


Inspired by immunity

By developing programs that mimic some of the functions of the immune system, computer scientists are tackling problems from fighting fraud to controlling robots. Erica Klarreich investigates.

Computer science has a healthy tradition of stealing nature's good ideas. Programmers have realized that evolution can be simulated using 'genetic algorithms', which drive a computer system towards a problem's most effective solution. The brain has inspired 'neural network' programs that are the basis of many attempts to develop artificial intelligence. But a third biological system — immunity — has until recently attracted little attention.

Computer scientists are now addressing this neglect. In the past few years, they have designed immune-inspired algorithms that might find a wide range of uses, such as detecting fraudulent financial transactions, controlling robots, and even distinguishing cancerous tumours from benign ones.

Our immune system has many desirable attributes. It is versatile and efficient, coping quickly with a diverse array of possible threats. Immunologists estimate that it can respond to



Programmed for success: the functions of white blood cells can easily be translated into algorithms. This montage shows T-cells attacking a large cancer cell (red), superimposed over a computer chip.

10,000 trillion different molecular signatures, or antigens. It is both decentralized and tolerant of errors, so that a few malfunctioning cells, or the loss of part of the system, will not lead to catastrophe. It can also learn to recognize the antigens of specific pathogens and remember them for the future.

Crucially for computer scientists, many of its functions can be translated easily into simple algorithms. "The immune system is rich with inspiration that is mostly untapped at this point," says Jonathan Timmis of the University of Kent in Canterbury, UK, who has devised programs that apply attributes of immunity to pattern-recognition problems.

Although the field is new, several immune-inspired programs are being readied for commercial use. Company 51, a software firm in Redwood City, California, plans to start marketing an immune-inspired system to protect computer networks from hackers and destructive 'worm' programs by the middle of this year. By 2003, Britain's post office, Consignia, could be using an immune-based fraud-detection system to sniff out suspicious transactions.

Fraud detection and computer security are natural candidates for applying the immune system's principles. "There's a pretty direct analogy between the problem of security and the problem the immune system faces in the body," says Stephanie Forrest of the University of New Mexico in Albuquerque, who was one of the first researchers to consider applying the attributes of immunity to computer security in the early 1990s.

Unlike conventional security systems, in which experts try to predict the kinds of

breach that may occur, the immune system starts with little information about the nature of potential invaders. Its ability to respond from a position of ignorance is one of its main attractions for computer-security experts, as the precise forms of computer viruses, hacker attacks or fraud can be hard to predict, and are often only subtly different from benign programs or behaviour.

"Patterns of fraud are often difficult to spot, as a clever fraudster tries to hide his activity, perhaps posing as an employee," says Peter Bentley, who studies immune-inspired security systems at University College London. He notes that pathogenic viruses behave in a similar way, often carrying proteins that mimic those in our own body. "Our immune system has solved this problem, so we'd like to use its processes to solve the problem in fraud detection," he says.

Production line

The immune system constantly reinvents itself, producing 10 million new lymphocytes — white blood cells — every day. The acquired immune system, the part that builds up tailored responses to specific pathogens, has evolved an array of techniques to make these lymphocytes home in on dangerous pathogens. Each lymphocyte either carries receptor proteins that can bind to specific antigens, or produces antibodies with similar specificity.

As lymphocytes are produced, a process of random genetic shuffling results in the receptors or antibodies being put together in subtly different ways, generating a diverse collection of cells that are each specific to a particular



Stephanie Forrest and Peter Bentley draw on the immune system's tricks to fight computer fraud.

antigen. Some receptors may bind to the body's own antigens, so to avoid sending out cells that will trigger such 'autoimmune' attacks, the cells undergo a process called negative selection. In this procedure, each newly made cell is matched up against the body's own molecular signatures, and any lymphocytes that react with them are either destroyed or their activity is suppressed.

The remaining lymphocytes course through the body. Those that fail to interact with an antigen die after a few days. But cells that do interact are stimulated, triggering an immune response. They immediately begin dividing, producing more cells that can detect the same antigen, in a process called clonal selection.

In the case of lymphocytes that produce antibodies, there is a second step which refines the immune response. As cells divide after being stimulated, random mutations alter the antibodies produced. Some of these antibodies will bind to the antigen more effectively — the cells that produce these antibodies subsequently divide more rapidly, and the response becomes increasingly precise.

Lifelong grudge


Once the immune system has learned to recognize a particular pathogen, this knowledge is not lost. Some of the cells become 'memory' cells, surviving for long periods and becoming reactivated in response to subsequent attacks by the same pathogen — this is how vaccination works.

Security experts are now trying to mimic these processes. In a typical immune-inspired system, transactions are represented as strings of data, and fraud detectors — equivalent to the immune system's receptors or antibodies — are other strings of data that are compared to entries in the transaction database, or rules that decide whether entries fall into a particular category. One rule, for instance, might recognize any transactions that occur between 2.00 and 2.30 am and involve \$10,000–15,000. Another might be stimulated by transactions in which the buyer enters a London address but a postal code from outside the city.

To build an army of detectors that ignore



Scott Thayer (left) and Surya Singh hope to extend their computer models to real robotics.



This 'minesweeping' model — blue dots are robots and red dots represent mines — uses similar principles to the innate immune system.

normal behaviour but catch fraudulent transactions, the system starts by putting together the components of its rules randomly, just as the immune system creates receptors or antibodies by random genetic shuffling. To weed out detectors that are stimulated by normal transactions, the system then sends the new detectors to 'live' for a period of time in a database of transactions that are known to be legitimate; any detector that recognizes one of these is thrown away.

Once the detectors that recognize normal transactions are removed, the remaining detectors are sent out to monitor the system. Just as lymphocytes that are not stimulated by an antigen die within days, detectors that do not encounter matching transactions are eventually killed off, to make way for rules that seem to match more closely the types of fraud that actually arise.

In a process that is analogous to clonal selection and the refinement of antibody responses, detectors that are alerted by unusual transactions make copies of themselves, adding tiny random mutations that could potentially improve their detective ability. Successful detectors remain in the system as 'memory cells', so that if the same kind of fraud is detected again, the system will respond swiftly. In this way, it learns to recognize common patterns of fraud. If, for example, a certain fraudster typically makes transactions late at night, the system will evolve a large number of detectors that are stimulated by late-night transactions.

Most immune-inspired security systems are similar in basic outline, but all incorporate subtly different attributes of the immune system. "There are just so many wonderful things we've got here," says Richard Overill of King's College London, who leads a team that is working to detect fraudulent post-office transactions. Overill's system, which his team began testing with Consignia this month, concentrates on clonal selection and immunological memory. "Artificial immune systems have the potential to be extremely smart and flexible," he says.

Other researchers are focusing not on the immune system's ability to identify abnormal behaviour, but on its skill in coordinating huge numbers of autonomous cells to create complex responses without a central organizing body. The immune system does this chiefly through local interactions and chemical signals, largely mediated by proteins secreted from its cells. "The immune system has no boss," says Lee Segel, who models immune responses at the Weizmann Institute of Science in Rehovot, Israel. "It has a trillion cells all running around, with no one telling them what to do, but still doing a pretty good job."

This attribute makes the system an ideal model for robotics projects that involve thousands or even millions of small, autonomous robots, argue Scott Thayer and Surya Singh of the Robotics Institute at Carnegie Mellon University in Pittsburgh, Pennsylvania. They are designing robotic systems to clear minefields, carry out search-and-rescue missions, or chart unexplored territory. So far, their work involves computer simulations, rather than real robots — the technology to produce thousands of cheap robots is still under development.

Rapid response

Unlike the fraud-detection researchers, who have focused on acquired immunity, Thayer and Singh have based their program on the 'innate' system, which springs into action at the first moment of an infection. Innate immunity is triggered by receptors that can only recognize pathogens in broad classes, but which respond faster than the more specific receptors and antibodies of the acquired immune system. Usually, they trigger a local inflammatory response, and the lymphocytes involved send out chemical signals to summon other cells.

In Thayer and Singh's minesweeping simulation, virtual robots initially patrol the minefield randomly. When one discovers an area with a concentration of mines, it broadcasts a signal to the other robots, some of which migrate to the area. In the immune system, the strength of secreted chemical signals decreases with distance. To mimic this effect, the virtual robots assign less importance to signals that emanate from far away. This ensures that only those in the vicinity of



Risk-reducing robots: could computerized minesweepers one day replace military experts?

▶ the signalling robot are recruited, while far-away robots continue to monitor other parts of the minefield.

This simple mechanism easily outperforms traditional 'raster' scans, which divide a search area into a grid and comb it cell-by-cell, row-by-row. Thayer and Singh's model is particularly effective for 'moving' minefields, in which the mines can jump, foiling raster scans by moving into previously cleared areas.

Thayer and Singh are working to incorporate the acquired immune system's ability to learn into their program. Although robots cannot copy themselves, if a robot finds itself to be especially well adapted to defusing a particular type of mine, it could broadcast the details of its program to nearby robots, duplicating its successful software, rather than its hardware.

Mutant minesweepers

The researchers also want to include mutation, allowing robots to change randomly the way they carry out operations, or the relative importance they assign to different signals; successful mutations could then be broadcast to others. For instance, a mutation might make a robot check a mine twice before summoning its neighbours. "As with the immune system, because there are so many robots, you can afford to risk two or three in trying out new behaviour," says Singh.

Computer scientists are even appropriating outdated immunological theories. In 1974, Niels Jerne of the Basel Institute for Immunology in Switzerland, who later won a Nobel prize for his theoretical work, introduced the idea of 'immune networks' to try to explain how the immune system maintains its memory of an antigen. Jerne proposed that certain antibodies, termed 'anti-antibodies', bind to other antibodies, while 'anti-anti-antibodies' recognize these, and so on. All of these antibodies, Jerne argued, suppress or stimulate one another in a way that is held in equilibrium until the original antigen comes along again and disturbs the network, triggering an immune response.

The theory has since fallen from favour — although antibodies that recognize other antibodies do exist, there is no strong evidence that they function in self-regulating immune networks as Jerne proposed. Nevertheless, computer scientists have taken inspiration from the idea to develop systems for discovering patterns in vast amounts of data. "From a computer-science point of view it's a lovely theory," says Timmis.

He is developing a pattern-recognition program that could carry out tasks such as finding relationships in people's Internet shopping habits or sifting through epidemiological data to identify risk factors for disease. His data-mining program uses virtual antibodies that cluster into networks, much as Jerne suggested. The program takes a small, random sample of a data set as its

starting 'antibody' population; the full data set is the collection of 'antigens'. If the program is looking for patterns in Internet shopping habits, each antibody and antigen might be a consumer profile that consists of a list of, say, 100 products together with details of whether or not the consumer bought each.

First, the antibodies are compared to each other, and ones that match closely are linked together to form a network. Two different profiles might be considered to match if they agree in, say, at least 50 of the entries. This matching threshold is set before the program starts running, and the level chosen affects the strength of the patterns found.

Once the antibodies have been linked into a network, they are compared to the antigens in the full data set. Antibodies that are not similar to enough antigens are removed, whereas antibodies that match many antigens make copies of themselves, introducing small mutations as they replicate. These mutated copies are then incorporated into the network, linked to similar antibodies. The entire process is repeated many times so that the network evolves clusters, each of which represents some pattern in the original data set.

If, for example, Internet shoppers who buy books online also tend to buy music, the original set of antibodies might contain a few consumer profiles that include both books and music. As these profiles are similar to each other, they will be linked together in the network, and as they are also similar to many profiles in the full data set of antigens, every time they are compared to the antigens they will be stimulated to make mutated copies of themselves. These copies are then linked in the network to the original profiles that include books and music, and gradually a cluster of antibodies forms. "Often the clusters will find obvious patterns that you knew already, but sometimes you get real nuggets of information," says Timmis.


Breast test

Timmis has tested his algorithm on the Wisconsin Breast Cancer Database, which contains patient profiles and data from their breast-tumour biopsies — documenting the sizes of clumps of cells, uniformity of cell sizes and shapes, and so on. The program learned to distinguish between samples from malignant and benign tumours.

Immune networks may prove to be more flexible learning tools than brain-inspired neural networks, which learn patterns by changing the strength of connections between simulated neurons. "With a standard neural network you have to define the network structure first, and the network is only as good as its initial design," says Timmis. By contrast, immune networks allow their components to change, as well as the connections between individual components, giving added plasticity.



Jonathan Timmis's 'immune network' program identifies patterns in huge data sets by evolving clusters of successful antibodies. This enables it, for example, to distinguish tumour types (below).



Researchers working in the field expect to garner many more computational insights from mimicking the immune system — particularly given that our biological understanding is advancing all the time. Forrest notes that immunologists have recently become very interested in the innate system. "We modellers are far behind," she says.

Some biologists argue that computer scientists may also provide insights for immunologists. "I think it would be beneficial to see more interplay," says Alan Perelson, a mathematical biologist who studies problems in immunology at the Los Alamos National Laboratory in New Mexico. "If the computer-science people have to go outside of 'normal' immunology, they might find something we've missed."

In the meantime, computer scientists have plenty of work to do in exploiting well-understood attributes of the immune system. "Nature has been perfecting this for so many years that even if we get it just one-tenth right, it's often better than what we could come up with on our own," says Bentley. ■

Erica Klarreich has just finished an internship with Nature.

Fraud detection

♦ www.dti-mi.org.uk/newweb/cifd.html

♦ www.icsa.ac.uk/Projects/link.html

Robotics

♦ www.ri.cmu.edu/pubs/pub_3845.html

Pattern recognition/data mining

♦ www.cs.ukc.ac.uk/people/staff/jt6

JONATHAN TIMMIS