An Application of Immunocomputing to the Evolution of Rules for Ecosystem Management

Marco A. Janssen and Daniel W. Stow
Department of Spatial Economics
Vrije Universiteit, Amsterdam, The Netherlands
m.janssen@feweb.vu.nl & dstow@feweb.vu.nl

Abstract - This paper discusses the evolution of rules between people for the management of ecosystems. Four aspects of the rules are discussed: coding, creation, selection and memory. The immune system provides us a useful metaphor to relate these four aspects into a coherent framework. We sketch a framework for a computational model to study the evolution of rules for the management of ecosystems.

INTRODUCTION

In the emerging field of immunocomputing, also known as artificial immune systems, computer scientists apply system characteristics of immune systems to other fields such as computer security and pattern recognition [4, 9]. In this paper we focus on a different use of immunocomputing – as a metaphor for the evolution of rules between people. We focus on a particular type of interaction between people, namely the management of common-pool resources. Common-pool resources are characterised by the fact that it is costly to exclude individuals from using the resource, and the fact that the benefits consumed by one individual subtract from those available to others. Common-pool resources can be found in many problem areas, but we mainly refer to ecosystem management problems, such as the management of fish stocks, forests and ground water reservoirs.

How to manage common-pool resources has been a topic of interest for many scholars. The so-called commons dilemma was made famous by [6] as "the tragedy of the commons". The traditional view is that when there is no regulation, individuals can ignore the costs their decisions impose on others. This leads to over-exploitation of the common resource. A cooperative action would allow more efficient use of the resource, but cannot be derived without central control or privatisation of the resource.

However, empirical evidence from field research and laboratory experiments does not support this view [23]. There are many examples where people dependent upon common-pool resources have organised themselves to achieve much higher outcomes than is predicted by the conventional theory [21]. Laboratory experiments show that communication is a crucial factor to derive cooperative behaviour. Furthermore, the ability of the participants to determine their own monitoring and sanctioning system is critical for sustaining cooperative behaviour [23]. However, the evidence for self-organisation of institutions is anecdotal. No formal explanation for it exists. Our goal is to develop such a formal model, based on models of the immune system.

Institutions can be considered as rules that shape human interactions. To understand self-organisation of institutions we have to understand the evolution of rules. This is different from the evolution of most organisms. The fitness of most organisms is related to the number of offspring. The fitness of a rule depends on whether it is used within a population of agents, which might be related to its functionality. An immune system seems to be a more suitable metaphor since it contains a large variety of effective responses to pathogens, creates new responses and remembers successful responses. The success of a response, like the success of a rule, is related to its functionality.

In this paper we discuss an immunocomputational framework to study the evolution of rules. We are especially interested in how rules are *encoded*, how new rules are *created*, how effective rules get *selected* and how rules are *remembered*. We discuss the immune system perspective of ecosystem management in more detail in [12] and [14].

THE IMMUNE SYSTEM

The immune system maintains the health of the body by protecting it from invasion by harmful pathogens, such as bacteria, viruses, fungi, and parasites. These pathogens are the cause of many diseases, so it is necessary to detect and eliminate them rapidly. The immune system also remembers successful responses to invasions and can re-use these responses if similar pathogens invade in the future. For the purposes of this paper, we will mainly discuss the coding, creation, selection and memory of immune system responses. Our discussion is based on the descriptions of [27] and [8].

The adaptive part of the immune system consists of a class of white blood cells called lymphocytes, whose function is to detect pathogens and assist in their elimination. The surface of a lymphocyte is covered with a large number of identical receptors. On the surfaces of pathogens are epitopes. The more complementary the structures of receptor and epitope are, the more likely they will bind together. Recognition occurs when the number of bound receptors on a lymphocyte's surface exceeds a certain threshold.

The immune system must maintain a diverse repertoire of responses because different pathogens must be eliminated in different ways. To achieve this, the immune system constantly creates new types of responses. These are subject to selection processes that favour more successful responses and ensure that the immune system does not respond to self-proteins. A memory of successful responses to pathogens is maintained to speed up future responses to those and similar pathogens.

The generation of new responses is done by a pseudo-random process of DNA recombination. The DNA used to create lymphocyte receptors consists of a number of libraries,

0-7803-7282-4/02/$10.00 ©2002 IEEE

each containing a number of gene segments. A new DNA string is assembled by picking a random segment from each library and joining these segments together. The resulting DNA is then used to make the receptor. If the DNA does not make a valid receptor the lymphocyte commits suicide, because it is useless without a receptor.

Lymphocytes are subject to two types of selection process. Negative selection, which operates on lymphocytes maturing in the thymus (called T-cells), ensures that these lymphocytes do not respond to self-proteins. This is achieved by killing any T-cell that binds to a self protein while it is maturing. The second selection process, called clonal selection, operates on lymphocytes that have matured in the bone marrow (called B-cells). Any B-cell that binds to a pathogen is stimulated to copy itself. Thus, B-cells are selected for their success in detecting non-self. The copying process is subject to a high probability of errors (“hypermutation”). The combination of copying with mutation and selection amounts to an evolutionary algorithm that gives rise to B-cells that are increasingly specific to the invading pathogen.

During the first response to a new pathogen the immune system learns to recognise it by generating new responses and selecting those that are successful, as described above. This response is slow and the organism will experience an infection. If the same or similar pathogens invade in the future the immune system will respond much more quickly because it maintains a memory of successful responses from previous infections.

There are several theories of how immune memory is maintained. One is that successful B-cells become long-lived memory cells that remain in the body in a dormant state until re-infection occurs. Another is that memory cells are not long-lived, but the immune system is constantly being stimulated by low levels of persistent pathogens, either left over from the infection or from subsequent invasions. This ensures that memory cells continue to produce descendants that can deal with future infections. Yet another theory is based on evidence that lymphocytes bind to each other as well as to pathogens. These cells can be described as a network, which dynamically maintains memory using feedback mechanisms [15]. If something has been learnt, it will be remembered if it continues to be reinforced by other parts of the network.

In the following four sections we discuss the coding, creation, selection, and memory of rules in social ecological systems, and the similarities and differences between rules and immune system receptors.

CODING OF RULES

To understand the emergence of rules, we must understand how rules are encoded. The creation of novel structures in any domain always takes place within the constraints of a generative system [2]. For example, English grammar and vocabulary is the generative system used for creating novel English sentences. We need a genetic structure of rules, just as DNA is the genetic structure used to generate new responses in the immune system. A useful starting point may be the grammar of institutions described in [3]. This grammar provides a framework to generate structural descriptions of institutional statements using a syntax of five components.

Comparing the proposed grammar of institutions with the encoding of lymphocyte receptors, we see that there are some interesting similarities. The overall structure of both can be described as a string of slots, into each of which are fitted components of a certain type. In both cases, each type of component is drawn from a library of possible variations and the number of variations to choose from varies between the different types of component.

CREATION OF RULES

Some researchers studying human creativity have found the notion of conceptual spaces useful [2, 24]. The dimensions of a conceptual space are defined by the generative system underlying the domain of interest. The grammar of institutions, combined with a mechanism for generating new institutions from it, defines a conceptual space containing the institutions that conform to the grammar. Exploration of this space creates novel rules that can enter the selection phase.

Generation of new responses in the immune system is by random recombination of the genetic material. Each combination is tested for validity – if it produces a valid receptor then the lymphocyte can enter the selection process. Similarly, an institutional statement may be invalid in a number of ways. Vital components may be missing, an individual component may contain an impossible or inconsistent value, or two or more components may be inconsistent with each other. At some point in the creation of a new rule there must be a test of its validity.

There are some significant differences between the ways new lymphocyte receptors and new institutional statements are created. Creating new rules at random seems like a costly process. The immune system can afford to do this because it contains so many millions of cells. Social groups do not contain as many agents as this or maintain such a large set of rules. However, people can reduce costs by ignoring vast areas of the space of possible rules.

Similar points are made in [24]. Evolution searches the conceptual space blindly – it cannot manage its search, it simply happens. Evolution’s (and the immune system’s) main weapons are time and parallel search. Human inventors do not have the time to search blindly through the possibilities, nor do they have the same capacity for parallel search that evolution has. Instead they are able to manage their search by following gradients of promise, ignoring large areas that are not cost-effective to search, changing the grain of the search, and shifting their starting point to a different area of the space.

Another problem with random recombination is that it cannot create completely novel components, only new arrangements of existing components. In the creation of institutional statements we may sometimes want to create new components and add them to those available for recombination.

SELECTION OF RULES

Before proposed rules become effective, that is, before becoming a social norm or a law, a selection process tests the rules. Within an immune system lymphocytes are subject to selection processes that favour more successful responses and ensure that the immune system does not respond to self-proteins. In an immune system, recognition occurs when the number of bound receptors on a lymphocyte's surface exceeds a certain threshold. In a social system, a rule may become effective when enough agents start to use the rule, or when enough votes are collected for a collective-choice.

The question regarding the selection of rules is whether enough support can be derived for a new rule. The ability of a group to support a new proposed rule can be considered to be dependent on social capital. Social capital comprises relations of trust, reciprocity, common rules, norms and sanctions, and connectedness in institutions [25].

Trust can be defined as the belief in reciprocity of another agent. A trustor will provide something of value to the trustee, but will expect something back later. A crucial element of trust is to recognise trustworthiness of others. In small groups one may know the reputations of all other agents. In larger groups one may use symbols to signal trustworthiness, such as membership of certain organisations, a university degree, or a uniform.

The existence of norms in a group that place group interests above those of individuals give individuals the confidence to invest in collective activities, knowing that others will do so too. Reciprocity and trust are important social norms which can be developed in a group [22]. Another important norm is to agree on sanctions for those who break the rules. Social norms can be developed during repeated interactions, but can decay easily by cheating.

Social capital reduces the costs of cooperation. Collective choice rules will only be selected when there is a sufficient level of social capital. In a population of distrust, selfishness and individualism, cooperative arrangements are unlikely to emerge and rules will not be selected.

REMEMBERING RULES

The memory of a society can take many different forms. These may be formal, such as laws and constitutions, or informal, such as taboos, rituals and religions. A useful starting point for looking at memory is the discussion of Traditional Ecological Knowledge in [1], which identifies a wide range of ecosystem management practices found in local and traditional societies and discusses the social mechanisms behind these practices. These include mechanisms for the generation, accumulation, and transmission of knowledge, the structure and dynamics of the institutions in which ecological knowledge is embedded, and mechanisms for cultural internalisation. Many of these mechanisms are relevant to the question of memory. Some – such as taboos, regulations, social and religious sanctions and folklore – can probably best be viewed as specific items (or collections of items) of memory. Others – the role of knowledge carriers, stewards, or wise people – emphasise the locations where memory is held. Finally, there are the processes that maintain memory – the transmission of knowledge between generations, community assessments of available resources, and rituals or ceremonies that serve as mechanisms for cultural internalisation.

We are interested in whether any of the mechanisms for maintaining social memory are similar to mechanisms in the immune system. In our overview of the immune system we presented three theories of how immune memory is maintained: long-lived memory cells, re-stimulation by pathogens, and immune networks.

Memory cells are analogous to individual items of memory, such as individual laws and taboos. An individual rule will become an item of memory if it is successful enough. This analogy is somewhat unsatisfactory because it does not take account of the processes involved in maintaining memory. If rules are not written down then they must be stored in people, and they must be able to survive the deaths of the people carrying them. Therefore the processes that transmit memory between people are crucial for the survival of rules that are not written down.

Another possible analogy to the memory cell theory is the revival of old knowledge and management practices in response to a resource crisis. A few examples of such revivals are given in [1]. Like memory cells, the knowledge being revived has lain dormant for a while before being reactivated in response to a disturbance. However, strong institutions and traditions are necessary for such revivals [1]. If these are not present economic incentives may be necessary. By contrast, memory cells respond to pathogens they recognise automatically, so the memory cell theory may not be able to say much about the circumstances that are necessary for the revival of dormant knowledge.

The theory that memory is maintained by continual re-stimulation of memory cells is more promising. In this theory, cells are short-lived and it is the descendants of the original memory cells that respond to future infections. This is similar to the process of intergenerational transmission, which ensures that memory can survive the deaths of the individuals that store it. In the immune system, memory is transmitted between generations as long as the memory cells are re-stimulated – i.e. as long as the information is relevant. Similarly, cultural change or persistence of memory depends on its continued relevance to the current context [20]. What this theory of memory cannot explain is the survival of memory that is no longer relevant.

The immune network is analogous to memory that is produced and maintained by a social network, which in effect covers almost all of the types of memory we are interested in. No single item of memory exists in isolation from other items, nor does any location or person holding items of memory exist in isolation from other locations or people.

0-7803-7282-4/02/$10.00 ©2002 IEEE

Furthermore, all of the processes that change and maintain memory take place in the context of a social network.

Immune networks could also be used to model networks of rules. When rules are created there is always a possibility that they conflict with pre-existing rules in some way. In an immune network, something is remembered if it continues to be reinforced by feedback from other parts of the network. If there is conflict between a rule and other parts of the network then the support that this rule receives will be weaker, making it more likely to be forgotten. An immune network model may help us to determine under what circumstances a new rule will be successfully integrated into a system of pre-existing rules.

COMPUTATIONAL MODELS OF IMMUNE SYSTEMS

As discussed above, and in more detail in [14], the evolution of rules within a social-ecological system has interesting similarities with immune systems. In this section we will provide an overview of the relevant formal models in theoretical immunology and immunocomputing that may be useful for understanding the evolution of rules.

Coding and creation

Many of the models of Artificial Immune Systems use a binary string encoding for antibody receptors and antigens [7, 9]. The reason for using such a representation is that it is easy to understand and analyse. It is also easy to define a simple matching rule to calculate the binding strength between receptor and antigen [9]. An alternative would be a variable-based symbolic representation, as suggested in [11]. This allows the use of concepts from logic such as *and*, *or*, *if-then-else*, etc. However, it may prove difficult to use.

There is a possible compromise between these two solutions, suggested by [17], in which the genotypes are binary strings and the phenotypes are rules written in a simple logic. Because the matching is done between rules and antigens, rather than binary strings and antigens, concepts from logic can be used. However, evolution acts on binary genotypes, which are generally simpler to evolve than logic expressions. It is also possible to define sophisticated mutation operators that can make rules more specific or more general [17], which would not be possible if the phenotype was a binary string.

Probably the simplest and most useful model of the creation of lymphocyte receptors is that of [7]. In this model, lymphocyte receptors and antigens are represented as 64-bit binary strings. The gene libraries used to construct the receptors consist of four libraries, each containing eight 16-bit segments. To construct a receptor, a segment is chosen at random from each library and these segments are joined together. When a genetic algorithm was used to study the effects of evolution on the libraries, they tended to self-organise so as to provide the best possible coverage of antigen space (and therefore high fitness scores).

Selection

The immune system implements two types of selection – negative selection and clonal selection. These two processes have tended to be dealt with separately by researchers in computer science, with negative selection being used mainly in the field of computer security [9], and clonal selection being used for machine learning [5].

The negative selection algorithm used by [9] consists of randomly generating detectors and then exposing these to self-patterns. If a detector matches one of these patterns while it is maturing it is killed. This closely resembles negative selection of T-cells in the thymus. When detectors are mature they can monitor a computer or network of computers for unusual activity (non-self patterns) and are unlikely to respond to normal activity (self patterns) because the negative selection process ensures they are tolerant of it.

Some computer security researchers have recognised the potential of clonal selection as a learning process to enhance the negative selection algorithm. For example [17] presents a clonal selection algorithm that uses a negative selection operator to ensure that the evolved detectors do not match self-patterns. This algorithm was applied to machine learning problems, evolving detector sets which achieved a high rate of detection of non-self while minimising detection of self. The best results were achieved when the evolved detector set contained a good balance of general and specific detectors. The general detectors can efficiently detect a large number of related antigens while the specific detectors can detect more unusual cases not covered by the general detectors.

Another model of clonal selection which has been applied to complex machine learning tasks is presented in [5]. At each iteration the model selects the best individuals in the population, based on their affinity to the antigen, and clones them with mutation. A certain number of low affinity individuals are also replaced in each iteration. Clonal selection is essentially a Darwinian evolutionary process, but there is an interesting difference between clonal selection and genetic algorithms. Whereas genetic algorithms tend to converge on a global optimum solution, the clonal selection algorithm evolves many local optima solutions [5, 11]. This is partly because clonal selection places more emphasis on the generation of diversity, both in the creation of new lymphocytes and in the relatively high rate of mutation. Another reason is that the immune system's memory can hold the diversity of responses that have been generated while responses to new disturbances are being evolved. This ability to generate and maintain a diversity of responses to different disturbances is one of the reasons why models of the immune system are an appealing analogy for the evolution of rules.

Memory

Artificial immune systems often seem to model memory in the simplest way possible, using long-lived memory cells. For example, in [9], successful detectors (those that match a certain number of non-self patterns) become memory detectors with a greatly increased life span. These memory

0-7803-7282-4/02/$10.00 ©2002 IEEE

detectors also need less stimulation to be activated, speeding up the secondary response to a previously seen antigen.

The immune network has also received much attention from computer scientists. A model of the immune network is used to solve machine learning problems in [11]. In this model the stimulation level of a B-cell is affected by its affinity to other B-cells in the network as well as its affinity to the antigen. Each antigen is presented to the network in a random location and a certain percentage of the B-cells local to this point is selected to respond to it. The worst 5% of these B-cells are killed and replaced by newly created B-cells. New cells can also be created when a B-cell binds to an antigen. When a new cell is placed in the network it is linked to the two B-cells to which it has the greatest affinity, resulting in the emergence of regions containing similar B-cells.

The memory of this system is maintained because cells that hold the memory of a particular antigen are continually stimulated by other cells in the network. Because B-cells are continually being deleted and replaced, the system as a whole can adapt to a changing environment by forgetting little used items of information. This will only occur when the cells holding these items lose the feedback from other parts of the network, so the system is not too quick to forget information when it is disturbed.

This model has an interesting advantage over neural network models [11]. Neural networks represent what they have learnt in a way that is very difficult to analyse, whereas the immune network represents what it has learnt explicitly – each cell is a representation of a learnt response to an antigen – making analysis of the results much easier.

A COMPUTATIONAL FRAMEWORK FOR THE EVOLUTION OF RULES IN A SOCIAL ECOLOGICAL SYSTEM

We will now sketch a possible framework to simulate the evolution of rules in a social ecological system. When it is disturbed, a well functioning system should be able to generate new rules that are effective to prevent severe consequences. For example, when a new management practise leads to over-harvesting, a healthy social ecological system should detect the problem in an early phase, create informal or formal rules to reduce the harvesting, and be alert for similar problems in the future.

Coding and creation of rules

The rules can be encoded as a bit-string using a similar encoding scheme to that suggested by [17]. The string might consist of parts that are created in different libraries in line with the different components of the grammar of institutions [3]. The first priority for future research is to try to develop the encoding scheme of [17] into one which can represent all of the possible institutions that the grammar can.

Creation consists of selective drawing from the space of possible rules. This means that probabilities for a new rule are not uniform. Due to experience and setting priorities, some rules might have a larger probability of being created. This process can be simulated by neural networks, hill climbing or genetic algorithms.

The coding of rules should be consistent. This means that rules can only be successfully created when the coding of a rule meets certain exogenous constraints of consistency. Such constraints can be absolute such as physical constraints, or can change in time, for example related to social norms.

Social capital

To derive sufficient amount of support for a new rule, it is important to build up enough social capital for a timely response. The question is how to formalise the elements of social capital and their dynamics.

Agents are assumed to be part of social networks, which formalise social interactions. Such social networks can be static, like small-world networks [29], or they can be the result of social interactions [16, 26]. In the latter case, agents start to have random interactions, but previous contacts reinforce social interactions. Furthermore, probabilities of social interactions can be related to the social structure of the network. For example, you may have a larger probability of meeting a friend of a friend, than a random other person.

Another aspect that can restrict social connectivity is the mutual trust between agents. How do agents recognise trustworthy others? One approach is to rely on the reputations of players. In [19], for example, agents keep track of image scores of individuals, where the image scores represent the degree of cooperative actions of the agent in the past. So, when an agent meets another agent, it derives information about its past performance of cooperation. This information is used by the agent to decide whether to cooperate or not. A drawback of such an image score is the assumption of perfect knowledge.

An alternative approach is the use of signals. Within the area of understanding the evolution of cooperation, the use of signals refers to models of “tagging” that have been used to explore behaviour in the iterated Prisoner’s Dilemma. Agents are able to “recognise” one another via an observable symbol, a tag, and based on this observed signal the agent may decide whether or not to interact. The use of tags for the study of social interactions was suggested by [10]. In [28] agents recognise one another and base refusals to play in a Prisoner’s Dilemma on past experience. In [18] agents are not identified from previous interactions, but instead it is assumed that symbols signal trustworthiness, such that they provide information about agents with whom no previous contacts have been experienced. Similarly, in [13] genetic and cultural symbols are used as ways to signal trustworthiness, where specific cultural symbols can be agreed on during interactions among the agents.

The recognition of trustworthy and untrustworthy agents could be likened to the recognition of self and non-self in the immune system. However, unlike in the immune system, recognition of an untrustworthy (non-self) agent does not result in destruction of the agent, but in avoidance of

0-7803-7282-4/02/$10.00 ©2002 IEEE

cooperative activities with that agent, because of the risk of getting defection as a response.

Memory

Memory of rules can be embedded in individuals like memory cells, or can be embedded in social networks like immune networks. When memory is stored in an individual, this individual receives the code of the rule and knows when to use the rule. When memory is stored within a social network, the interactions between the agents generate the code of the rule and the conditions in which to apply the rule. In both cases memory is limited, and needs to be maintained actively. This occurs by re-stimulation of the memory. When a problem occurs repetitively the memory storage is reinforced frequently. In case a situation happens only rarely, but can have severe consequences, it can be worthwhile to put energy in training the memory. This can be done by external stimulation such as celebration days, monuments, rituals, and taboos.

CONCLUSION

In this paper we have described an immunocomputational framework for the evolution of rules for ecosystem management. The next phase will be the development of computational models. But what can such models contribute to the understanding of institutions for environmental management? One of the possible topics of research is to understand in which situations what types of clusters of rules emerge. Do differences in ecosystem dynamics lead to different types of rules emerging?

We are entering a rather unexplored and exiting area of research. The anecdotal evidence that self-governance of common-pool resources is possible might be studied by formal approaches based on artificial immune systems.

ACKNOWLEDGEMENTS

The authors thank Leandro de Castro, Elinor Ostrom, Jeroen van den Bergh, and the participants of the Institutional Analysis and Development seminar course (Fall 2001) at Indiana University for their helpful comments on an earlier version of this paper. We also thank the participants of a seminar at the Universite Libre de Bruxelles. Support of the European Union (contract nr. IST-2000-26016) is gratefully acknowledged.

REFERENCES

[1] Berkes, F., J. Colding and C. Folke (2000). Rediscovery of traditional ecological knowledge as adaptive management. *Ecological Applications*, 10: 1251-1262.
[2] Boden, M.A. (1994). What is creativity? In M.A. Boden (ed.), *Dimensions of Creativity*, Cambridge, MA: MIT Press, pp. 75-117.
[3] Crawford, S.E.S. and E. Ostrom (1995). A Grammar of Institutions. *American Political Science Review*, 89 (3): 582-600.
[4] Dasgupta, D. (ed.) (1999). *Artificial Immune Systems and their Applications*, Berlin: Springer-Verlag.
[5] De Castro, L.N. and F.J. Von Zuben (2000). The clonal selection algorithm with engineering applications. In A.S. Wu (ed.), *Proceedings of the 2000 Genetic and Evolutionary Computation Conference (GECCO) Workshop Program*, Las Vegas, Nevada, pp. 36-37.
[6] Hardin, G. (1968). The tragedy of the commons. *Science*, 162: 1243-1248
[7] Hightower, R, S. Forrest and A.S. Perelson (1995). The Evolution of Emergent Organization in Immune System Gene Libraries. In *Proceedings of the Sixth International Conference on Genetic Algorithms* (Eshelman, L.J. ed.), pp. 344-350, San Francisco, CA: Morgan Kaufman.
[8] Hofmeyr, S.A. (2001). An Interpretative Introduction to the Immune System. In I. Cohen and L. Segel (eds.), *Design Principles for the Immune System and other Distributed Autonomous Systems*. Oxford University Press, pp. 3-26.
[9] Hofmeyr, S.A. and S. Forrest (1999). Immunity by design: An artificial immune system. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, San Francisco, CA: Morgan Kaufmann, pp. 1289-1296.
[10] Holland, J.H. (1993). The Effect of Labels (Tags) on Social Interactions. Santa Fe Institute Working Paper, 93-10-064.
[11] Hunt, J.E. and D.E. Cooke (1996). Learning using an artificial immune system. *Journal of Network and Computer Applications*, 19: 189-212.
[12] Janssen, M.A. (2001). An Immune-system Perspective on Ecosystem Management. *Conservation Ecology,* 5(1): 13. [online] http://www.consecol.org/vol5/iss1/art13
[13] Janssen, M.A. and E. Ostrom (2001). Critical factors that foster local self-governance of common-pool resources. Unpublished manuscript.
[14] Janssen, M.A. and D.W. Stow (2001). Ecosystem Management and the Evolution of Rules: An immune system perspective. Submitted to *Ecological Economics*.
[15] Jerne, N.K. (1973). The immune system. *Scientific American*, 229(1): 52-60
[16] Jin, A.M., M. Girvan, and M. E. J. Newman (2001) The Structure of Growing Social Networks. Santa Fe Institute Working paper 01-06-032
[17] Kim, J. and P. Bentley (2001). Towards an artificial immune system for network intrusion detection: An investigation of clonal selection with a negative selection operator. *Congress on Evolutionary Computation*, Seoul, Korea.
[18] Macy, M. and J. Skvoretz (1998). The Evolution of Trust and Cooperation Between Strangers: A Computational Model. *American Sociological Review*, 63: 638-660
[19] Nowak, M.A. and K. Sigmund (1998). Evolution of indirect reciprocity by image scoring, *Nature*, 393: 573-577.
[20] Olick, J.K. and J. Robbins (1998). Social Memory Studies: From ‘collective memory’ to the historical sociology of mnemonic practices. *Annual Review in Sociology*, 24: 105-140.
[21] Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action*, New York: Cambridge University Press.
[22] Ostrom, E. (2000). Collective Action and the Evolution of Social Norms. *Journal of Economic Perspectives*, 14(3): 137-158.
[23] Ostrom, E., R. Gardner and J. Walker (1994). *Rules, Games. & Common-Pool Resources*, The University of Michigan Press.
[24] Perkins, D.N. (1994). Creativity: Beyond the Darwinian paradigm. In M.A. Boden (ed.), *Dimensions of Creativity*, Cambridge, MA: MIT Press, pp. 119-142.
[25] Pretty, J. and H. Ward (2001). Social Capital and the Environment. *World Development*, 29: 209-227.
[26] Skyrms, B. and R. Pemantle (2000). A dynamic model of social network formation. *Proceedings of the National Academy of Sciences*, 97 (16): 9340-9346.
[27] Sompayrac, L. (1999). *How the Immune System works*, Blackwell Science.
[28] Stanley, E.A., D. Ashlock and L. Tesfatsion (1994). Iterated Prisoner’s Dilemma with Choice and Refusal of Partners. In C.G. Langton (ed.), *Artificial Life III*, Addison-Wesley: Reading, MA., pp. 131-176
[29] Watts, D.J. (1999). *Small Worlds: The Dynamics of Networks between Order and Randomness*, Princeton University Press, NJ.