ELSEVIER

Biophysical Chemistry 72 (1998) 223–230

Biophysical Chemistry

Multiple attractors in immunology: theory and experiment

Lee A. Segel

Department of Applied Mathematics and Computer Science, The Weizmann Institute of Science, Rehovot 76100, Israel

Revision received 12 January 1998; accepted 13 February 1998

Abstract

This selective survey discusses the relative merits of various modeling approaches in immunology that exhibit multiple attractors, and also assesses the ability of the different models to contribute to deeper biological understanding. The first topic is global anti-idiotypic network models, which, like Hopfield neural network models, exhibit a large number of steady states that are identified with memory. It is shown that a 'reverse engineering approach' to T-cell vaccination for autoimmunity, featuring steady states corresponding respectively to 'normality', 'vaccination' and 'disease', is able to spur new experiments, in spite of the model's deliberate neglect of almost all biological detail. Mention is made of several other T-cell models that feature bistability for Th1 or Th2 dominance, or for activation and unresponsiveness.

Keywords: Idiotypic networks; Autoimmunity; T-cell vaccination; T-helper cells; Allergy; Falsifiability

1. Introduction

It is no wonder that the immune system is complex, for its principal job is defense against a wide variety of pathogens with short generation times. Thus, pathogens have evolved a remarkable variety of strategies to coexist with the verte-brate host, and the host in turn has coevolved a sophisticated set of counterstrategies. For modelers, complex systems are 'good for business' since understanding such systems generally requires many models in order to comprehend a variety of emergent properties at several levels of organization [1]. Here we survey a few representative non-linear models of immune action. We focus on the predicted appearance of multiple steady states and how such predictions relate to the biology of immunity and autoimmunity.

2. Multiple steady states in large immune networks

In biology, the archetype of models exhibiting multiple steady states is the Hopfield neural network [2,3] and its descendants. In the Hopfield model each of the myriad steady states represents a memory. Retrieval of a memory from a rough reminder of the memory state (associative memory) is modeled by the passage of the network to a given steady state from anywhere within that state's domain of attraction. Neural nets can perform biologically-interesting computations such as 'learning' to distinguishing coding from non-coding sequences of DNA [4]. Nonetheless, many biologists are skeptical as to whether representing a neuron by a switch captures enough of its nature even to provide overall qualitative insights. As has been said, a neuron is more like a chip than a switch [5].

0301-4622/98/$19.00
PII S0301-4622(98)00136-7

Well before Hopfield, Jerne [6] advocated the importance of the analogy between neuronal networks and the massive network of interacting immune cells. Due to Jerne's eminence (he won the Nobel prize) and because of the intrinsic attractiveness of his idea, many theorists were motivated to construct increasingly complex models of immune networks. There is memory in immunology too – one doesn't get measles twice. Biologists almost unanimously ascribe memory to special 'memory cells' [[7], Ch. 9]; theorists emphasize the alternative that network attractors can provide a form of memory.

As reviewed by Perelson and Weisbuch [8], the prototypical immune network is composed of N interacting clones of antibody-secreting B-cells b_i, i = 1,2,.., N. (Terminology: *antigens* generate an immune response. *Antibodies* are molecules that bind to antigens and help lead to their elimination.) In the basic network model these cells have a source (of strength m) in the bone marrow, a death rate d, and a maximum proliferation rate p. How close the proliferation rate approaches its maximum depends on the excitation h_i received by cells of clone i. This yields

$$\frac{db_i}{dt} = m - db_i + pb_i f(h_i) \tag{1a}$$

Exactly as in Hopfield neural networks, the *influence field* h_i is assumed to be expressible as a linear combination (via constants J_{ij}) of the populations of all the cell types:

$$h_i = \sum_j J_{ij} b_j \tag{1b}$$

The excitation function is taken to be

$$f(h_i) = \frac{h_i}{\theta_1 + h_i} \frac{\theta_2}{\theta_2 + h_i} \tag{1c}$$

where θ_1 and θ_2 are constants. We take $\theta_1 \ll \theta_2$, in accordance with various observations and certain theories, so that ln $f(h)$ is a symmetric bell-shaped function of the field h.

Further simplification renders the above model much more tractable. It is assumed that clone b_1 is the only one to react with an invading antigen (concentration γ). In the simplest case, the antigen does not reproduce and thus disappears at a rate proportional jointly to γ and the B-cell concentration b_1:

$$\frac{d\gamma}{dt} = -k\gamma b_1 \tag{2a}$$

Additional assumptions are that exactly c clones interact with the 'root' clone b_1, and that each of those clones interacts with exactly c further clones, etc. (Such a network organization is called a *Cayley tree* by physicists.) All non-zero interaction coefficients are taken to have the same value, one. This gives

$$h_1 = cb_2 \tag{2b}$$

$$h_i = b_{i-1} + (c-1)b_{i+1} \quad (i \geq 2) \tag{2c}$$

Let us consider the special case $N = 2$. There are two types of clone: b_1 reacts with the antigen and b_2 reacts with b_1. The latter reaction is possible because portions of the antigen receptors of b_1 are molecules unique to the b_1 clones that the immune system can recognize, so that they can be viewed by the b_2 clones as antigens.

Similarly, b_1 can recognize b_2. Since in Greek 'idio' means 'own, personal, private, peculiar, distinct' (Oxford dictionary) such networks are called *idiotypic* and a b_1–b_2 interaction is called *anti-idiotypic*.

Analysis of the model given by Eqs. (1) and (2) when $N = 2$ reveals that there are typically three possible stable steady states; see Fig. 1. (In all of these

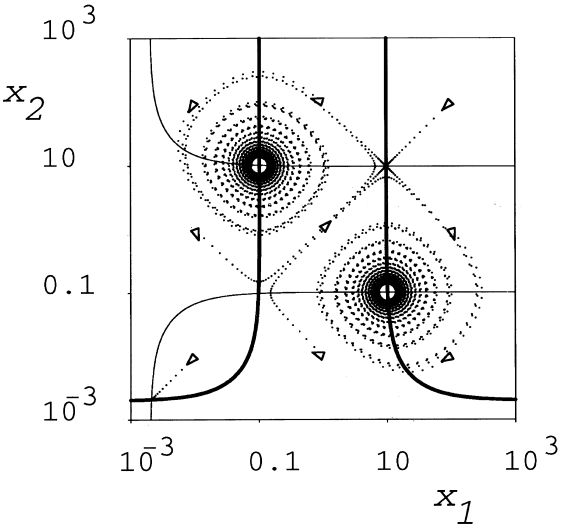

Fig. 1. (Fig. 15, Perelson and Weisbuch [8]. Reproduced with permission.) Virgin state V (10^{-3}, 10^{-3}), immune state I (10, 10^{-1}) and tolerant state T (10^{-1}, 10) of model (1) with simplifications (2).

states, the antigen concentration is zero; antigen serves to stimulate the network and then is always eliminated.)

(a) State V. Antigen stimulation is very small or non-existent. Only the root clone b_1 has a non-negligible concentration, balancing the source m and death d. This is the *virgin state*, in which the immune system is essentially unstimulated.

(b) State I. Clones b_1 and b_2 both have significant concentrations (considerably above the virgin level), with b_1 much larger than b_2. To understand this situation, recall that the excitation function $f(h)$ is a bell-shaped function of the field h. When excitation is below a critical value h_c (where f is maximal) an increase of the field results in an increase of excitation. When h exceeds h_c, further increase of h decreases excitation. These two ranges of fields are termed *stimulatory* and *suppressive*, respectively. In state I the clone b_1 has a high population and hence induces a suppressive field on b_2, which has a low population. The field on b_1 is excitatory. State I is called *immune* because if antigen is injected into a system in state I then elimination of the antigen is much faster than it would be if the initial state were virgin. This model can thus explain the basic phenomenology of immune memory. An initial injection of antigen shifts the system from the virgin state V to the immune or memory state I. A second injection of antigen is rapidly eliminated.

(c) State T is the mirror image of state I: b_1 is small and b_2 is large. This state is called *tolerant* since clone b_1 is quite suppressed and is thus only able to react weakly to antigen. There may be a bearing here on the fact that the immune system is tolerant to self. (Most, but not all, anti-self clones are eliminated by various mechanisms, but anti-self clones are known to be present in normal people – yet they generally do no harm.)

Some results of analyzing the simplest immune network model are sketched. Among the early modeling of this type was work by Hoffmann [9] and de Boer [10]. A great deal of further research has been done; see the review by Perelson and Weisbuch [8]. Arbitrary values of N have been considered and the restriction to Cayley trees has been lifted. Antibody secretion, ignored in Eq. (1) and (2), has been explicitly incorporated. Other frameworks for modeling have been explored, notably the computer-oriented 'bit-string' approach and the integro-differential equations of the shape-space approach. The 'metadynamics' approach deals with ever-shifting sets of equations, as very small clones are deleted and new clones enter. Networks of automata have been examined. Among the further issues considered are the size and the connectivity of the network as it emerges in the newly born. Estimates of the 'memory capacity' of the network have been made. Investigations have been carried out to determine conditions permitting and forbidding the 'percolation' of an initially-localized disturbance throughout a network.

Theorists (with some exceptions [11]) love the intricacies of immune network theory. Yet, for a variety of reasons, most experimentalists are disenchanted. A colorful expression of this attitude appears in an article by Langman and Cohn [12] which asserts that 'the complete idiotype network is an absurd immune system.. masked behind colorful metaphors, fetching analogies with the nervous system and cunning neologisms'. Langman and Cohn assert that either 'the idiotype network fragments into many sets of non-overlapping pairs.. or connectivity is a complex web.. and the network is tied in a Gordian knot.' As to the 'Gordian knot', a variety of studies have shown that signals need not percolate wildly throughout essentially infinite models of immune networks; a priori, infinite networks (or very large finite networks) pose no problem of principle. As to the likelihood that interactions are linked into id and anti-id pairs, it is indeed true that direct observations of a chain of idiotype anti-idiotype anti-anti-idiotype network activities have typically been performed under artificial conditions. Moreover, models of the type we have been discussing assume the presence of crucial auxiliary signals that typically come via so-called 'T-helper cells'. It is thus crucial that the interaction of B- and T-cells be included in large network models. A good start on this difficult problem has been made by Carneiro et al. [13,14]. There is also an interesting contribution by A. Neumann and A. Perelson (unpublished).

In spite of wide-spread skepticism among some experimental immunologists concerning anti-idiotypic networks, a recent major textbook cautiously concludes that 'while the elements of an immune network exist in mature immune systems and can be experimentally manipulated, their actual role in repertoire

control remains speculative' [7; p. 12:20]. In favor of the importance of large-scale networks is accumulating experimental evidence that an extensive network of self-reacting antibody-secreting cells is found in normal individuals [15,16]. Note that the network is held to be primarily self-reacting, not antigen-driven as is the case in most network models. Among theoretical works, particularly relevant here is the approach to networks exemplified by Ref. [17]. Both theory and experiment are concerned with immune states that typically continually vary with time; it appears that more complex attractors than steady states are involved. Indeed, oscillatory and/or chaotic-seeming oscillations have been found in several network models. It remains to be determined whether the observed fluctuations are mainly network generated or whether they are principally an expression of stochastic forcing of the network from without. If the former is the case then the interesting situation will have arisen wherein an important natural dynamical system has been demonstrated to feature complex aperiodic attractors.

Especially noteworthy are experimental findings concerning the autoimmune disease 'lupus' (systemic lupus erythematosus). It appears that network connectivity patterns differ in diseased and normal individuals, for 'multivariate statistics..distinguished the two groups of donors, and demonstrated a larger dispersion and wider time-dependent variations in the patient population, compared with the healthy controls' [18]. Similarly, the putative autoimmune disease idiopathic thrombocytopenic purpura was 'associated with generalized alterations of antibody repertoires that may be characteristic enough to allow for diagnosis' [19]. Thus there is strong evidence that a global network view of the self-reacting clones not only provides a useful perspective for understanding immune physiology, but also may have important medical applications. Little network modeling is focused in this direction; more should be.

3. Small immune networks via reverse engineering

Continuing our earlier comparison, we recall that Jernean immune networks are cousins of Hopfield neural networks, both of which are controversial attempts to take a global overview of network complexity. In neurobiology, reacting to worries about oversimplification, theorists have constructed more complex neuronal models but so far the price paid, usually, is restriction to an 'oligonetwork' containing relatively few cells. Much has nonetheless been achieved that is of undoubted relevance to many biological applications where such oligonetworks play key roles. (See [20–22].) Analogously, there is increasing evidence that preformed oligonetworks play major roles in immunology [23–25]. Theorists have correspondingly influenced the course of experiments with models of small immunological networks. To demonstrate this I will first examine an attempt by myself and E. Jäger to abstract essential features of a phenomenon in immunology via a method that we termed 'reverse engineering' [26]. Much of the presentation of this topic closely follows material that appears in Ref. [27].

Of particular interest to us was the use of 'T-cell vaccination' to combat autoimmune diseases. The relevant experiments concern mice and rats that sometimes can be induced to exhibit diseases that are close in symptoms and cause to human afflictions, for example, multiple sclerosis (MS). By inoculation of suitable doses of appropriate T-cells, certain strains of rodents can be driven into an MS-like disease called EAE (experimental autoimmune encephalomyelitis). Smaller doses do not induce disease. In fact the animals are vaccinated, in the sense that if the smaller dose is later followed by the standard disease-giving dose no disease develops.

Deliberately ignoring almost all biological detail, Jäger and I attempted to construct the simplest-possible mathematical model that would exhibit the observed phenomenology. This we did by exhibiting a variety of differential equation pairs, all of which had three stable steady states that we could identify respectively with the *normal*, *vaccinated* and *diseased* states of the mouse. One of the differential equations described the population dynamics of the 'bad guy' effector cells E that bring about disease, and one described the 'good guy' regulator cell population R that regulates the proliferation of the 'bad guys'. Conversely, the effectors influence the proliferation of the regulators. The following is an example of the equations:

$$\frac{dE}{dt} = 0.01 + \left(3.5 - 0.5R + \frac{100E^2}{25 + E^4}\right) \quad (3a)$$

$$\frac{dR}{dt} = 0.01 + R(-0.1R + 0.02E^3 - 0.33E^2 + 1.3E + 1) \quad (3b)$$

In another application of 'reverse engineering' to T-cell vaccination [28], models reproducing the phenomenology were constructed wherein *disease* was a transient, not a steady state (as is in fact the case in EAE but not in autoimmune diabetes for example). 'Tristability' is replaced by bistability, with *normal* and *vaccinated* as coexistent stable steady states.

Modelers derive satisfaction from constructing equations whose solutions reproduce some complex phenomenology, but biologists do not regard this as a meaningful achievement. What gives the work biological significance is the fact that the models have a life of their own, and thus predict more phenomena than they are set up to reproduce. In particular, various versions of the present model predicted that although a high dose of effector cells gives disease, an even higher dose might lead to vaccination. Our biological colleagues found this prediction simultaneously counterintuitive and attractive (since such a general model produced the prediction). They performed the relevant experiments – and verified the predictions [29].

Fig. 2 depicts some details of the theory and the experiments. In the figure, E and R represent populations of effector and regulator cells. The points labeled N, V and D represent 'normal', 'vaccinated' and 'diseased' stable steady states of the system, with respective domains of attraction hatched, clear, and dotted. Heavy lines delineate five representative solution trajectories. (1) A sufficiently large addition of effectors E to the normal state leads to disease. (2) A smaller effector dose leads to the vaccinated state. (3) When the standard disease-giving dose is given to a vaccinated animal, no disease results and the system returns to the vaccinated state. (4) Surprisingly, a very large dose of effectors leads to the vaccinated state. (5) An even larger dose leads to the normal state. A very small dose of effectors gives a state of the system that is within the domain of attraction of the normal state and thus leads to a return to that state (not shown).

The insert to Fig. 2 shows experimental results for autoimmune diabetes in female non-obese diabetic mice. Disease (hyperglycemia) is monitored 1 week after injection of the diabetogenic C9 clone. (The number in parenthesis indicates the number of mice used to obtain an experimental point.) The major predictions of the model are verified, notably Section 4.

4. Other small network models with multiple attractors

Despite the interesting experiments spurred by the reverse engineering approach, it must not be thought that this is normally the preferred approach to modeling complex systems. Reverse engineering is just one weapon in the theoretician's armory. Usually there is no substitute for the conventional strategy of close examination of the experimental findings and subsequent construction of models that are firmly based on these findings. (Even here, the use of reverse engineering reminds the modeler that reproduction of considerable phenomenology is no guarantee that a detailed model is correct – different details may lead to the same set of predictions.) In particular,

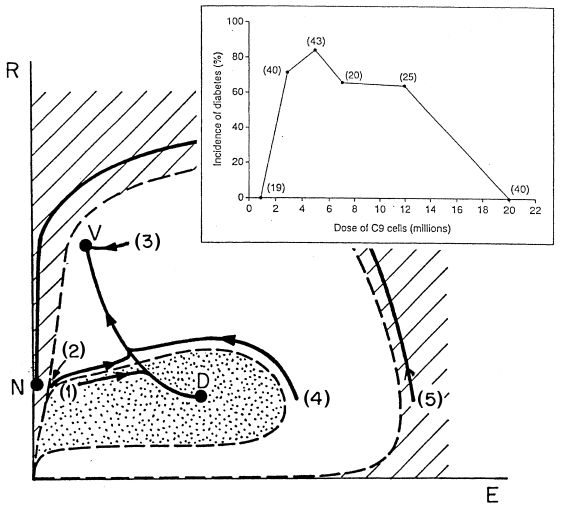

Fig. 2. (Composite of Figs. 2 and 3 from Segel et al. [29], with permission of Elsevier Science). Main figure: phase portrait for a 'reverse engineering' model of T-cell vaccination. Insert: experimental results. For explanation, see text.

Refs. [30] and [31] contain rather different fairly detailed models that both yield similar mathematical structures (i.e. similar phase planes) and hence the same overall predictions as the reverse engineering models for T-cell vaccination.

A good contrast to the super-phenomenological reverse engineering models is the careful study of Th1-Th2 cross regulation in Ref. [32]. These two types of T-helper cell secrete two different sets of cytokines and thereby direct radically different immune responses. Some of the cytokines are cross-inhibitory: Th1 secretions suppress Th2 and vice versa. The basic model here consists of a non-linear system comprising a dozen ordinary differential equations for the concentrations of various cell types and cytokines. (Although detailed, the model of course makes a number of simplifications.) The fact that the time scale of chemical changes is much shorter than that of cellular changes is the basic observation that makes possible reduction of the basic model to four differential equations together with a number of auxiliary algebraic equations. Analysis of the model reveals large parameter domains where bistability appears. Depending on the initial conditions, at steady state either Th1 cells or Th2 cells are present, but not both. For other parameter choices an all-Th1 or all-Th2 state results, regardless of initial conditions. Thus, the model demonstrates that a sharp choice between Th1 and Th2 can be made, and delineates many factors that influence the choice.

The final model in [32] consists of differential equations for the Th1 and Th2 cells, for antigen, and (in essence) for a population of antigen-presenting cells that activate the Th cells. The accompanying algebraic equations provide formulae for the mutual interaction of the four populations whose change is described by the differential equations. These formulae embody many underlying biological assumptions. However, quite different biology might lead to similar formulae and hence to similar qualitative and even semi-quantitative behavior of solutions to the four differential equations. These considerations have been exemplified in two models based on the description in Ref. [32] of Th1-Th2 interaction [31].

The line of thought just presented is hardly novel, but it bears continual emphasis. The line between the reverse engineering approach and conventional modeling is not as sharp as it first appears. It must always be kept in mind that a number of rather different underlying modeling assumptions can give rise to similar final results.

An application of the Th1-Th2 dichotomy has been made to allergy, over-response of the immune system to foreign substances. Ref. [33], apparently the first allergy model, attempts to explain why the standard immunotherapy treatment for allergy often fails, while even when it succeeds, treatment can take years [34]. The model is based on the description in Ref. [32] of Th1-Th2 interaction, for allergy is associated with Th2 dominance. An additional matter that is important for allergy is the fact that Th2 cells are favored in lymph nodes that drain mucosal surfaces, while Th1 responses seem to predominate in the spleen and non-mucosal lymph nodes. If we focus on pollen allergies, inhalation of the allergen introduces it into mucosal compartments but immunotherapy provides small antigen doses in non-mucosal compartments. Hence, exchange between mucosal and non-mucosal compartments was incorporated into the model.

Numerical exploration of the two-compartment allergy model indicates that its overall behavior is analogous to the underlying model of Ref. [32]; in particular, steady states are exclusively Th1 or exclusively Th2. According to the model, only for certain parameter domains does immunotherapy indeed switch the dominant state from Th2 (allergy) to Th1 (no disease). In particular, if intercompartmental communication is too weak then the compartments operate essentially independently and injection in one compartment has no long term effect on the other. If communication is too strong, the two compartments are effectively merged and again treatment fails. The required intermediate strength of communication between the mucosal and non-mucosal compartments is in fact rather weak, which accounts for the length of time taken by successful immunotherapy to effect a cure.

Another type of model is not for a small network of interacting cells, but rather for a small network of intracellular chemicals. Indeed, within any cell there is in fact a complex network of messengers that originate from the binding of various receptors, and that trigger other messengers. Interaction of all these messengers eventually results in the turning on or off of suitable genes. For T-helper cells, two major states

can result, *activation* (cytokine production or proliferation) or *unresponsiveness* (anergy). In Ref. [35] these states are alternative stable steady states of a dynamical system. The model is of logical type (concentrations are 1 or 0, high or low) with characteristic time delays for each transition. Tyrosine phosphorylation is a key regulator. A major feature of the results is a demonstration, in accord with experiment, that whether a given ligand activates or inhibits depends crucially on the delays.

Related earlier work [36,37] dealt with regulation of antibody production. A multiplicity of steady states was obtained for models whose core is a negative feedback loop between helper and suppressor T-cells, together with autocatalytic feedbacks of these two populations on themselves. The two cited papers contain insightful comparisons between logical and continuous (differential equation-based) analyses, both of which were used in attacking the same fundamental problem.

5. Concluding remarks

In surveying immunological models with multiple attractors, I have tried not only to sketch the theoretical aspects of the models but also to examine their relevance, or lack of it, to experiment. Further discussion of the results from the reverse engineering model will illustrate some of the subtleties involved in such an examination.

In the first submitted version of the reverse engineering paper [29], after showing that the model simulated salient experiments involving T-cell vaccination we wrote 'At this point the reader may wonder why he or she should bother to read about equations that merely reproduce what we can know from experimental results already described in the literature. But hold on. The chief value of the model to experimentalists is that it leads to intuitively bizarre predictions that are experimentally feasible; it makes the unthinkable thinkable and so do-able.' From the journal we received a strong objection, which asserted that the value of a model is that it makes the thinkable testable and insisting that the unthinkable is just that, unthinkable. We thus changed the offending assertion to the milder form: 'The chief value of a model is to formulate predictions that can prompt experimentalists to do experiments that they might not have otherwise done.'

Note that the 'predictions' of the reverse engineering model are weakly based. Details of the phase-plane behavior illustrated in Fig. 2 are essentially guessed; these details are not derived from the solution of equations based on specific and reasonable biological hypotheses concerning the cell populations treated in the model. In particular, just as likely, a priori, as the shape of the dotted domain of attraction in Fig. 2 is another much more elongated shape in which this region forms a belt, stretching far to the right of the steady state D that represents disease. If the domain of attraction of D had such a belt shape, then the model would predict the unsurprising result that larger and larger addition of effectors to a normal animal would all result in disease.

The real world is subtle. Maneuverability was a concomitant of David's weakness compared with the mighty giant Goliath; daring vanquished overconfidence. The reverse engineering model provides another example of strength in weakness. The model's very lack of verifiable biological assumptions makes it clear that a wide variety of particular implementations might yield a phase plane of the form of Fig. 2. The essence of what is required is a sufficiently strong regulatory response to high effector levels. The model points out that if such regulation is present then certain counterintuitive consequences result.

Experiments showed that these unexpected consequences in fact occur. Popper [38] is the prophet of falsifiability as the criterion for a good scientific theory (although, as is so often the case, the prophet's opinions are much more shaded than the stark dogmas of his followers). By contrast, Segel et al. [29] asserted that 'especially for conceptual models such as ours, the value of a model may lie not so much in its falsifiability as in the experiments it makes thinkable.' Such opinions are not limited to immunology. For example, in a consideration of models in the earth sciences, it was concluded that 'the primary value of models is heuristic' [39]. However, if positivism and falsifiability are often not central in biology and if making the unthinkable thinkable is frequently the true goal of mathematical modeling, then this makes the modelers' challenge much more difficult.

Acknowledgements

This work was supported by the US-Israel Binational Science Foundation (BSF), grant 95–00526. I am grateful to A. Perelson for useful suggestions.

References

[1] L.A. Segel, Complexity 1 (1995) 18.
[2] J.J. Hopfield, Proc. Natl. Acad. Sci. USA 79 (1982) 2554.
[3] J.J. Hopfield, Proc. Natl. Acad. Sci. USA 81 (1984) 3088.
[4] R. Farber, A. Lapedes, K. Sirotkin, J. Mol. Biol. 226 (1992) 471.
[5] T. Poggio, C. Koch, Sci. Am. 256 (1987) 46.
[6] N.K. Jerne, Ann. Immunol. (Inst. Pasteur) 125C (1974) 373.
[7] C.A. Janeway, Jr., P. Travers, Immunobiology, Blackwell, Oxford, 1994.
[8] A.S. Perelson, G. Weisbuch, Rev. Mod. Phys. 69 (1997) 1219.
[9] G.W. Hoffmann, in: S. Levin (Ed.), Systems Theory in Immunology, Lecture Notes in Biomathematics, Springer, Berlin, 32 (1978) 239.
[10] R.J. De Boer, in: A.S. Perelson (Ed.), Theoretical Immunology, Part Two, SFI Studies in the Sciences of Complexity, Addison–Wesley, 1988, p. 265.
[11] Z. Grossman, Theories of Immune Networks, Springer, Berlin, 1989, p. 38.
[12] R.E. Langman, M. Cohn, Immunol. Today 7 (4) (1986) 100.
[13] J. Carneiro, A. Coutinho, J. Faro, J. Stewart, J. Theor. Biol. 182 (1996) 513.
[14] J. Carneiro, A. Coutinho, J. Stewart, J. Theor. Biol. 182 (1996) 531.
[15] A. Nobrega, M. Haury, A. Grandien, E. Malanchere, A. Sundblad, A. Coutinho, Eur. J. Immunol. 23 (1993) 3851.
[16] A. Coutinho, M.D. Kazatchkine, S. Avrameas, Curr. Opin. Immunol. 7 (1995) 812.
[17] F.J. Varela, A. Coutinho, J. Stewart, in: W. Stein, F.J. Varela (Eds.), Thinking About Biology, SFI Studies in the Sciences of Complexity, Vol. III, Addison–Wesley, 1993, p. 1.
[18] A. Ayouba, C. Ferreira, A. Coutinho, Scand. J. Immunol. 45 (1997) 408.
[19] A. Sundblad, C. Ferreira, A. Nobrega, M. Haury, E. Perreira, F. Padua, A. Coutinho, J. Autoimmun. 10 (1997) 193.
[20] E. Marder, L.F. Abbott, Theory in motion Curr. Opin. Neurobiol. 5 (6) (1995) 832.
[21] J.M. Weimann, P. Skiebe, H.G. Heinzel, C. Soto, N. Kopell, J.C. Jorge-Rivera, E. Marder, J. Neurosci. 17 (1997) 1748.
[22] E. Av-Ron, H. Parnas, L.A. Segel, Biol. Cybern. 69 (1) (1993) 87.
[23] I.R. Cohen, in: H. Atlan, I.R. Cohen (Eds.), Theories of Immune Networks, Springer, Berlin, 1989, p. 6.
[24] I.R. Cohen, D.B. Young, Immunol. Today 12 (4) (1991) 105.
[25] E. Mozes, in Y. Shoenfeld (Ed.), The Decade of Autoimmunity, Elsevier, Amsterdam, 1998, in press.
[26] L.A. Segel, E. Jäger, Bull. Math. Biol. 56 (1994) 687.
[27] L.A. Segel, Curr. Sci. 73 (1997) 929.
[28] L.A. Segel, E. Jäger, J. Biol. Syst. 3 (1995) 441.
[29] L.A. Segel, E. Jäger, D. Elias, I.R. Cohen, Immunol. Today 16 (2) (1995) 80.
[30] J.A.M. Borghans, R.J. de Boer, Proc. R. Soc. Lond. B259 (1995) 173.
[31] E. Jäger, Ph.D. Dissertation, Dept. of Applied Mathematics and Computer Science, Weizmann Institute of Science, Rehovot, 1996.
[32] M.A. Fishman, A.S. Perelson, J. Theor. Biol. 170 (1994) 25.
[33] M.A. Fishman, L.A. Segel, Bull. Math. Biol. 58 (1996) 1099.
[34] L.M. Lichtenstein, Sci. Am. 269 (1993) 117.
[35] M. Kaufman, F. Andris, O. Leo, in: D. Beysens, G. Forgacs (Eds.), Dynamical Networks in Physics and Biology, Springer, Berlin, 1997, in press.
[36] M. Kaufman, J. Urbain, R. Thomas, J. Theor. Biol. 114 (1985) 527.
[37] M. Kaufman, R. Thomas, J. Theor. Biol. 129 (1987) 141.
[38] K.R. Popper, Conjectures and Refutations, Basic Books, New York, 1962, p. 253.
[39] N. Oreskes, K. Shrader-Frechette, K. Belitz, Science 263 (1994) 641.