detects cancer more accurately than existing ones, but also acts more like an intelligent colleague than a black box. Instead of relying on feature extraction, Dr Tourassi's technique works by comparing the images taken by a radiologist with a large collection of normal and abnormal mammograms held in a database. This still requires algorithms, but of a different type. She and her colleagues have developed template-matching algorithms to compare the intensity and distribution of pixels in different images. They have also created decision algorithms to determine, after it has been compared with the entire database, whether a region of interest on a mammogram is normal or cancerous.

The diagnosis of the cases in the database has been confirmed, either by biopsy or by long-term follow-up, so there is no doubt about their details. If a new mammogram is similar to known cases of breast cancer, this would give reason for suspicion. This is exactly how a radiologist relates a case to those he saw in the past.

Making a clean breast

Dr Tourassi has found that her system can reliably distinguish tumour masses from normal tissues, and has a lower rate of false positives than systems based on feature extraction. Also, crucially, it can explain to a radiologist how it reached its decision by showing him similar mammograms in the database. The radiologist is then in a better position to decide whether the computer's judgment is valid.

The knowledge-based system has another bonus, too. As mammograms of new cancer cases are added to the database it is looking at, it will become cleverer—just as radiologists and physicians become more experienced and skilful as they come across more patients. This is in contrast to feature-based CAD systems, which cannot adapt to new cases unless their algorithms are suitably modified. However, there is a potential problem in the long run. The knowledge-based CAD system has a huge demand for computing power and, as the database grows, it will get slower and less efficient.

Dr Tourassi, however, has been thinking about this problem. She suspects that by using only the most informative mammograms, it might be possible to keep the size of the database within reasonable limits. To decide which ones to select, she turned to a branch of science called information theory. This theory says that the amount of information in a system can be measured in terms of its entropy.

Entropy is actually a measure of disorder. In the context of image processing, it is an indication of the complexity of an image. For example, an image that is all black or all white has zero entropy; an image of a chessboard, which contains an equal number of regularly distributed light and dark pixels, has low entropy; images with more varied distributions of many intensity levels of pixel have high entropy and are considered more informative.

Entropy is the basis of a standard indexing strategy for images and is often calculated automatically when they are put into a database. It is therefore easy to rank a collection of images by their entropy. In a pilot study, Dr Tourassi and her colleagues tested whether a subset of high-entropy mammograms would work as effectively as using the whole database.

Their entire database contained 2,318 mammograms from the Digital Database of Screening Mammography collected at the University of South Florida, but they varied the size and content of the collection of images examined by their algorithms, by either randomly selecting a subset of mammograms or choosing those with high entropy. They measured the utility of the different sets of images by taking each mammogram within it in turn, and testing it against all the others.

Their conclusion was that testing against only the 600 most informative mammograms was as effective as using all 2,318 images. And the system took less than three seconds per query—far faster than a radiologist could manage. A mammogram-reading machine that not only mimics but surpasses human perception and can explain its diagnosis would truly be a girl's best friend. ■

Evolution

Stooping to conquer

How to become multicellular

VIEWED from humanity's lofty heights, single-celled creatures are the scum of the earth. In reality, though, almost all living things are unicellular, and with good reason. Multicellularity requires most of the cells in a body to make the supreme Darwinian sacrifice, by giving up reproducing. This helps the few cells specialised for reproduction to do so more effectively. Since the self-sacrificing cells have the same genes as the specialised reproducers they are, in effect, reproducing collaterally. But it is still a hard trick to pull off, and it has not happened often.

One creature that has managed the trick—separately from plants, animals and fungi, who are the real experts in the field—is an alga called *Volvox carteri*. An adult *Volvox* consists of around 2,000 body cells, whose job is to move the organism around using their flagella, and 16 cells capable of reproducing.

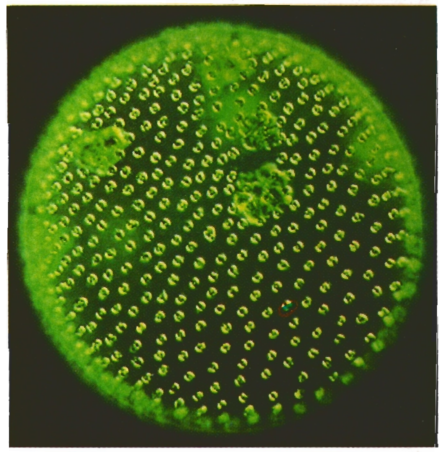

A perfect example of division of labour

Clearly, *Volvox* evolved from similar algae that exist only as single cells, but until now the genetics of the process have been obscure. However, as they report in *Molecular Biology and Evolution*, Aurora Nedelcu of the University of New Brunswick and Richard Michod of the University of Arizona think they have worked out what happened. In doing so, they have shed light on the type of process that eventually resulted in human beings.

The gene that stops the body cells of *Volvox* reproducing is called *regA*. It works by suppressing the production of proteins needed to make new chloroplasts in a cell. Chloroplasts are the structures in which photosynthesis happens. Without an adequate supply of them a cell cannot grow big enough to divide.

What Dr Nedelcu and Dr Michod did was to look for an antecedent of *regA* in a single-celled creature, in order to find out what its job was. By searching for genetic sequences similar to *regA* in the burgeoning databases of genes that now exist, they found one in a unicellular alga called *Chlamydomonas reinhardtii*. In *Chlamydomonas*, the gene only gets switched on when environmental conditions are poor—for instance, when sunlight or nutrients are scarce. That keeps *Chlamydomonas* from wasting resources during hard times, increasing its likelihood of reproducing over the long haul.

Apparently, a mutation in this gene—or, more probably, in the control system that activates it—gave rise to *Volvox*. When that happened, individual cells could turn their reproductive capacity on or off according to the function they served in *Volvox*'s body. Though this often doomed an individual cell's own genes, it gave rise to a co-operative organism that successfully passed those genes on indirectly.

What genetic change allowed multicellular animals to come into existence remains obscure. But *Volvox* shows the sort of thing to look for: a gene that stops reproduction in single-celled creatures and has been co-opted to do a new job. ■