Autonomous Distributed Control in the Immune System Using Diffuse Feedback

JASON BROWNLEE
Technical Report 070329A
Complex Intelligent Systems Laboratory, Centre for Information Technology Research,
Faculty of Information and Communication Technologies, Swinburne University of Technology
Melbourne, Australia
jbrownlee@ict.swin.edu.au

***Abstract*-How does the immune system identify a pathogen and regulate the resources and severity of an immune response without a master controller, a guiding goal, or a guiding global performance measure? A feedback and control theory called diffuse feedback (diffuse information network) proposed by Lee Segel is discussed which addresses this complex and multiple-faceted question. The result is a theory and set of principles that address the autonomous control of a distributed population of decision-making agents through diffuse information signals and spatial self-organization.**

Keywords- Autonomous Decentralized System, Immune System, Diffuse Feedback, Spatial Self-Organization

I. INTRODUCTION

The immune system is a completely distributed system (to a first approximation) that autonomously provides and maintains a measured and localised immune response to an antigen. *How might the immune system achieve this highly complex functionality?* This question was addressed by mathematician and theoretical immunologist Lee Segel in his proposition of a 'diffuse feedback' or 'diffuse information network' theory. In elucidating this theory, Segel proposes that the principles of this specialised form of information-dissemination and feedback-control observed in the immune system may provide a new path for the design of autonomous and distributed artificial intelligence systems.

Segel's work began with the definition of 'diffuse feedback' in the context of the spatial self-organization and autonomous control of immune response in the immune system [4,13]. In this work a series of diffuse feedback (differential equation) models were proposed which were later expanded upon in an immunological context [14]. The specialised form of feedback control system was compared to tradition feedback control [5], and evolutionary algorithms were used to solve aspects of the differential equation model [1]. Diffuse information was proposed as the medium for feedback [6], and the theory was discussed at length [7-9]. The work culminated in a book by Segel and colleague Cohen [12] that focused on the principles of the theory. The book presented autonomous and distributed principles from immunology and similarly realted biological systems. Considering the theory, ones thoughts cannot help but wonder into those of the spatial models of the immune system. This was Segel's background in theoretical immunology (for example his exploration of the 'shape space' paradigm [3,11,16]). Although not directly related to the theory, he provides a high-level summary of such spatio-temporal theoretical models in [10].

This work reviews the diffuse feedback and diffuse information network theory of the immune system. Section II discusses possible goals of the immune system from a high-level, which Segel suggests is dubious. He proposes an evolutionary perspective and a set of 'low-level' (diffuse) goals that may have overlapping and or contradictory properties. Section III discusses the meat of Segel's diffuse feedback theory of the immune system. This section covers the diffuse information, sensing, feedback and control aspects of the theory, and briefly summarises (qualitatively) Segel's two exemplar models. Finally, section IV summarises Segel's comments on the use of the theory in the design of autonomous and decentralised systems.

II. IMMUNE SYSTEM GOALS

What is the goal of the immune system? One may impose the principle function of the immune system is to limit damage to the host by pathogens. How does the system measure it performance towards this goal? If the system is minimizing damage to tissue then should the average damage be minimized? Should the maximum damage be minimized (to avoid fatal attacks)? Should their be some combination? Should some tissues or organs be protected more than others?

Evolution is not an extremalizing process, there is no good reason to think that evolution has 'really' optimized anything. The immune system is one bodily system that has adapted to detect and neutralise pathogens, potentially extending a hosts lifetime and increasing its reproductive 'fitness'. The goal of evolution is not to maximise individual fitness, although such a model is useful in dealing with the complexities of evolution. The same approach may be used in a general approach when dealing with the immune system.

In addition to evolving without a specific high-level goal, and the immune system *acts* without a high-level goal. Many elements make up the immune system, and many different specific defence strategies, from which the correct one must be selected for a given pathogen.

Further, pathogen themselves have strategies for overcoming various aspects of defence in the immune system, such as sabotaging immune signalling. Making the wrong 'tactical' decision has consequence. A response may damage self-tissues, it may divert much needed resources away from a more pressing invasion, or it may blunt the effectiveness of the immune response.

In summary, defining a high-level goal for the immune system fail because (1) it is hard to define a suitable objective, (2) the system was constructed without one, and (3) the system operates without such a high-level goal control system. The imposed high-level goal of 'effective pathogen killing' may be approximated by a series of low-level goals:

1. It is good for cells in the immune system to kill harmful pathogens (perhaps proportional to virulence).
2. It is bad for cells of the immune system to harm self-tissues.
3. It is good for the immune system to minimise energy expenditure.

Unlike the initial high-level, long-term goal, these approximations are low-level and short-term. In addition, they are spatially localised, and as such require distributed localised performance measures. In locally seeking sub-goals, no optimum is sought. Different cells may deal with different sub-goals, and an overall solution is provided. This provides a flexible, and self-tuning, rather than a hard-wired response.

How is local performance assessed? How does feedback promote the sub-goals? How might such a 'bottom-up' control system work in the immune system?

III. Diffuse Feedback in the Immune System

Traditional feedback control systems work by measuring the difference between a control signal and an actual signal and act to reduce that difference. Any influence that flows from the state of the system to modify the base actions of system components. A classical example is that of a temperature control system that seeks temperature homeostasis. Ambient temperature is measured with a sensor and compared to a control temperature, the system then attempts to correct for any discrepancy between the two. In this example, there is a single clearly defined goal, a clearly defined way for measuring relative performance to that goal, and a well-defined method for improving performance. The immune system has the physiological short-term goals of minimizing damage to tissues by pathogens, and minimizing the damage to tissue by its own effector systems. Addressing the first goal may violate the second, and achieving the second goal may result in the take-over by an invading pathogen. These goals are *overlapping* and *contradictory*.

The sub-goals of the immune system are achieved through the actions of the immune effectors exhibiting locally good and bad tendencies.

Diffuse feedback is an alternative to classical feedback. It is a generalization of classical sensor-based feedback to a situation with multiple-sensors, many interacting 'plants' (agents or decision makers), and multiple-progress indicating signals that can enhance the attainment of multiple conflicting goals. Diffuse feedback in the immune system is the tuning of the behaviour of the system components in order to achieve better performance within the context of the systems sub-goals. The principles of diffuse feedback are:

1. Diffuse information from many sensors modifies the actions of many cells to improve the performance with respect to the sub-goals.
2. Different cells of the immune system are affected differently by the same information.
3. Different types of information, and different weightings of the same set of information elements are employed by a given cell when it decides to act (affects the intensity of a cells action).
4. Cellular actions governed by this feedback may include proliferation, migration, signalling, effector function, and cell death.

The goals are diffuse, the feedback is diffuse, and the information provided as feedback is diffuse, but *what is diffuse*? Segel describes diffuse as distributed and dispersed effecting many decision-making agents in space and time. Information, goals, and feedback are inaccurate, blurred, fuzzy, and non-specific.

The components of the immune system are divided into effectors and signals. Cytokines provide a signalling mechanism in the immune system. Cytokines are molecules that regulate the response, and are conventionally viewed as a 'command network'. One cell may be triggered to release many different types of cytokines, and there is a 'many-to-many' relationship between cytokines and the different receptors they may bind to. A cytokine typically affects several functions, and each function may be affected or require several cytokines. Rather than a command network, Segel proposes that the signalling by cytokines may provide a Diffuse Information Network (DIN). The principles of such a network are as follows:

1. The diffuse information network provides information about the state of the immune system, the pathogens, and the host. In particular, when effector cells perform a function, they 'advertise'. For example this may facilitate cell recruitment resulting in the inflammation response.
2. Information is often coded not in the form of a single signal, but rather in a collection of signals (spectrum, vector, or set of properties).
3. Information generated at some point is typically effective only near that point (localised), and only for a limited period of time (although some forms of information may spread for great distances).
4. The association of two pieces of information generates a new information.

Diffuse information and diffuse feedback allow the immune system to continuously-monitor performance toward the two conflicting goals. As a control strategy, it provides an autonomous and distributed feedback strategy for dealing with uncertainty and a way of addressing multiple conflicting goals. Credit assignment

(the allocation of reward and punishment) of actions is achieved through reinforcement within a geographical proximity (neighbourhood). Reward is allocated for the destruction of pathogens, and punishment may be the absence of reward (implicit negative selection) such as in the case of damaging self. Feedback may let the best response take over, facilitating trial-and-error 'in the field'.

The following summarises the immune response in the context of diffuse information and diffuse feedback:

- There is a broad-spectrum dominant initial response (perhaps biased by evolution).
- The immune system may be considered as addressing a number of sub-goals; it senses and monitors progress towards these sub-goals.
- The diffuse information network provides feedback regarding progress towards sub-goals and state of the system and the host.
- The information provides diffuse feedback, the system adjusts its behaviour, fine-tuning at the lowest level.

This example of a diffuse feedback system shows how a distributed population of autonomous decision-makers may harmonize conflicting goals. Finally, it is worth noting that the proposed control system may fail. Using problems with the immune system as an example, the system may choose the wrong action and respond too aggressively (such as in the case of allergies), strongly respond against self-tissues (as in autoimmune disease), and even the inflammation process may damage tissues if too many cells are recruited to the cause.

Segel, provides two examples of how such a diffuse feedback system can (1) improve the performance of effector cells, and (2) amplify effective effectors. The specifics of the differential equation models are omitted for brevity, opting for qualitative descriptions. See [14] for a complete dissection of the models, or [8,13] for a non-mathematical treatment.

A.*Optimizing Effector Performance*

This is a static model (no clonal selection or cell proliferation) and the system must control the performance of effectors (size of the response). The response is regulated in the context of the conflicting goals: to maximise destruction of pathogens (P) and minimise destruction of self (h). These goals are overlapping in that the effector cells (E) have an effect, the release of a noxious chemical (N), which both kills pathogen and harms self.

Model A is a simple hard-wired response (no information model) where the secretion rate (s) of N by E is fixed in proportion to a single pathogen virulence in an effort to minimise harm to self (h). The model is limited because different pathogens may have different virulences, requiring a case-by-case tuning of the secretion rate (s) of N. A fixed secretion rate may result in damage to self by N if the pathogen virulence is lower than expected (overreact), or damage to self by P if the virulence is higher than expected (under-react).

Model B extends the hard-wired model by making the secretion of N adaptive based on monitoring feedback of local performance (an information model). The model adds additional chemical signals: a harm indicator chemical (H) which provides information of damage to the host (by any means), and a kill indicator chemical (K) that provides information of the amount of pathogen killing. Each provides sensor information as to the local environmental state in the context of the two overlapping and conflicting goals. Combinations permit the regulation of the secretion rate s of N providing an indication of when dangerous pathogens are being killed (requiring an increase in response), and alternatively when the immune system is harming self (requiring a down-regulation of the response).

B.*Optimizing Effector Choice*

This is a spatial scenario involving two effector types E and F, where E is better at combating pathogen P. There are two spatial compartments and cells and chemicals may travel freely between the compartments. A kill indicator chemical (K) is released that indicates how many pathogens have been killed. Proliferation of an effector type is proportional to the amount of K. How can the effector-blind signal be used to select the better effector type E? These models provide an example of how the interpretation of the same chemical signal in two different locations in space can result in the selection of one effector over another.

Model C requires that the effector types be initially grouped together such that effector types E are abundant in the first compartment, and effector types F are abundant in the second compartment. The result is that more K is produced in the first compartment than the second compartment, and given that effector proliferation is proportional to K, effector type E proliferates faster than F. The model is dependant on its starting condition, which results in an initial reinforcement (positive feedback) of the better effector cells, although, becomes less efficient with time as the chemical K diffuses between the compartments. The model could become more efficient through the use of 'sharp selection', that is short-range activation and long-range inhibition (see [2], page 6 and 12).

Model D is an extension of model C, except the K chemical which controls the proliferation rate is delivered to an effector cell by another (helper cell) A cells. A cells are attracted to the chemical K (chemotaxis), thus preferentially move towards larger concentrations of K. In this scenario, the build-up of K in the first compartment leads to an increased proliferation of E effectors, which is maintained given the attraction and maintenance of A cells from both compartments. This both enhances the proliferation of better effectors, and implicitly inhibits the proliferation of less effective effectors.

C. *Other Examples*

Segel developed the theory to describe the feedback-control mechanisms in the immune system, thus the immune system is an emblematic (cellular) example. He suggests that diffuse feedback systems are rampant in

biology; other examples in the body may include regulation systems such as temperature control, and the respiration systems. Segel [5] enumerates a number of examples including chemotaxis which describes the movement of bacteria along chemical gradients towards nutrients and away from hazards, the metabolic network (molecular) that provides the raw material and energy for physiological processes in the body, and an artificial example of a checker-playing system. In [12] these examples are again highlighted in addition to ant colonies (community), and gene expression (the genomic regulation network).

IV. Autonomous Distributed System

In the preface of [12], Segel and Cohen describe an 'autonomous system' as a system without a boss where distributed agents make decisions 'autonomously'. Segel [4] proposed that the immune system, as evaluated in the context of diffuse feedback, may be a prototype for an autonomous distributed system. The theory may be used in the design of a 'bottom-up' artificial intelligence, and initiate a paradigm of distributed intelligent systems. He advocated that insights from the behaviours of the immune system are useful both in understanding other biological systems, and are applicable to non-biological distributed autonomous systems. Segel considered two important principles in this regard: the importance of geography (IV.A), and the importance of communication (IV.B).

A. *Spatial Organization*

Selection (with the right information) can drive spatial dispersion of groups of identical cells through positive reinforcement and inhibition. Spatial organization can allow non-specific signals to select specific cells to contribute to short-term and localised system goals and may be required in an immune response given the turbulent and stochastic nature of pathogen arrival. For example: the lymph nodes facilitate the exposure of lymphocytes to samples of pathogen such that useful effector mechanisms may be generated and dispersed throughout the host organism.

Segel cites the relevance of Maes [15] discussion of situated agents and the benefits of spatial organization. Specifically how task selection by the agents results from the emergent effects of activation-inhibition dynamics amongst tasks. Rational control strategies emerge in a distributed manner by parallel local interactions between simple sets of modules, ultimately providing a way of selecting between competing goals.

B. *Communication*

Feedback is information communicated amongst components of the system. The immune system has various communication mechanisms, not limited to the cytokines, which inspired the diffuse information network conceptualisation. Communication between cells (extra-cellular) is achieved through chemicals and receptors to specific chemicals. Chemical secretion and detection may result in information communication cascades between different cell types. Some chemicals are blind and diffuse, and others may be directed and localised messages. Intra-cellular communication refers to the processes within a cell that occur, such as after activation of the cell through a molecule binding with its receptor. Intra-cellular process may result in various differentiation and proliferation actions. A final form of communication is indirect, such as the migration of cells the directed movement of cells in homing and recruitment.

C. *General Method*

The immune system does not have a well-defined goal, has trillions of individual cells and hundreds of signal types that act concurrently, and uses a proliferation function (positive feedback) to select effector densities. Diffuse feedback provides an approach to forge an improve immune response without global knowledge of what a good response is. It provides a general improvement method and a way to control large collections of leaderless agents to perform complex or intelligent tasks.

Segel [5] defines a general 'bottom-up' method for guiding a distributed system to improving performance without a guiding goal or guiding global performance measure:

1. **Elementary Sub-Goals**: Compile a list of performance goals, do not worry about contradiction or overlap and keep goals general so that they can be reached by many paths.
2. **Install Sensors**: Install sensors that give information about progress towards sub-goals and information on the system and its environment (really, sensors should be local in space and time).
3. **Monitor Performance**: Do the best you can to outline a plan that can provide progress on the various goals. The plan needs to be 'sloppy' (flexible) such that several different options can be used to improve a given performance.
4. **Combine with Action**: Design some way that feedbacks can be used to improve resources allocation (like effector choice).

Acknowledgments

Tim Hendtlass for his patience and for providing useful feedback on drafts of this paper

References

[1] C. Warrender , S. Forrest, and L. Segel, "Effective Feedback in the Immune System," *Genetic and Evolutionary Computation Conference Workship Program,* pp. 329-332, 2001.

[2] Hans Meinhardt. *Models of biological pattern formation*, London ; New York: Academic Press, 1982.

[3] L. A. Segel and A. S. Perelson, Shape space: an approach to the evaluation of cross-reactivity effects, stability and controllability in the immune system *Immunology Letters*, vol. 22, pp. 91-99, Aug, 1989.

[4] Lee A. Segel , "The immune system as a prototype of autonomous decentralized systems," *IEEE International Conference on Systems, Man, and Cybernetics, 1997. 'Computational Cybernetics and Simulation,* Orlando, FL, USA, pp. 375-385, 1997.

[5] Lee A. Segel , Diffuse feedback from diffuse information in complex systems *Complexity*, vol. 5, pp. 39-46, 2000.

[6] Lee A. Segel , Controlling the immune system: Diffuse feedback via a diffuse informational network *Complexity in Biological Information*

Processing, Novartis Foundation Symposium, vol. 239, pp. 31-51, 2001.

[7] Lee A. Segel . Diffuse feedback from a diffuse informational network: in the immune system and other distributed autonomous systems. In: *Design Principles for the Immune System and Other Distributed Autonomous Systems*, ed. Lee A. Segel, I. R. C. New York: Oxford University Press, 2001.pp. 203-226.

[8] Lee A. Segel, How can perception of context improve the immune response? *Autoimmunity and Emerging Disease*, vol. pp. 169-191, 2001.

[9] Lee A. Segel, How does the immune system see to it that it is doing a good? *Graft*, vol. 6, pp. 15-18, 2001.

[10] Lee A. Segel, Some Spatio-temporal models in immunology *International Journal of Bifurcation and Chaos*, vol. 12, pp. 2343-2347, 2002.

[11] Lee A. Segel and A. S. Perelson, Computations in Shape Space: A New Approach to Immune Network Theory *Theoretical Immunology, Part Two, SFI Studies in the Sciences of Complexity*, vol. pp. 377-401, 1988.

[12] Lee A. Segel and Irun R. Cohen. *Design Principles for the Immune System and Other Distributed Autonomous Systems*, USA: Oxford University Press, 2001.

[13] Lee A. Segel and Ruth Lev Bar-Or. Immunology Viewed as the study of an Autonomous Decentralized System. In: *Artificial Immune Systems and Their Applications*, ed. D. Dasgupta. Germany: Springer, 1999.pp. 65-88.

[14] Lee A. Segel and Ruth Lev Bar-Or , On the role of feedback in promoting conflicting goals of the adaptive immune system *The Journal of Immunology*, vol. pp. 1342-1349, 1999.

[15] Patti Maes, Situated Agents Can Have Goals *Designing Autonomous Agents*, vol. 6, pp. 49-70, 1990.

[16] R. J. de Boer, L. A. Segel, and A. S. Perelson, Pattern formation in one- and two-dimensional shape-space models of the immune system *The Journal of immunology*, vol. 155, pp. 295-333, 1992.