Artif Life Robotics (2001) 5:46–57

ORIGINAL ARTICLE

Elpida S. Tzafestas

Aging agents

Received: October 12, 1998 / Accepted: May 21, 2002

Abstract We are adopting Brooks and Wiley's view of evolution as an irreversible process capable of producing increasingly greater complexity at higher organizational levels. We start from the assumption that the evolutionary force is intrinsic in the living system, and is in reality a continuous senescence function leading gradually and unavoidably to death. We are therefore seeking a senescence function that favors social rather than solitary agents in terms of longevity, without prespecifying in detail the agent's life span. We show that a senescence function relying on negative (destructive) feedback links from metabolism to genetic program conforms with these specifications. We also show that senescence should affect all the regulation parameters of the agent, and that the system remains nonmanipulable and unpredictable as far as its life span is concerned. This senescence function favors the more "cognitive" agent models (the ones having additional regulation loops), and thus the emergence of organizations of a higher order that have more elaborate social relations.

Key words Aging · Adaptivity · Sociality · Self-regulation · Autonomy · Emergence

Introduction

Why age?

Artificial life is concerned with the study of "life as it could be, instead of life as we know it".[1] One major research theme is the evolution of structures and organizations. Here, we are adopting Brooks and Wiley's[2] view of evolution as an irreversible process capable of producing increasingly greater complexity at higher organizational levels: "Because we know of no natural hierarchical configurations not produced by an irreversible process, we are logically compelled to accept the proposition that the natural hierarchy of form we see is the result of historical and irreversible processing operating on discrete units" (p. 82). *Our assumption is therefore that the evolutionary force is intrinsic in the living system, and is in reality a continuous senescence function leading the living system irreversibly, gradually, and unavoidably to death.* Such a senescence function may be a true evolutionary force, or else the significant constraint of evolution,[3] and may lead to the emergence of higher-order organizations if it favors sociality between individuals, i.e., if the life span of the social individual is longer than that of the asocial one. In this way, a multicellular organization will have a "selective advantage" over a unicellular one, a society of multicellular agents will have a selective advantage over a single agent, and so forth. Of course, the senescence function should be coupled with a mechanism of detection and reinforcement, or stabilization, of "emergent" properties for higher-order organizations to appear. In addition, the whole process should be shown to be recursive, in the sense that the senescence function will be the same across organizational levels, while the exact form of the emergent structures will vary according to the level. In what follows, we will only be concerned with the first part of this hypothesis, i.e., in finding a senescence function that favors agent societies rather than individual agents; the transition from level to level and the evolution of organizations will be treated in a further study.

What might a senescence function look like?

Our first concern during the development of a senescence model is to maintain a degree of similarity with biological reality, but with the aim of *understanding* this reality rather than using it blindly. The basic model should therefore give rise to a senescence phenomenon that does not appear once the agents have reached a certain age, but that is instead a continuous process during the agent's life, that leads him gradually and unavoidably to death:

E.S. Tzafestas (✉)
Intelligent Robotics and Automation Laboratory, Electrical and Computer Engineering Department, National Technical University of Athens, Zographou 15773, Greece
Tel. +30-1-772 2489; Fax +30-1-772 2490
e-mail: brensham@softlab.ece.ntua.gr

". . . the decline in old age would be the price to pay for the vigor in youth. On the one hand, the aging speed would increase due to factors tending to favor vigor in youth. On the other hand, it would decrease due to other forces tending to delay the deleterious effects. The equilibrium between those opposite forces would hence adjust the aging speed and the life span. (Jacob,[4] p. 93, translated by us)."

Or, in Kanungo's words[5] (p. 267)

"senescence or aging should not be viewed as an isolated and independent phase in the life span of organisms, but should be considered together with development and adulthood phases. These earlier phases may not only influence the organism's longevity but also the rate, duration and mode of its senescence."

The senescence force (or function) should therefore be modeled as a developmental force having the properties listed below.

- Senescence should favor social relations in the sense that social agents live longer than asocial ones.
- Senescence should be continuous and ensure the agent's death.
- However, senescence should not prespecify the agent's life span; the agent should have a limited possibility to act in a way which will relatively extend its life.
- Senescence should be the result of the coupling of two opposite forces, one of which delays while the other one accelerates the agent's death. The agent should therefore search for the proper compromise between these forces that will allow it to live longer.

Note that such a senescence force, relying on the coupling of opposite forces that do not prespecify the agents' life span, leads to a system not directly controllable and manipulable, as will be shown below.

Senescence model

An agent has two "behavioral" parts: its *metabolism* (or its body) that is responsible for the "consumption" of messages coming from the outside (and is therefore responsible for the self-regulation of the agent in front of the world), and its *genetic program*, that determines and codes what the metabolism does. Dyson[6] drew the analogy between metabolism and hardware, and between genetic program and software, to develop his theory of a double origin of life. According to his view, the genetic program (or software) passes from generation to generation (and the agent cannot explicitly control or modify it during its life; it is his *inert* part). On the other hand, the metabolism (or hardware) is important only to a particular individual (it is therefore the part where learning is allowed). For instance, all the environmental and social adaptation functions are essentially metabolic functions, but are regulated by the agent's genetic program. The agent's *physiology* is then the mode of coupling between its genetic program and metabolism. The designer of an artificial agent works at the program level to try to adapt it to the specific features of the implementation medium that constrain the possible forms of the metabolism.

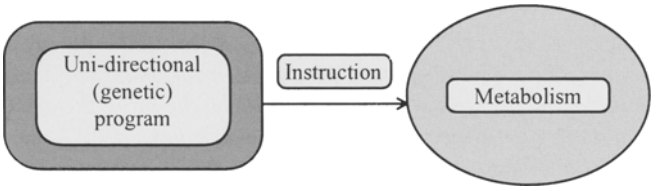

Fig. 1. The agent model without senescence

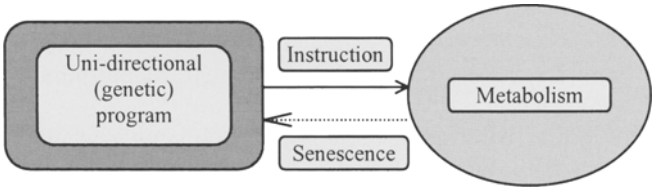

Fig. 2. The model of the senescent agent. In what follows, the *broken lines* represent the senescence feedback

In order to find a senescence model which meets the given specifications, we imagined an agent physiology such that the genetic program and the metabolism of the agent are coupled with negative feedback links of the following type: each time an agent metabolizes, it is because its program instructed it to do so, and it does so in a direction that makes it more adaptive to its environment, but every metabolic action acts in a progressively destructive manner upon the program, and therefore upon the future possibility of metabolizing. The senescence function is therefore a negative feedback link added to the instruction of the metabolism by the agent's program (this situation is represented schematically in Figs. 1 and 2). The simplest model complying with this idea is that of unidirectional (linear) programs with threshold effects that determine the agent's potential to metabolize. The metabolic rate, which depends not only on the program but also on the environment in which the agent is situated, is responsible for the senescence rate, and is self-catalytic. This idea of negative self-catalysis has been inspired by browsing the biological theories of senescence,[7] and more specifically by the view of senescence not as a single mechanism, but as a whole set of mechanisms acting at multiple organizational levels and probably interconnected, so that practically every existing theory is partially true[8]: senescence has to be programmed but also not programmed, continuous but also catastrophic at times, etc. More precisely, Zs.-Nagy's membrane hypothesis of aging[9] and Kanungo's gene regulation theory[5] attempt a synthesis of experimental data and existing theories under a common level-independent principle, i.e., progressive membrane deterioration and progressive gene expression deterioration, respectively.

As well as Dyson's theory,[6] the importance of metabolism is also a foundation of the theory[10] that free-radicals have a deleterious effect on the DNA of somatic cells, as well as on proteins and other essential molecules, and whose accumulation depends on the metabolic rate. For example,

Rusting[8] (p. 91) notes that the mean life span of various species is inversely proportional to their metabolic rate. *Senescence is therefore the wear of the program–metabolism connection, and this wear is unavoidable: the more the agent metabolizes, the more its program deteriorates.*

Usually, the genetic program of an agent is represented as a sequence of genes, each acting as a parameter of the (metabolic) behavior of the agent. An implementation of the senescence model above can be derived in one of two ways:

- either by acting directly on a gene's value every time the gene fires (is used in the agent's context to determine its behavior);
- or indirectly by acting on the gene expression mechanism every time the gene fires.

In both cases, the results are equivalent, and the gene values change in a way that makes their future effect less and less pronounced. The exact way the gene values change (directly or indirectly) depends on the details of the particular phenomenon modeled. However, as will be shown later, those gene modifications should follow some not so simple rules in order to conform to the specifications of the previous section.

In the implementation that follows, we regard simple numeric parameters of the agent's behavior as genes, whose values change every time they are used. The genes and their functions are described as they are encountered. The concept of gene modification in relation to genetic representation and senescence will be discussed again later in view of the results obtained below.

Application domain

To implement those ideas, we developed a senescence model for a problem that is apparently simple, but sufficiently abstract to allow for the observation of complex enough phenomena and serve as a validation basis. A population of agents is initialized in a closed space; each agent has an individual goal and a possibility to "socialize." Furthermore, an agent may encounter obstacles that, as their name indicates, will distract it from its goals. Eventually, the agent will meet its goal and stop, i.e., it will die; the multiagent system will be dead when all its component agents die.

To simulate the system, the notions of space, goal, and obstacle have been instantiated as follows. The closed space is a closed 2D grid. The goal is a point in that grid toward which the agent heads, obstacles are "inert" objects that the agent destroys, and sociality relies on the range of perception of other agents (in what follows, we will often refer to this range as the sociality factor, or simply sociality). Sociality distracts the agent from its goal by inducing it to "lunatic behavior": the agent turns around itself without moving and, eventually, when a time-out expires, it advances (a step forward) but in a random direction, since it has been perturbed. As far as arbitration between those three behaviors is concerned, the destruction of obstacles takes priority over the social distraction, which in turn takes priority over the heading toward the goal; arbitration is thus static. This setup is summarized in Figs. 3 and 4.

The view of sociality as a distraction is based on the following observation: an agent has its own goal point that it tries to reach. If it encounters other agents that also have their own goals, it is distracted from its goal to participate in social activity, and so it achieves its goal later and lives longer. In an engineering or design context, sociality as a distraction would have an operational meaning: the global problem, that of the multiagent system, would be better solved with sociality/distraction than without.

We have implemented and compared three successive behavioral models. The first is a "reactive" model (reactive

Basic behavioral model

1 ***Distraction***. *If there is an obstacle in the agent's position, then the agent destroys it.*

2 ***Social behavior***. *If there are other agents in the perception range (sociality), then the agent is perturbed and after a timeout it migrates elsewhere.*

3 ***Goal-direction***. *If its goal is met, then the agent stops, otherwise it heads toward this goal.*

- *There may be from 1 to 5 obstacles in the same place (this corresponds to obstacles of varying degree of difficulty).*
- *The agent's perception range (its sociality) takes a value between 1 and 5.*
- *The social distraction timeout takes a value between 1 and 10.*
- *The goal-point is randomly initialized in the grid.*

Fig. 3. Basic behavioral model. For all subsequent aging models, the environmental and behavioral parameters have been tuned in a 40 × 40 world in a way that will amplify the observed phenomena, such as distraction. A few sporadic experiments with other parameter settings showed that those phenomena did not change *qualitatively*

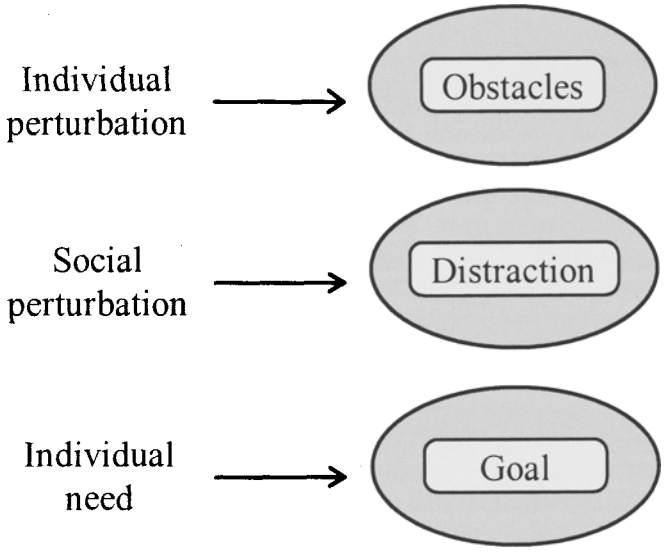

Fig. 4. Abstract behavioral model, i.e., a motivational system with static arbitration between the three tasks of an agent. The goal is a point in space (in a 2-dimensional grid), obstacles are objects whose presence delays the agent, and distraction induces a random change of heading

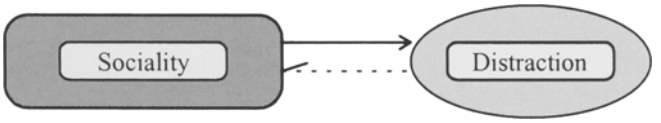

Fig. 5. Reactive model. The agent's sociality, i.e., the range of social perception, is negatively catalyzed

as far as its primary need is concerned), the second is an "autonomous" model (motivated and senescent), and the third is a "cognitive" (self-regulated) model; these three models are presented below. Senescence is continuous; it intervenes to the physiological level as a negative feedback link, and it mainly affects sociality. In fact, we will see that the system dies slowly owing to the decay of the agents' sociality: since sociality decreases, the inter-agent relations degrade and the system loses its identity as a coherent system; it decomposes. The sociality parameter simultaneously expresses selfishness and tolerance: it expresses selfishness since the agent is distracted from its goal, and it expresses tolerance since it allows the agent to live longer than if it were not social.

We will show that senescence should affect all the regulation parameters, and that the system remains nonmanipulable and unpredictable as far its life-span is concerned. We will also show that senescence favors the more cognitive agent models (the ones having additional regulation loops), and thus it favors the emergence of organizations of a higher order that have more elaborate social relations.

As well as theoretical and simulation work, we also had the system observed by volunteers (mostly artificial intelligence researchers) to identify the relation between the real mechanism underlying the behavior of the simulated agents and the conclusions drawn through observation to respond to the question "what does the system do and why do agents die?" As a result, it was found that all observers would come to some sort of top-down conclusion, such as "the agents are trying to create configurations with high agent concentrations in some areas," while they remained unable to describe how the agents would act toward such a goal. Furthermore, they appeared systematically to ignore behavioral regularities (for instance they would find a straight walk pretty dumb) and underestimate the value of possible manipulations of the simulated system, such as adding, deleting, or moving agents, etc. Overall, they would be misled by the complexity and diversity of the emergent behavior of the system toward elaborate, complex, and nonmechanistic explanations of individual agent behavior, and away from simple, deterministic models such as the one presented here. While we will not describe in detail the observation procedure and the results obtained, we will occasionally repeat some telling remark from an observer whenever this is useful to the discussion.

The aging mechanism

A "reactive" model

The first senescence model is founded on the observation that if an agent's sociality is constant over time, the multiagent system may be "immortal," i.e., it may stay alive for ever. This may happen if two agents have neighboring goals *and* the sociality of either one of the two is sufficiently high, i.e., superior to the distance between the two agents' goals. The solution appears to be a mechanism that makes sociality decay gradually. The first idea was to make sociality fade "naturally," i.e., for a reason intrinsic to the agent and independently of its activity. However, this would lead to a "preprogrammed death" phenomenon, and, as we have already explained, we are interested in senescence models where agents are mortal *without* the conditions and the context of death being preprogrammed. We have therefore introduced a feedback link between the effects of sociality (distraction) and sociality itself: every time the agent socializes, i.e., every time it is distracted from its goal, its sociality decays by a small proportion less than 1 (the time-out decays by the same factor). The idea of a time-out that makes sociality decay and that decays itself in the same way is consistent with the specification of senescence as a self-catalytic phenomenon that is observable, but not directly controllable (cf. the discussion below). This first senescence mechanism is presented in Figs. 5 and 6.

Observation and experimentation with this system have shown that it is fairly manipulable: if we move an agent away from its goal, it will always try to return to it, despite the presence of obstacles and of other agents, but also despite all these manipulation trials. Successive manipulations will lead to a decay of sociality close to 0, so that the agent will no longer be sensitive to the presence of other agents. However, since the manipulations do not directly affect the

Fig. 6. Basic aging model (reactive model)

Reactive model

Program : system of sociality decay (otherwise, adaptation)

Timeout or social distraction

$\Rightarrow$ *Sociality := sociality*(1-f_s) & timeout := timeout*(1-f_s)*

- *Genes : f_s, timeout, sociality*
- *Initial values : $f_s \in (0.05,0.15)$, timeout(0) $\in (5,10)$, sociality(0) $\in (1,10)$*

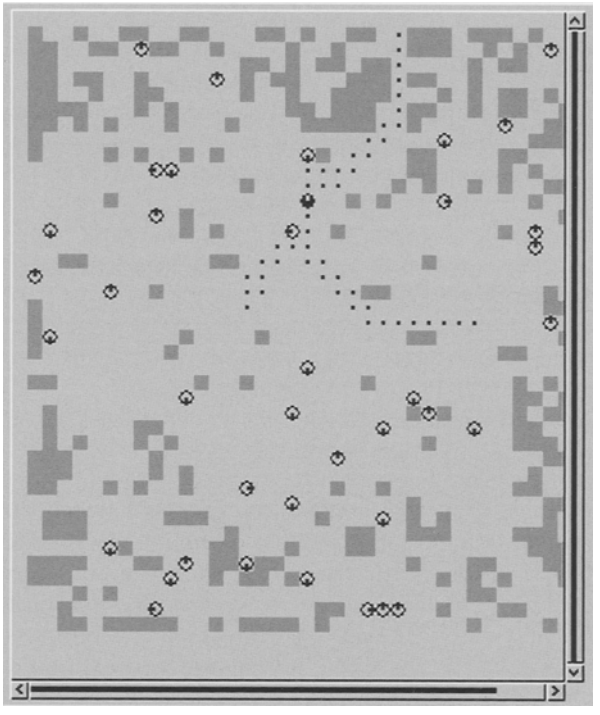

Fig. 7. Manipulability in the reactive model. The black agent's original trajectory (denoted by a series of *small dots*) is the one on the left, and the other two took place after manipulation, i.e., repositioning of the agent in a different position. The agent always returns to its initial goal. The obstacles are visualized as dark

agent's goal, it remains manipulable by an external agent who, after observing the system, has discovered the agent's goal-position (this is visualized in Fig. 7, which gives the agent's trajectory after multiple manipulation trials). I discovered that the system was manipulable in this sense when some of the volunteer observers asked me to create specific agent configurations that they presumed would be interesting according to their previous observations; one of these configurations involved a single agent alone in a world (the observer remarked in this case, "it isn't interesting, this agent follows a very regular path, so when an agent is alone, he is stupid"). In short, *the system is reactive as far as its goal is concerned*; what appears to be necessary is a feedback at the level of the goal, and this has to be dynamic.

An "autonomous" model (motivated and senescent)

Having a dynamic goal necessitates a goal-revision mechanism. For the basic behavioral model (Fig. 3), goal revision means to adopt a new goal position in space. The important question is, *when does an agent have to revise its goal?* Following the observations in the last section, goal revision is necessary if there has been a lot of manipulation trials. Since an agent has no means of detecting that it has been manipulated (unless it possesses a gyroscope, an odometer, or other sensor of a similar type), the only criterion it can use is a temporal one: with a time-out, the agent revises its goal. It also appears natural to say that an agent will revise its goal on attainment in an effort to live longer. In any case, we run the risk of again having the same manipulability problem on a metalevel if we do not introduce a goal-revision criterion that is true in the beginning and that gradually becomes false.

We have therefore introduced an additional variable – called *adaptivity* to show that it expresses the agent's possibility to adapt or revise its goal – and a threshold to which it is compared. If adaptivity is superior to this threshold, the agent can revise its goal, otherwise it cannot. Feedback affects the adaptivity parameter to ensure the nonmanipulability of the system. As a negative force acting on adaptivity, we have introduced a temporal adaptivity degradation factor. Finally, for the senescence mechanism to be self-catalytic, we also had to introduce two parameters of adaptivity regeneration on goal attainment, or whenever there is social distraction. In this way, *the two parameters of adaptivity and sociality are coupled in a double loop*: if the agent is adaptive, then it discovers a new goal and it runs the risk of encountering other agents and being distracted. If the agent is distracted, its adaptivity rises. This double loop is a reinforcement, or self-catalytic, loop, since by itself it leads to a system where adaptivity constantly increases. However, we have already seen that the agent's activity leads to a decrease in sociality, which will, in turn, prevent adaptivity from increasing. If, on top of these, there is a temporal degradation mechanism, adaptivity will eventually decrease, and sometimes it will show local maxima. An agent complying with this senescence model is therefore more autonomous than the previous one (since the user cannot keep it alive artificially by persistent manipulations) because it shows a relative au-

tonomy as far as its goals are concerned. Its problem is how to socialize enough in order to live a long time, but not socialize too much, otherwise it runs the risk of not achieving its goals in time and dying prematurely. The model is summarized in Figs. 8 and 9.

As expected, the agents in this model are not as manipulable as those in the previous model. Figures 10 and 11 show the trajectory of an agent and the effect of multiple manipulation trials. We observe that after persistent manipulation trials, an agent may revise its goal or die, and that the more an agent has "worked" during its life, i.e., the more goals it has pursued, the less manipulable it is and the more susceptible to death (as if it were exhausted).

Adaptivity is the crucial parameter that shows if an agent is "alive" or not. Sometimes, the maximum adaptivity is at the agent's birth, so that the agent remains alive by "inertia," as if it were condemned right from the beginning (and in a sense it is). Otherwise, the adaptivity curve of a multiagent system (which is computed as the mean of the adaptivity variables of the individuals) will be bell-like, whereas the adaptivity variables of the individuals may show one or more local maxima (Figs. 12–14). Adaptivity is also a measure of the manipulability of an agent in a social context: if a moving agent passes near a dead agent whose sociality is high enough and whose adaptivity is close enough to its threshold, the latter may be resurrected and set off toward a new goal. Consequently, a system may be dead while its components would be very alive in another

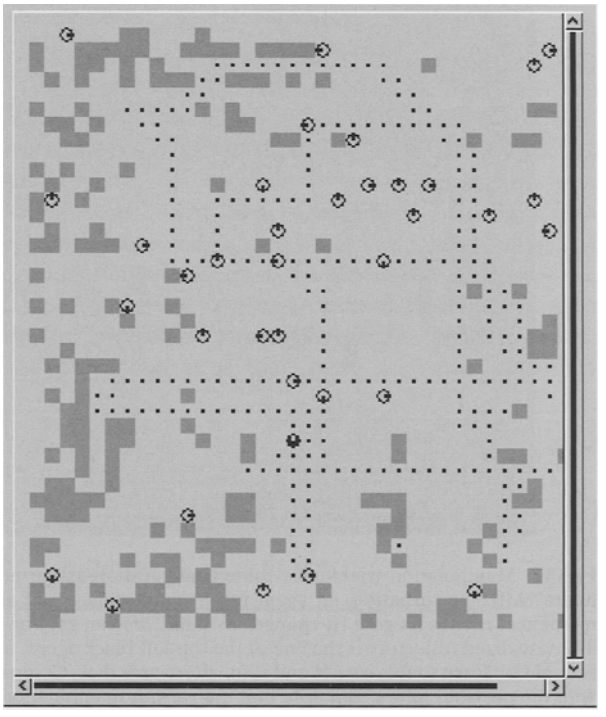

Fig. 10. Manipulation trial in the autonomous model. The original trajectory of the black agent passes through multiple goals (as can be seen from the changes of heading in the trajectory)

Fig. 8. Autonomous model. The adaptivity of the agent is negatively catalyzed. The positive reward links are not visualized

Fig. 9. Autonomous model (motivated and senescent)

Autonomous model (motivated and senescent)

Program

If the goal is met or timeout, then the agent sets a new goal

(if adaptivity ≥ threshold, then (it is adaptive) new goal),

otherwise death (death is "simulated" as goal=current_position)

Temporal decay	:	$adaptivity := adaptivity * f_0$
Regeneration whenever a goal is met	:	$adaptivity := adaptivity * (1+f_1)$
Regeneration during distraction	:	$adaptivity := adaptivity * (1+f_2)$
Decay of f_1 on goal revision	:	$f_1 := f_1 * f_0$,
Decay of f_2 on distraction	:	$f_2 := f_2 * f_0$
Decay of timeout on goal revision	:	$timeout := timeout * f_0$

- *Genes : f_0, f_1, f_2, timeout, sociality, adaptivity, threshold*
- *Initial values : $adaptivity(0) \in (0.5,1)$, $timeout(0) \in (50,100)$, $threshold \in (0.1,0.2)$, $f_0(0) \in (0.75,0.95)$, $f_1(0) \in (0.5,0.9)$, $f_2(0) \in (0.1,0.3)$*

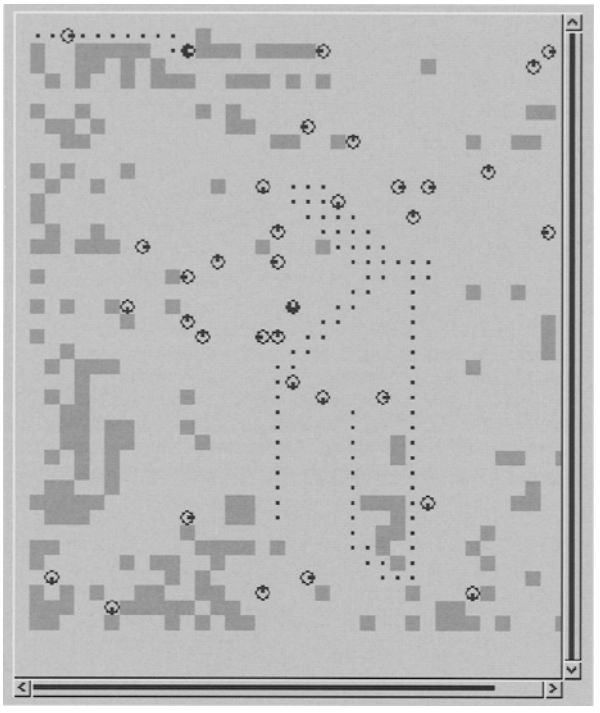

Fig. 11. Manipulation trial in the same configuration as the previous figure. After perturbation, an agent retries many times, and at some moment it revises its goal (it changes its mind, we can say) or it dies. The visualized trajectory is the one of the top-left black agent, which is brought by force to this corner and soon afterwards dies. Compare this with the previous figure, and note that the various manipulations have slightly perturbed the overall system; e.g., the black agent of Fig. 10 is now situated in the middle of the grid

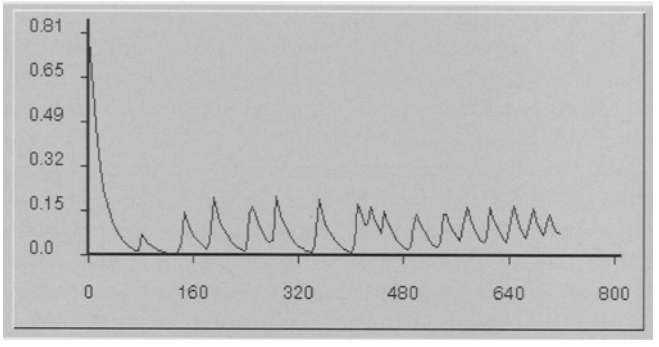

Fig. 12. The adaptivity curve (the mean adaptivity across agents) is bell-like. One can say that the agent's adaptivity curve shows its developmental history: the system develops up to the point of maximal vigor, and then declines until its death. As far as sociality is concerned, it decreases continuously owing to the temporal decay mechanisms and the negative feedback loops

social context. (In a real organism, many cells are individually alive when the organism dies.[11])

This brings us to another question: can we (and *should* we) intervene in a dying system to increase its adaptivity? This looks possible only in the first stages of the system's (and the agent's) life, where one can, for instance, replace an agent by a "younger" one or inject new agents. However, the more the system ages, the greater the number of agents that would need such a "replacement." What is even more

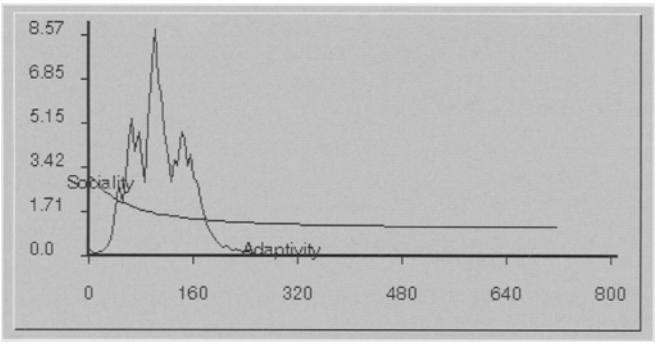

Fig. 13. An agent's adaptivity shows small jumps during each "regeneration"; this agent's parameters at death are adaptivity = 0.0918, sociality = 0.941, time-outs = 0, revision trials = 16, regenerations = 15

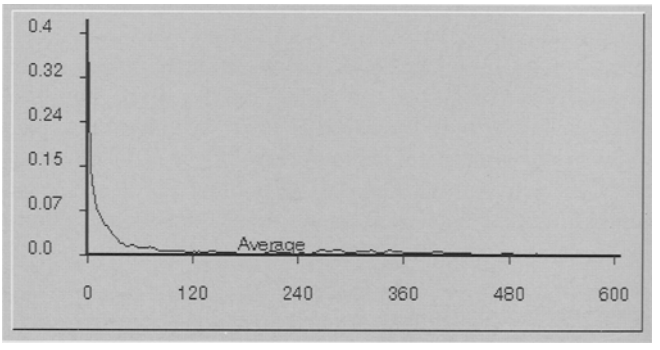

Fig. 14. Sometimes the adaptivity decreases continuously from the beginning, with the global maximum being at $t = 0$ (this is the case of a system of 100 agents that died at $t = 514$)

important, is that if social relations were "manufactured" by the agents themselves, such a renewal of a system of agents would be possible in principle but impossible in practice (unless we had access to the whole network of social relations between agents, as well as to the details of their goal-revision mechanism, in which case we would do better to intervene drastically and cut once and for all those deleterious feedback links). Note, finally, that the goal revision mechanism induces some sort of diversity in the spatial relations that in turn determine the social ones.

A "cognitive" model (self-regulating)

The model of the previous section relies on a network of coupled functions that make certain physiological parameters of the agent decay. An agent could live longer if it had the possibility to "delay" the degradation of those parameters with the aid of a mechanism having a "regenerative" effect on them. Such a regenerative mechanism can be seen as a "cognitive" mechanism, i.e., one that possesses knowledge, or in our case metaknowledge of the agent upon itself. Delay of the degradation may be achieved in the case of the previous model through regulation of the same parameters, i.e., through self-regulation of the agent itself. In an effort to show that the relation between individuality and sociality is conserved when one ascends regulation

Fig. 15. Cognitive model (self-regulating)

Cognitive model (self-regulating)

1 Individuality regulation

Goal distance regulation system

Goal met $\Rightarrow$ *distance := distance*(1+r_D)*

Timeout $\Rightarrow$ *distance := distance*(1- r_D)*

new goal's distance := distance(1+noise)*

2 Sociality regulation (restoration)

Sociality factor regulation system

with a regulation (adaptation) window w_S

If sociality < eigen_sociality,

sociality := sociality + r_S(eigen_sociality-sociality)*

r_S := r_S(1-dr_S) (r_S decay)*

and w_S := w_S(1-dr_S) (w_S decay)*

- *Genes : As before plus distance, r_D, noise, w_S, r_S, dr_S, eigen_sociality,*
- *Initial values : As before plus $r_D \in (0.1,0.2)$, $r_S \in (0.2,0.5)$, noise $\in (-0.1,0.1)$, eigen_sociality $\in (2,4)$, $w_S \in (20,30)$, $dr_S \in (0.01,0.05)$*

levels, we have implemented and compared two self-regulation mechanisms: a mechanism of regulation of the distance of the new, revised goal, and a mechanism of sociality regulation.

The first is a mechanism of regulation of the agent's individuality, and is based on the following observation: if an agent follows a very distant goal, it runs the risk of getting too distracted with respect to its goal, and therefore it will time-out prematurely and have to revise its goal (as a side effect, the sociality time-out itself would decrease). Regulating the distance of the new goal means that the agent will gradually choose closer or farther goals according to whether the goal revision is due to a time-out or to a goal achievement, respectively. This mechanism tends to regulate the agents' behavior in a way that will avoid too frequent socialization.

The second mechanism regulates sociality directly, and is based on the observation that the system's sociality always decreases, and this is why the system dies. Thus, we can think of direct regulation of sociality, but an absolute criterion concerning the direction of regulation (increase or decrease) cannot be found, since there is no means of knowing whether an agent has socialized too much or too little, or to estimate the appropriate value for the sociality factor in a particular configuration. We have therefore adopted the alternative of regulation according to an endogenous reference value (an eigen-sociality): since the problem is that sociality decreases, we will try to restore it if its value falls below the reference value. The two models are summarized in Fig. 15.

Fig. 16. Cognitive model 1. The distance regulation is independent of the agent's activity and is not affected by the senescence mechanism

Individuality regulation

Individuality regulation (cognitive model 1, Fig. 16) is regulation of the distance at which the new goal is chosen during revision. At the beginning of the simulation, this distance is chosen randomly between 0 and the maximal possible distance in the grid. Unlike the previous models, this one does not necessitate a negative feedback link affecting the regulation parameters (here the r_D parameter) because the goal distance is not a parameter whose decrease (or increase) would explicitly be connected to the senescence function; distance regulation is just a means to regulate the sociality indirectly.

This regulation mechanism has been introduced to show that the agents managing to live longer are those that are more adaptive in their search of new goals, i.e., those that, in one way or another, "learn" from their failures (for instance, because they choose successive goals close to one another). However, this regulation mechanism did not lead to a considerable increase in life span as compared with the previous model: actually, the performance has fluctuated above and below that of the previous model. Why? One

reason is that individuality plays only a minor role in an agent's survival in a society. However, the most important reason is that this regulation may have an effect only in the case where it would make sense to *choose* the direction of distance regulation, and thus of sociality adaptation, i.e., to *choose* between an increase or a decrease of sociality. On the other hand, the next model relies on a direct regulation of sociality, and it will be compared to this one later. For the time being, it is sufficient to note that indirect sociality regulation via distance regulation has had an unstable performance, because the new goal distance by itself does not induce a specific degree of sociality. The justification is that there remains an uncontrollable parameter in this model, which is the actual configuration of other agents: choosing a closer goal in order to restrict the consequences of sociality makes no difference if all other agents are nearby.

Sociality regulation

We implemented this direct sociality regulation mechanism (cognitive model 2, Fig. 17) to show that, compared with the previous models, it leads to a longer life span for the multiagent system.

Comparative results and evaluation

Table 1 gives the comparative results obtained with all four models for various agent population sizes. Except for the instability of the goal distance regulation mechanism, the other results are as predicted: regulation makes the system live longer.

The agent's physiological network is presented in Fig. 18. Its essential elements are the negative instruction and feedback loops (a few secondary positive feedback links and the internal loops of the "programmatic" parts are not visualized).

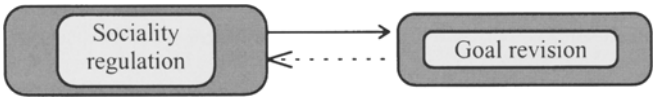

Fig. 17. Cognitive model 2. The sociality regulation has to be a senescent mechanism as well

Discussion

Time and knowledge

The critical parameter of the senescence model is time: the dynamics of the metabolic rate of agents is responsible for the system's aging rate, *although the two rates are not the same, and each one is the result of a complex system of interactions*. The self-catalyzed feedback loops have the consequence that the dynamics remain irreversible, and the system is sure to die despite the unpredictability of its emergent behavior. This means that the dynamics and the lifespan of the system cannot be deduced solely from the agent's parameters: outside an interaction context, an observer cannot say much, and the same agent in two different contexts may give rise to different phenomena/behaviors. In this sense, *time as a design parameter does not have the same "semantics" as time as an observation parameter*. Time

Table 1. Comparative table of life spans in the three models (in a 20 × 20 world). The results are average values over 20 simulations per case

	Reactive model	Autonomous	Cognitive 1	Cognitive 2
10 agents	172.4	245.9	241.85	339.25
20 agents	225.6	274.4	280.85	882.95
30 agents	255.95	271.05	292.1	1146.95
40 agents	293.3	280.15	291.3	1355.55
50 agents	298.35	300.65	308.4	1440.3

Fig. 18. Cognitive model: complete physiological network. *The global adaptivity of the agent is negatively catalyzed*

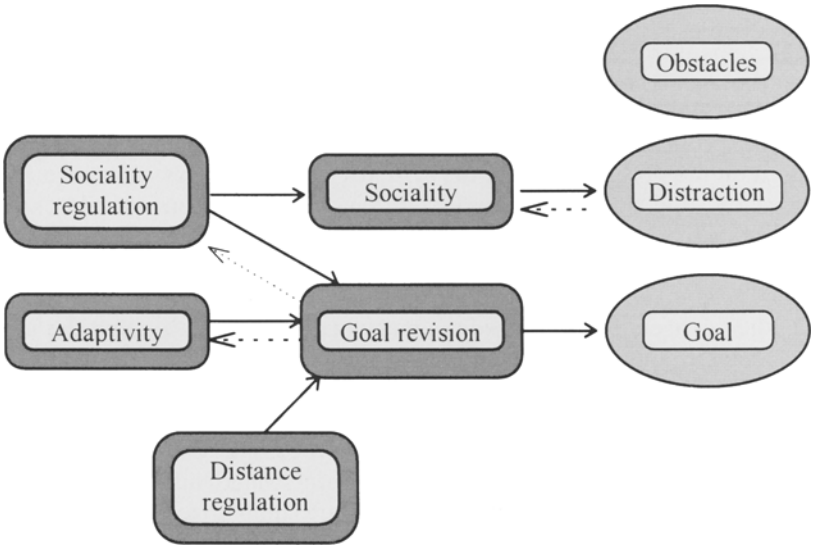

as an observation parameter may have more than one sense: as observers of the system, we can *choose* to observe the life-spans, or certain cycles generated by the emergence of various forms. However, for the agent itself, time as a design parameter has one single meaning, i.e., its survival potential, which is not independent either from the other behavioral parameters of the agent or from its interaction with its environment and with the other agents. An agent with a high time-out is expected to live a long time provided that its perception range allows it to socialize enough but not too much. As stated elsewhere, in a system founded on a network of feedback links between metabolism and program, the life-span is neither prespecified nor specifiable. Furthermore, the life-span *should not* be specifiable, otherwise an external observer could artificially manipulate the system and extend or shorten its life-span at will. We can reverse this conclusion to say that *what is not specified cannot be directly controlled, and therefore the behavior of a truly autonomous system must show a degree of emergence.* Once more, an external manipulator agent maintains only a limited intervention freedom: I managed to double the life-span of a system by injecting brand new agents from time to time at selected positions. However, this freedom results from the system's own potential without being connected in any explicit way to the manipulator's intentions.

We also saw that the system's life-span may increase if the agents have the possibility to modify their own dynamics of interaction with the world and with the other agents. The important result is that, despite this possibility, the life-span of the system is never specified in advance. Our view of time as a parameter of the aging process is consistent with the current trend in aging research: ". . . we need to make time an independent rather than a dependent variable in our analyses. Instead of using the calendar to measure aging, we need to be able to use the changes in important physiological variables to measure aging." (Arking[12], p. 11)

Regulation and feedback

We have shown that the life-span of the system increases if agents possess implicit or explicit mechanisms of sociality regulation, which is the crucial parameter of senescence. The parameters of those self-regulation mechanisms have been defined as senescent since this ensures the system's mortality. We managed to show the occasional immortality of systems whose regulation parameters are not negatively catalyzed by the same senescence feedback links (occasional immortality means that all systems having this behavior are not necessarily immortal).

First, Fig. 19 shows the sociality and adaptivity curves for a cognitive-type system without negative feedback on the sociality restoration factor; this system has become immortal. Next, Fig. 20 shows the behavior of a system of the same type with negative feedback on the restoration factor, but without negative feedback on the restoration window. This system is mortal, but the decrease in sociality is linear, and therefore after a certain point it is more "predictable." A third system of the same type with feedback at all levels is visualized in Fig. 21, which shows a sociality evolution typical of self-catalytic systems, and more specifically a phenomenon of avalanche or catastrophe. We have also shown that a second-type system (autonomous model) without a decrease in factors f_1, f_2, and *time-out* can also become immortal, but this time its immortality shows not as sociality

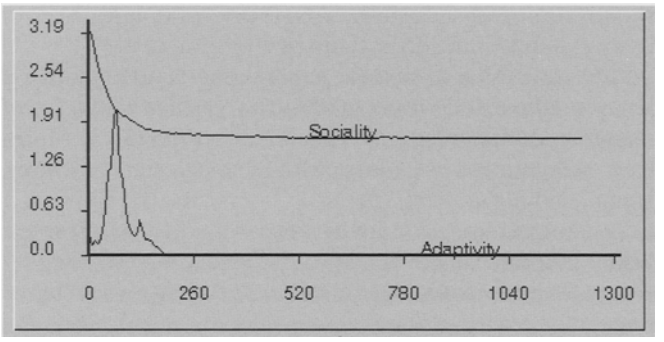

Fig. 19. Demonstration of the need to act on the regulation parameters. The adaptivity and sociality curves of an "immortal" (cognitive type) system that stabilized to a sociality value of 1.65 with an adaptivity that converges to 0. This result is due to the absence of negative feedback on the sociality restoration factor

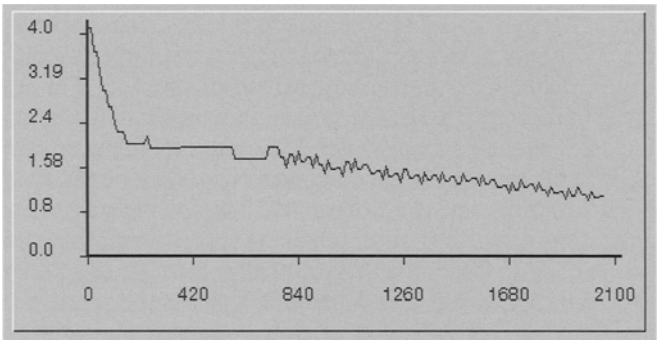

Fig. 20. But *all* the regulation parameters ought to be affected by a senescence link. This is the curve of an agent's sociality in a "cognitive" system with negative feedback on the sociality restoration factor, but not on the restoration time-out. The restoration trials are therefore of uniform frequency, so that the sociality decrease is nearly linear

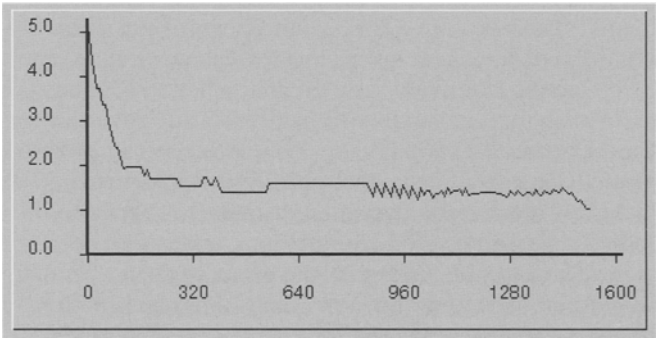

Fig. 21. An agent's sociality in a "cognitive" system with feedback at all levels. The sociality restoration trials are not of the same frequency, and sociality shows a final "abrupt" fall, i.e., a catastrophe or an avalanche effect (since the restoration factor decrease is self-catalyzed through the decrease of the restoration window)

stabilization with adaptivity converging to 0, but as adaptivity stabilization with sociality converging to 0.

The conclusion from these experiments is that *a system is senescent if the parameters of its highest regulation level (and therefore its highest cognitive level) are decreasing*. Moreover, self-catalysis as defined, with recursive feedback links, is not a simple embellishment of the mechanism to bring it closer to the biological reality: self-catalysis is the simplest aging mechanism (the preprogrammed, monotonic, decrease being excluded because of predictability and relative manipulability). *Self-catalysis is therefore a necessity, and in its absence the system would run the risk of becoming immortal.*

Genetic representation and senescence

In view of the above experiments and discussion, it is interesting to observe that all of the parameters we have qualified as "genes" are subject to one-way modification except the adaptivity parameter, that is subject to complex modification defined by three rules instead of one (see Fig. 9). Adaptivity is, precisely, the crucial parameter that characterizes the agent's vitality and determines its life-span. An agent with genetically encoded adaptivity may be manipulated by direct intervention at the level of the adaptivity gene, and so a truly autonomous senescent agent should not encode this parameter genetically. Instead, adaptivity should be a variable of the metabolic part of the agent, but of course, the initial value of this parameter may still be genetically encoded. More generally, *all parameters that undergo complex modifications during an agent's life should be implemented as metabolic variables rather than be genetically encoded in order to ensure maximal autonomy.*

One can still intervene in the other genes and indirectly influence the agent's life span, but this is not a safe method according to the above discussion, i.e., results are not guaranteed. The deeper reason for this phenomenon is that the behavior of the agent, as well as its actual life span, are a result of the interplay between the genetic program and metabolic "action" modulated in its enclosing environment. Furthermore, *the various genes exhibit complex interactions through the agent's body*: the expression of a particular gene yields different results at different moments because of its indirect dependence on the previous expression of other genes (for example, in Fig. 9, the effect of f_1 on adaptivity will differ depending on how many times f_2 has already fired, and vice versa).

Conclusion

We have presented a senescence model using feedback between the metabolism and the genetic program of an agent, and three successive models that increase the life expectancy of the population. It has been shown that every additional regulation level leads to an increase in the system's life-span, and that this span is neither prespecified nor deducible from the phenomena. It has also been shown that in this social system, sociality regulation is much more important than individuality regulation, since sociality is itself much more important than individuality. To ensure the agent's mortality, the senescence feedback should affect all the parameters of the highest regulation level.

This mode of coupling between genetic program and metabolism favors the emergence of nested regulation loops that increase the system's life span. We have not described *how* these loops or regulation levels could emerge or evolve, we have simply demonstrated their "selective" advantage in terms of longevity. Table 2 recapitulates the major points of the approach, the application, and the results obtained.

Table 2. Recapitulation of the approach and the results

The problem	A model of continuous senescence that ensures the death of the agent without prespecifying its life span This senescence model should favor social rather than solitary agents
The application	An abstract problem of goal direction with sociality that shows as distraction from the goal Goal revision Regulation of the distance from goal, or of sociality
The solution	Senescence = negative feedback from the metabolism to the genetic program
The conclusions	• Reactive model ⇒ manipulability • Reactive model vs. motivated (autonomous) model vs. self-regulating (cognitive) model • Regulation of the distance from the goal is less efficient than sociality regulation • Each subsequent regulation level increases life span • Regulation of the distance from the goal ⇒ unstable performance • The senescence feedback should affect all the parameters of the highest regulation level • Some parameters should not be genetically encoded

As immediate further work, we plan to use this model in systems with more than one type of agent, which is equivalent to a multicellular organism with differentiated cells. We intend to perform two studies, one on reproduction, where we will try to reproduce the classic results of Hamilton[13] at a suborganism level, and another on the emergence of cancer in such systems. We have had some initial clues[14] that when the number of agents becomes unusually high, the system dies sooner, thus exhibiting a cancer-like behavior. The case of differentiated multiagent systems would be a much better vehicle for this type of study.

As well as such theoretical considerations, this study of senescence is also of some practical importance: since such a senescence function is thought to favor social relations at an arbitrary organizational level, we intend to use it as a learning force at the cellular level within an artificial animal. We[15] have already reported that a network of cells self-organizes spontaneously in cases of individual cell failures, and tries to "discover" new social interactions. A natural degradation force, such as a senescence function, might induce "natural" failures by perturbing the network, and

could maintain self-organizational activity at its maximum, thus maintaining the activity of the network (and keeping the agent alive). As a side effect, the new social interactions within the cellular network might give rise to what an external observer would perceive as learnt behaviors. Note that this view of cellular learning as the emergence of new social structures within the cellular agent implies that *learning, too, is an irreversible process* and that *memory is not cyclic*. The cellular agent will then learn not because of any externally imposed forces, but for an intrinsic reason, that of living longer.

Acknowledgments I wish to thank Bernard Victorri for useful suggestions. I also thank the volunteers who observed the system: Bruno Bouzy, Effie Cassoura, Gilles Corron, Philippe Darche, Michel Futtersack, Zahia Guessoum, Philippe Laublet, Laurent Magnin, Dominique Pastre, Geber Ramalho, Maria Rifqi, Raphaël Rispoli, Pierre-Yves Rolland, Nikitas Sgouros, Costis Zikidis.

Part of this work was completed while the author was a PhD candidate in the LAFORIA–IBP (Laboratoire Formes et Intelligence Artificielle, Institut Blaise Pascal, Université Pierre et Marie Curie, Paris, France) under the supervision of Prof. Jacques Ferber. An earlier version of this article has appeared elsewhere.[14,16]

References

1. Langton CG (1988) Artificial life. In: Langton CG (ed) Artificial Life. Proceedings of an Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems. Addison-Wesley, Redwood City
2. Brooks DR, Wiley EO (1988) Evolution as entropy: Toward a unified theory of biology, 2nd edn. University of Chicago Press, Chicago
3. Stewart J (1994) Un système cognitif sans neurones: les capacités d'adaptation, d'apprentissage et de mémoire du système immunitaire. Intellectica 1(18):15–43
4. Jacob F (1981) Le jeu des possibles. Essai sur la diversité du vivant. Éditions Fayard, Paris
5. Kanungo MS (1994) Genes and aging. Cambridge University Press, Cambridge
6. Dyson F (1985) Origins of Life. Cambridge University Press, Cambridge
7. Warner HR, Butler RN, Sprott RL, et al. (eds) (1987) Modern biological theories of aging. Raven Press, New York (Aging series, vol 31)
8. Rusting RL (1992) Why do we age? Sci Am, December, pp 86–95
9. Zs.-Nagy I (1994) The membrane hypothesis of aging. CRC Press, London
10. Kirkwood TBL, Holliday R (1979) The evolution of ageing and longevity. Proc R S London, Ser B, Biol Sci 205:531–546
11. Rosen R (1992) Cells and senescence. In: Bittar EE (ed) Developmental biology. JAI Press, Greenwich, pp 191–203
12. Arking R (1998) Biology of aging. Sinauer, Sunderland
13. Hamilton WD (1966) The moulding of senescence by natural selection. J Theor Biol 12:12–45
14. Tzafestas ES (1996) Aging agents. LAFORIA Research Report 96/01
15. Tzafestas ES (1995) Compromising algorithmicity and plasticity in autonomous agent control systems: the autonomous cell. J Intell Syst in press (earlier version appeared as LAFORIA Research Report 95/33, December)
16. Tzafestas ES (1995) Vers une systémique des agents autonomes: des cellules, des motivations et des perturbations. Thèse de Doctorat de l'Université Pierre et Marie Curie, Paris