On Immune Inspired Homeostasis for Electronic Systems

Nick D. Owens[1], Jon Timmis[1,2], Andrew J. Greensted[1], and Andy M. Tyrell[1]

[1] Department of Electronics, University of York, UK
{ndlo100,jt512,ajg112,amt}@ohm.york.ac.uk
[2] Department of Computer Science, University of York, UK

Abstract. Many electronic systems would benefit from the inclusion of self-regulatory mechanisms. We strive to build systems that can predict, or be aware of, imminent threats upon their specified operation. Then, based on this prediction, the system can alter its operation or configuration to circumvent the effects of the threat. In this position paper, we discuss the role of the immune system can play in serving as inspiration for the development of homeostatic engineered systems, through the development of an immune inspired extensible architecture. We outline the major requirements for such an architecture, and discuss issues that arise as a result and propose possible solutions: things are never as simple as they first appear.

1 Introduction

The biological term Homeostasis coined by Cannon [1] refers to an organisms ability to maintain steady states of operation in a massively changing internal and external environment. Engineering homeostasis in electronic systems is a challenging endeavour. There have been many attempts at employing various mechanisms to endow certain systems with homeostasis, for example the Unix operating system [4] and robotics [5]. Within a biological context, it is generally accepted that organism homeostasis is an emergent property of the interactions between the immune, neural and endocrine system. Taking this view, work in [6] discussed mechanisms inspired by the neural and endocrine systems and how these might be exploited in the context of robotic systems. However, there is a great deal of complexity issues when one examines the interactions of these three systems, therefore we have decided to focus on a single subsystem, the immune system, in an attempt to get a handle on the inherent complexity. In practice, and as it will be seen in this paper, it is almost impossible to draw lines between each of these systems, in particular the immune and endocrine systems, as there are so many types of interactions at so many different levels. In this position paper we examine the issues involved in creating a general extensible architecture for homeostasis for use in electronic systems that will endow homeostatic properties on engineered system. However, this is not a simple task and this position paper serves as a discussion on the issues regarding designing such an architecture.

L.N. de Castro, F.J. Von Zuben, and H. Knidel (Eds.): ICARIS 2007, LNCS 4628, pp. 216–227, 2007.

2 Biological Homeostasis and Biological Homeostatic Control Systems

The processes that encompass homeostasis are best understood by looking to the original definition [1]:

> The coordinated physiological processes which maintain most of the steady states in the organism are so complex and so peculiar to living beings involving, as they may, the brain and nerves, the heart, lungs, kidneys and spleen, all working cooperatively that I have suggested a special designation for these states, *homeostasis*. The word does not imply something set and immobile, a stagnation. It means a condition a condition which may vary, but which is relatively constant.

A present day reference on homeostasis, Vander's Human Physiology [2] acknowledges the stability provided by homeostasis is due to interactions of the immune, neural and endocrine systems; and that homeostasis also occurs individually within each one of these systems. Vander [2] opens with a chapter on homeostasis to provide a context to the whole book, the chapter describes homeostasis in terms of *homeostatic variables* and *set points* of those variables. Examples of homeostatic variables in the human body are blood glucose levels, or body temperature. The set points are the steady states (not necessarily equilibria) at which the system attempts to maintain these variables. Vander notes that over a given time period there may be massive variability in homeostatic variables, there is a rise in blood glucose after a meal, for example. But, if a time-averaged mean across that time period is taken and compared with consecutive time periods the behaviour is far more predictable. The control of the of homeostatic variables and set points is performed by *homeostatic control systems*. They are predominantly feedback systems, most often negative feedback, but positive feedback does also occur. Vander [2] supplies a list of general properties homeostatic control systems, which is reproduced here.

1. Stability of an internal environmental variable is achieved by balancing inputs and outputs. It is is not the absolute magnitudes of the inputs and outputs that matter, but the balance between them.
2. In negative feedback systems, a change in the variable being regulated brings about responses that tend to move the variable in the direction opposite the original change — that is, back toward the initial value (set point).
3. Homeostatic control systems cannot maintain complete constancy of any given feature of the internal environment. Therefore, any regulated variable will have a more-or-less narrow range of normal values depending on the external environmental conditions.
4. The set point of some variables regulated by homeostatic control systems can be reset — that is, physiologically raised or lowered.
5. It is not always possible for homeostatic control systems to maintain constancy in every variable in response to an environmental challenge. There is a hierarchy of importance, so that the constancy of certain variables may be altered markedly to maintain others at relatively constant levels.

Vander [2] discusses a number of other issues and examples of homeostatic control systems in context of the human body, we abstract these to produce a further list of desirable properties of homeostatic control systems.

1. *Prediction.* Vander [2] determines this as feed-forward regulation. In response to an environmental change the homeostatic control system manipulates the internal environment in order to avoid a deviation from a set point before it has happened.
2. *Innate and Adaptive Response.* The homeostatic control system is built up of innate and adaptive *reflexes* which are used to bring homeostatic variables back to set points. The innate reflexes are involuntary, unpremeditated and unlearned, and are instigated in response to a particular stimulus, internal or external. As one would imagine, adaptive reflexes are learned to correct unforeseen deviations from set points. Vander also states that all reflexes, innate or adaptive, are subject to further learning.
3. *Acclimatisation.* Although encompassed by both adaptive responses and re-setting of set points, it is an important enough property in its own right. It represents the ability for a set point to semi-permanently change in response to semi-permanent change in the environment. To aid explanation we take the analogy in [2] of a runner who is asked to run for 8 consecutive days in a hot room (a room hotter than the runner's normal environment). Details of the runner's sweating are recorded. By the 8th day the runner starts to sweat earlier and in far greater quantities than the 1st day, this allows to the runner to limit the deviation of the temperature homeostatic variable from its set point. The 'sweating' homeostatic set point has acclimatised to the new environment. When the runner returns to running in the original environment the set point will, over a number of days, acclimatise back to the original.

3 The Immune System for Homeostasis

First, it is worth noting why we are attempting to construct artificial homeostatic systems using solely the immune system, apparently ignoring the neural and endocrine systems. All three systems are necessary for human homeostasis and none of the systems are singularly sufficient. When one investigates the immune system, it is clear that the endocrine system is inextricably linked to the immune system. Immune cytokine networks share many of the same functional properties of the endocrine system, and are in effect considered part of the endocrine system: therefore we are not ignoring the role of the endocrine system. There is clear evidence to suggest immune, neural and endocrine interactions [3], however, as previously mentioned we have excluded the neural system from our studies, as by doing so reduces the level of complexity that we are dealing with and allows us to focus our efforts at exploring the role of the immune system in body maintenance, a view held by some theoretical immunologists [7] and [9].

3.1 Cohen's Immune System

Cohen believes the role of the immune system is to repair and maintain the body. As the removal of pathogen is beneficial to the health of the body, defence against pathogen can be seen as just a special case of body maintenance. In order to achieve body maintenance, the immune system must select and regulate the inflammatory response according to the current condition of the body. This condition is assessed through the Co-respondence of both the adaptive and innate immune agents, which are required to recognise both the presence of pathogen (non-self antigen) and the state of the bodys own tissues (self antigen). The specificity of the immune response, therefore, is not just the discrimination of danger [11] or the distinction of self/non-self, but the diagnosis of varied situations, and the evocation of a suitable response. Degeneracy is a concept central to Cohen's ideas and is discussed in terms of immune receptors and cytokine networks. Degeneracy is defined as [20]: "The ability of elements that are structurally different to perform the same function of yield the same output." Cohens maintenance role of the immune system requires it to provide three properties: Recognition: to determine what is right and wrong, Cognition: to interpret the input signals, evaluate them, and make decisions. Action: to carry out the decisions.

3.2 Grossman's Tunable Responses

Grossman [8] sees the immune system as a system which reacts to perturbations, to changes in the environment rather than the specificity of any particular pathogen. Grossman's view is constructed around immune cells having with tunable activation thresholds, which control proliferation, differentiation and choice of effector function. The activation thresholds are tuned to a cell's recent excitation history (its interactions and interaction affinity) [8], a change in the environment will cause a change in the cells excitation. The rate and size of perturbation with respect to the cell's history is what ultimately determines the response of the cell. Grossman believes control of immune response (e.g. severity of attack, tolerance and memory) emerges out of the dynamics of a population of tunable cells [9]. Grossman provides a simple mathematical model for the tuning of cells in [8] and outlines potential biological evidence for the tuning in [10].

3.3 Appropriateness of Immune Inspiration

Cohen's and Grossman's theories concern the immune system as a whole, their arguments relate to interactions producing behaviour rather than analysis of immune machinery. They view the immune system in terms of maintenance and tolerance rather than attack and defence of invading pathogens. The concepts of Cohen and Grossman have commonalities with the concepts of biological homeostasis and homeostatic control systems, section 2. They provide immune theories that would seem to provide excellent inspiration for construction of a homeostatic control system. There are certainly conflicts between the two models of the immune system, Grossman requires some immune receptor specificity [8],

and Cohen prefers to do away with specificity altogether [12]. However there is definitely common ground and opportunity to combine the two, in fact Cohen does precisely this in [12] to produce a model to describe T Cell behaviour.

4 An Architecture for Artificial Homeostasis

We will now discuss some of the issues that arise in attempting to construct an architecture for homeostasis based the notion of homeostatic control systems provided. Again, it should be noted that this is a position paper, we discuss potential problems and propose some tentative solutions.

It is our intention that if this architecture can be developed then systems built adhering to the rules of the architecture will then have the properties of a homeostatic system. Consequently the system should have an innate level of homeostasis and then adapt and acclimatise to the environment it is placed within. We now define some terms within our system, and say that the system is comprised of:

1. *Sensors.* These can sense from the environment.
2. *Homeostatic Variables and Set Points.* For the system to maintain homeostasis we must define what it means for the system to be homeostatic. The intention is that homeostatic variables are evaluated by functions on the sensors, and other internal variables of the system. Each of these are associated with a priority, to represent the importance of certain variables over others as mentioned in section 2.
3. *Actuators.* These can act to manipulate the environment.
4. *Homeostatic Responses.* Similar to the homeostatic variables and set points, the system requires innate methods to correct deviations from homeostatic set points. The responses would make use of the systems actuators.
5. *Tasks.* These describe the behaviour of our system.
6. *Homeostatic Control System.* This maintains the homeostatic variables at their set points, while allow the system to complete its tasks.

4.1 Splitting the Problem: Breaking into a Homeostatic System

Imagine that we have constructed a homeostatic control system, which is able to maintain homeostasis given the Sensors; Actuators; Homeostatic Variables, Set Points and Responses; and Tasks. The problem now is: can we sensibly and tractably split up a system into these components? Although choice of sensors, actuators and tasks are ultimately specific to problem domain of the system, there are still general considerations. It is important to understand the purpose of the system, this may seem obvious, but it raise some interesting issues. For example: is it more important for the system to complete its task, or is survival of the system (it not becoming irreparably damaged) more important? Is there only a single system performing a given task, or are there many systems? Therefore, is losing one or two systems an acceptable cost in order to complete the task? If this

is the case the homeostatic control system could be a little more cavalier with the choices of homeostatic response, this should be reflected in innate definitions of the system.

Similarly there are some general points on choice of innate homeostatic variables, set points and responses. An interesting way to split the problem is to observe how the problem is split in biological systems. There is a natural hierarchy to biological systems, the homeostasis of an individual is maintained by the immune, neural and endocrine systems. The homeostasis of these systems is maintained by systems internal to them, and then the homeostasis of those systems is maintained by systems internal to them, and so on.

The homeostasis of artificial systems can be broken up in a similar manner. At the population; at the individual; at the tasks; the physical components of a system. Problem splits can be both logical and physical, clearly splitting by tasks represents a logical split whereas by components represents a physical split. There is no constraint on splits being entirely logical or entirely physical, one can imagine a system split both logically and physically with the split problem represented as a graph. Nodes in the graph would represent *homeostatic units*, and edges of the graph represent a communication of homeostasis information between homeostatic units of the system. This raises the questions: what information should be shared between homeostatic units and how should this information be shared? It is reasonable to envisage homeostatic variable deviation information propagating through this graph, but sharing of homeostatic responses and correlations is less obvious. Imagine a 4-wheeled robot with a homeostatic architecture split so that each of the four is controlled by four homeostatic units. Each wheel is mechanically identical and has identical sensors. One of the four discovers a correlation or response that is useful in predicting and avoiding flat tyres, how can this information be propagated to the other three wheels? The difficulty arises that each homeostatic unit is by intention self-organising; there is a black box element to these homeostatic units. We don't not know what coding scheme has be adopted in each of the four homeostatic units, to communicate the information we would have to isolate the information that deals with new discovery along with the other information it depends upon. Then a mechanism would be required to translate from the coding scheme of the discovering unit to the coding scheme of the other three units. We should note that we are not suggesting these steps be literally implemented, but that they hopefully can emerge as part of an appropriately immune inspired algorithm.

We return, briefly, to discussing logically splitting the problem and suggest a temporal heuristic may be useful. Systems are required to operate over a variety of time scales, it is a property of general systems including computational, psychological and social [17]. For example, parts of our system may need to respond on time scale t_1 and other parts may need to respond on timescale t_2. The difference between t_1 and t_2 may be so pronounced that from the point of view of t_1 operations on timescale t_2 occur instantly and atomically, and from the point of view of t_2, t_1 looks constant. This property is noted in Burns et al. [17] who take inspiration from psychology and real time systems to develop a

formal framework of *timebands*. It can be used to describe and prove relationships between entities interacting at different timescales. No proof or analysis of interactions between timescales is needed here, in fact such analysis flies in the face of the purpose of our architecture. However, these time scales still provide a natural way to break up the system into homeostatic variables with timescales in common sharing hierarchical levels. It is worth noting that this hierarchy does not necessarily impose a order of importance, or that lower levels in the hierarchy must be completed before higher levels. It is just an intuition with which to split the problem and to characterise interactions between homeostatic variables and responses.

There is an issue concerning the granularity with which homeostatic set points are defined. The intention is to allow the architecture to discover an emergent method of maintaining homeostasis, as a consequence we do not want the set points to be defined with too fine a granularity. For example, homeostatic variables could be represented by a real numbers and if the set points are set at specific real target values, then the homeostasis of the system is very prescriptive. The solution to maintaining homeostasis reduces to an error minimisation problem, and biologically inspired techniques, although still useful, are no longer necessary. However, this overlooks the fact that it may be very hard to assign fixed target values to the set points to achieve the desired behaviour. Moreover, systems where this is a straight forward task are not the systems we are interested in bestowing with homeostatic properties. Set points should be defined in a fuzzier sense: intervals, minimisations, maximisations etc. For example a component in the system may have an operational temperature range in $[-10, 50]$ degrees Celsius, this range is an obvious choice for a set point.

A final point of discussion in the choice and the assignment of homeostatic variables is their priorities. In section 2 it is noted that there is a hierarchy of importance of homeostatic variables and the homeostasis of some variables may be sacrificed to maintain the homeostasis of others. Homeostatic variable priorities are a natural way to represent this property. The assignment of priorities to homeostatic variables will necessarily constrain the behaviour of the system. An incorrect intuition about priorities may lead to system failure in the worst case and unnecessarily restrictive behaviour in the best case. We suggest that priorities are assigned very carefully and potentially sparingly, perhaps assigning many homeostatic variables the same priority. The concept of the homeostatic control system sacrificing certain homeostatic variables in favour of others can still exist in a system with many homeostatic variables sharing the same priority. Many homeostatic variables can be allowed to deviate slightly from their set points to avoid a large deviation on a separate homeostatic variable.

Choosing innate responses to correct deviations from homeostatic set points is subject to similar issues as that of choosing homeostatic variables. We would like the homeostatic control system to discover good choices for homeostatic responses, so we do not necessarily need sophisticated innate responses, the hope being the homeostatic control system will discover them. A simple heuristic for choice, then, is all that is needed. We suggest a greedy response: each homeostatic

variable is associated with the response that will best correct that variable, if perturbed, regardless of effect of that response on the rest of the system. For example in an autonomous robot system, the temperature homeostatic variable may be associated with a response to turn on the fan (pull the variable back towards its set point) and to turn off all motors (eliminate a potential source of the perturbation).

4.2 Homeostatic Control System

Before discussing how an immune inspired homeostatic control system may be constructed, it is useful to examine the issues that arise in attempting to maintain homeostasis on our new definition of a system. Immune inspiration or not our system is beginning to fit into Cohen's [7] three stage system: *Recognition*, sensors sense the environment which position the homeostatic variables. *Cognition*, based on the sensor values and homeostatic variables in relation to their set points we must decide on homeostatic responses to take. *Action*, the homeostatic responses act to move homeostatic variables back towards set points.

It is worth noting that linking a homeostatic response to the correction of a homeostatic variable is perhaps sufficient for an innate response, but it is definitely not sufficient for adaptive responses. To ease discussion we will define a new term: *homeostatic error* for a homeostatic variable, this is simply the distance (by whatever distance metric we care to choose) of a homeostatic variable from its set point. To reiterate, now with the new parlance, it is not sufficient to link adaptive homeostatic responses to homeostatic error. The homeostatic error is only a context with which to understand the sensor values.

We return to the properties of homeostatic systems given in section 2, and determine what our system must do given our current definitions:

1. *Arbitration* — given homeostatic variables with homeostatic error and the system's current understanding of the homeostatic responses, the control system must arbitrate between the possible responses to best maintain homeostasis.
2. *Correlation of Sensor and Homeostatic Error* — The system must correlate the sensor conditions under which homeostatic errors tend to arise.
3. *Response Learning* — The system must improve on innately supplied homeostatic response choices, by learning responses for the correlations learned in the previous step. This includes recovery from the failure of responses, if an actuator used in a response fails the system must use a different combination of actuations to achieve the desired response.
4. *Prediction* — Using the correlations and responses learned the system should predict the movement of homeostatic variables to avoid homeostatic error wherever possible.
5. *Acclimatisation* — This represents the systems ability to change its correlations and learned responses as the environment changes. This highlights the issue of *overspecificity*, the system becoming too specific to a given environment and failing when the environment changes. We want our environments change and so the system must be able to acclimatise to any new environment that arises.

An interesting way to think of the acclimatisation and the avoidance of over-specificity is in terms of robustness. Specifically robustness to changes in the environment, this can be examined in information theoretic terms using the mutual information [18], I, between the system S and the environment E:

$$I(S : E) = H(S) - H(S|E) \tag{1}$$

$H(S)$ is the entropy of the system and $H(S|E)$ is the conditional entropy of the system with respect to the environment. This conditional entropy term can be thought of as the amount information that is in the system that is not correlated to anything in the environment [18]. This represents an information excess, it allows the system to change neutrally (a change that does not effect the mutual information) without changing the systems behaviour. By a similar argument the environment can change neutrally, without the system noticing, if E changes in a way not correlated in S. The system and its behaviour can be robust to both internal and external changes. The problem now becomes one of striking the appropriate level of mutual information, to allow for acclimatisation but not overspecificity. This can be answered by examining the rate of change of the environment with the rate of change of the mutual information, that is the ability of the system to keep up with a changing environment. Insight is given observing by the specificity of the system to the environment at a given time, if the system too specific (overspecificity, or overtraining) a change in the environment will seem large from the point of view of the system, and so a large rate of change and the system may not be able to keep up. If the system is unspecific for the environment a change in the environment will seem small, there will be a small rate of change. However if the system is too unspecific changes in the environment will not cause appropriate acclimatisation and incorrect system behaviour.

Artificial systems have considerations of homeostatic control beyond the those that are obvious from biology. Sensors can fail and their ability to sense the environment to a specified range and resolution can degrade. The consequence of sensor degradation could cause an incorrect calculation of homeostatic error which in turn could cause an incorrect homeostatic response. Catastrophe can occur, a failure of the system to fulfil its purpose, in our terms: either failure of system survival or failure of the system to complete its task. Imagine a situation where the homeostatic response caused by an incorrect error perpetuates the true homeostatic error, which causes positive feedback and dramatic increases of homeostatic error to the point of catastrophe. Sensor failure is a far more malicious failure than response failure. Although it is impossible to completely eliminate the possibility catastrophe, with appropriate system design it is possible to reduce the probability of a catastrophic collapse. There is a natural physics provided by the environment, for example in a robot, it may be possible to estimate robot temperature as a function of motor speeds rather than using a thermometer. This highlights a need for multiple methods of evaluating the same homeostatic variable, this is an extension of an already identified desirable property of our system. Currently our system must correlate sensor data and homeostatic errors to better identify when homeostatic errors occur and

when they are likely to occur. Now, we would like our system to correlate sensor data and homeostatic error in order to produce multiple methods of evaluating homeostatic variables. This is providing degenerate methods of evaluating homeostatic variables, this turns out to be exactly what we want.

Degeneracy and Redundancy. We have defined degeneracy in section 3.1. In contrast *redundancy* is characterised by multiple identical structures performing the same function [20]. A further important distinction is that although under certain conditions degenerate structures can perform the same function, under different conditions they can perform different functions. Degeneracy and redundancy are analysed in information theoretic terms in [19], the formulation is in terms of mutual information between subsets of inputs and outputs. They establish that [19]:

> A degenerate system, unlike a fully redundant one, is thus extremely adaptable to unpredictable changes in circumstances.

They stress that a degenerate system must have a degree of functional redundancy,there is scale between a fully redundant systems (everything performs the same function) and fully independent systems (a one to one mapping, everything performs different functions), degeneracy lies in the middle. Tononi et al. [19] also comment that degeneracy and complexity go hand in hand, and systems with high degeneracy have the potential for high complexity. Appropriate degenerate structures would seem to provide precisely the robustness for the balance between specificity and acclimatisation discussed earlier. Consequently degeneracy is a property that we would like at many levels of our systems. It makes a case for our systems being sensor rich (mimicking biology), and not just sensor redundant, to allow for sensor noise, degradation and failure; and to be able to represent the complex relationships between degenerate sensors.

Returning to the more general topic of our homeostatic control system; biological systems develop their homeostatic control through evolution species over many generations. Similar ideas are possible artificially if there have a population of identical systems maintaining homeostasis and sharing information on homeostasis. But, in the case of a single system how is performance evaluated? Imagine all homeostatic variables are at or within homeostatic set points, but a more efficient and preferred system operation exists and to get to this operation requires movement into homeostatic error. Is this movement ever possible? The movement clearly requires a homeostatic variable or task linked with a concept of efficiency, however it is likely that this will have a lower priority than other homeostatic variables. A lower priority than the homeostatic variables that are required to briefly delve into error to allow the efficiency error to drop. We do not provide an answer here, more a comment from Cohen [7] "what works, works.", if the total homeostatic error is at zero the system is successful, if the behaviour is not the desired one then the homeostatic variable definitions must be changed.

5 Towards an Immune Inspired Solution

There are themes that arise in the discussion of artificial homeostasis which clearly link to the immunological theories. The co-respondance of [7] and [9] combined with the tuning of [8] and the cytokine networks of [14], to provide the arbitration. The degeneracy of [7] providing the robustness through appropriate mutual information between the environment and the system. The population dynamics of [7] and [10] giving rise to necessary decisions, proliferation and memory.

Although there are many conflicting theories on the immune systems operation and function, there is less argument to the immune system's function on a basic mechanical level. We demonstrate analogies between the function of components of our homeostatic system and the function of some immune components. To begin with, there is analogy between the lymphatic system, and our homeostatic unit graph, the units and the lymph nodes being distributed locations in which the problem and the solution is determined. Clearly both systems have an innate/adaptive divide. At the innate level there are analogies between antigen presenting cells (APCs), macrophages, the context they present and the the innate variable evaluations and response. At the adaptive level there is analogy between the T Cells, B Cells and there proliferation and differentiation decisions and the adaptive recognition and response of the homeostatic system. These analogies serve little purpose in the design of an immune inspired solution, but are used to demonstrate that inspiration is needed from the whole immune process, not one individual part or concept.

5.1 Moving from Inspiration to Algorithms

Much of our inspiration is born from the arguments made Cohen [7] and Grossman [8], however a number of the arguments are neither scientifically verified nor their workings fully understood. We suggest taking a principled approach in the form of the conceptual framework [13], which may allow us to build immune algorithms which are correctly biologically grounded, better theoretically understood and ultimately more successful. The first step is to model mathematically and computationally the biology as it is explained by [7] and [8], to gain a better understanding of the biological processes and interactions involved. Distance along this road has already been travelled with modelling of degeneracy in both biological and artificial immune systems [19], [15] and [16]; with mathematical modelling of cytokine networks [14]; and with model combining some of the out of Cohen and Grossman [12]. This leaves much of the work in understand to theories put forward by Grossman: [8], [9], [10], which we believe has the potential to be some of the most fruitful immune inspiration we have outlined.

6 Conclusions

We feel that it is indeed possible to build an immune inspired architecture for homeostasis. The greatest chances of success lie through: proper understand of

the homeostasis problem; appropriate choice of immunological inspiration; and considered understanding of the biology behind the inspiration.

References

1. Cannon, W.B.: The Wisdom of the Body. Norton, New York (1932)
2. Widmaier, E.P., Raff, H., Strang, K.T.: Vander's Human Physiology: The Mechanisms of Body Function, 10th edn. Mc-Graw Hill, New York (2006)
3. Besedovsky, H.O., Del Rey, A.: Introduction: immune-neuroendocrine network. Front. Horm. Res. 29, 1–14 (2002)
4. Somayaji, A., Forrest, S.: Automated Response Using System-Call Delays. In: 9th USENIX Security Symposium (2000)
5. Brooks, R.: A Robust Layered Control System for a Mobile Robot. IEEE Journal of Robotics and Automation 2(1), 14–23 (1986)
6. Neal, M., Timmis, J.: Timidity: A Useful Mechanism for Robot Control? Informatica 27(4), 197–204 (2003)
7. Cohen, I.R.: Tending Adams Garden: Evolving the Cognitive Immune Self. Elsevier Academic Press, Amsterdam (2000)
8. Grossman, Z., Paul, W.E.: Adaptive Cellular Interactions in the Immune System: The Tunable activation threshold and significane of subthreshold response. PNAS 89, 10365–10369 (1992)
9. Grossman, Z.: Cellular Tolerance as a Dynamic State of the Adaptable Lymphocyte. Immunological Reviews. No. 133 (1993)
10. Grossman, Z., Paul, W.E.: Autoreactivity, dynamic tuning and selectivity. Current Opinion in Immunology 13, 687–698 (2001)
11. Matzinger, P.: The Danger Model: A renewed sense of Self. Science 296, 301–305 (2002)
12. Hershberg, U., Solomon, S., Cohen, IR.: What is the basis of the immune system's specificity? Mathematical Modelling and Computing in Biology and Medicine, 377–384 (2003)
13. Stepney, S., Smith, R.E., Timmis, J., Tyrrell, A.M., Neal, M.J., Hone, A.: Conceptual Frameworks for Artificial Immune Systems. Int. J. Unconventional Computing 1(3), 315–338 (2005)
14. Hone, A., van den Berg, H.: Modelling a Cytokine Network. In: FOCI 2007 (2007)
15. Andrews, P.S., Timmis, J.: A computational model of degeneracy in a lymph node. In: Bersini, H., Carneiro, J. (eds.) ICARIS 2005. LNCS, vol. 4163, pp. 164–177. Springer, Heidelberg (2005)
16. Mendao, M., Timmis, J., Andrews, P.S., Davies, M.: The Immune System in Pieces: Computational Lessons from Degeneracy in the Immune System. In: FOCI 2007 (2007)
17. Burns, A., Hayes, I.J., Baxter, G., Fidge, C.J.: Modelling Temporal Behaviour in Complex Socio-Technical Systems. Techinical Report YCS-2205-390. Department of Computer Science, University of York (2005)
18. Weeks, A., Stepney, S., Polack, F.A.C.: Neutral Emergence: a proposal. In: Symposium on Complex Systems Engineering, RAND Corporation (2007)
19. Tononi, G., Sporns, O., Edelman, G.M.: Measures of Degeneracy and Redundancy in Biological Networks. PNAS 96, 3257–3262 (1999)
20. Edelman, G.M., Gally, J.A.: Degeneracy and complexity in biological systems. PNAS 98(24), 13763–13768 (2001)