Forum: How Complexity Helps to Shape Alloimmunity

Charles G. Orosz, Stephanie Forrest, Steven Hoffmeyr, Irun R. Cohen and Lee A. Segel

Graft 2001; 4; 365

FORUM: How Complexity Helps to Shape Alloimmunity

Introduction

Charles G. Orosz, Ph.D.
Professor, Depts. of Surgery, Pathology, and
Molecular Virology/Immunology/Medical Genetics
The Ohio State University College of Medicine
Rm 355 Means Hall
1654 Upham Drive
Columbus, Ohio, USA 43210
Tel.: 614.293.3212
Fax: 614.293.4541
email: orosz-1@medctr.osu.edu

Over the last 50 years, we have learned much about the basic elements of the immune system, but the exact mechanisms by which these elements come together to eliminate viruses, allografts and parasites from the body remain conjectural. Why is this? Why has the prolonged uncovering of massive amounts of information regarding the basic constituents of the immune system been unable to provide us with a clear understanding of how the immune system operates to maintain the integrity of self? Further, why is this question so rarely asked?

In general, one could argue that the growing gulf between immunologic information and immunologic understanding stems from at least two broad philosophic issues. The first issue involves the scientific method. Under question is not the premise of the scientific method that proscribes the hypothesis-and-test approach to problem solving. What is questionable is the premise that the whole of a problem is but the sum of its parts, and that an understanding of the whole can be achieved by an analysis of sufficiently dissected parts. Clearly, some questions cannot be answered using this approach. For example, if the question is: what constitutes a watch? or how does a watch work?, then the dismantling of a watch would be quite informative. However, if the question is: what is time?, than dismantling a watch will probably provide little information of value.

A correlate of this premise is that really large scientific problems require really laborious dissections. Problems like the function of the immune system may require investigation by generations of scientists before understanding reveals itself. Hence, the current lack of understanding simply reflects the state of the art: inadequate dissection, more analysis needed. This concept permits scientists to labor at dissective investigations without any concern for overall understanding. They toil to find the answer, but remain unconcerned when they do not. In time, a critical mass of information will finally be obtained by the research community, and from that will materialize the underlying pattern of the many parts, leading to the final solution of the problem. In this situation, it is the acquisition of information that becomes the goal, while the acquisition of understanding recedes from attention.

The second philosophic issue involves acceptable scientific approaches to the study of complex systems. The perplexing panorama of accumulated immunologic facts clearly illustrates that immunology is an enormously complicated system. So complicated, indeed, that it has defied human understanding for decades. Further, it is continuously made more complicated by the relentless analytic efforts of investigators. This is the catch-22 of modern science. To grapple with this complexity, working immunologists routinely oversimplify immunology. They envisage simple mechanisms and simple paradigms, upon which they hang complexity as decoration. The fact is, humans are ill-equipped to deal with raw, unbuffered complexity. To study complex systems, they either carve them up into small pieces, or simplify them beyond recognition. In either case, they first destroy a complex system's complexity in an effort to study its function. The operating principle is that complexity is merely a cosmetic feature which contributes little of importance to the function of complex systems. But what if complexity is an integral, formative element of the immune system? It is possible that the immune system, for example, has not only learned to cope with complexity, but has learned to employ complexity to its advantage. While intuition suggests that complex systems are inherently problematic and prone to error, biologic systems have not only maintained their complexity, but built upon it. Perhaps complexity provides more advantages than liabilities to the immune system.

What has allowed the conceptual approach that permits us to dismiss the complexity of the immune system? What is responsible for the unquestioning belief that the whole of a system is no more than the sum of its parts? What if one of the triumphs of life is that the whole of the organism is far more richly endowed than any collection of its isolated parts would ever suggest? What if complexity is the key that has made the immune system continually victorious over most of its adversaries?

Why has the prolonged uncovering of massive amounts of information regarding the basic constituents of the immune system been unable to provide us with a clear understanding of how the immune system operates to maintain the integrity of self?

If complexity is integral to immune function, than the continued disregard of complexity constitutes an underappreciated, debilitating barrier to our understanding of immunity. Then we should consider the possibility that the time-honored, analytic approach, while informative, is insufficient to

solve this problem. If so, we will have to find new and unusual scientific approaches that directly address the complexity of the immune system. This is the spirit of this forum. As might be suspected, addressing this complexity is a huge challenge. We must ask how the immune system can be studied in the context of its own complexity? In the first article, Stephanie Forrest demonstrates how computers can provide a tool for the study of immunity as a complex adaptive system. We must ask how complexity helps to shape the design of the immune system. In the second article, I suggest that complexity permitted the immune system to self-organize in ways that have been best exploited by social insects, ways that contrast dramatically with human organizational schemes. We must ask what capabilities complexity contributes to the immune system. In the third article, Irun Cohen suggests that complexity has allowed a new, sophisticated feature of lymphocyte function to emerge: cognition of the ever-changing interface between an individual and the external environment. In the fourth article, Lee Segal suggests that this cognitive function is also used to guide the successful development of immune responses. Finally, we must ask what disadvantages the immune system endures for embracing complexity. What has been the price?

In the first article, Stephanie Forrest describes how she used principles of the natural immune system to engineer an artificial immune system that protects computer networks from attacks by "dangerous non-self". Her immune program has features that correspond to thymic function, T cell receptors, MHC molecules, peptide processing and presentation, costimulation and immune memory. Indeed, she was surprised at "how many features of the natural immune system... [she was]... forced to incorporate in order to achieve acceptable performance of the artificial immune system." Dr. Forrest is among the first to abstract and field-test the working principles of the immune system. Further, she has demonstrated that various aspects of the immune system can be assembled in simulation for study with a computer. Such studies would preserve the complexity of the immune system while it is studied, rather than disregard it. In an effort to develop protection for computers, she has pioneered a radical, new approach to studies of the immune system. Inadvertently, she may have developed a better understanding of the immune system than have most immunologists.

In the second article, I suggest that our understanding of immune system function is hampered by our humanization of immunology. We have modeled our concepts of immune function around human organizational patterns. However, the organizational schemes of social insects, which are quite different than those of humans, may better reflect how leukocytes relate to one another and accomplish tasks. If so, there are fundamental errors in our basic concepts regarding immune system organization and function. For example, social insects have no leaders, and neither may the immune system. The article provides arguments for why we should abandon our humanized organizational schemes of immunity and consider the organizational schemes of social insects as models for leukocyte behavior.

IMMUNOLOGIC COMPLEXITY

- How can the immune system be studied within the context of its own complexity?
- How does complexity help to shape the design of the immune system?
- What new capabilities does complexity contribute to the immune system?
- What disadvantages must the immune system endure for embracing complexity?

In the third article of this forum, Irun Cohen develops the concept that the immune system's job is not just to protect the body from infection, but to "keep the body in working order despite daily wear and tear." This is an extension of concepts presented in his recent book, *Tending Adam's Garden*, in which he describes the immune system as a cognitive element of the body. Cohen now argues that the immune system is constantly busy keeping up with the changing world and its impact on the body. He wonders how the immune system knows when it is doing the right thing. He suggests that this requires a dialogue between the immune system and the body, and he introduces the fascinating concept of "immune linguistics," a process by which various immune elements are used by the cognitive immune system to facilitates this dialogue.

This theme is explored further by Lee Segel in the fourth article of the forum. Dr. Segel examines a question that has long interested him: "How does the immune system know that it is doing a good job?" As a mathemetician with an interest in mathematical modeling of the immune system, he shows how different immune elements could provide feedback information on the effectiveness of a developing immune response. With this article, he continues to define the working principles of immunoinformatics, a concept that was first introduced in a Santa Fe Institute symposium proceedings edited by Segal and Cohen, called "Design Principles for the Immune System and other Distributed Autonomous Systems." Immunoinformatics is the study of how the immune system generates, posts, processes and stores information. In this forum, both Segal and Cohen describe why information is needed for effective immune function, and thus provide the conceptual framework for immunoinformatics.

Overall, the concepts provided in this forum step well outside of the traditional bounds of immunology, and describe some new perspectives on some old problems. These newer approaches, when used in conjunction with the continued analysis of immune system components, may help to narrow the widening gulf between the generation of new immune information and the development of a better understanding of immune system function.

Contributors

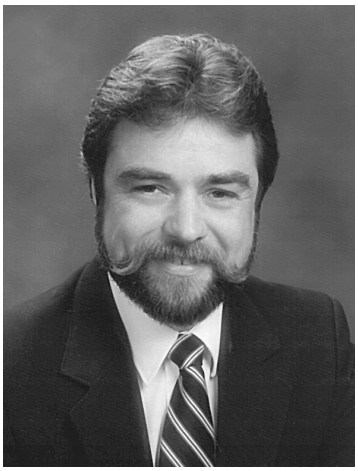

Charles G. Orosz, PhD, is Professor in the Departments of Surgery, Pathology and Molecular Virology/Immunology/Medical Genetics. He is the Director of the Clinical Histocompatibility Laboratory, and the Director of the Transplant Sciences Program. His research interests include: a) immunohistologic, immunologic and molecular biologic analyses of graft-reactive lymphocyte redistribution/function in murine cardiac, kidney or aortic allograft recipients; b) mechanisms of tolerance induction in murine allograft recipients; c) mechanisms controlling the development of transplant vascular sclerosis and other features of chronic rejection in murine cardiac or aortic allografts; d) methods of clinical immune monitoring in transplant patients; and e) application of "Complexity Theory" to the study of immunologic functions.

Related References

1. Orosz CG. Networked alloimmunity: A brief reexamination of some basic concepts in transplantation immunobiology. **J Heart Lung Transplant 1996; 15:1063-1068**.
2. Orosz CG, Sedmak DD. Concerns regarding the current paradigm for chronic allograft rejection. The chronic graft rejection paradigm. **Transpl Immunol 1997; 5:169-172**.
3. Orosz CG. If you think that you know the answers, you probably don't appreciate the nature of the problem. ASTS Hume Lecture, May 14, 1998. **Graft 1998; 1:175-180**.
4. Orosz CG. Complexity and Transplantation. In: Stuart FP, Abecassis M, Kaufman DB eds. Organ Transplantation. **Georgetown, TX: Landes Bioscience, 2000; 418-432**.
5. Orosz, C.G. An introduction to immuno-ecology and immuno-informatics. **Oxford Univ Press, 2000; 119-146.**

Stephanie Forrest, PhD, is a Professor of Computer Science at the University of New Mexico in Albuquerque and a member of the External Faculty at the Santa Fe Institute. Her research interests cover a wide spectrum of parallel adaptive systems, including genetic algorithms, immune systems, computer security, and emergent computation.

Related References

1. Forrest S, Genetic algorithms: principles of adaptation applied to computation. **Science 1993; 261:872-878**.
2. Forrest S, Hofmeyr S, and Somayaji A. Computer immunology. **Communications of the ACM 1997; 40(10):88-96**.
3. Smith DJ, Forrest S, Ackley DH, et al. Variable efficacy of repeated annual influenza vaccination. **Proc of the National Academy of Sciences (PNAS) 1999; 96:14001-14006**.
4. Oprea M and Forrest S. How the immune system generates diversity: pathogen space coverage with random and evolved antibody libraries. **Proc of the Genetic and Evolutionary Computation Conference (GECCO), Morgan-Kaufmann 1999; 1651-1656**.
5. Maley CC and Forrest S. Exploring the relationship between neutral and selective mutations in cancer. **Artificial Life (in press)**.

Contributors

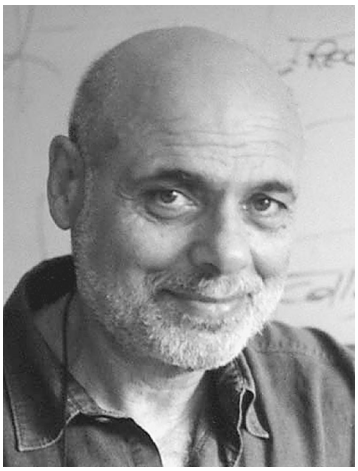

Irun R. Cohen, MD, is the Mauerberger Professor of Immunology, and the Director of the Robert Koch-Minerva Center for Research in Autoimmune Diseases in the Department of Immunology at The Weizmann Institute of Science in Rehovot, Israel. His research interests focus on autoimmunity, T cells, hsp60, type-1 diabetes mellitus, specific immune modulation and complex systems.

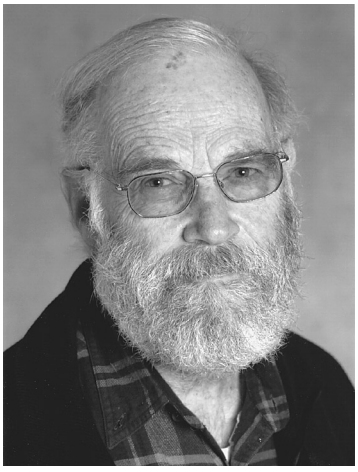

Lee Segel, PhD, is an Emeritus Professor in the Department of Computer Science and Applied Mathematics of the Weizmann Institute and an External Faculty Member of the Santa Fe Institute. He is editor-in-chief of the Bulletin of Mathematical Biology. His research centers on a variety of problems in cellular neurobiology, hematapoetics and immunology.

Related References

1. Parnas H, Segel LA, Dudel J, et al. Autoreceptors, membrane potential and the regulation of transmitter release. **Trends in Neurosci 2000; 23:60-68**.
2. Lev Bar-Or R, Maya R, Segel LA, et al. Generation of oscillations by the p53-Mdm2 feedback loop: a theoretical and experimental study. **Proc Natl Acad Sci USA 2000; 97(21):11250–11255.**

A

Engineering an Immune System[1]

Stephanie Forrest and Steven Hofmeyr

Introduction

The immune system is highly complex, and many researchers believe its complexity prevents us from ever knowing definitively what purpose each individual component serves. As they point out, the forces of natural selection don't guarantee perfect or minimal solutions, simply ones that work well enough to increase survivability. We know that many simpler immune systems exist in lower animals, and these systems seem to work just fine. Thus, so the argument goes, the human immune system is not optimized, and ascribing specific purpose to the various components and mechanisms is foolhardy.

Even worse, there is still debate about what role the immune system plays for the body. The traditional view that the immune system's primary job is to distinguish self from nonself[1] has been challenged in many places.[9] Likewise, the view of the immune system as a "danger detector" is controversial.[11] And, more recently, some have advocated viewing the immune system as a means of preserving homeostasis for the body.[2] These three views are not mutually exclusive, but they emphasize different aspects of immune-system function and suggest not only that the system itself is complex, but so is its functional role in the body. By contrast, there is little debate that the primary function of the heart is to circulate blood. In spite of these controversies, it seems reasonable to assume that one of the immune system's main roles in the body is "protective."

(continued on p. 370)

B

Immunity as a Swarm Function

Charles G. Orosz

The Humanization of Immunity

In characterizing the function of the immune system, immunologists have often borrowed the human organizational schemes employed by their armies, businesses or orchestras. Indeed, cooperative activity among large numbers of humans seems to require a decision-making leader, a pyramid of lieutenants who implement the decisions, and hordes of underlings who do the actual work. Apparently, Immunologists presume that lymphocytes have the same organizational requirements. This is but one of many humanized presumptions about lymphocyte function. For example, an individual human life is inherently important and worth protecting at all cost. The maximal effectiveness of a human undertaking is associated not only with outcome but also with efficiency and cost. Waste is considered an unacceptable use of effort and resources. In the same vein, immunologists presume that individual leukocytes are important and do not function or die uselessly, and that immune resources are not squandered on unnecessary tasks. These presumptions are made, in part, because immunologists are human, and know of few, if any, reasonable alternatives. However, several insect organizational schemes have now become fairly well-delineated. These are highly effective, yet quite unlike the organizational schemes used by humans. Thus, they provide intriguing alternative organizational schemes for leukocytes that deserve some consideration.

(continued on p. 373)

C

Immunity, Set Points, Reactive Systems and Allograft Rejection

Irun R. Cohen

Set Points

Most physiological systems are regulated with the help of set points. Set points hatch homeostasis. The body's thermoregulatory system, for example, aims at a core temperature of 37° C; above that temperature, a center in the hypothalamus activates peripheral vasodilatation, shunts blood to the skin, and turns on the perspiration until the resulting loss of heat returns the body temperature to 37° C. A core temperature below the 37° C set point triggers peripheral vasoconstriction and, if that does not suffice to raise the temperature, shivering. IL-6 and IL-1 make you feverish by stimulating the hypothalamus to raise the thermal set point above 37° C. The kidneys and lungs use set points to regulate body pH; blood volume and osmolarity are regulated by the kidneys and the hypothalamus using set points; cardiac output has a set point that varies with metabolism; alpha and beta cells in the islets regulate the concentration of blood glucose by set point and so on with other endocrine functions. Some set-point systems are not yet well characterized; the amount of body fat and the numbers and types of bacteria residing in skin, mouth and bowel are examples. But we do not have to know the mechanism by which a set point works; we can recognize a set point by the way it works

The set point of a system is the system's "goal" (in a metaphorical sense; real goals are the products of minds). The set point, in functional terms, defines the circumstances in which the system-in-charge need make

(continued on p. 376)

D A continued

How Does the Immune System See to It That It Is Doing a Good Job?

Lee A. Segel

In his essay[1] Irun Cohen asks "How does the immune system know when it is making the right response, when it is doing the right thing?" My own answer follows.

Feedback Toward Improving Immune Response

My first point is that there is no single characteristic of "doing the right thing." Aspects of making an appropriate response to the ever shifting challenges with which the immune system is faced are numerous, somewhat ill-defined, overlapping and often even contradictory. For example, in an appropriate response dangerous pathogens are killed and harm to self is avoided. These aspects are in partial conflict, because inflammation harms self (for example to kill intracellular pathogens the immune system kills cells of the host). As Cohen stresses, other aspects of the immune system involve maintenance tasks ranging from assisting in wound healing to helping control tissue regeneration.

Cohen[1] holds that the immune system "aim(s) at representing a part of the world" and it uses the resulting internal images to learn and adjust various reaction programs, each of which "amounts to a functional image of a stimulus that elicits a response." True, I think, but not the whole truth. A fuller truth takes into account more explicitly the way the immune system copes with uncertainty by a generalized feedback. This feedback aims not to optimize immune response, but rather to improve it in real time with the aid of suitable sensory input.

(continued on p. 379)

If we don't know exactly what the immune system does, and if we believe that many of its components are redundant or unnecessary, then how should we go about understanding the immune system? The point of this article is to suggest that an engineering perspective might be helpful. That is, if we set out to engineer a protective system that operates successfully in an environment with some of the same constraints as those faced by the immune system, what components would we need and what would we need them for? To what extent would these components resemble the natural immune system?

Building a Computer Immune System

A suitable domain for such an exercise is that of computer security. A computer security system should protect a computer or network of computers from unauthorized intruders, which is similar in functionality to the immune system protecting the body from invasion by foreign pathogen. Further, a computer security system should protect against insider attacks, malfunctioning software (analogous to misbehaving cells) and other internal errors, maintaining the computer within normal operating tolerances. Because of the compelling similarity between the computer security problem and the problem of protecting a body against damage from internally and externally generated threats, we designed an artificial immune system to protect computer networks based on immunological principles, algorithms and architecture.[7,8]

Defining Self

The security of computer systems depends on such activities as detecting unauthorized use of computer facilities, maintaining the integrity of data files, responding to "denial-of-service" attacks and detecting and eliminating computer viruses. We view these protection problems as instances of the more general problem of distinguishing self (legitimate users, uncorrupted data and so on) from dangerous nonself (unauthorized users, viruses and other malicious agents). Just as the natural immune system evolved to monitor certain observables in the body (e.g., peptides or heat-shock proteins), so must an artificial immune system be designed to monitor certain aspects of a computer.

We briefly describe an intrusion-detection system built to protect networked computers, although we have studied other problems, including computer virus detection[6] and host-based intrusion detection.[5] Our intrusion-detection system treats self as being synonymous with normal behavior of a local area network (LAN). This dynamic view of self is quite different from that taken by typical anti-virus software, which looks for changes in files stored on hard disks. The distinction is roughly analogous to that between gene products and genes themselves. For example, if the natural immune system had evolved to inspect directly the genomes of all cells for irregularities, we would have a system more closely analogous to anti-viral software. Instead, the immune system typically monitors gene products.

Our equivalent of an organism is a LAN of computers. TCP/IP is the most common connected communication protocol used on the internet, and the behavior of our model organism can be characterized by its TCP/IP connections, or datapath triples.[10] A datapath triple is comprised of a source address, destination address and communicating program, and this information completely specifies a network connection. Our equivalent of a "peptide" in this environment is a binary string representing the datapath triple. All normally observed and acceptable connections, both those within the LAN and those connecting the outside world to the LAN, form the set of self patterns, and all others (potentially an enormous number), form the set of nonself patterns.

Architecture

Network traffic is monitored by a set of detectors on each computer in the LAN.

A continued

Each detector consists of a binary string (analogous to a receptor on a lymphocyte), and a detection event is a partial match between a detector string and a datapath string (analogous to the binding between a receptor and an epitope). The partial matching is implemented by a threshold-based rule, for example, two strings match if they have more than a given number of bits in common. Partial-matching can be thought of analogously to the affinity of a receptor to a ligand.

Detectors have a finite lifetime and follow a lifecycle reminiscent of the lifecycle of immune system cells, such as T cells and B-cells. Initially, a detector is created with a randomly generated receptor string and remains immature for a certain period of time. During this maturation period, the detector is compared to all occurring datapath triples (network connections), and if the detector matches any triples, it "dies" and is replaced by a detector with a new, randomly generated receptor string. This is analogous to negative selection of thymocytes in the thymus. If the detector survives the maturation period without matching anything, it becomes mature, and future matches raise an alarm, indicating that a potentially dangerous network connection pattern was detected. A mature detector has a finite lifespan, after which it is replaced by a new, immature detector.

Negative selection is an important feature of our system, because it allows for distributed detection. Once censored by the negative-selection process, each detector can function independently of other detectors, that is, without communication between detectors or coordination of multiple detection events. This is because each detector covers part of nonself. Thus, a set of detectors can be split up over multiple sites, which will reduce the coverage at any given site but provide good system-wide coverage. To achieve similar coverage using detectors which match against self would be computationally inefficient.

Lessons Learned

We have described the basis for our artificial immune system. In experiments, this simple start was not sufficient to create an effective security system. We had to add several more immunological mechanisms to address limitations in the basic model. Because these mechanisms were added to solve specific problems, we believe that we can offer some novel perspectives on why these mechanisms might exist in the immune system.

Tolerance and the reduction of autoimmunity. In our negative-selection algorithm, we assumed that detectors would be exposed to a comprehensive sample of self during their maturation period. This was not always the case, and the problem is exacerbated by legitimate changes to self. In the network, legitimate changes to the profile of self patterns might occur when a computer is added or dropped from the network, when a new user is added to the network or when new software is installed. These new self patterns caused unacceptably high false positive rates analogous to autoimmune reactions.

In experiments, this simple start was not sufficient to create an effective security system. We had to add several more immunologic mechanisms to address limitations of the basic model.

To reduce the resultant false positives, we implemented a mechanism similar in spirit to the costimulation that a B-cell must receive from a T-helper cell. In our case, the second costimulatory signal was provided by a human observer. If a mature detector does not receive this signal within a given period after matching (typically 24 hours), it will die. Consequently, detectors responsible for false positives are automatically eliminated, whereas a human intervenes to confirm true positives. This allows the system to adapt to incomplete or evolving definitions of self. It also allows for shorter maturation periods in the negative-selection phase, and hence a higher ratio of mature to immature detectors. Costimulation in the natural immune system presumably has similar benefits, as well as protecting against potentially inappropriate somatic mutations.

The form of tolerization we have described is similar to peripheral tolerance in the immune system, in that mature detectors are still subject to tolerization through costimulation. This form of tolerization is essential when the description of self is gradually changing, or when it is impossible to accumulate a complete description of self in a single location. However, it can be very inefficient. One way to address this inefficiency is to accumulate as many self strings as possible in a single safe location and store them, so that new detectors can be compared against all strings in the store. In this way, detectors can be generated that are more likely free from nonself contamination and they can be generated more efficiently, reducing the need for lengthy maturation periods. This is the equivalent of central tolerization in the thymus, and illustrates that the ideal system may be one in which as much tolerization as possible is carried out centrally, with peripheral tolerization being used to address the issues of perpetual novelty, incomplete descriptions of self and somatic mutation. There have been debates in immunology about the relative roles of tolerization in the thymus and peripheral tolerization; our experiences suggest that both are essential, because they play complementary roles.

Finite lifetimes and long-lived memory cells. The perfect detection system would have a single detector for every nonself

A continued

string, and the detector would match only that nonself string and no other. However, such perfect detection would require as many detectors as there are nonself strings, an infeasible number. The first mechanism for overcoming these resource limits is generalization: because of partial matching, each detector can detect a subset of nonself strings. As detectors become more general, a single detector can match a larger subset of nonself, and so fewer detectors are needed. However, as the generality increases, so the ability to make precise discriminations decreases, diminishing the detection abilities of the system. Further, as the probability of an immature detector matching self increases, it takes longer to generate mature detectors that are tolerized to self. Hence, the generality of detectors is limited.

It is likely that a detector set that is limited both in generality and numbers will fail to detect at least some important nonself patterns. To combat this problem, we introduced a second mechanism, dynamic coverage: the detectors in the detector sets are continually changing. This dynamic coverage is a consequence of randomly generating detectors coupled with finite lifetimes. Dynamic coverage ensures that an attacker cannot repeatedly exploit the same gaps in coverage. There is a trade-off here, however. As the lifespan of an individual detector decreases, the potential for exploiting a hole in the coverage decreases, but a detector spends relatively more of its life in the maturation phase, where it is not contributing to detection at all. We expect the immune system to be subject to a similar trade-off, one which will govern the optimal lifespan of lymphocytes and other cells.

However, not all of our detectors have a finite lifespan. Those detectors that have detected anomalies and received human-mediated confirmation enter a competition where the best-matching detectors become memory detectors. Memory detectors are analogous to long-lived immune memory cells, in that they have much extended lifespans and lower thresholds of activation. Memory detectors greatly enhance detection of previously seen attacks by automatically extracting and encoding signatures of attacks. The efficacy of this mechanism reflects the efficacy of secondary responses in the immune system.

MHC and diversity. Generalization is an important tool in a resource-limited environment: if each detector can match a subset of nonself patterns, fewer detectors are needed. However, generalization introduces potential discrimination errors in the form of holes: a hole is a nonself string for which no valid detectors can be generated.[3,4] That is, any possible detector that could match patterns in the hole would also match some patterns in self. As the generality of detectors increases and the specificity decreases, so the potential for holes also increases. Holes are problematic. Once discovered, they can be continuously exploited by an attacker, because holes are a consequence of the structure of the self set and cannot be overcome by dynamic coverage.

We were surprised at how many features of the natural immune system we were forced to incorporate in order to achieve acceptable performance of our artificial system.

A solution that proved to be effective at reducing the overall number of holes is multi-representation—different representations are used for different detectors. One way of achieving this is for each detector to have a randomly generated permutation rule, according to which all datapath triples are permuted before being matched against the detector. This effectively changes the structure of the self set for each detector, with the result that different detectors will be subject to different holes. Consequently, where one detector fails to detect a nonself triple, another may succeed. Multi-representation was particularly effective at reducing the number of holes when the nonself patterns were similar to self patterns.

Similar to our artificial system, the immune system also faces problems of limited resources, and appears to use both generalization and dynamic coverage. Generalization is a consequence of the fact that a monoclonal lymphocyte can bind to a set of structurally similar peptides, which is analogous to partial matching. It is not unreasonable to assume that this generalized detection also results in holes, and if so, pathogens will evolve away from detection towards the holes. We speculate that molecules of the major histocompatibility complex (MHC) implement a form of multi-representation. Each different type of MHC can be regarded as a different way of representing a protein (depending on which peptides it presents); in effect, the immune system uses multiple representations of proteins. Hence, varying the MHC varies the holes that exist. This is illustrated by the existence of diseases, such as leprosy, that are strongly affected by MHC types. This perspective on MHC can give us insights into the evolution of MHC.

Conclusion

In summary, we were surprised at how many features of the natural immune system we were forced to incorporate in order to achieve acceptable performance of our artificial immune system. Studies such as these can help shed light on the question of what role different components and mechanisms play in the natural immune system, and they can provide a partial answer to those who argue against teleological explanations of the immune system. In the long run, understanding what role different components play and why they evolved will

A continued

B continued

help us design more effective and robust interventions and therapies.

Acknowledgments. The authors gratefully acknowledge the support of the National Science Foundation (grants IRI-9711199, CDA-9503064, and ANIR-9986555), the Office of Naval Research (grant N00014-99-1-0417), Defense Advanced Projects Agency (grant AGR F30602-00-2-0584) and the Intel Corporation. Many of the ideas described in this paper were developed during many years of collaboration with Alan Perelson and Lee Segel.

References

1. Burnet, FF. The Production of Antibodies. 2nd edition. **London: Macmillan, 1949**.
2. Cohen IR. Discrimination and dialogue in the immune system. **Immunology 2000; 12:215-219**.
3. D'haeseleer P, Forrest S, Helman P. An immunological approach to change detection: Algorithms, analysis and implications. In **Proceedings of the 1996 IEEE Symposium on Computer Security and Privacy. IEEE Press: 1996**.
4. D'haeseleer P. An immunological approach to change detection: Theoretical results. In **Proceedings of the 9th IEEE Computer Security Foundations Workshop. IEEE Computer Society Press: 1996**.
5. Hofmeyr S, Somayaji A, Forrest S. Intusion detection using sequences of system calls. **Journal of Computer Security, 1998; 6:151-180**.
6. Forrest S, Perelson AS, Allen L, et al. Self-nonself discrimination in a computer. In **Proceedings of the 1994 IEEE Symposium on Research in Security and Privacy, Los Alamitos, CA. IEEE Computer Society Press: 1994**.
7. Hofmeyr SA, Forrest S. Architecture for an artificial immune system. **Evolutionary Computation Journal 2000; 8(4):443-473**.
8. Hofmeyr SA. An immunological model of distributed detection and its application to computer security. **PhD thesis, University of New Mexico, Albuquerque, NM: 1999**.
9. Matzinger P. Tolerance, danger, and the extended family. **Ann Revs Immun 1994; 12:991-1045**.
10. Mukherjee B, Heberlein LT, Levitt, KN. Network intrusion detection. **IEEE Network 1994:26-41**.
11. Vance RE. Cutting edge commentary: A Copernican revolution? Doubts about the danger theory. **J Immunol 2000; 165: 1725-1728**.

The Dehumanization of Immunity

Social insects have societies that are careless about the needs or rights of individuals. Indeed, individuals are considered to have no particular importance and to be expendable. As far as the insect society is concerned, each individual is heartless, brainless, and relatively worthless. There is no recognized individuality and no individual thought or creativity. Interestingly, there is also no overt leadership. Each individual operates independently, guided by a set of internal rules and a set of signals from the local environment. Thus, individuals are slaves to the situation, and their behavior develops in the absence of free will or choice. In such situations, the same task may be redundantly performed by numerous individuals. Waste and error are acceptable in insect societies.

Despite this, large numbers of individuals within insect societies can perform highly complex tasks that benefit the colony, such as foraging, nest building and brood care. For ants, bees and termites, these colony-oriented cooperative undertakings are so sophisticated that they amaze humans, who have considered them to be manifestations of a hive mind[1] or swarm intelligence. Humans believe that intelligence is required for complex or creative undertakings, so social insects must harbor some form of intelligence. In his recent book,[2] Eric Bonabeau defines swarm intelligence as "the emergent collective intelligence of groups of simple agents". Bonabeau describes how insect swarm functions have been mathematically modeled by several individuals, and how these models can generate results that are remarkably similar to the events that occur in nature. The presumption in these models is that a social insect colony is "a decentralized problem solving system, comprised of many relatively simple interacting entities". One important characteristic of these systems is self-organization, defined by Bonabeau as "a set of dynamical mechanisms whereby structures appear at the global level of a system from interactions among its lower level components". The rules specifying these interactions among the system's constituent units are executed on the basis of purely local information, without reference to the global pattern, which is an emergent property of the system rather than a property imposed upon the system by an external ordering influence". In other words, major projects, like nest building, are undertaken without a blueprint or a project foreman. The labor force is a large set of pre-programmed biologic robots that work inefficiently but effectively to perform small tasks as they are permitted by the evolution of the project. What emerges from this reflexive mass activity is an unique structure with defined functional features that is remarkably well integrated into the local environment.

It is important to remember that even human leukocytes are inhuman, and it seems reasonable to suggest that the immune system may construct its responses to a given problem, such as resistance to invasion by an unidentified agent, in the same manner as social insects. Through the process of inflammation, the immune system localizes temporary clusters of leukocytes to sites of possible invasion, even while it transports peptide samples from the inflammatory site to permanent leukocyte cluster sites, that is, the lymph nodes. Indeed, leukocytes always operate en masse, as leukocyte swarms. According to the social insect model, each leukocyte in such swarms would operate independently, based on its perception of signals in its local environment, including those provided by other swarm members. The fact that individual leukocytes constitute the functional unit of immunology is well-appreciated by immunologists. However, leukocytes must swarm to function, and these swarms always contain several interacting subtypes of leukocytes (T cells, macrophages, NK cells, etc.). Thus, it appears that there are at least two functional units of immunity. The second unit is the self-organized leukocyte population, which

B continued

might be called a leukoswarm, a leukoflock or the leukocluster, and refers to a grouping of leukocytes that has functionally significant properties which cannot be reduced to the level of the individual leukocytes. Examples of this would be the T cell/antigen presenting cell (APC) cluster, which produces cytokines, or the T cell/ B cell/APC cluster, which produces antibodies. The existence of these leukocyte clusters is well appreciated by immunologists, but their role as a second, higher order functional unit of immunity is not. This is one contribution of complexity to the design of the immune system.

Despite the emergent functions of leukocyte swarms, every individual within the swarm has an internally (genetically) defined set of behavior patterns, causing them to respond in a reproducible way to a given set of environmental stimuli. It is interesting to note that this is the basic principle behind in vitro studies, that is, when the experimental conditions are accurately reproduced, the pattern of leukocyte behavior observed under those conditions will also be reproduced. The social insect model would further suggest that each immune response would evolve to fit the current, local conditions. There would be no response blueprint, and no constraint to develop a response pattern like one that would develop either elsewhere in the same individual, or anywhere in another individual. Further, effective responses could occur in the complete absence of immune leadership, and there would be no need to postulate a leadership role to any one component of the immune response, including T cells or macrophages. This is somewhat contrary to the current perspective of modern immunology, which, at best, continues to debate whether macrophages or T cells are in control of the immune response.

One argument against this view is that most immunologists do not believe that leukocytes behave predictably, especially in vivo. Given the principle of self-organization, this unpredictability could result from the fact that activated leukocytes modify their environment (express/secrete bioactive mediators) as features of their genetically encoded responses to specific stimuli. This is comparable to the secretion of pheromones by insects. As with insects, these environmental modifications cause changes in the behavior of other local leukocytes, which are constantly monitoring the environment for behavior-shaping cues. When they respond to these new cues, they produce bioactive agents which again remodify the environment. Thus, the environment of an immune response is constantly evolving, and leukocyte behavior seems unpredictable because it is constantly changing as an immune response evolves. This suggests that an immune response may not be a consequence of numerous leukocytes instructed by regulatory cytokines to adopt a similar, synchronous behavior, as is currently believed (Th1/Th2 hypothesis). Rather, it may be the emergent consequence of numerous independent interactions among large numbers of leukocytes and their similarly evolving microenvironments throughout sites of inflammation. If so, it would be difficult to appreciate immune responses without studying the influence of leukocytes on the local microenvironment, or the influence of the local microenvironment on the behavior of leukocytes. This adds critical importance to bioactive agents like chemokines and provisional matrix molecules (such as fibronectin), both of which appear to be necessary for effective immune responses.[3,4] Indeed, it suggests that the extracellular matrix, and not a particular leukocyte subset, may be an important coordinating element of an immune response. If so, immunologists may have to learn to read the matrix, as do leukocytes.

An Unique Feature of Immunity

The humanization of leukocyte function diminishes regard for at least one important feature of leukocyte function. Leukocytes often exhibit "selective perception" of signals in their environment, due to their ability to display different patterns of signal receptors. In humans, all senses operate concurrently and continuously, so there is no comparable situation in the experience of individual humans. Leukocytes may respond to a given set of environmental cues not only by producing bioactive mediators, but also by modulating their expression of receptors for specific environmental signals. Thus, there are two variable components to the leukocyte perception system: the panorama of signals present in the environment and the differential distribution of signal receptors among the leukocytes. This "double kaleidoscope" system permits the extraction of useful information from pools of mixed signals, and facilitates the ordered perception of environmental signals. The key here is "useful information", which may differ significantly from cell type to cell type, and even for a single cell at different stages of differentiation. Given that the key element in the system is the individual cell acting in defined ways to defined sets of stimuli, swarms of leukocytes acting in this manner tend to confuse the observer.

In some ways, an immune response is similar to events that would unfold if several different plays were acted out simultaneously by separate casts on the same stage. To avoid confusion among the players, each would have to wear a headphone that connects him only to the other members of his own cast. The risk of a miscue delivered inadvertently by a member of another cast is significantly reduced by the selective reception signals. Chaos would occur if the headphones were removed, or traded among the different casts. Given the selective reception system, each player receives the necessary cues, and responds appropriately by delivering his lines, thus providing the necessary cues to others of his cast. In this manner, each play unfolds appropriately and independently of the others. As the individual plays end, the members of the cast leave without affecting the progression of the oth-

B continued

er plays. Of course, to the audience the event looks and sounds like chaos until the last play is left on stage. It would help the audience immensely if they were provided with a set of headphones.

This metaphor has important implications. What immunologists are lacking are the appropriate sets of headphones, so they remain confused by the many concurrent agendas expressed by the different cells involved in any given immune response. Further, the presence of particular environmental signals is not, of itself, informative to the outside observer. Rather, the value of the signal depends completely on the presence of cells which can recognize and respond to it. Thus, investigators should pay much more attention to the receptors displayed by the individual cell types involved in an immune response, and less attention to the panorama of signals that are present in the environment. This is contrary to the current trend of investigations regarding cytokines, chemokines and growth factors.

Swarm-Like Features of Immunity

The question remains whether immune responses are self-organizing leukocyte swarm functions. According to Bonabeau, self-organization relies on four basic ingredients. The first is positive feedback, which operates to permit the creation of new structures. The second is negative feedback, which operates to stabilize new structures. The third ingredient is amplification of fluctuation (random error, etc), that operates to permit the discovery of new solutions. The fourth ingredient is multiple interactions, which operate to reinforce the other ingredients. Immunology has numerous examples of each of these ingredients, with the possible exception of the third, amplification of fluctuation. This apparent lack may be somewhat artificial, rather than biologic. Most investigators design their experiments to maximize the reproducibility of the outcome, and few have intentionally explored the fluctuations that occur in their experimental systems. Nevertheless, such fluctuations are common and well-appreciated by most investigators.

Bonabeau also suggests that self-organized phenomena exhibit several characterizing properties. One is the creation of "new spatiotemporal structures within an initially homogenous medium". Pathologists can vouch that pro-inflammatory immune responses involve the characteristic remodeling of tissue histology by infiltrating leukocytes. The second characteristic of self-organized systems is multistability, the coexistence of several alternative stable states. "Because structures emerge by amplification of random deviations, any such deviation can be amplified, and the system converges to one among several possible stable states, depending on initial conditions." This may be the fundamental principle behind the experimental data that has led to the Th1/Th2 hypothesis of T lymphocyte function. The third characteristic is the "existence of bifurcations (at which the behavior of the system changes dramatically) when some parameters are varied." Again, there are instances of this in immunology. For example, when low to moderate levels of lipopolysaccharides (LPS) are detected by the immune system, they effectively potentiate the immune response[5] and help to re-establish physiologic homeostasis, but when the levels of LPS become very high, they have the opposite effect, typified by the onset of multi-organ failure.[6]

In general, it appears that experimental observations would support the hypothesis that leukocytes employ inhuman, social insect-like organizational schemes to accomplish their tasks. Such schemes are characterized by self-organization and the absence of overt leadership. Indeed, clustered leukocytes may exhibit swarm intelligence. This would endow them with a set of operating options that differs significantly from those that would be available if human organizational schemes were employed. Thus, it seems that an appreciation of swarm intelligence may promote a better understanding of leukocyte function than is currently available.

Swarm-like features of immunity

- Positive feedback - which operates to permit the creation of new structures
- Negative feedback - which operates to stabilize new structures
- Amplification of fluctuation - which operates to permit the discovery of new solutions
- Multiple interactions - which operate to reinforce the other features

References

1. Kelly K. Out of Control. Reading, MA: **Addison-Wesley Publishing Co, 1994:1-472**.
2. Bonabeau E, Dorigo M, Theraulaz G. Swarm Intelligence. From Natural to Artificial Systems. **New York, NY: Oxford University Press, 1999**.
3. Gutierrez-Ramos JC, Lloyd C, Kapsenberg ML, et al. Non-redundant functional groups of chemokines operate in a coordinate manner during the infllammatory response in the lung. **Immunol Rev 2000; 177:31-42**.
4. Coito AJ, Brown LF, Peters JH, et al. Expression of fibronectin splicing variants in organ transplantation. A differential pattern between rat cardiac allografts and isografts. **Am J Pathol 1997; 150:1757-1772**.
5. McGhee JR, Farrar JJ, Michalek SM, et al. Cellular requirements for lipopolysaccharide adjuvanticity. A role for both T lymphocytes and macrophages for in vitro responses to particulate antigens. J **Exp Med. 1979; 149:793-807**.
6. Mayeux PR. Pathobiology of lipopolysaccharide. **J Toxicol Environ Health 1997; 51:415-435**.

Continued

no further adjustment. The right pH, the right blood volume, the right cardiac output, the right O^2 concentration mean that the system-in-charge of the particular operation need not change its output. A system at its set point may rest; for example, the climate control system in your hotel room mercifully becomes quiet when there is no need to add or remove heat from the room; the air temperature is at the 20° C set by the management. Alternatively, a system at set point may not stop its activity, but only persist in a steady state: the heart beat, for example. Reaching the set point means that all is well; the best of all possible states has been reached. One may view a set point as an attractor of the system[1]. Does the brain operate using a set point? I hope not.

An Immune Set Point

Does the immune system operate using a set point? I don't think classical immunology ever raised the issue explicitly, but according to classical clonal selection theory (CST) thinking,[2] the immune system would have to have a single set point, a point of non-reactivity. The CST would propose that the immune system achieves stable non-reactivity in two different ways during its lifetime. During its ontogeny, the immune system strives to become inactive by having all its (self-)reacting cells commit suicide; after ontogeny, the immune system strives to become inactive by killing or removing, not its cells but rather all the (foreign) antigens its cells can see. The logic of the CST is based on the notion that the immune system acts like a reflex; what it sees, it attacks.[3] Inactivity means, therefore, not seeing. Self-tolerance results from killing the cells that see, and defense against the foreign results from killing that which can be seen. The difference is due to developmental timing. Defense against the foreign is well served by such a simple logic, and the persistence of the CST as the ruling paradigm of immunology owes much to the esthetic simplicity of its viewpoint; the immune system is proposed to operate with a clearly defined goal.

Immune Maintenance

The cognitive paradigm of immunity, in contrast to the CST, takes account of the fact that the immune system is not only a defense system, but also a system charged with maintenance of the body; autoimmunity, according to this view, is a necessity.[1,3-5]

Characteristics of self-organized systems

- Creations of new spatio-temporal structures within an initially homogenous medium
- Multistability - the coexistence of several alternative stable states
- Existence of bifurcations at which the behavior of the system changes dramatically when some parameters are varied

The continuous grind of existence takes its toll continuously; entropy increases automatically, without respite.[1] To keep the body in working order despite daily wear and tear, the immune system is continuously called upon to help heal wounds, regenerate tissues, stimulate new blood vessels, fashion connective tissue, kill infected or aberrant cells, remove waste, and tend to other maintenance jobs. Simplistically, one might say that the immune system looks to the self, sees what needs mending, makes repairs, and kills what is beyond repair. Many maintenance functions are activated and regulated by cytokines and other immune system molecules in an ongoing dialogue between the body tissues and the immune system.[1,6] Rejecting the foreign is viewed by the cognitive paradigm only as an aspect of immune self-maintenance. Maintenance means attending to the self, not ignoring it.[7]

Immune maintenance is carried out by processes we call inflammation; inflammation has been defined as the dynamic processes set into motion by injury that lead to healing.[8] The body is continuously in need of maintenance, and the immune system is continuously busy regulating the expression of inflammation; fortunately, most of the time inflammation is beneficial, subclinical, localized and minimal.

It was once assumed that maintenance was carried out by the innate arm of the immune system; the adaptive T cells and B cells were assumed to be oblivious of the self and responsive to the foreign only.[2] That convenient division of labor is not operative; the adaptive immune system is filled with self-recognizing lymphocytes,[7] and autoimmune T cells are demonstrably involved in body maintenance.[4] Is there a set point for maintenance?

No Immune Set Point

If the immune system were responsible only for protecting the body against pathogens, then one could argue that an immune set point could be freedom-from-pathogens,[9] or avoidance-of-danger;[10] the system could be seen by such reasoning as inactive until activated by a dangerous pathogen. Rejection of the pathogen would allow the system to rest. Rejection of the pathogen could then function as the system's set point; no pathogen, no immune response. However, if the immune system is responsible for body maintenance, then the immune system can never be allowed to rest.[1] Maintenance is a continuous activity. Immune maintenance can stop only when

Continued

there is nothing to maintain, when there is no life left in the body.

The immune system, like the brain, spends a lifetime just keeping up with the changing world. The immune system, like the brain, has to improvise; these systems react to whatever impinges on the body. Constantly reacting systems can have no set point; there is no pre-determined optimum in responding to the unpredictable vagaries of the world.

Immune Self-Organization

Constantly reacting systems like the brain and the immune system don't aim to rest at set points; they adapt. Adaptation means being able to update a response to a world that changes. An immune response is like a key to a particular lock; each immune response amounts to a functional image of the stimulus that elicited the response.[1] Just as a key encodes a functional image of its lock, an effective response encodes a functional image of its stimulus; the stimulus and the response fit each other. The immune system, for example, has to deploy different types of inflammation to heal a broken bone, repair an infarction, effect neuroprotection,[4] cure hepatitis or contain tuberculosis. Each aspect of the response is a functional representation of the challenge.

Self-organization allows a system to adapt, to update itself in the image of the world it must respond to.[1,6] No set point would serve either immunity or thinking. The immune system, like the brain, does not aim at homeostasis; both of these reactive systems aim at representing a part of the world.

One of the key differences between set-point physiological systems and reacting systems like the immune and nervous systems is that the latter two organize themselves through experience with the world; they change with the times because they remember the times; they have a history.[1] Set-point systems, in contrast, only have their genetically determined set points as guides; they do not learn from somatic experience. The immune system responds to whatever infects, is infected, breaks, dies, overgrows, tears or rots; and the system learns and remembers the lesson. The immune system even responds to itself.[7] The processes of learning and memory constitute self-organization.[1,6]

Feedback Regulation

Set-point systems work by striking a balance between positive activation and negative feedback; they need no other control mechanism. The body's thermoregulation is a clear example; the system is activated to produce heat by any temperature below the 37°C set point, and it is activated to dissipate heat at any temperature higher than the set point (positive feedback). Likewise, the production of heat is shut off by a higher temperature and the dissipation of heat is shut off by a lower temperature (negative feedback). The positive and negative feedbacks guarantee attainment of the set point; the set point functions as the metaphoric goal of the set-point system. The system is regulated by the negative and positive feedbacks centered around its set point.[11]

The immune system, like set-point systems, does act in response to both positive and negative feedbacks;[11] activation acts like positive feedback and suppression acts like negative feedback. Inflammation, for example, positively feeds back on itself; inflammation molecules up-regulate the expression of major histocompatibility complex (MHC), cytokines, chemokines and other molecules, which, in turn, up-regulate the immune response to generate more inflammation, and so forth. The inflammatory cascade, like a fire, has the capacity to spread out of control. So inflammatory reactions, similar to the coagulation cascade, have to have negative feedback signals built in. Natural negative feedback is just beginning to be appreciated, but it is clear that inhibitory signals can be generated by the action of pro-inflammatory enzymes on the extra-cellular matrix,[12] and on cytokines such as IL-2.[13] Thus, reactive systems, like set-point systems, do have feedback controls; but the function of feedback differs in the two systems. Set-point systems use feedback exclusively to navigate to their set points; reactive systems do not have set-point "goals" to reach. Reactive systems use feedback to adapt their internal images, to change their patterns of response in accord with the patterns of signals emanating from their worlds of interest.[1]

Updating the Response

But how does the immune system or the brain know when they are making the right response, when they are doing the right thing? The point is that the immune system, like other reactive systems, has no set point to tell it what to aim for. The immune system responds to continuously changing patterns of signals.[1] The responses are local at particular lymph nodes and tissue sites, and they are also global in the blood, lymph, spleen, bone marrow and elsewhere. The immune system modifies inflammation by adjusting its response to dynamic changes in the signals presented by the inflamed tissue. The immune system, in essence, modifies the tissue, and the tissue, in turn, modifies the immune response. Reactive systems don't aim for set points; they aim for dialogue.[1,5,6] The immune system, like the brain, cannot know when it is doing the right thing; the immune system only modifies its response to the body as the body responds to it. The body and the immune system are like a married couple.

Internal Languages

As I have discussed elsewhere, the immune dialogue is a chemical abstraction; the processing and presentation of antigens, the receptors that see the antigens, the innate accessory signals that activate the response phenotype, the mix of cytokines and antibodies and other molecules come together as a string of signals that constitute a chemical language.[6] An antigen may function as the subject of an immune sentence, while

Continued

the innate accessory signals function as predicates that determine what the system will do with the antigen subject.[1,6] Abstract languages are essential to internal images. In fact, abstractions of reality are the substance of internal images. The molecules of immunity are simultaneously the language of immune communication and a functional representation of the stimuli to which the immune system responds. We can illustrate this notion metaphorically; a verbal description of a tree is a communication about a tree and, at the same time, is a abstract representation of a real tree in symbols legible to the brain. Similarly, a processed peptide in the context of an activated antigen presenting cell can be both an immune communication to certain clones of T cells and the immune representation of an infection or an injury. A processed peptide is not an infectious agent or a sick cell; a processed peptide is only a sign, a representation. Systems that construct representations must use abstract languages, be they verbal, electronic, or molecular, for carrying out the construction. The string of chemical signals that the immune system exchanges with the tissues, the immune-body dialogue, is self-adjusting.[1] The rules of this dialogue constitute a system of immune linguistics. The individual cells and molecules that make up the immune system do not know, and cannot know, when each is doing the right thing. The right thing emerges from a collective and dynamic interaction of autonomous cells and molecules. But that complexity is in need of a long discussion (see ref 1).

Allograft Rejection

Writing a piece for publication in *Graft* invites some comment on graft rejection. Why are allografts destroyed so energetically? In former times, immunologists were wont to say that the immune system existed to distinguish between self and non-self,[2] and the allograft was the epitome of non-self inviting destruction. More recently, some immunologists have taken a more sophisticated view and note that that the immune system is not so tuned to rejection of the foreign, but is more bent on fighting danger, irrespective of self or non-self.[10] Allografts, however, pose no threat; why does the immune system attack them? Or to use the metaphor of language, what image does an allograft communicate to the host that commands its rejection?

The molecules of immunity are simultaneously the language of immune communication and a functional representation of the stimuli to which the immune system responds.

It has always intrigued me that, barring pre-formed antibodies, an allograft tends to be more vigorously rejected than is a blatantly foreign xenograft. So it is not the foreigness per se of the allograft that so irritates the host's immune system. Barring a graft-versus-host reaction, there is no threat inherent in an allogeneic MHC; so an allogeneic MHC poses no danger to its host. So what is the compelling message of an allogeneic tissue to host immune cells? I propose that the concept of immune maintenance,[1,4] might be worth considering.

An allogeneic tissue, in contrast to a totally foreign entity, bears a great many signals of the host itself. These graft signals are seen as self by almost all host innate immune receptors and by many antigen receptors (at least those of host B cells). The relatively few differences in MHC-presented peptides constitute signals that call for maintenance; the allograft looks like aberrant self-tissue in need of maintenance. Now the maintenance program, as we discussed above, leads the immune system to repair what can be healed (so allografts undergo, for example, scar formation and angiogenesis) and to destroy what resists healing. Since the allograft cannot resist expressing its allelic genes, the allograft is read by the host immune system as a tissue continuously calling for maintenance. The allograft cannot change its aberrant signaling. The allograft is perceived as chronically sick self. And so the intractable allograft finally gets rejected; it's only a case of extreme maintenance. Or to use the language metaphor, the chemical words of the allograft communicate signals in the language of the host immune system; the allograft's words are familiar enough; it's their intractable accent that irritates unto death.

References

1. Cohen IR. Tending Adam's Garden: Evolving the Cognitive Immune Self. **San Diego: Academic Press, 2000**.
2. Burnet FM. Self and Not-Self. **Cambridge: Cambridge University Press, 1969**.
3. Cohen IR. The cognitive principle challenges clonal selection. **Immunol Today 1992; 13:441-444**.
4. Schwartz M, Cohen IR. Autoimmunity can benefit self-maintenance. **Immunol Today 2000; 21:265-268**.
5. Cohen IR. Discrimination and dialogue in the immune system. **Semin in Immunol 2000; 12: 215-219; 269-271; 321-323**.
6. Atlan H, Cohen IR Immune information, self-organization and meaning. **Int Immunol 1998; 10:711-717**.
7. Cohen IR. The cognitive paradigm and the immunological homunculus. **Immunol Today 1992; 13:490-494**.
8. Florey HW. General Pathology. **London: Lloyd-Luke Ltd, 1970:22**.
9. Medzhitov R, Janeway CA. How does the immune system distinguish self from nonself? **Semin Immunol 2000; 12:185-188**.
10. Matzinger, P. Tolerance, danger, and the extended family. **Ann Rev Immunol 1994; 12:991**.
11. Segel LA, Bar-Or RL. On the role of feedback in promoting conflicting goals of the adaptive immune system. **J Immunol 1999; 163:1342-1349**.
12. Lider O, Cahalon L, Gilat D, et al. A disaccharide that inhibits tumor necrosis factor-a is formed from the extracellular matrix by the enzyme heparanase. **Proc Natl Acad Sci USA 1995; 11:5037-5041**.
13. Ariel A, Yavin EJ, Hershkoviz R, et al. IL-2 induces T cell adherence to extracellular matrix: Inhibition of adherence and migration by IL-2 peptide generated by leukocyte elastase. **J Immunol 1998; 161:2465-2472**.

D continued

Dispersed Feedback Copes with the Unpredictable

Much of immune challenge and response is effectively unpredictable. Consider for example the matter of the generation of dominant epitopes by antigen presentation. This involves numerous fluctuating molecular level processes such as interaction with chaperones to protect the major histocompatibility complex (MHC), directing intracellular traffic, cleavage by restriction enzymes, selection of peptides to bind to the MHC and removal of degradation products from the MHC groove. The result is that the "choice of the favored determinant is essentially aleatory... and requires empirical test...".[18] Other unpredictable matters include the chance gene rearrangements that yield the operational spectrum, at any given time, of antibodies and T cell receptors, the random choice of host genotype, the shifting spectrum of pathogen mutations, the function-altering polymorphisms in cytokine genes[12] and the impingement of unforseeable environmental fluctuations. Even macroscopic organ systems might well provide a degree of uncertainty, for the general possibility of chaos in complex dynamics is now well-accepted. More concretely, evidence has been obtained that features of fractal-like dynamics underly aspects of the complex variability in the healthy human heartbeat[7] and also in human gait.[5] Further, certain forms of life-threatening pathology, such as heart failure, are associated with a breakdown of these multiscale fluctuations and the emergence of excessively periodic and highly predictable behavior.[3]

Engineers know how to deal with uncertainty, by **feedback**. The rotating-weight governor on steam engines is an old example of monitoring performance and reacting to the monitoring to reduce deviation from some set-point; if the engine rotates too fast then the swinging out of the weights partially opens a valve and thus decreases the steam pressure. Too slow rotation is countered analogously. More sophisticated engineering feedback came quite late, in the cybernetic age of the mid-twentieth century.

As a rule biology is way ahead of engineering, and the rule holds for feedback too. Indeed, Cohen[1] mentions several examples of physiological feedback where sophisticated adjustments reduce deviations from evolutionarily established set points.

If evolution can establish **classical feedbacks** toward set points, then the ubiquity in biology of redundancy and variability lead one to expect that evolution can handle systems where set points are ill-defined, numerous and contradictory. I believe that this expectation is indeed confirmed by the presence in biology of complex control systems that employ what one can term **dispersed feedback**. Dispersed feedback uses information from a variety of sensors to improve in some sense the variegated aspects of immune performance. There is no optimization.

Dispersed feedback uses information from a variety of sensors to improve in some sense the variegated aspects of immune performance.

Dispersed Feedback in Metabolism

Chemotaxis provides an example of dispersed feedback.[14] Metabolism provides an even better example. To see this consider just the single instance provided by the glycolytic enzyme PFK, which has at least five regulatory sites. There is up-regulation of PFK activity by AMP and (indirectly) by glucose as well as inhibition by ATP, H+, and citrate. Biochemistry textbooks explain the functional advantages of the multiple regulation.[19]

- Up-regulation of energy is promoted by a combination of a high level of the ATP precursor AMP and a low level of ATP itself.
- Glycolysis supplies precursors such as citrate for reactions that synthesize enzymes, signalling chemicals and structural elements. High levels of citrate signal abundant precursors.
- Via excessive production of lactate, glycolysis can bring about a harmful drop in blood pH (acidosis), hence the down-regulation by H+.

Multiple aspects of an appropriate metabolic response to changing environments are evident: supplying energy; providing enzymes, signalling chemicals and structural elements in suitable quantities and avoiding pH extremes. It is also evident that these aspects overlap and conflict, and that the degree of effort invested in achieving one aspect will effect the amount of resources that can be allocated to another aspect. It seems quite clear that evolution has honed the ramified regulation of the metabolic machinery in such a way as to allow metabolism to respond appropriately to unpredictably shifting demands on the structure and function of the organism. In general, the multiple conflicting aspects of appropriate physiological response reflect the multiple conflicting influences on the complex organism's survival that are exerted by the physical and biological environment.

Immunology: Feedback Exploits Information to Improve on an Initial Broad Spectrum Response

Here is how I believe that dispersed feedback works to improve performance of immune system goals. When the immune system is faced by a challenge, it first replies rather reflexively with a broad spectrum response, for example with a mixture of antibodies, with Th1 *and* Th2. How this response is triggered will not concern us here; suffice it to say that the trigger can be one or a combination of factors such as detecting conserved microbial constituents,[8] exceeding a tunable activation threshold,[4] sensing danger[11] or sensing tissue destruction.[6]

Feedback modifies the immune response in the light of information that the immune system collects on how well it is performing, and on the general physiological state of the organism. Of central relevance here is what Orosz[12] aptly terms immunoinformatics: how

D continued

the immune system generates, posts, processes and stores information. Aspects of immunoinformatics appear in a number of the following paragraphs.

Information Molecules Indicate the Appropriateness of Various Aspects of Immune Response

The primary sources of information for dispersed feedback are molecules that can indicate the appropriateness of different aspects of the immune response of the system. Two categories of such molecules are **kill chemicals** and **harm chemicals**. A kill chemical (K) provides evidence that a pathogen has been destroyed. A harm chemical (H) indicates that damage is being done to the host. But is the host damage being done by the pathogens, in which case the immune response should be elevated? Or is the host damage being done by the immune system, in which case the immune response should be damped? What might be termed the **principle of association** distinguishes between these two alternatives. If a cell's receptors simultaneously sense high levels of the pathogen (P) and of H, then it is likely that there is considerable harm (H_P) due to the pathogens. If H is high and P is low then the inference should be that H_I, the harm due to the immune system, is high while H_P is relatively low. In another use of this idea of "guilt by association" extensive killing of dangerous pathogens (large K_{DP}) can be signalled by the simultaneous sensing of considerable pathogen killing (large K) and considerable pathogen damage (large H_P). In brief, $K_{DP} = KH_P$.

A general H must be generated by host damage. Epitopes of host hsp would be good candidates. A chemical that has been shown to provide evidence of host harm, and indeed downgrades inflammation, is a trisulfated disaccharide fragment from the inflammation-induced cleavage of extracellular matrix by heparanase.[10]

Is there evidence for the presence of K? K can be identified if it fulfills the following requirements.

- It is far more prevalent in pathogens than hosts.
- It is an intracellular molecule (for then its presence indicates that the host has been destroyed).
- It is essential to the pathogen (otherwise pathogen mutation will replace the molecule in question).
- There is evidence that the presence of the molecule modifies the immune response, presumably via a suitable receptor.

Candidates for K include N-formyl peptides, palindromic DNA sequences, endotoxins, and mycolic acid.[15] Note the important difference between mycolic acid, an intracellular constituent of cell walls in gram-negative bacteria, and lipopolysaccharides (LPS), an extracellular constituent of the same cell walls. The information content of ligation of the LPS receptor by LPS is "potential damage is present, from gram-negative bacteria". Ligation of CD1 by mycolic acid should propagate the message "gram-negative bacteria are being destroyed."

KEY CONCEPTS

- Sensors record the status of multiple, overlapping goals.
- Sensor-receptor ligation produces the cytokine in-information network.
- The same information differently drives different cells and different actions of the same cell to improve immune system performance.

How Information Is Used to Improve Performance

The title of this essay asks how the immune system "sees to it" that it is doing a good job. Implicit here is the assertion that the immune system "sees" or, more generally, "senses," which it certainly does. It senses many things and processes the information. But it is not enough to know how good a job the system is doing; for survival the system must see to it that the information it has is used to improve performance. The immune system can employ dispersed feedback to do this in at least two ways. Dispersed feedback can select appropriate effectors among the varied options available to the system and it can utilize information to improve the performance of a given effector.

Let us consider (briefly) the possibility of using feedback to improve effector performance. (See ref. 15 for a discussion of effector selection.) How can the system be driven toward better performance of twin, conflicting, aspects of immune response—the positive aspect of enhancing K_{DP} and the negative aspect of causing harm H_I by the immune system to the self? A simple possibility is that the system can operate according to the following prescription (where *a*, *b*, and *c* are constants):

$$\text{Cellular action} = \frac{aK_{DP}}{1 + bH_I + cK_{DP}} \qquad (1)$$

All other things being equal (H_I fixed), cellular action will be more intense the more it leads to K_{DP}. (Cellular action increases when K_{DP} increases, in just the same saturating way that increasing substrate concentration increases the velocity of a Michaelean reaction.) All other things being equal (K_{DP} fixed), cellular action will be more and more strongly damped the more it damages self (cellular action decreases when H_I increases, in just the same way as increasing inhibitor concentration damps the velocity of a Michaelean reaction).

D continued

Cellular actions include proliferation, motion (both random and chemotactic), signalling and effector functions. Formula 1 gives an example of how these actions can be affected by information on how the system is doing. It is to be expected that a piece of information such as H_1 will not act directly on cellular action, but rather indirectly, via cytokines. (For example binding of the kill chemical CpG to intracellular receptors in macrophage up-regulates the secretion of Th1 cytokines.[20,21] The functional message is translated into a cytokine profile. **A cardinal immunoinformatic principle is that the same information differently affects different actions of different cells**; for example, knowledge that gram-negative bacteria are in the vicinity should enhance the production of complement and down-regulate cytotoxicity. In Formula 1, this would mean that the constants *a*, *b*, and *c* in fact depend on cytokine levels in different ways for different cell types and in different ways for different actions of the same cell type (for example in different ways for inducing B cells respectively to proliferate and to secrete antibody).

The classical view of cytokines is that they form a **command network** (for example, IL-6 commands a switch from IgM to IgG). I believe that it is more profitable to regard cytokines as forming an **informational network** (for example TGF-β, PGE-2 and PAF are secreted when scavenger receptors are ligated,[2] giving the message "apoptosis is occurring"). The focus on cytokines as providing information renders understandable the initially puzzling findings that each immune function is effected by a variety of cytokines (the immune system is simultaneously adjusting its overlapping and contradictory aspects) and that a given cytokine has many functions (the same information differently effects different cells and different cellular functions). Viewing cytokines as an informational network means that there should be more research emphasis on determining what receptor ligations lead to cytokine secretions.

Remarks on Graft Rejection

Here are two connections between the picture I have been trying to develop and ideas of Orosz[12] on graft rejection.

The notion that immune response is initially broad-spectrum and later honed according to performance effectiveness is consistent with findings that there are multiple mechanisms of acute allograft rejection and with the suggestion that immune resources are stockpiled at an inflammatory site, such as a graft site, for possible later use.

The classical view of cytokines is that they form a command network….I believe that it is more profitable to regard cytokines as forming an informational network.

Acute allograft rejection displays different patterns in different tissues reflecting the different homeostatic agendas of the different tissues. The "competing network agendas"[12] of the tissues and the immune system can be regarded as an instance of the necessity to harmonize the diverse aspects of immune response.

In his essay, Cohen[1] suggests that an allograft "looks like an aberrant self-tissue in need of maintenance". Our view of cytokines as an informational network suggests that the physician may not be wise to counter graft rejection by using cytokines to command a diminishment of "harmful maintenance" for commands have numerous unforeseen side effects. Rather cytokines should be used as an information modifier, in order somehow to change the internal image of the allograft from aberrant to normal self-tissue.

Final Remarks

In my view, it is desirable to employ the term feedback not, as is common, in the broad sense of "interaction" but rather in the more focussed sense of **reaction to information concerning response appropriateness**, appropriateness with respect to set points (classical feedback) or appropriateness with respect to a set of overlapping and conflicting aspects of response (dispersed feedback).

Dispersed feedback and checking response appropriateness via sensors are intimately interrelated. If a biological system has a sensor then it is probably worthwhile to regard the sensor as having evolved to sense the effectiveness of one or more aspects of homeostatic responses, or to sense useful general information on the state of the system. Via intracellular processing, dispersed biological feedback integrates information from membranal sensors to promote fitness by enhancing the appropriateness of the overlapping and contradictory aspects of system response. "Promoting fitness" is too lofty an aim to be sensed. But it is inherent in my definition of the aspects of a physiological response that the effectiveness of a given aspect of response can be sensed by a cell. Indeed, "the individual cells and molecules... do not know... when each is doing the right thing".[1] But the cells can know from their sensors when they are doing a better thing or a worse thing and they can modify their behavior in light of this knowledge. Furthermore a given cell can provide information to other cells on its own action and on its reading of the general state of the system; the cell collective can employ this information to select cell subsets that more appropriately help promote the ever-shifting spectrum of actions required of an appropriate response. Thus the cells' reaction to their sensors' information is responsible for the fact that not "the right thing" but a somewhat better thing usually "emerges from a collective and dynamic interaction".[1]

In this essay I have concentrated on how dispersed feedback can use information from a variety of sensors to strengthen positive aspects of the immune response and to diminish negative aspects. Another whole

D continued

chapter, whose investigation has barely begun, concerns how dispersed feedback can contribute to an appropriate balance among the many overlapping physiological systems, of which the immune system is just one, that contribute to an organism's health and reproduction. Hard information is just beginning to become available on the fitness costs of immune response, both innate responses in invertebrates and adaptive responses in vertebrates.[13] Common sense indicates that indeed fitness requires different emphases on different physiological systems at different times and circumstances. The nervous system would seem an obvious candidate for sensing the performance of the various systems and using the results to improve overall response. Indeed, evidence is accumulating in favor of powerful nervous system influence on the immune system, for example by neuropeptides that act via specific T cell receptors to influence T cell secretion.[9]

Acknowledgments. I am grateful to I. Cohen, A. Goldberger and C. Orosz for useful reactions to a preliminary version of this essay.

REFERENCES

1. Cohen IR. Immunity, set points, reactive systems, and allograft rejection. **Graft 2001; 4(5):369-378**.
2. Fadok VA, Bratton DL, Konowal A, et al. Macrophages that have ingested apoptotic cells in vitro inhibit proinflammatory cytokine production through autocrine/ paracrine mechanisms involving TFG-β, PGE2, and PAF. **J Clin Invest 1998; 101:890-898.**
3. Goldberger AL. Non-linear dynamics for clinicians: Chaos theory, fractals, and complexity at the bedside. **Lancet 1996; 347:1312-1314**.
4. Grossman Z, Paul WE. Adaptive cellular interactions in the immune system: The tunable activation threshold and the significance of subthreshold responses. **Proc Natl Acad Sci USA 1992; 89:10365-10369**.
5. Hausdorff JM, Purdon PL, Peng C-K, et al. Fractal dynamics of human gait: stability of long-range correlations in stride interval fluctuations. **J Appl Physiol 1996; 80:1448-1457**.
6. Ibrahim M, Chain B, Katz D. The injured cell: The role of the dendritic cell system as a sentinel receptor pathway. **Immunol Today 1995; 16:181-186**.
7. Ivanov PCh, Amaral LAN, Goldberger AL, et al. Multifractality in human heartbeat dynamics. **Nature 1999; 399:461-465**.
8. Janeway CA, Jr. The immune system evolved to discriminate infectious nonself from noninfectious self. **Immunol. Today 1992; 13:11-16**.
9. Levite M. Neuropeptides, by direct interaction with T cells, induce cytokine secretion and break the commitment to a distinct T helper. **Proc Natl Acad Sci USA 1998; 95(21):12544-12549**.
10. Lider O, Cahalon L, Gilat D, et al. A disaccharide that inhibits tumor necrosis factor is formed from the extracellular matrix by the enzyme heparanase. **Proc Natl Acad Sci USA 1995; 92:5037-5041**.
11. Matzinger P. Tolerance, danger, and the extended family. **Ann Rev Immunol 1994; 12:991-1045**.
12. Orosz C. An introduction to immuno-ecology and immunoinformatics. In Segel LA, Cohen IR, editors. Design principles for the immune system and other distributed autonomous systems. **London: Oxford University Press; 2001. p. 125-150**.
13. Read AF, Allen JE. The economics of immunity. **Science 2000; 290:1104-1105**.
14. Segel LA, Diffuse feedback from diffuse information in complex systems. **Complexity 2000; 5:39-46**.
15. Segel LA. Diffuse feedback from a diffuse informational network. In Segel LA, Cohen IR, editors. Design principles for the immune system and other distributed autonomous systems. **London: Oxford University Press; 2001. p. 203-226**.
16. Segel LA and Lev Bar-Or R. On the role of feedback in promoting conflicting goals of the adaptive immune system. **J Immunol 1999; 163:1342-1349**.
17. Segel LA, Cohen IR, editors. Design principles for the immune system and other distributed autonomous systems. **London: Oxford University Press; 2001**.
18. Sercarz E. Distributed, anarchic immune organization: semi-autonomous golems at work. In Segel LA, Cohen IR, editors. Design principles for the immune system and other distributed autonomous systems. **London: Oxford University Press; 2001. p. 241-260**.
19. Stryer L. Biochemistry, 3rd ed. **New York:W.H. Freeman, 1988**.
20. Van Uden J, Raz E. Immunostimulatory DNA and applications to allergic disease. **J Allergy Clin Immunol 1999; 104(5):902-910**.
21. Wagner H. Bacterial CpG DNA activates immune cells to signal infectious danger. **Adv Immunol 1999; 73:329-368**.