Proceedings of the 2007 IEEE Symposium on
Foundations of Computational Intelligence (FOCI 2007)

Modelling a cytokine network

(Special session: Foundations of Artificial Immune Systems)

A. Hone* and H. van den Berg†
*Institute of Mathematics, Statistics & Actuarial Science, University of Kent, Canterbury, UK
Email: anwh@kent.ac.uk
†Warwick Systems Biology, University of Warwick, UK
E-mail: hugo@maths.warwick.ac.uk

***Abstract*—Artificial immune systems have taken much of their inspiration from theoretical models of immune processes, in particular the process of clonal selection in populations of B cells and T cells. Here we focus on a generic dynamical model for the interaction of cytokines with immune cells, which we refer to as an artificial cytokine network (ACN).**

I. Introduction

Mathematical models form a vital part of theoretical immunology. Mathematics is essential for making quantitative statements about the response of the immune system to antigens. Nonlinear dynamical systems [14], [15], and especially systems of nonlinear differential equations [9], appear throughout the physical sciences, and they have also been used successfully in population biology and ecology to describe the growth and spread of plant and animal populations [12]. The use of nonlinear differential equations has further been exploited to model the interactions of populations of lymphocytes with virus or antigen populations [13], [16], and dynamical systems have been used in many other theoretical models in immunology.

The emerging field of artificial immune systems has been inspired by mathematical models of immune interactions, and in particular network models of lymphocyte populations (see for example [16] and references). Jerne introduced the notion of an *idiotypic network* of immune cells that are able to recognize one another as well as antigen, while Farmer et al. suggested that the adaptive responses of immune networks might provide useful paradigms for machine learning. Subsequently, immune network algorithms [2] were based on the idea of replacing nonlinear differential equations with analogous discrete or iterative schemes, together with the inclusion of stochastic effects due to mutations. Many other artificial immune system algorithms, such as CLONALG (due to de Castro and von Zuben [3]), IMMALG and optIA (due to Nicosia and Cutello [6], [7]), were based on simplified mechanisms for clonal selection. In the mean time, Jerne's idiotypic network theory had already been discredited by some immunologists. However, the biological implausibility of a theoretical model does not necessarily preclude its usefulness in a computational context.

We believe that mathematical analysis, and especially the theory of dynamical systems, has a whole host of useful tools to offer to practitioners of bio-inspired computing. By translating abstract processes into precise mathematical models, it should become much easier to decide exactly what the natural immune system has in common with an artificial immune system (AIS). Mathematical models of immunological processes can be adapted to AIS in order to provide objective measures of their performance, and further to work out how they can best be controlled. In a recent review of theoretical techniques [1], we gave a simple model for the dynamics of an artificial immune response, and indicated how this could be analysed using ideas from optimal control theory [17]. AIS algorithms based on clonal selection result in algorithms that are stochastic rather than deterministic, usually described in terms of populations of "cells" that evolve between different states according to probabilistic rules. In that case, they have many properties in common with genetic algorithms and other kinds of evolutionary algorithms. Thus it makes sense to model them in terms of Markov chains: exact Markov chain models of certain AIS optimization algorithms have begun to be developed fairly recently [4], [5]. These models allow one to prove convergence, but more detailed analysis is required to derive more useful properties such as rates of convergence; the methods developed in [11] should be helpful in this regard.

The development of AIS has been inspired by abstract models of immunological processes. In this article, we aim to provide further inspiration for computer scientists by outlining a new model for the interaction of signalling molecules (cytokines) with immune cells. In turn, this cytokine network model could also provide insights for biology.

II. Cytokine Networks

The allocation of computational resources to the processes of an artificial immune response is critical to AIS performance. The urgency of the job load, the tasks at hand, the specific efficacies of the available response mechanisms, the likelihood of discovering more efficacious responses (e.g. a higher-affinity receptor), as well as demands on CPU time by processes outside the AIS: these are all factors that determine how much should be allocated to the available AIS and developing better responses.

The natural immune system regulates this allocation by means of a system of various immune cells that mutally influence each other's activities via hormone-like intercellular messenger molecules called cytokines. Cytokines stimulate proliferation of various immune cells, with different immune cell types responding to different cytokines. Such immune

cells include the effector cells (those which carry out the response) as well as the cells that produce cytokines themselves. Indeed, immune cells often have a dual role, producing cytokines in addition to an effector function. The activities modulated by the cytokines produced by one cell include the production and secretion of cytokines by other immune cells. These interactions form a lymphoid endocrine system called the *cytokine network*. Each cell type in this network is characterised by its own subset of the cytokines that it can secrete, as well as its own subset of cytokine receptors that govern its activity. The cytokine network integrates stimuli from a variety of sources (e.g. distressed cells, effector cells, naive response-precursor cells), and the cytokines produced by the network regulate the development and growth of the responding immune cells.

From a computational point of view, the cytokine network has an input of information about the state (extent and severity) of the disease, the state (extent and efficacy) of the ongoing responses, and an output that governs proliferation of selected effector cells as well as the organization of new responses (e.g. antigen presentation, germinal centre reaction). Both input and output are encoded by the concentrations of the various cytokines. From a modelling point of view, the complexity of cytokine networks poses considerable challenges. Moreover, the cytokine network operates both locally and more globally, and is thus intermediate between a paracrine system and an endocrine system. However, in what follows we ignore all such spatial aspects for the sake of simplicity.

Below we give a mathematical specification of the cytokine network which emphasizes it as a computational paradigm. To highlight this aspect, we refer to it as an *artificial cytokine network* (ACN). The ACN is one example of a computational system inspired by biological para-/endocrine systems. As will become clear, the ACN has much in common with the associative memory models studied by neural network theory [10]. This is not very surprising since the ACN is likewise a system that matches a vector representing a given situation to a vector representing a (hopefully suitable) response. However, there are a few interesting points of contrast. The analogues of "synaptic weights" in the ACN are non-changing. However, the ACN is in some sense a superposition of a number of associative memory structures, with the relative contributions of these structures changing on a second, slower time scale.

III. Artificial cytokine network

For our model of the cytokine network, we consider an intercellular medium in which n distinct chemical species of cytokines diffuse and are well mixed. We can then define cytokine concentrations $u_1, \ldots, u_n$. The cytokines are produced by cytokine-producing cells, of which there are m types. Cytokine production by a cell of any one of these types depends on external stimuli $s_1, \ldots, s_r$ (i.e. signals arising outside of the cytokine network) as well as the cytokines themselves. The density of cell type ℓ in the medium is denoted as v_ℓ. We thus have the following kinetics:

$$\dot{u}_k = \sum_{\ell=1}^{m} \psi_{\ell k}(u_1, \ldots, u_n, s_1, \ldots, s_r) v_\ell - \nu_k u_k, \tag{3.1}$$

$$\dot{v}_\ell = \Big(\varphi_\ell(u_1, \ldots, u_n, s_1, \ldots, s_r) - \mu_\ell\Big) v_\ell, \tag{3.2}$$

where

$$k = 1, \ldots, n, \quad \ell = 1, \ldots, m;$$

the function $\psi_{\ell k} > 0$ expresses the effect of the cytokines and external stimuli on the production of cytokine k by a cell of type ℓ; $\nu_k > 0$ is the rate of degradation of the kth cytokine; $\varphi_\ell > 0$ expresses the effect of the cytokines and external stimuli on the proliferation rate of a cell of type ℓ; and $\mu_\ell > 0$ is the death rate of cells of type ℓ. There is often a separation of time scales between the dynamics of the cytokines, equation (3.1), and the dynamics of the cytokine-producing cells, equation (3.2). The system as a whole is a functional, mapping the external stimuli into a cytokine profile,

$$s_1(t), \ldots, s_r(t) \mapsto u_1(t), \ldots, u_n(t)$$

where the latter directs the immune response.

When the stimuli evolve slowly (i.e. much slower than the typical timescale $1/\mu_\ell$ given by the inverse decay rates of the cells), and when there is just one cell type ($m = 1$), this mapping is quasi-static, in the sense that the stimuli at time t uniquely determine the cytokine profile at that moment in time (up to transient behaviour). In that case, the cytokine network then behaves essentially like a look-up table. However, when $m \geq 2$, this look-up table will itself evolve over time, depending on the *history* of the stimuli. Moreover, certain rapid changes in the stimuli may precipitate sudden transitions to a different look-up table.

To illustrate these points more concretely, consider the following specification:

$$\psi_{\ell k}(u_1, \ldots, u_n, s_1, \ldots, s_r) = \overline{\psi}_{\ell k} S\Big(\textstyle\sum_{i=1}^{n} w_{\ell k i} u_i - \tilde{\theta}_{\ell k}\Big) \tag{3.3}$$

with

$$\tilde{\theta}_{\ell k} \stackrel{\text{def}}{=} \theta_{\ell k} - \sum_{j=1}^{r} \tilde{w}_{\ell k j} s_j \tag{3.4}$$

where $\theta_{\ell k} > 0$ represents the *stimulation threshold* of cell type ℓ as regards production of cytokine k, and $\overline{\psi}_{\ell k} > 0$ represents the maximum cell-specific secretion rate of cytokine k by a cell of type ℓ. The function S is monotonically increasing, with $S(x) \in [0, 1]$ for all $x \in \mathbb{R}$, with $\lim_{x \to -\infty} S(x) = 0$ and $\lim_{x \to +\infty} S(x) = 1$; an example is $S(x) = 1/(1+\exp\{-x\})$. The parameters $w_{\ell k i}$ and $\tilde{w}_{\ell k j}$, which may be negative, zero, or positive, characterize how the production of cytokine k by cell type ℓ is affected by stimulation by cytokine i or external stimulus j.

To gain an insight into the dynamics of this model, consider first the case $m = 1$ (just one cell type), $v_1(t) \equiv \overline{v}$ (timescale separation) and with the $\tilde{\theta}_{1k}$ fixed for all cytokines k. Also,

let us take S to be the Heaviside step function, i.e.

$$S(x) = 0, \quad x < 0, \tag{3.5}$$
$$= 1, \quad x \geq 0; \tag{3.6}$$

the quantities $\tilde{\theta}_{1k}$ behave as 'crisp' thresholds for this choice. Let $(\overline{u}_1, \ldots, \overline{u}_n)$ be a (quasi-) stationary point which solves $\dot{u}_k = 0$ for all k - in that case we approximate that v is a constant, which does not change appreciably over the response timescale of the cytokines. There are *at most* 2^n such stationary points, since there are n different sigmoid functions ψ_{1k} for $k = 1, \ldots, n$, and each of these can take either the value zero or the positive value $\overline{\psi}_{1k}$ at the steady state, but only provided that this is compatible with the corresponding threshold for the sigmoid. We consider the region around this stationary point bounded by the hyperplanes that are the locus of $\sum_{i=1}^{n} w_{\ell k i} u_i = \tilde{\theta}_{\ell k}$; stationary points do not generically lie on such a hyperplane. This region is the basin of attraction of the asymptotically stable point $(\overline{u}_1, \ldots, \overline{u}_n)$, since we have either $S = 1$ or $S = 0$ throughout this region for all k.

When S is a smooth sigmoid function (such as $S(x) = 1/(1 + e^{-x})$), the situation becomes more complicated. However, we will in general be able to define regions around the stationary points in which $\sum_{k=1}^{n} (\dot{u}_k)^2$ satisfies the properties of a Lyapunov function (see chapter 10 in [9]). The weight parameters, that characterize cell types, are fixed in this model. Thus "learning" in the classical neural network sense only takes place over the much slower evolutionary timescale on which novel cell types arise.

IV. Numerical results

In this section we present some preliminary numerical results for the dynamics of the system (3.2). We concentrate on the special case when there are only two cytokines ($n = 2$) and one cell type ($m = 1$). The external stimuli are also combined into a single quantity $s_1 = s(t)$. In that case, the system can be rewritten in the form

$$\begin{aligned} \dot{u}_1 &= \psi_1(u_1, u_2, s)\, v - \nu_1\, u_1, \\ \dot{u}_2 &= \psi_2(u_1, u_2, s)\, v - \nu_2\, u_2, \\ \dot{v} &= (\phi(u_1, u_2, s) - \mu)\, v. \end{aligned} \tag{4.1}$$

We assume that ψ_j, $j = 1, 2$ are given by sigmoid functions (denoted S) as before, and so take

$$\psi_j = \overline{\psi}_j\, S(\sum_{k=1,2} W_{jk}\, u_k - \tilde{\theta}_j), \qquad j = 1, 2;$$

in fact we fix $S(x) = 1/(1 + \exp\{-x\})$. Since there is the freedom to rescale u_1 and u_2, we can set $\overline{\psi}_1 = 1 = \overline{\psi}_2$ without loss of generality. We further suppose that the stimulus s encourages the proliferation of cells, as does the presence of the cytokine corresponding to u_2, while the cytokine measured by u_1 instead tends to decrease the overall growth rate of cells (which can lead to rapid cell death). The latter assumptions suggest that (as a first approximation) a suitable form for the function ϕ should be

$$\phi(u_1, u_2, s) = s\, u_2 \exp(-\gamma u_1), \qquad \gamma > 0.$$

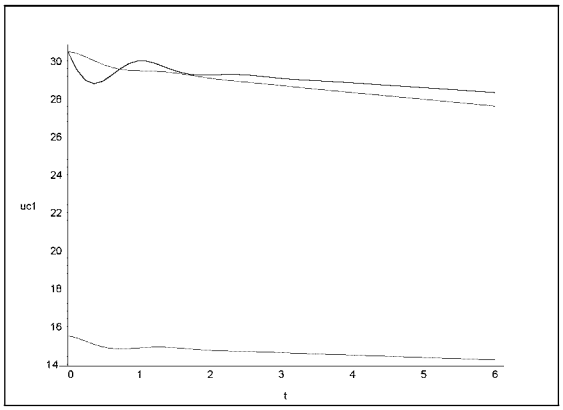

Fig. 1. *Solutions of the system (4.1) for constant stimulus $s = 1.5$, starting from initial values $u_1 = 15.5$, $u_2 = 30.5$, $v = 30.5$.*

For all the numerical results presented here, we have taken parameter values as follows:

$$\nu_1 = 1, \quad \nu_2 = 2, \quad \mu = 10, \quad \gamma = 0.1$$

Also, we have fixed constant threshold values $\tilde{\theta}_1 = 6$, $\tilde{\theta}_2 = 11$, and the interaction matrix (that determines the regions where the sigmoid functions are close to 0 or 1) is taken as

$$\begin{pmatrix} W_{11} & W_{12} \\ W_{21} & W_{22} \end{pmatrix} = \begin{pmatrix} -1 & 1 \\ 1 & 0 \end{pmatrix}.$$

To begin with, in Figure 1 we have plotted the case where the stimulus is constant ($s = 1.5$), and the cytokine and cell densities are started at initial values close to a quasi-stationary state, which can be calculated by approximating S by a Heaviside function. To be more precise, we solve the equations $\dot{u}_1 = \dot{u}_2 = \dot{v} = 0$ to get the steady states for this simplified choice of S, and these serve as suitable initial data. The topmost curve with the initial wiggle is v, while the adjacent curve is u_2, and u_1 is at the bottom. A gradual decay is initially observed, but as the time increases further there is a sudden rapid drop (corresponding to a threshold being crossed), and thereafter all three quantities decay rapidly to zero. (This ultimate decay is not shown in the Figure, however.) Thus in this case the constant stimulus means that the quasi-stationary state does not persist.

In Figure 2 we have plotted the case where the stimulus is constant, and the cytokine and cell densities are started at initial values close to a different quasi-stationary state (whose approximate value is found in the same way as before). The initially topmost curve for v goes through an inflection point, while u_2 is eventually the slowest to decay: in this case rapid (exponential) decay towards zero of all three quantities sets in almost immediately. Once again the point $(0, 0, 0)$ is an attracting fixed point. Thus in this case once more the constant

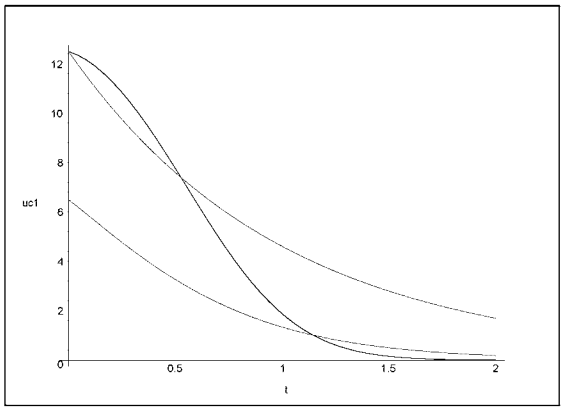

Fig. 2. *Solutions of the system (4.1) for constant stimulus* $s = 1.5$*, starting from initial values* $u_1 = 6.5$*,* $u_2 = 12.5$*,* $v = 12.5$*.*

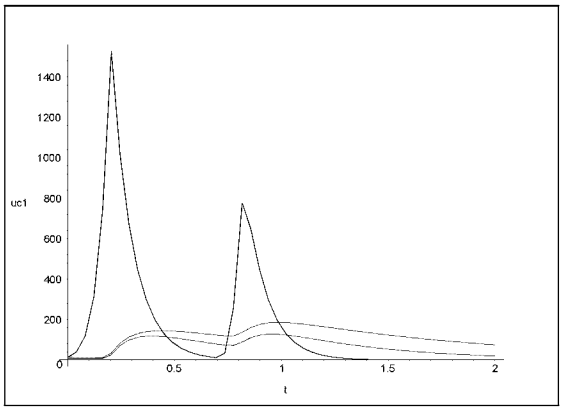

Fig. 3. *Solutions of the system (4.1) for two different stimuli applied between times* $t = 0, 0.2$ *and between times* $t = 0.7, 0.9$ *respectively. starting from initial values* $u_1 = 6.5$*,* $u_2 = 12.5$*,* $v = 12.5$*. All profiles are shown together here, with the cytokine profiles at the bottom.*

stimulus means that the (approximate) quasi-stationary state does not persist.

In order to get a different sort of behaviour (for this choice of network functions and topology) we have solved the same system but with two different stimuli applied at different times. To begin with, we take the choice

$$s(t) = 6\,H(0.2 - t)\,H(t) + 700\,H(0.9 - t)\,H(t - 0.7)$$

where H denotes the Heaviside function, corresponding to a weak stimulus between $t = 0$ and $t = 0.2$, followed by a much stronger one between $t = 0.7$ and $t = 0.9$. The surprising effect is that the weaker initial stimulus produces a much stronger output in terms of the proliferation of the cytokine-producing cells, as can be seen from the topmost curve (the profile of v) in Figure 3. This is not so surprising given that we have a nonlinear system i.e. the outputs are not directly proportional to the inputs.

To see the effect of history on the system, we apply the same two stimuli but now with larger intervals between them, for we take the choice

$$s(t) = 6\,H(0.2 - t)\,H(t) + 700\,H(1.1 - t)\,H(t - 0.9)$$

so that now the much stronger stimulus is applied between $t = 0.9$ and $t = 1.1$. This time the the stronger second stimulus produces a stronger response than the first, although not proportionately so, in terms of the increased concentration of cytokine-producing cells, i.e. the topmost curve (the profile of v) in Figure 4. The cytokine responses can be observed from the lower curves.

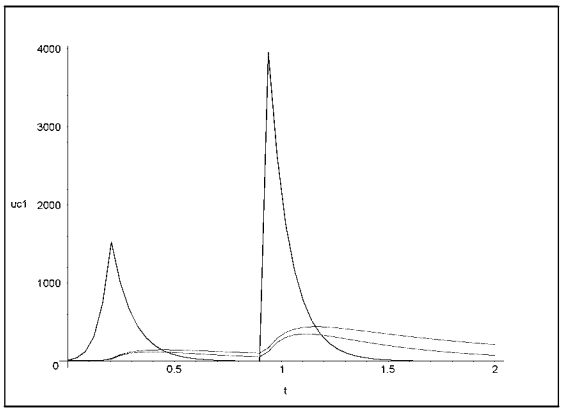

Fig. 4. *Solutions of the system (4.1) for two different stimuli applied between times* $t = 0, 0.2$ *and between times* $t = 0.9, 1.1$ *respectively. starting from initial values* $u_1 = 6.5$*,* $u_2 = 12.5$*,* $v = 12.5$*. All profiles are shown together here, with the cytokine profiles at the bottom.*

V. Conclusions

Even with only two cytokines and one type of cell, the ACN model defined by (3.2) shows a rich variety of different behaviours. For autonomous ODE models with three dependent variables, it is known that there is already the potential for chaotic behaviour (as in the famous Lorenz system), with strange attractors, so if we increase the number of variables then there are even more possibilities. A great deal of further analysis (both exact and numerical) will be needed to understand how best to choose the parameters and specify the interactions (i.e. how to fix the topology of the network). Ideally, to exploit this network for computation, some way to control the response of the system would be required (perhaps using the optimal control ideas outlined in [1]). Nevertheless, we hope to gain a better understanding of this model in future, and compare it with existing models of cytokine networks. One

fruitful direction is to make a comparison with differential equation models for neural networks such as in [8], which suggest that it is possible to use a training phase for the ACN, in order to adjust parameters in the interaction matrix for the cytokines to suitable values.

ACKNOWLEDGMENT

Andy Hone is grateful for the support of an EPSRC Springboard Fellowship.

REFERENCES

[1] H. van den Berg and A. Hone, Mathematical analysis of AIS dynamics and performance, in *In Silico Immunology*, D. Flower and J. Timmis (eds.), Springer (2006) at press.
[2] L.N. de Castro and J. Timmis, *Artificial Immune Systems: A New Computational Intelligence Approach.* Springer-Verlag (2002).
[3] L.N. de Castro and F.N. von Zuben, Proc. IEEE Comp. Soc. Press, SBRN '00, vol. **1** (2000) 84-89.
[4] E. Clark, A. Hone and J. Timmis, Springer LNCS **3627** (2005) 318-330.
[5] M. Villalobos-Arias, C.A. Coello Coello, O. Hernandez-Lerma, Springer LNCS **3239** (2004) 226-235.
[6] V. Cutello and G. Nicosia, Springer LNCS **2527** (2002) 361-370.
[7] V. Cutello, G. Nicosia, M. Pavone and J. Timmis, *An Immune Algorithm for Protein Structure Prediction on Lattice Models*, IEEE Trans. Evol. Comp. **10** (to appear).
[8] J.J. Hopfield, Proc. Natl. Acad. Sci. USA **81** (1984) 3088-3092.
[9] D.W. Jordan and P. Smith, *Nonlinear Ordinary Differential Equations*, 3rd edition, Oxford: Oxford University Press (1999).
[10] A.J. Meade & H.C. Sonneborn, *Adv. Eng. Software* **27** (1996) 213–225.
[11] B. Mitavskiy, J.E. Rowe, A. Wright and L.M. Schmitt, Exploiting Quotients of Markov Chains to Derive Properties of the Stationary Distribution of the Markov chain associated to an Evolutionary Algorithm, preprint (2006).
[12] J.D. Murray, *Mathematical Biology*, Springer (1993).
[13] M.A. Nowak & R.A. May, *Virus dynamics*, Oxford: Oxford University Press (2000).
[14] Ott, E.: *Chaos in dynamical systems*, Cambridge: Cambridge University Press (1993).
[15] H.-O. Peitgen & D.H. Richter, . *The Beauty of Fractals: Images of Complex Dynamical Systems.* Springer-Verlag (1986).
[16] A.S. Perelson & G. Weisbuch, *Rev. Modern Phys.* **69** (1997) 1219–1267.
[17] E. Pinch, *Optimal Control and the Calculus of Variations*, Oxford: Oxford University Press (2002).