

Self-Awareness and Control in Decentralized Systems

Melanie Mitchell

Computer Science Department
Portland State University
P.O. Box 751
Portland, OR 97207-0751
mm@cs.pdx.edu

In *Working Papers of the AAAI 2005 Spring Symposium on Metacognition in Computation*. Menlo Park, CA: AAAI Press, 2005.

Abstract

How can self-awareness emerge in a distributed system with no central control? How can such awareness feed back in a decentralized way to control the system's behavior? Many people have written about how self-awareness might come about in the brain. In this paper, I examine mechanisms for self-awareness and control in two other decentralized biological systems: the immune system and ant colonies. I then attempt to isolate some principles common to both systems. Finally, I discuss ways in which these mechanisms can serve as inspiration for the design of artificial intelligence systems with sophisticated abilities for distributed self-awareness and self-control.

Introduction

A primary mystery of cognitive science is how a unified sense of consciousness emerges in the brain, a system consisting of billions of cells with no central control. "Consciousness" is a difficult term to define, but it is clear that the brain has what we might call "adaptive self-awareness": information about the global state of the system, which feeds back to adaptively control the actions of the system's low-level components. This information about the global state is distributed and statistical in nature, and thus is difficult for observers to tease out. However, the system's components are able, collectively, to use this information in such a way that the entire system appears to have a coherent and useful sense of its own state.

Defined in this way, self-awareness is not unique to the brain. In this paper, I discuss two other biological systems that exhibit adaptive self-awareness: the immune system and ant colonies. In these systems, as in the brain, a major purpose of self-awareness is to control the low-level components of the system in order to insure that appropriate resources get to the right place at the right time. For example, a global response of the brain to visual input ("What I'm seeing looks like trouble") can feed back to control head position, focus of attention, and eye movement, in order to get

a better look at what is worrying the system at a high level. Likewise, the immune system is able to monitor the current state of the body in order to manufacture and spatially place an appropriate set of antibodies for the current environment of pathogens. Ant colonies are able to sense the state of the environment so as to allocate appropriate proportions of individual ants to different tasks.

In all these cases, information must continually be obtained from the environment and used as it is obtained to effect appropriate resource allocation throughout the system, all without the benefit of central control. How do decentralized systems continually collect and use this information?

Here I sketch some mechanisms that have been identified for achieving this in immune systems and ant colonies. I then attempt to abstract some common principles, and discuss how these principles can inspire mechanisms for self-awareness in artificial intelligence systems with decentralized architectures.

Self-Awareness in the Immune System

The vertebrate immune system is a distributed system with multiple adaptive functions, many of which are not yet well understood. Several immunologists have proposed that the immune system should be viewed as an information-processing network that has some "cognitive" properties, including pattern recognition, memory, and self-awareness (Cohen 1999; Forrest & Hofmeyr 2001; Segel 2001d; Orosz 2001).

Segel, in particular, has argued that the immune system must have a way to "see to it that it is doing a good job", in the absence of any central control (Segel 2001d). He has proposed that the immune system's self-awareness comes about via several layers of what he calls "diffuse feedback" (Segel 2001b). Orosz (2001) has proposed a similar set of processes in the context of his work on "immunoinformatics". In this section I present a simplified sketch of one set of immune system mechanisms, those that effect so-called "affinity maturation", in which the immune system adaptively increases its specificity for detecting foreign pathogens such as viruses or bacteria. I will then discuss the ideas of diffuse feedback and how it relates to these mechanisms.

Detection of Pathogens

Pathogens are detected by the immune system via the collective actions of trillions of lymphocytes (white blood cells) that continually circulate in the blood and tissues. A lymphocyte's surface is covered with identical receptors that bind to a particular range of molecular shapes. Within this range, there are particular shapes to which these receptors bind most strongly and others to which they bind more weakly.

In order to "cover" the huge space of possible pathogen shapes in a reasonable way, the population of lymphocytes in the body at a given time is enormously diverse: each individual lymphocyte recognizes a range of shapes that differs from the range recognized by other lymphocytes in the population. When a lymphocyte is born, a novel set of identical receptors is created via a complicated randomization process in the lymphocyte's DNA. Because of continual turnover of the lymphocyte population (about 10^7 new lymphocytes per day) and thus continual introduction of new receptors, the space of possible pathogen shapes is well covered in a dynamic way.

This means that any foreign pathogen is almost certain to encounter a lymphocyte whose receptors will match its shape to at least some degree. When this happens, a bond is formed. The better the match ("affinity"), the stronger the bond. If the number of strongly bound receptors on a lymphocyte exceeds a threshold, and the lymphocyte gets "go-ahead" signals from other lymphocytes with similarly bound receptors, the lymphocyte is activated. Once activated, the lymphocyte secretes antibody molecules, which bind to and neutralize the pathogens, and mark them for destruction by other immune system cells.

The activated lymphocyte then migrates to a lymph node, where it divides rapidly, producing large numbers of daughter lymphocytes. These daughters are subject to a very high rate of mutation, which results in some modifications in their receptors. These mutants are tested on pathogen molecules that are captured in the lymph node. The mutants that do not bind die after a short time.

The ones that bind to the pathogen are unleashed into the bloodstream, where some of them encounter and bind to pathogen molecules, in some cases more strongly than did their mother lymphocyte. These activated daughters themselves migrate to lymph nodes, and create mutated offspring. This cycle continues, with the best-matching new lymphocytes themselves producing the most offspring cells. A Darwinian process ensues, evolving large numbers of new lymphocytes with increasingly better affinities for the pathogen.

This process of detection and destruction typically takes a few days to weeks to eradicate the pathogen from the body. Other mechanisms prevent the immune system (in most cases) from mistakenly attacking the body's own molecules. For more details on affinity maturation and other immune system mechanisms, see Hofmeyr (2001).

Diffuse Feedback

Because the immune system itself, in the process of attacking pathogens, causes inflammation and other harm to tissue, Segel (2001b) and Orosz (2001) have argued that the

immune system must balance an impetus to kill foreign pathogens with the obligation to prevent harm to the body as much as possible. Segel has speculated that this balance requires an additional level of "self-watching". He hypothesizes that specialized signals in the form of molecules are created by three types of situations: harm to the body due to pathogens, the immune system's successful killing of pathogens, and self-harm to the body due to immune system actions. These signals communicate such harm or success to other immune system cells. The "pathogen-harm" and "successful-kill" molecules should up-regulate immune system activity in the region in which they are produced, while the "self-harm" molecules should down-regulate such activity. Up-regulation might mean, for example, speeding up the cell-division process in activated lymphocytes; down-regulation might mean slowing it down.

While Segel's hypothesized signals have not been definitively identified, he proposes plausible candidates. For example, it is known that trisulfated disaccharide molecules are produced by inflamed tissue—a sign of self-harm—and that they also down-regulate the immune response in the neighborhood of such tissue.

How is such up- or down-regulation achieved? Segel's hypothesized signaling molecules would affect the concentration levels of *cytokines*—chemicals that mediate all immune system functions, and can have the desired up-or-down-regulation effects. Cytokines are known to form a signaling network: all immune cells have receptors for cytokines and are affected by them. Cytokines are currently the object of intense study in immunology: "More than 10,000 papers a year are currently being written [concerning cytokines], yet 'practically nothing is known about the behaviour of the network as a whole'" (Segel 2001a). Segel proposes that it is the spatial-temporal concentration of cytokines that provides a "diffuse feedback" to the immune system and cause it to up- or down-regulate by affecting the actions of individual lymphocytes at different spatial locations. Diffuse feedback is a statistical effect of spatially heterogeneous concentrations of cytokines that are sampled locally by lymphocytes.

Immune System Summary

The sections above described two levels of feedback mechanisms:

- Lymphocytes sample their local environment of pathogens by forming bonds to molecules indicating a pathogen's presence. The stronger the affinity, the more daughter lymphocytes are created, in a Darwinian cycle that leads to greater and greater affinity.
- It is speculated that molecules are produced that indicate harm to the body from pathogens, successful killing of pathogens, and harm to the body from immune-system actions. These molecules affect local cytokine distributions that in turn serve to up- or down-regulate immune activity as appropriate.

These are two examples of a set of possible complex feedback mechanisms that together create a distributed self-awareness in the immune system. This self-awareness allows the immune system to generate and use the right lym-

phocytes at the right place at the right time in order to benefit the entire system, in a continually changing environment of pathogens and in the face of multiple, often conflicting requirements (e.g., kill pathogens, but do not harm the body).

Foraging and Task Allocation in Ant Colonies

Analogies have often been made between ant colonies and the brain (e.g., Hofstadter 1979, "...Ant Fugue"); both are composed of relatively simple elements (neurons, ants) from which emerge larger-scale, intelligent and adaptive behavior that appears to exhibit self-awareness. Two examples of such behavior in ant colonies are the ability to optimally and adaptively allocate resources (ants) in foraging for food, and the ability to adaptively allocate ants to different tasks as needed by the colony. Both types of behavior are accomplished with no central control, via mechanisms similar to those described above for the immune system.

Foraging

In many ant species, foraging works roughly as follows (Bonabeau, Dorigo, & Theraulaz 1999). Foraging ants in a colony set out moving randomly in different directions. When an ant encounters a food source, it returns to the nest, leaving a pheromone trail. When other ants encounter a pheromone trail, they are likely to follow it. The greater the concentration of pheromone, the more likely an ant will be to follow the trail. If an ant encounters the food source, it returns to the nest, reinforcing the trail. In the absence of reinforcement, a pheromone trail will dissipate. In this way, ants collectively build up and communicate information about the locations and quality of different food sources, and this information adapts to changes in these environmental conditions. At any given time, the existing trails and their strengths form a good model of the food environment discovered collectively by the foragers.

Task Allocation

Task allocation is another way in which an ant colony regulates its own behavior in a decentralized way. Gordon (2002) has studied task allocation in colonies of Red Harvester ants. Workers in these colonies divide themselves among four types of tasks: foraging, nest maintenance, patrolling, and midden (refuse sorting) work. The numbers of workers pursuing each type of task adapts to changes in the environment. Gordon found, for example, that if the nest is disturbed in some small way, the number of nest maintenance workers will increase. Likewise, if the food supply in the neighborhood is large and high quality, the number of foragers will increase. How does an individual ant decide which task to adopt in response to nest-wide environmental conditions, even though no ant directs the decision of any other ant and each ant only interacts with a small number of other ants?

The answer seems to be that ants decide to switch tasks both as a function of what they encounter in the environment and as a function of their rate of interaction with ants performing different tasks. For example, an inactive ant—one not currently performing a task—who encounters a foreign object near the nest has increased probability of doing

nest-maintenance work. In addition, an inactive ant that encounters many nest-maintenance workers entering and leaving the nest will also have an increased probability of adopting the nest-maintenance task; the increased activity in some way signals that there are important nest maintenance tasks to be done. Similarly, a nest-maintenance worker who encounters many foragers returning to the nest carrying seeds will have an increased probability of switching to foraging; the increased seed delivery signals in some way that a quality food source has been found and needs to be exploited. Ants are apparently able to sense, through direct contact of their antennae with other ants, what task the other ants have been engaged in, by perceiving specific chemical residues associated with each task.

Gordon points out that the proposed mechanism of task switching based on interaction rate can explain an initially puzzling finding of her experiments. She found that perturbing a nest by placing toothpicks around it increased the number of nest maintenance workers, who proceeded to assist in removing the toothpicks. However, this increase was more reliably seen in older colonies (ones that have produced more generations of ants) than in younger colonies. The individual ants in each type of colony were the same age, and ants presumably do not have the ability to pass on complex information about tasks and task switching to their offspring. But it does turn out that older colonies are larger than younger colonies, and in larger colonies ants relying on individual interactions for information would be able to obtain better statistics about the proportion of ants currently doing each task. Thus one would expect a stronger effect in larger colonies, which is what Gordon observed.

Ant Colony Summary

Similar to the discussion of feedback in the immune system, the sections above describe two kinds of feedback mechanisms:

- Foraging ants sample their local environment, responding to food or pheromone signals from other ants. If a food source is found, a forager creates or reinforces the pheromone trail leading to the source. If a trail is found, a forager will have follow it with probability related to its strength.
- Ants sample the local environment of other ants, adopting new tasks according to information gained about what other ants are currently focusing on.

As in the immune system, these examples illustrate how global information about the current state of the colony is represented in a distributed way over a population of ants, and how that information is used to direct the behavior of the individual ants to the benefit of the entire colony.

What is Really Meant by "Self-Awareness" in These Systems?

Some readers may disagree with my use of the term "self-awareness" to describe the kinds of adaptive feedback mechanisms sketched above. A devil's advocate might argue that "self-regulation via multiple feedback systems" would be a

better description, since the immune system and ant colonies are not *really* conscious in the way the brain is. Perhaps this is true (and I think it probably is) but then the devil's advocate must explain, precisely, what the difference is. In the first paragraph of this paper I defined self-awareness as information contained in a system about its global state that feeds back to adaptively control the system's low-level components. Given this definition, both the immune system and ant colonies have some degree of self-awareness. It is not clear where, exactly, the "self" is located in these systems, but neither is it clear for the brain. If there is something more to self-awareness in the brain, it needs to be elucidated.

Some of the most useful discussions I have seen along these lines are in the works of Hofstadter and Dennett (Hofstadter 1979; Hofstadter & Dennett 1981; Dennett 1991). Obviously, these issues are worth exploring for anyone interested in metacognition.

Four Principles of Self-Awareness and Control in Decentralized Systems

The following is a list of four principles abstracted from the above sketches of immune systems and ant colonies. While the principles are given here in a serial list, in reality they overlap and interact. I believe that these principles carry over to intelligence, self-awareness, and self-control in other decentralized systems. Broadening their reach—and their number—is the subject of future work.

1. Global information is encoded as statistics and dynamics of patterns over the system's components.

In the immune system, for example, the spatial distribution and temporal dynamics of lymphocytes can be interpreted as a representation of the continually changing population of pathogens in the body. Similarly, according to Segel's speculations, the spatial distribution and dynamics of cytokine concentrations encode global information about the immune system's success in killing pathogens and avoiding harm to the self.

In ant colonies, the colony's food environment is represented, in a dynamic way, by the statistical distribution of ants on various trails, and the colony's overall state is represented by the dynamic distribution of ants performing different tasks.

But *who* actually interprets this representation, and performs the corresponding adaptations of the system? It is the collection of cells or ants themselves, taking actions depending on statistics they sample locally, via other cells or ants they encounter, or via the local chemical environment. In order for statistics to be reliably interpreted, actions are taken based on thresholds: a lymphocyte undergoes affinity maturation only if its receptors have a sufficient number of strong bonds, plus additional verification signals from other cells. A nest-maintenance ant switches to foraging only if it has encountered a sufficient number of other successful foragers.

2. Randomness and probabilities are essential.

Both immune systems and ant colonies use randomness and probabilities in essential ways. The receptor shape

of each individual lymphocyte has a randomly generated component, so as to allow coverage by the population of many possible pathogenic shapes. The spatial distribution of lymphocytes has a random component, due to the distribution of lymphocytes by the blood stream, so as to allow coverage of many possible spatial distributions of pathogens. The detailed thresholds for activation of lymphocytes, their actual division rates, and the mutations produced in the offspring all involve random or noisy aspects.

Similarly, the movement of ant foragers has random components, and these foragers are attracted to pheromone trails in a probabilistic way. Ants also task-switch in a probabilistic manner.

It appears that such intrinsic random and probabilistic elements are needed in order for a comparatively small population of simple components (ants or cells) to explore an enormously larger space of possibilities, particularly when the information to be gained is statistical in nature and there is little a priori knowledge about what will be encountered.

3. The system carries out a fine-grained, parallel search of possibilities.

Many, if not all, complex systems in biology have a fine-grained architecture, in that they consist of large numbers of relatively simple elements that work together in a highly parallel fashion.

Several possible advantages are conferred by this type of architecture, including robustness, efficiency, and evolvability. One additional major advantage is that a fine-grained parallel system is able to carry out what Hofstadter (1995) has called a "parallel terraced scan". This refers to a simultaneous exploration of many possibilities or pathways, in which the resources given to each exploration at a given time depends on the perceived success of that exploration at that time. The search is parallel in that many different possibilities are explored, but is "terraced" in that not all possibilities are explored at the same speeds or to the same depth. Exploration is "on-line" in the sense that information is used as it is gained to continually reassess what is important to explore.

The fine-grained nature of the system not only allows many different paths to be explored, but it also allows the system to continually change its exploration paths, since only relatively simple micro-actions are taken at any time. Employing more coarse-grained actions would involve committing time to a particular exploration that might turn out not to be warranted. In this way, the fine-grained nature of exploration allows the system to fluidly and continuously adapt its exploration to the information it obtains.

For example, at any given time, the immune system must determine which regions of the huge space of possible pathogen shapes should be explored by lymphocytes. Each of the trillions of lymphocytes in the body at any given time can be seen as a particular mini-exploration of a range of shapes. The shape ranges that are most successful (i.e., bind strongly to pathogens) are given more

exploration resources, in the form of mutated offspring lymphocytes, than those shape ranges that do not pan out (i.e., lymphocytes that do not bond strongly). However, while exploiting the information that has been obtained, via the process of affinity maturation, the immune system continues at all times to generate new lymphocytes that explore completely novel shape ranges. Thus the system is able to focus on the most promising possibilities seen so far, exploring these most quickly and to the most depth, while never neglecting to explore new possibilities. This is closely related to Holland's general principles for optimally balancing "exploitation" and "exploration" in adaptive systems (Holland 1992). Orosz's (2001) principle of "dynamic engagement" for the immune system also requires a fine-grained architecture.

Similarly, ant foraging uses a parallel-terraced-scan strategy: many ants initially explore random directions for food. If food is discovered in any of these directions, more of the system's resources (ants) are allocated, via the feedback mechanisms described above, to explore those directions further. At all times, paths are dynamically allocated exploration resources in proportion to their promise (the amount and quality of the food that has been discovered at those locations). However, due to the large number of ants and their intrinsic random elements, unpromising paths continue to be explored as well, though with many fewer resources.

Note that the redundancy inherent in fine-grained systems such as the immune system and ant colonies allows the system to work well even when the information available is only statistical in nature. Redundancy allows many independent samples of information to be made, and allows fine-grained actions to be consequential only when taken by large numbers of components.

4. The system exhibits a continual interplay of bottom-up and top-down processes.

The notions of "bottom-up" and "top-down" processes are often used to describe aspects of cognition: for example, perception is an interplay between what is perceived "bottom-up" by the retina (and other early vision systems) and the "top-down" expectations derived from concepts activated by these bottom-up sensations.

A similar interplay can be found in the immune system: bottom-up explorations are carried out by a continual patrol of lymphocytes with different receptors, collectively prepared to approximately match any pathogen (what Segel, 2001c, has called "a broad spectrum initial response"), while top-down explorations consist of focused lymphocytes, which, when activated by a match, create offspring that are variations of themselves in order to zero in on a particular pathogen shape.

Likewise ant foraging consists of bottom-up explorations by ants moving at random, looking for food in any direction, and top-down explorations in which ants follow existing pheromone trails.

As in all adaptive systems, maintaining a correct balance between these two modes of exploration is essential. Indeed, the optimal balance shifts over time. Early explo-

rations, based on little or no information, are largely random, unfocused, and bottom-up. As information is obtained and acted on, exploration gradually becomes more deterministic, focused, and top-down, in response to what has been perceived by the system. Exhibiting such an appropriate shift in exploration mode has been hypothesized to be a general property of adaptive and intelligent systems (Holland 1992; Hofstadter & Mitchell 1994).

Implications for Artificial Intelligence

The four principles listed above, along with other general principles abstracted from the study of decentralized complex adaptive systems, can be a guide in designing artificial intelligence systems with decentralized architectures that have sophisticated abilities for pattern perception and self-awareness.

In fact, these principles guided the design of the Copycat system, developed by Douglas Hofstadter and myself (Hofstadter & Mitchell 1994; Mitchell 2001), and its successor, Metacat, developed by Hofstadter and James Marshall (Marshall & Hofstadter 1998; Marshall 2002).

Copycat and Metacat are programs that perceive patterns and make analogies in the domain of letter strings, such as "If **abc** changes to **abd**, what is the analogous change to **ijjkk**?" The purpose of both these projects was to develop general mechanisms of high-level perception, analogy-making, and self-watching that would be extensible to many domains. Metacat, in particular, is able to monitor its own processing so as to improve its perceptual abilities, and is considerably more sophisticated than Copycat.

Limitations on space prevent a detailed description of Copycat and Metacat here; interested readers should consult the references given above. However, the philosophy behind the programs can be summarized by the following principles for modeling perception, which closely follow the principles abstracted above.

- (1) The perceptual process must be fine-grained, diverse, redundant, and decentralized.
- (2) Perception is guided by "fluid" concepts, which are themselves shaped as the perceptual process unfolds.
- (3) The perceptual process proceeds as an interplay of bottom-up modes (driven by stimuli from the environment) and top-down modes (driven by expectation, prior knowledge, biases, and what has already been discovered). This interplay is not preprogrammed, but is an emergent effect of the collective actions of low-level components of the system.
- (4) The perceptual process shifts over time from being highly parallel, random, and bottom-up, to being more focused, deterministic, and top-down. As in (3), this shift is not pre-programmed, but rather is an emergent effect of collective behavior in the system.
- (5) The perceptual process must have a means of "self-watching"—monitoring its own state and progress—that feeds back to affect behavior. In Copycat, this is implemented by a computational "temperature". This is similar to the ideas proposed by Segel described above concerning self-watching in the immune system, as well as ideas of Orosz concerning "variable connectivity" (Orosz 2001). In

Metacat, self-watching is a more explicit and sophisticated part of the system.

I believe that the work on Copycat and Metacat is the first time these principles have been stated together in this general format, and actually tested in computer models.

While Copycat and Metacat are limited to perceiving analogies in a simple domain, they are based on quite general principles. I believe that similar principles, gleaned from mechanisms that produce adaptive behavior and self-awareness in natural systems, will be necessary components of future artificial intelligence systems with sophisticated metacognitive abilities.

Conclusion

In this paper I have sketched some ideas about self-awareness and control in the immune system and in ant colonies, attempted to abstract some general principles, and outlined the possible importance of these principles for designing artificial intelligence systems with decentralized architectures. This discussion, of course, has only scratched the surface of some very difficult and complex issues.

Much has been written about how information in the brain, including information about its own state, is represented in a distributed, statistical, and dynamical fashion. One of the most important research challenges for cognitive science—and more generally, the sciences of complex systems—is to elucidate common principles of such representation of information and its self-interpretation more broadly across decentralized systems.

Acknowledgments This work has been supported by a grant from the J. S. McDonnell Foundation. My thinking on these issues has been greatly influenced by many people, in particular, Douglas Hofstadter, John Holland, Stephanie Forrest, Daniel Dennett, and Lee Segel.

References

- Bonabeau, E.; Dorigo, M.; and Theraulaz, G. 1999. *Swarm Intelligence*. New York: Oxford University Press. Chapter 2, pp. 25-31.
- Cohen, I. R. 1999. *Tending Adam's Garden: Evolving the Cognitive Immune System*. San Diego, CA: Academic Press.
- Dennett, D. C. 1991. *Consciousness Explained*. Boston: Little, Brown and Co.
- Forrest, S., and Hofmeyr, S. A. 2001. Immunology as information processing. In Segel, L. A., and Cohen, I. R., eds., *Design Principles for the Immune System and Other Distributed Autonomous Systems*, 361–387. New York: Oxford University Press.
- Gordon, D. M. 2002. The regulation of foraging activity in Red Harvester ant colonies. *The American Naturalist* 159(5):509–518.
- Hofmeyr, S. A. 2001. An interpretive introduction to the immune system. In Segel, L. A., and Cohen, I. R., eds., *Design Principles for the Immune System and Other Distributed Autonomous Systems*, 3–26. New York: Oxford University Press.
- Hofstadter, D. R., and Dennett, D. C. 1981. *The Mind's I: Fantasies and Reflections on Self and Soul*. New York: Basic Books.
- Hofstadter, D. R., and Mitchell, M. 1994. The Copycat project: A model of mental fluidity and analogy-making. In Holyoak, K., and Barnden, J., eds., *Advances in Connectionist and Neural Computation Theory, Volume 2: Analogical Connections*. Norwood, NJ: Ablex. 31–112.
- Hofstadter, D. R. 1979. *Gödel, Escher, Bach: An Eternal Golden Braid*. New York: Basic Books.
- Hofstadter, D. 1995. *Fluid concepts and creative analogies*. New York: Basic Books. chapter 2, 97–126.
- Holland, J. H. 1992. *Adaptation in Natural and Artificial Systems*. Cambridge, MA: MIT Press. First edition, 1975.
- Marshall, J., and Hofstadter, D. 1998. Making sense of analogies in Metacat. In Holyoak, K.; Gentner, D.; and Kokinov, B., eds., *Advances in Analogy Research: Integration of Theory and Data from the Cognitive, Computational, and Neural Sciences*. Sofia, Bulgaria: New Bulgarian University.
- Marshall, J. 2002. Metacat: A self-watching cognitive architecture for analogy-making. In *Proceedings of the 24th Annual Conference of the Cognitive Science Society*, 631–636. Mahwah, NJ: Lawrence Erlbaum Associates.
- Mitchell, M. 2001. Analogy-making as a complex adaptive system. In Segel, L. A., and Cohen, I. R., eds., *Design Principles for the Immune System and Other Distributed Autonomous Systems*. New York: Oxford University Press. 335–359.
- Orosz, C. G. 2001. An introduction to immuno-ecology and immuno-informatics. In Segel, L. A., and Cohen, I. R., eds., *Design Principles for the Immune System and Other Distributed Autonomous Systems*. New York: Oxford University Press. 125–149.
- Segel, L. A. 2001a. Controlling the immune system: Diffuse feedback via a diffuse informational network. In Bock, G., and Goode, J., eds., *Complexity in Biological Information Processing*. Chichester; New York: John Wiley and Sons. 31–44. Novartis Foundation Symposium.
- Segel, L. A. 2001b. Diffuse feedback from a diffuse informational network in the immune system and other distributed autonomous systems. In Segel, L. A., and Cohen, I. R., eds., *Design Principles for the Immune System and Other Distributed Autonomous Systems*. New York: Oxford University Press. 203–226.
- Segel, L. A. 2001c. How can perception of context improve the immune response? In Steinman, L., ed., *Autoimmunity and Emerging Diseases*. Jerusalem: Center for the Study of Emerging Diseases.
- Segel, L. A. 2001d. How does the immune system see to it that it is doing a good job? *Graft* 4(6):15–18.