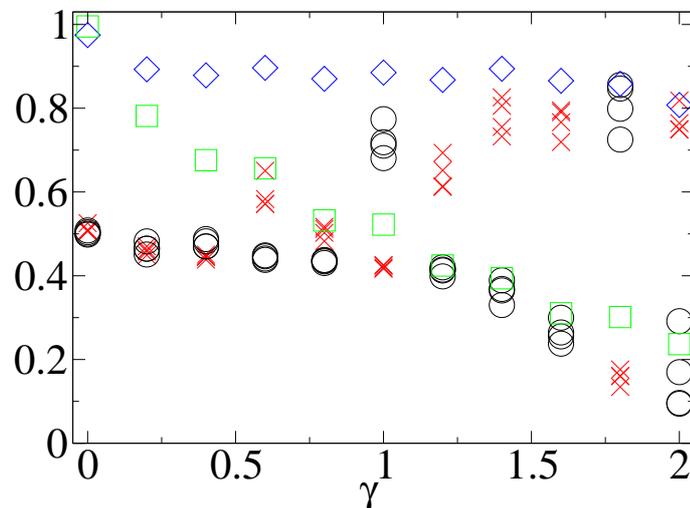


Figure 6.25: Averages over the last 100 generation as a function of γ : fitness (green squares), g (blue diamonds), l (red crosses) and r (black circles) moves, in the “collision” case, 4 species. For $\gamma \geq 1$ circles and crosses are clearly distinct, i.e. a single rule has emerged in all the species.

thus assume that *the agents following a traffic flow rule are able to influence the evolution of the other ones*.

Also in this case we can see two zones, the first one corresponding to a traffic rule that follows a trajectory of minimal distance, while for higher values of γ a traffic rule that follows a longer path (we refer again to figure 6.19) emerges in all the species. For low values of γ some of the simulations have not shown the emergence of a clear traffic flow rule in the first 1000 generations. Since the simulations with a single family lead to a traffic flow rule for the same values of γ , we can think that the absence of this rule in the 4 species case is due to a greater difficulty to develop it simultaneously for many families. The transition to a clear emergence of a traffic flow rule seems to happen around $\gamma = 1$.

The difference between the *direction independent* and *collision* case is due to the fact that in the first case only the geometry, and not the direction, of trajectories matters, and thus families can develop different rules, while in the second case the flow is optimised only if all the agents move in the same direction.

6.6 Conclusions

We have shown that a quite general indirect communication system can emerge between agents moving in a mobility system, in order to optimise traffic flow. The communication system can solve both the time independent problem of finding the paths on which a higher velocity is possible, and the time varying problem of avoiding traffic jams. Communication emerges together with a traffic flow rule that depends on the geometry of the mobility system, as the result of a form of “*evolutionary memory*”. This rule emerges as a *symmetry breaking* choice between paths of equal length when the occurrence of traffic jams is quite low, while it emerges as an actual choice of a longer but quicker path when the occurrence of traffic problems is very high. We have also shown that, if the motion of the agents depends on the direction of the motion of the other agents located in the same site, as expected for a realistic mobility system, the evolution of a group of agents, and thus their ability to develop a traffic flow rule, can influence the evolution of other agents even if there is no exchange of genetic information.

6.6 Conclusions

Bibliography

- [1] E. Bonabeu, “*Control Mechanism for Distributed Autonomous Systems: Insights from the Social insects*”, in L.A. Segel, I.R. Cohen (eds.), “*Design Principles for the Immune System and Other Distributed Autonomous Systems*”, Oxford University Press (2001).
- [2] E. Bonabeu, M. Dorigo and G. Theraulaz, “*Swarm Intelligence*”, Oxford University Press (1999).
- [3] P. P. Grassé, “*La théorie de la Stigmergie: Essai d’interprétation du Comportement des Termites Constructeurs*”, *Insect. Soc.* 41-80 (1959).
- [4] J. L. Deneunburg, S. Aron, S. Goss and J. M. Pasteels, “*The Self-Organizing Exploratory Pattern of the Argentine Ant*”, *J. Insect Behaviour* 3: 159-168 (1990).
- [5] R. Beckers, J. L. Deneunburg and S. Goss “*Trails and U-Turns in the Selection of a Path by the Ant Lasius niger*”, *J. Theor. Biol.* 159, 397-415 (1992).
- [6] M. Dorigo, V. Maniezzo and A. Colorni “*Positive Feedback as a Search Strategy*”, *Tech. Rep. No. 91-016*, Politecnico di Milano, Italy (1991).
- [7] M. Dorigo “*Ottimizzazione, Apprendimento Automatico ed Algoritmi Basati su Metafora Naturale*”, Ph.D. Dissertation, Politecnico di Milano, Italy, (1992).
- [8] Y. Nakamichi and T. Arita, “*Effectiveness of Emerged Pheromone Communication in an Ant Foraging Model*”, *Proceedings of the Tenth International Symposium on Artificial Life and Robotics*, 324-327 (2005).

BIBLIOGRAPHY

- [9] K. O . Stanley and R. Miikkulainen, “*Evolving Neural Networks through Augmenting Topologies*”, *Evolutionary Computation* 10(2), 99-127 (2002).
- [10] F. Gomez and R. Miikkulainen, “*Solving Non-Markovian Control Tasks with Neuroevolution*”, *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, pages 1356-1361, Morgan Kaufmann, San Francisco, California (1999).
- [11] http://en.wikipedia.org/wiki/Driving_on_the_left_or_right
- [12] <http://users.pandora.be/worldstandards/driving%20on%20the%20left.htm>
- [13] D. Helbing, L. Buzna, A. Johansson and T. Werner, “*Self-Organized Pedestrian Crowd Dynamics: Experiments, Simulations, and Design Solutions*”, *Transportation Science*, Vol. 39, No 1, 1-24 (2005).
- [14] D. Helbing, P. Molnár, I. J. Farkas and K. Bolay, “*Self-organizing pedestrian movement*”, *Environment and planning B: Planning and design*, v. 28 361-383 (2001).
- [15] K. Ando, H. Oto and T. Aoki “*Forecasting the flow of people*”, *Railway Res. Rev.* 45(8) 8-13 (1988).
- [16] I. D. Couzin and N. R. Franks, “*Self-Organized Lane Formation And Optimized Traffic Flow In Army Ants*”, *Royal Soc. Proc. Biological Sc.* (2002).

Conclusions

The purpose of this thesis has been the study of artificial life systems starting from the dynamics of basic constituents. In particular we focused on systems whose fundamental units are individuals capable of perception and data processing, used as the basis of their decision mechanisms and thus of their ability to modify the environment. These systems have been studied using computer models (agent models) in which each individual is an independent unit (the system is distributed and autonomous). In particular we have studied models of the dynamics of agents moving in crowds or in large scale mobility systems (“urban areas”) and of cells in the immune system. Even if the fundamental components of these systems are extremely complex, usually their behaviour is introduced in computer models using simple local rules, and one of the purposes of our research has been to investigate how from these simple rules could emerge global self organising properties, also as the result of an evolutionary process.

After a brief introduction to the basic concepts of interest in the field of artificial life (chapter 1) and of the computational tools used in the thesis (chapter 2), in the chapters from 3 to 6 we exposed the research projects developed during the Ph.D. activity. This activity can be divided in two parts: in chapter 3 and 4 we have studied the possible connection between the microscopic dynamics and the macroscopic one in models with very simple interaction rules, while in chapters 5 and 6 we have used a more “genuine” Alife approach based on genetic algorithms and neural networks to study evolving and emergent properties of complex adapting systems.

In chapter 3 we modify a physical system (interacting charged particles) in order to make it resemble a perception based interaction (sight). More precisely, we have introduced a dependence on the relative positions and velocities in the forces between particles, in order to introduce a “cone of vision”. As a consequence of this dependence the interaction between particles

is non Newtonian, since if a particle (an “automaton”) sees another particle, but is not seen by the latter, the third law of dynamics does not apply. We have performed an analytical and numerical study of the two body problem, and then a numerical study of the statistical properties of the many body system. To prevent the strong energy dissipation in the many body system we introduced a “memory effect”, and studied the dependence of the equilibrium state on the control parameters of the model. We also verified that for a given range of parameters the equilibrium state of the system results in an ordered structure (a “crystal”) and proposed a model with attractive and repulsive forces in order to simulate a more realistic behaviour.

In chapter 4 we have studied a model inspired by the immune system. T cells, APC and antigens have been represented as agents moving on two overlapping discrete grids, and particular attention has been given to the T cell-APC interaction, simulated as happening on the surface of the cells. Also a differential equation based mean field theory has been proposed and the results between the microscopic agent model and the mean field theory have been compared.

In Chapter 5 we have developed an evolvable crowd dynamics model in which pedestrians avoid collisions in order to safely reach their goal. The model was based on the ability of agents to predict the actions of the others using a “Theory of Mind”, i.e. trying to predict the future actions of the others. The model presented emergent global self-organisation patterns, which resemble those of actual pedestrians. We also used this model to study the evolution of the “level” of Theory of Mind, i.e. we examined for which range of the control parameters in the model agents developed “recursive thinking”.

In chapter 6 we have studied a mobility system on a larger scale, using a discrete space-time representation. Agents are controlled by evolving neural networks, and the model presents the emergence of global and local strategies that optimise traffic flow. In particular, we studied the emergence of a communication system (inspired by pheromone communication in social insects) and the emergence of traffic flow rules.

All the models in this thesis are at a preliminary stage and should deserve more study. In particular, for what concerns the “gas of automata” model in chapter 3, the relevance with respect to more realistic situations, and thus the possible application to the study of actual crowd or animal behaviours, has to be investigated. Also the model on the immune system should be investigated more deeply, and both the space structure of the environment,

and the interactions between cells should be formulated in a more biologically sound way.

The crowd dynamics model in chapter 5 opens different directions of research. Given the ability of the model to reproduce actual self organising pedestrian behaviour, we intend to optimise it and to use it in realistic simulations of pedestrians dynamics (i.e. in environments with a complex topology reproducing actual urban areas or large buildings). We also intend to study more deeply the self organising properties of the model, with a particular attention to their dependence on density, comparing the results with those obtained using more simple models, in the attempt to understand which are the minimal requests to reproduce these patterns.

Finally, we would like to use the neural network based approach of chapter 6 to the physical setting of chapter 5, to verify if self organising patterns can result also from a completely emergent behaviour. Furthermore, a first approach to this problem has shown the emergence of traffic flow rules, similar to those of chapter 6.

Acknowledgements

Algunas veces gano, y otras veces pongo un circo y me crecen los enanos
Joaquín Sabina, “Que se llama soledad”

飲食男女

Important things in life

When I started writing this thesis I asked myself if I should have used the first singular person *I* or the plural *we*. I soon decided that it would have been more correct to use *we*. If you wondered to who that *we* referred, here I'll give you an answer.

First of all I have to thank my Ph. D. advisor, professor Giorgio Turchetti. Not only he has read and discussed all my works, but he has thought and developed the “gas of automata” model, of which I just cured the numerical treatment of the N body problem. Not surprisingly the mathematical insight of that part of the work is extremely higher than that of the rest of the thesis. I have then to thank professors Armando Bazzani, Bruno Giorgini, Sandro Rambaldi and Graziano Servizi for useful discussions and help. In particular professor Servizi has always provided a quick answer to all my questions concerning the C++ programming language and other informatics problems, besides being an excellent system administrator.

All the “young” members of our research group, Olivia Bernardi, Giuseppina Melchiorre, Carlo Benedetti, Marco Brambilla, Massimiliano Capriotti, Luca Cattelan, Enrico Lunedei and Luca Rossi have been extremely useful, first of all because their company made working less hard, and then for useful discussions. In particular Carlo Benedetti explained me the basics of his Coulomb oscillator model which I (mis)used for the N body automata problem.

Everything good in my immune system model is due to Fabio Luciani, while all the bad part is due to me... to the reader the task of understanding how

much Fabio has worked with me... is it sure that I would have liked to collaborate more with him.

Laura Mecozzi worked with me on the numerical study of the two automata problem as part of her undergraduate thesis.

Professor Takaya Arita has been so kind to give me hospitality in his artificial life laboratory at Nagoya University for three months, to discuss my research project and to read my papers even when I returned to Italy. I consider him as a “second Ph.D. advisor”. The “originally Alife” chapters of my thesis have been thought and developed in Japan. Reiji Suzuki gave me a lot of useful advices and probably saved my life carrying me and my 70 kilos of baggage to Nagoya airport. Masanori Takano explained me in detail is ToM model, and Yoshiyuki Nakamichi helped me to prepare the latex file for my AROB 2006 paper.

Artur Matos has helped me to find an apartment in Nagoya, to solve a lot of little and big problems during my stay in Japan, helped me with my research, introduced me to a pleasant Brazilian-Japanese-and-not-only company, tolerated my bad Portuguese, explained me the basics of Buddhism, invited me to 坐禪, suggested me to visit 三十三間堂 with its one thousand 十一面千手千眼觀音菩薩 golden statues, hosted me in his house when I went back to Nagoya... isn't that enough for a *muito obrigado*?

I also want to thank all the people in the lab, all the friends that I met in Nagoya, and even all 愛知県, that has hosted me many times and always made me feel at home. I think that Aichi deserves the wonderful 漢字 (love and “know”) in its name. お世話になりました.

Going back to Bologna, I want to thank INFN for hosting me for three years in the computer room, which turned to be “my” room, and in particular the system administrator Franco Martelli. The main problem of that computer room has been that for most of the time almost no one except me was using it... I thus want to thank all the people who spent part of their time during their graduation thesis (or Ph.D. activity) and saved me from unbearable boredom... hoping to remember every one: Federica Antonucci, Elena Canovi, Lara Giotto, Stefano Bernardi, Fabio Santandrea, and Fulvio Sbisà.

Elisabetta Moroni borrowed for me a book I needed very much at the library of Milano Bicocca University, I book that I lost a few days ago (... don't worry, I've found it again!)... going back to Bicocca I had a talk with Concetta Miccio, one of these talks that helps us going doing research (she knows

why...)

My parents have supported me (least of all, financially) during the last three years, as also during the previous ... ones (I'm writing on the day of my birthday, and these numbers are getting too large for me to be glad to see them on a piece of paper) and this thesis is due and dedicated to them. I'm also glad with them because they gave me a sister, Claudia (well, they actually gave a little brother to her, and I know she wasn't so glad at that time, but I hope and think that she is now), and I want to thank my grandparents who started (at least from my limited point of view) our family, and that are always with me even if they are not Here anymore, and all my relatives.

My cats: Micio James and Rocco.

(Some) people I like: Emaunuele Vercesi, Luigi Discalzi and his family (Marianna and Sofia), Davide Dagrada, Giacomo Guerra and his (brand new) wife Chiara, Marco Sabbadini, Matteo Carotta, Marco Codegoni, Gabriele Vipadi, Federico Carugo and Roberta Carbone (who are married and waiting to enlarge their family, looks like everyone but me has turned into a responsible adult), Marta Zanotti, Chiara Vassena, Helise Stabile (far away so close), Jasmine 辛桂芳 (so far away so close).

Thank to you, that I forgot to put in this list even if you deserved it, please forgive me, I'm very tired.

Obrigado

ありがとう

謝謝

Grazie