Immune Network Theory

Geoffrey W. Hoffmann

The University of British Columbia

For many years we have known that the immune system can sense the presence of foreign things in the body. The immune system reacts to foreign substances by making antibodies, and thus rids the body of intruders. We now know that the immune system also has the ability to sense its own components. It is an inward-looking system, in addition to being a defensive system against intruders. This perspective leads to a unified theory, called the symmetrical network theory, that is based on a small number of postulates, and covers many aspects of the regulation of the adaptive immune system.

Table of Contents

Outline

Immune network theory has been developed since the 1970s, and has been previously published in journals and in book chapters including conference proceedings. This monograph presents and extends the scope of a version of immune network theory called the symmetrical network theory.

There has been a problem for students and even many scholars of immunology of not being able to see the wood for the trees. The subject of immune system regulation is huge, and the selection of facts provided by conventional textbooks is not geared to understanding the immune system in terms of immune network theory. In a typical immunology textbook the network "idea" is summarily dealt with in a few paragraphs, and most of today's immunology graduates are not aware of the extensive experimental and theoretical literature that exists on network immunology.

Chapter 1 describes the general role of theory in biology and specifically in immunology. How do we build theories, and how we know when we have a good one? The pivotal role of mathematical modelling in theoretical biology is outlined. Another key component is to identify any paradoxes that exist within a reigning paradigm, and to seek and find ways to resolve them.

Chapters 2 to 7 introduce some basic concepts of immunology, the building blocks that the theorist has to work with. Much of this material can be found also in introductory texts on immunology. The main difference is that almost no more detail is given here than is necessary for the development of network theory in the subsequent chapters.

Chapter 2 begins with a thumbnail sketch of antibodies, the protein molecules that are highly diverse and can be highly specific in their interactions with substances that are foreign to the body (antigens).

Chapter 3 describes the clonal selection theory. This fundamental law of immunology describes how antibodies are produced by lymphocytes.

Chapter 4 describes some of the forms of the response of the immune system that can be observed following stimulation by various foreign antigens. This chapter concludes with the question of whether the many different immune response phenomena are the result of a corresponding diversity of mechanisms (which tends to be the conventional view) or whether a unifying model can be found that accounts for a wide range of diverse phenomena. The latter is the aim of immune network theory.

The specificity of antibodies and lymphocytes is the subject of Chapter 5. While the "exquisite" specificity of antibodies is a cliché, antibodies actually have variable parts called V regions with multiple sites for potentially binding to antigens, and antibodies are consequently multispecific.

The regulation of specific B cell and T cell responses is the most fascinating aspect of the immune system, and some of the main phenomena are described in Chapter 6. Of particular importance are the ways in which T cells can be shown to help and suppress immune responses. These are the most basic phenomena that a theory of regulation of the immune system needs to account for, so this chapter sets the stage for the later development of network theory.

While network regulation deals mostly with molecules that have V regions and hence are called specific, there are also some cells and molecules without V regions that play a role, and some of these non-specific mediators are described in Chapter 7.

Niels Jerne's network hypothesis is introduced in Chapter 8, and states that interactions between the V regions of antibodies (and similarly interactions involving other molecules and receptors with V regions) play a key role

in the regulation of the system. The hypothesis did not spell out many details of how this might occur; that came later with the explicit network models that have been developed. The hypothesis is firmly based on clonal selection, and involves clones recognizing each others' V regions, instead of being independent of each other. Mutual recognition of clones means that the immune system is introspective, with self-recognition being at least as important as the recognition of foreign things. The first explicit network model to be developed was formulated by Peter Richter. It is also discussed in Chapter 8.

Chapter 9 documents an important symmetry in the immune system, that was recognized in 1975 and led to the symmetrical network theory, which is the main focus of this book. "First symmetry" was the starting point for the theory that provides us with a way of fitting many pieces of the immunological puzzle into a coherent whole.

The symmetrical network theory is introduced in Chapter 10. This chapter describes how much of the important phenomenology can be understood in terms of this most complete version of network theory. The central role of mathematical modelling in finding a self-consistent theory is also made clear in this chapter. The clones specific for a given antigen include two main classes, namely those that have complementarity to the antigen and those that resemble the antigen to a greater extent than the extent to which they are complementary to the antigen. A two-variable mathematical model of the population levels of these clones simulates some of the main ideas of the theory. Non-specific accessory cells play a role in switching between stable states of the system.

Chapter 11 extends the two-variable mathematical model (two classes of clones) developed in chapter 10 to an N-variable model, and the question of how the added complexity can be compatible with the system being stable is explored. A central aspect of the system is unusual as a mathematical system. It is robustly stable even when there are a very large number of interconnected clones; in fact the stability increases with the complexity of the system. Most dynamical systems become less stable as they become more complex.

Chapter 12 describes a phenomenon that was discovered around the time that network theory was invented, namely "MHC (major histocompatability complex) restriction." T cells are selected to recognize the MHC molecules of the individual in addition to having specificity for foreign antigens, and thus they appear to have a dual specificity. The interpretation of this phenomenon is straightforward in the context of the symmetrical network theory.

Chapter 13 relates how, in 1982, one of the pieces in the puzzle appeared not to fit. This became known as the "I-J paradox". The I-J phenomenon had been discovered in 1976, and appeared to be central to an understanding of immunoregulation. The I-J paradox concerned what seemed to be a missing gene, a discrepancy between serological mapping of a presumed protein(s) called I-J and molecular genetic mapping. The discrepancy unfortunately caused many immunologists to lose interest in the entire area. It was shown in 1988 that the conundrum could be solved in terms of the symmetrical network theory, but by then the waves caused by the paradox had caused many immunologists to be sceptical of a large section of immunoregulatory phenomena and theory, including suppressor T cells, I-J and network theory, and they threw out the baby with the bathwater. It is time to revisit a largely forgotten story. In this chapter we describe the I-J paradox and show how it is important and can be resolved in the context of the symmetrical network theory.

Chapter 14 discusses antibodies that are found in the serum of animals that have been immunized with cells from animals of the same species, but of a different strain. When we take account of antibodies that are specific for receptors of foreign lymphocytes that recognize self antigens, an intriguing new symmetry emerges, that we call "second symmetry". This symmetry also has implications for our understanding of I-J.

Chapter 15 addresses the question of how the immune system distinguishes between self and non-self structures. The question of how the system can go wrong, and make an immune response against self components (autoimmunity) is explored.

Chapter 16 describes an idiotypic network model of HIV pathogenesis. In this model HIV triggers autoimmunity. The theory is based on the model of I-J described in chapter 13. The theory is furthermore based on the fact that HIV preferentially infects HIV-specific helper T cells. The theory predated the experimental demonstration that this is the case. The theory leads to a potential vaccine based strategy for the prevention of HIV infection.

Chapter 17 describes a more detailed idiotypic network models of regulatory T cells that includes the influence of self MHC class II on Th1, Th2, Ts1, Ts2 and Ts3 cells and serum IgG. Phenomena called I-J restriction and Igh restriction of suppressor T cells are described and explained in the context of the model. Serum IgG is ascribed a role in the maintenance of self tolerance. The inverse relationship between antibody immune responses and delayed type hypersensitivity is ascribed to competition between suppressor T cells and contrasuppressor T cells. The model is then extended to include all self antigens and IgM secreting and IgG secreting B cells.

Chapter 18 assesses what has been achieved in immune network theory, and where the field needs to go from here.

Chapter 1. Immunology: many facts, few theories

I don't believe any experiment
until it is confirmed by theory
-Arthur Eddington

Hypotheses, models and theories

The immune system is an ongoing major focus of medical research. This book is about a theory for understanding the way the immune system regulates itself. What is a theory? A theory can be defined in terms of the related terms hypothesis and model. A scientific hypothesis or postulate is an idea about how something works. A model is a set of hypotheses, that together may help us to understand how something works. A theory is a well-developed model.

The immune network hypothesis

In the early 1970s an eminent Danish immunologist working in Switzerland triggered a revolution in our understanding of the immune system. Niels Jerne proposed that cells and molecules of the immune system not only recognize foreign substances, but also recognize, respond to and are regulated by each other. It followed that we should regard the immune system as a network of interacting cells and antibodies.[1] This perspective is known as the idiotypic network theory, or more simply the immune network theory. The idea is that the cells of the immune system are functionally connected via variable components with enormous diversity (V regions), with each cell being connected to a small subset of the rest via interactions between V regions. Important properties of this system, including memory, are then properties of the network of cells as a whole, rather than of the individual cells. This was a revolutionary paradigm for immunology, and much progress has since been made towards understanding the immune system in these terms.

The hypothesis becomes a theory

The formulation of a set of explicit, self-consistent postulates, the development of mathematical models based on those postulates, and the garnering of experimental evidence supporting the postulates have in the meantime transformed the immune network hypothesis into an immune network theory. Immune network theory is as central to cellular immunology as is the clonal

[1]N. K. Jerne (1974) Towards a Network Theory of the Immune System, Ann. Immunol. (Inst. Pasteur), 125C, 373-389.

selection theory, its immediate predecessor. Towards the end of the book we will learn about a theory of AIDS pathogenesis that is based on immune network theory, and we can expect that network theory will be important in developing a better understanding also of other autoimmune diseases, of allergies and possibly even cancer. In typical courses on cellular immunology, however, network ideas are not treated in any depth, and the fact that specific regulation of the immune system is most easily understood in terms of network mechanisms is not widely appreciated. A pedagogic presentation on immune network theory has been lacking for some time, and this book is an attempt to provide one. Undergraduate immunology courses do not typically include immune network theory, but they should.

Mathematical modeling as a tool

Textbooks for the life sciences do not traditionally contain much mathematics, and many biologists have yet to concede that mathematics has an important role to play in biology. The relevance of mathematics for network immunology will become evident as we develop specific network models.

This book is intended to make immune network ideas widely accessible. It is written for biologists, physical scientists and readers of popular scientific magazines. While it contains some mathematics, the main ideas are presented in a way that can be appreciated also by someone who is not interested in the mathematics. Just as a theoretical biologist can learn the principles and results of immunology experiments without having the knowledge and skills of an experimental immunologist, experimental immunologists and others can learn the principles and results of computational experiments without having the knowledge and skills of a theoretical biologist. The mathematics is primarily a tool for visualizing the dynamics of the system and adds rigour to the development of the ideas. The main kind of mathematics used is calculus, in particular coupled non-linear ordinary differential equations. A reader will be able to understand the diagrams that come from the mathematics without necessarily working through the details of the mathematics.

We have learned a lot about immune network regulation, but there is still a long way to go. We will see a growing need in the future for research immunologists who are conversant with mathematical modeling as an investigative tool. The number of immunologists trained also in the physical and mathematical sciences is small. They will have an important part to play in the continuing elucidation of immune regulation mechanisms.

A model is a set of postulates about a system. Model building in theoretical biology can begin with a broad, relatively undifferentiated idea, that gains definition as we attempt to capture its essence mathematically. A mathematical model is a mathematical formulation of the set of postulates. Mathematical models are powerful tools for analyzing the consequences of particular sets of

postulates. A set of explicit postulates about a system is translated into the form of equations. The mathematical model typically reflects both non-linear complexities[2] and any simplifications inherent in the postulates. We can then study the behaviour of the mathematical model, and see whether it has properties similar to the biological system of interest. We can see whether the simplifications result in the loss of the properties we are looking for. On the other hand non-linearities in the model frequently give rise to unanticipated and even counterintuitive behaviour. If the model does not have desired (experimentally observed) properties, we know that there must be something wrong with the set of postulates we made, and we move on to a different, sometimes more complex set of postulates. In this way mathematical models are effective filters for ideas. Just like experiments, they can be used to determine which ideas (or combinations of ideas) might be correct and which ones are certainly wrong.

Experimental immunology is the starting point

Working on immune network theory requires learning the language of experimental cellular immunology. This book includes an introduction to that language. The treatment of the experimental basis of network immunology is however skeletal. Some experimental facts that are important for distinguishing between alternative models will be presented in detail, while many experimental systems (that have to do with network immunology but not directly with immune network theory as presented here) are not included. This treatment is intended to be a self-contained exposition of experimental facts with a theory that ties them together. Notwithstanding this, some readers will wish to supplement reading it with reading a conventional immunology textbook, for example "Introduction to Immunology" by John W. Kimball (Macmillan), followed by reading some of the original literature.

A problem for the immune network theorist is the plethora of published facts. About 8000 papers are published in immunology journals each year. Which of these are most relevant and which ones are of secondary importance? Which experimental results have been given the right interpretation? Most of the experimental papers that are important for network immunology are more

---

[2] Linear models are models in which the variables, and/or their rates of change, are linear sums of the other variables, for example
dx/dt = the rate of change of x = ax + by
dy/dt = the rate of change of y = cx + dy
is a pair of differential equations describing the rates of change in x and y as linear functions of x and y. Linear systems are readily solved and exhibit simpler behaviour than non-linear systems, that contain terms such as x^2 and xy^2. Models of interesting biological systems, including the immune system, are typically non-linear systems.

than two decades old. In this field, "older" does not necessarily mean "inferior". Facts are important or not in the context of specific theories. Theories are important or not in the context of specific facts. This cycle of mutual interdependence of facts and theories has to be broken somehow, and to this end we need criteria for evaluating theories. We will find that many old facts about immune system regulation are simple and important in the context of immune network theory.

Evaluation of competing theories

The following is a set of qualities that a theory of immune system regulation should have to be acceptable. This set provides a basis for systematically, and hopefully objectively, comparing competing theories.

(a) Simplicity. The simpler a theory, the better the theory. This criterion was enunciated in the fourteenth century by the theologian and philosopher, William of Ockham. His formulation was "plurality should not be assumed without necessity", which is known as "Ockham's razor." "Plurality" is here synonymous with "complexity". Ockham's razor is indeed the sharp knife that often separates a good theory from inferior theories. It may seem that the more complex we make a theory (that is, the more postulates), the more phenomenology we will be able to explain, and the more testable predictions we may be able to make. However, a theory with many postulates is inelegant, and it is unlikely to be unique. If for each phenomenon to be explained, we make one postulate, we can have a theory with very many postulates. As in the case of murder mysteries, many complex theories can account for a given set of facts, but typically only one simple theory accounts for the same set of facts. Compared with the simple theory, the complex theory looks far-fetched. It is satisfying to have a single idea or postulate leading to an understanding of multiple facets of a case. The simple theory is more readily disprovable if it is wrong, since there are less adjustable parameters that can be used to fit experimental results. A simple theory typically also has more predictive power than a complex theory, again because there are less adjustable parameters. Simple theories tend to be supported by simple experimental results.

There is an urgent need for simple theories of the immune system because much of experimental immunology is like stamp-collecting. Immunologists collect an enormous number of facts, but an underlying simple structure for the facts often fails to emerge.[3] The point of doing theoretical immunology is

[3] For examples of encyclopedic collections of facts see "Fundamental Immunology", W. E. Paul, Ed., 1984 (1st ed.) and 1989 (2nd ed.). Even the two chapters on idiotypic networks in these two volumes are collections of facts with no significant attempt to tie them to the rest of immunology. They are of course useful as collections of facts, but we need a unifying theory.

to find an underlying simplicity in a system that may at first appear to be complex.

(b) Scope. The broader the scope of a theory for a given level of complexity, the better. The scope of a theory is the number of phenomena explained. Not all phenomena are of equal importance; and an objective awarding of "points" can be tricky. There are four equations, called Maxwell's equations, that together describe a vast array of electromagnetic phenomena. This is an example of an extremely successful theory, both from the point of view of simplicity and scope.

(c) Predictions. We can test a theory by determining whether its predictions can be confirmed. We need to be careful, however, since more than one theory may lead to a particular prediction. A confirmed prediction does not therefore prove that we have found the right theory. On the other hand, if an experimental result clearly conflicts with a prediction of a theory, the theory is disproved. It is therefore possible to disprove theories, but impossible to prove them. Karl Popper was a philosopher of science who discussed questions of the provability and disprovability of scientific theories.[4] He saw verified predictions as failed attempts to disprove a theory. Theories gain credibility from such failed attempts to disprove them. On the other hand, if the prediction of a theory is highly specific (for example, the theory predicts an experimental curve should have a particular shape and a specified maximum value), and this prediction is confirmed experimentally, we justifiably become confident that the theory is on the right track. Several predictions of the symmetrical network theory have been confirmed, and there are no failed predictions that have disproved the theory. As the theory is fleshed out, more predictions emerge, and this volume includes 18 predictions of the symmetrical network theory, of which 3 have been confirmed.

(d) Resolution of Paradoxes. A dominant view of how a particular system works is called the paradigm (broadly accepted working hypothesis or theory) for the system. Facts that do not fit the prevailing paradigm are particularly important, because they may be windows to fundamental progress.[5] Such facts are called

---

[4] K. Popper, in The Logic of Scientific Discovery, 1934, chapter 1, section 6: "I shall certainly admit a system as empirical or scientific only if it is capable of being tested by experience. These considerations suggest that not the verifiability but the falsifiability of a system is to be taken as a criterion of demarcation ... It must be possible for an empirical scientific system to be refuted by experience."

[5] The most striking examples of paradoxes that led to revolutions in scientific understanding occurred in physics. Quantum theory was spawned by a conflict between theoretical expectations and experimental measurements of black body radiation. Relativity was born following a failure to find evidence of an ether as the medium of electromagnetic waves in the Michelson-Morley experiment, a paradox in the context of 19th century physics.

paradoxes. We cannot have a paradox without an accepted paradigm from which the result of an experiment can be predicted, and a paradox exists only in the context of such an accepted paradigm. A famous paradox of cellular immunology was the I-J paradox, a molecular biological result that contradicted results from classical immunogenetics (the genetics of the immune system). This paradox emerged in 1982; we will describe it in chapter 13, and see how it can be resolved within the network immunology paradigm. The experimental facts underlying a paradox are particularly important facts.

(e) Mechanistic basis. A model that simulates dynamical properties of a system, but is not explicit about underlying mechanisms, is called a phenomenological model. In contrast to this, a model that is based on specified components and mechanisms is a called a mechanistic model. A model with a mechanistic basis is typically superior to a phenomenological model, since it is more amenable to experimental tests. Phenomenological models are abstract constructs (typically mathematical constructs) with the desired dynamical properties, and can be useful as forerunners to a mechanistic theory in which the variables are given explicit interpretations in terms of specified physical components and processes.

(f) Rigour. We can have rigour in the extent to which invoked "facts" or postulates of the theory have been "proven" experimentally (disproof has failed, see *(c)*). The experimental findings have to be reproducible. We can also have rigour in the derivation of the properties of the model from its postulates, which may involve mathematical modeling. Obviously, the more rigour the better is the model.

(g) Robustness. If a model "works" (gives sought after behaviours) for only a very narrow range of model parameters, the question arises as to whether the physical or biological system being modeled has parameters in that precise range. If it works over a wide range of values of one of the parameters, we say the model is robust with respect to the value of that parameter. If it is robust with respect to (say) four parameters and requires a fifth variable to have a sharply specified value, this leads to a highly specific prediction that can be useful in validating the model. We test robustness of a mathematical model by determining the ranges of parameters that lead to the various types of behaviour for which the model is designed to account. These ranges can typically be determined also experimentally. The experimentally measured values of the parameters should then fall within those ranges. The non-linearities contained in many mathematical models of biological systems result in the models exhibiting the required behaviour over a satisfactorily broad range of parameter values. In such cases, even if only limited experimental data is available, the models are often considered to be robust.

(h) Aesthetics. Beauty in many cases is in the eye of the beholder, but a model or theory can also be seen as attractive from the perspective of a broad constituency. The aesthetic appeal of a theory is not independent of the criteria

already listed, especially simplicity and scope. The double helix was a beautiful solution to the problem of the molecular basis for the inheritance of genetic information.

Two networks

In 1983, at the International Immunology Congress in Kyoto, Japan, the immunologist Baruj Benacerraf (Nobel Laureate, 1980) pronounced that the immune network theory had been proven to be correct. In 1984 Niels Jerne was awarded the Nobel Prize, largely for his formulation of the network hypothesis. This was an important step in the level of recognition of network theory. At that time another type of complex system was however becoming important, and immunologists found themselves in a position of having to choose one of two complex systems on which to focus. The first system has to do with what is known as specific regulation of the immune system, and the second has to do with non-specific regulation. It was difficult or impossible to gain an understanding of both of these complex systems simultaneously. Specific regulation (also referred to as "antigen-specific regulation") involves changes in the system that typically affect its response to one particular substance (or "antigen") but not to another substance (a "control antigen"). Non-specific regulation, on the other hand, concerns regulatory mechanisms that do not involve discrimination between various antigens. The antigen-specific aspect of regulation ("adaptive immunity") involves the network of V (variable) regions, which, like the antigens, have a great diversity of shapes. Non-specific regulation involves a system of soluble regulatory molecules called lymphokines, and cell surface molecules called CD molecules, both of which do not have V regions. Lymphokines and CD molecules play a role in both adaptive immunity and a more primitive form of immunity that does not involve learning and memory, called "innate immunity".

Since the mid 1980s there has been a preference among experimental immunologists to concentrate on the lymphokine/CD network rather than the V region network. The V region network has continued to be studied by a relatively small band of devotees. Eventually we will need to understand both the V region network and the lymphokine/CD network, and interactions between the two. It has become clear that the lymphokine/CD network alone does not lead to a simple way of understanding the complexities of cellular and molecular immunology. This book focuses on the V region network (also called the idiotypic network), and we will find that a satisfying picture of many specific regulatory aspects of the immune system emerges from this perspective, with only limited reference to the lymphokine/CD network. This leads to the hope that an understanding of the complete system (V region network plus lymphokine/CD network) will eventually be reached via an initial understanding of the V region network.

Top-down and bottom-up approaches

Biologists and theoretical biologists[6] typically tackle the job of trying to understand the immune system in two different ways. Biologists traditionally have a "bottom-up" approach to unravelling how something biological works. This means they seek to understand the components first, then determine how the components could fit together to produce a functioning system. An example of the success of this approach was the determination of the structure of DNA, and the way in which the structure provides the basis for an understanding of genetics. From the double helical structure of DNA, with complementary base pairs, it was easy to see, at least in principle, how genetic information is stored and reproduced.

Theoretical biologists, physicists and physical chemists, typically use a "top-down" approach. They study a system as a whole, and may determine how it responds to small perturbations or other stimuli. They ask what minimal set of components would be needed to produce the observed behaviours, and if there is no direct evidence of such components, they feel free to postulate their existence. In chemistry, this worked well in the formulation and development of the atomic theory of matter. The law of equal proportions led to the postulates of the existence of atoms and molecules. For example, two moles of hydrogen plus one mole of oxygen react to make one mole of water (H_2O). In physics quarks are postulated to explain the properties of the fundamental particles in an optimally simple way. Postulating anything in the absence of direct evidence is anathema to most biologists. So it is perhaps not so surprising that the structure of DNA was not predicted by anyone using a top-down approach, even though it perhaps could have been.

The distinction between top-down and bottom-up approaches is related to a distinction between two kinds of experimental data, namely system component data and system response data. The bottom-up approach is based mainly on system component data, while the top-down approach is based mainly on system response data.

A purely bottom-up approach to immunology does not work because there is an enormous amount of component data already available about the immune system, and a core problem is to determine which data are of fundamental importance for understanding the system. It has been estimated that the

6 The designations "biologist" and "theoretical biologist" are chosen for want of better terms. Hypotheses are routinely formulated and tested by experimentalists without reference to professional theorists. In other words, many cellular immunologists are effectively their own theorists, and in many cases this works well. Nevertheless, an expanded role for theorists, who both understand the experimental side of the discipline and who appreciate the role of mathematical models, is needed.

science of immunology consists of more than a million facts and beliefs. Conservatively, 8000 or more papers are published per year, with an average of 5 or more display items (figures or tables). Each display item makes a specific point, reflecting a fact or belief. This means that more than 40,000 facts and beliefs are published per year, and after twenty-five years the scholar has a million of these at his or her disposal. Needless to say, the vast majority of these facts and beliefs are not retained in the consciousness of the community of immunologists. In this book we can obviously only touch the surface of that data. Data discussed here are selected on the basis of their relevance to an understanding of immunology in the context of network theory. A network theory emerges that is based on a small number of postulates, reflecting what are deemed to be the most important facts. The theory accounts for much of what we know about how the system behaves. It is called the "symmetrical network theory". It satisfies the criteria we have specified for a satisfactory theory.

The development of the symmetrical network theory utilized both the top-down and the bottom-up approaches. The theory illustrates how an understanding of system response data can be related to system component data. While a significant amount of progress to this end has been made, there is a vast amount of information available in libraries, that is unprocessed in the context of network theory, and much is still to be done. The extent to which the theory is correct will be determined by the extent to which the many predictions of the theory (see index) can be validated.