

intuitive mixtures to reconstruct objects of interest.

Learning techniques (the core of 'neural network' theory) have been used in various ways to find components of objects or images^{2,3}. Lee and Seung take a so-called unsupervised approach, the goal of which is to find structures — or 'basis functions' — that help to 'explain' a training set, in the sense that training examples can be faithfully reconstructed using appropriate combinations of the discovered basis functions. What distinguishes Lee and Seung's approach from other structure-finding algorithms is the way in which the basis functions must be combined to build target objects — these authors say that the combinations should be exclusively additive. Their approach reflects the idea that objects are most naturally described by the inventory of the parts they contain. Or, to look at things the opposite way, parts are precisely those entities that allow objects to be reassembled using purely additive combinations.

When Lee and Seung asked their NMF algorithm to decompose a set of faces, they found that the resulting basis functions were visibly different — and more part-like — than those found by conventional structure-finding algorithms such as principal components analysis (PCA) or vector quantization (VQ), although see ref. 4. Unlike NMF, PCA decomposes images into conceptually opaque additive and subtractive combinations of basis functions. The VQ technique represents objects using a set of whole-object prototypes, similar to the way in which suitcases are identified at an airline's lost-baggage counter. In other words, PCA finds parts that are not very intuitive and VQ ignores part structure altogether.

What is the importance of Lee and Seung's contribution? Most compelling is the argument that NMF's non-negativity constraint is well matched to our intuitive ideas about decomposition into parts. Second is the fact that the resulting basis functions differ in an interesting way from those generated without the non-negativity constraint. And third is the elegance and simplicity of the algorithm.

But there are caveats too. The authors rely on subjective appraisal — the look and feel — of their results to support the claim that NMF is better at finding the underlying component-based structure of complex objects than, say, PCA. And although Lee and Seung find that the NMF basis functions learned from faces seem to consist of "several versions of mouths, noses and other facial parts", another observer might describe many of these same basis functions as partial renderings of several face parts in unintuitive combinations (see Fig. 1 on page 789). Moreover, a simple threshold applied to PCA or VQ basis functions could considerably reduce the differences between these more

'global' basis functions and the more 'local' NMF ones.

Another question is how NMF basis functions relate to 'parts' as conventionally defined⁵. Given that the basis functions depend only on coactivation of input elements, a blob in a basis function extracted from a set of facial images could arise from the collusion of nose pixels in one image, upper-lip pixels in another and shadow pixels in a third. That is, unrelated entities, or 'causes', could be lumped within the same so-called NMF 'part'. But this display of poor partsmanship is less a commentary on NMF than it is a consequence of the input format (raw pixel values) and simplified representational scheme (weighted sums). Together these variables make it difficult for NMF, or any comparable scheme, to discover the deeper semantic structure that is normally associated with a parts-based representation.

Lee and Seung draw part of the inspiration for their non-negativity constraint from biology — that is, they point out that the firing rates of neurons are never negative, and that synaptic strengths do not change sign. But is this germane? Receptive fields in the brain often come in opponent pairs, where the positive and negative ranges of a variable are represented by the positive firing rates of two separate neurons (for example, 'on'- and 'off'-centre cells). Moreover, individual neurons, which are presumed to be the feature detectors in our own visual systems, typically show one or more kinds of opponency within their receptive fields. So a stimulus that is excitatory in one part of a receptive field may be inhibitory in another. In short, positive and negative quantities seem to be intermixed throughout our perceptual systems.

Despite the caveats, Lee and Seung's work is clear and thought-provoking, bearing on a scientific question of great importance — how to extract meaningful patterns from a complex world. To boot, their NMF procedure will be a practical tool for structure-finding that will probably have a significant effect on future work in the area. ■

Bartlett W. Mel is in the Department of Biomedical Engineering, University of Southern California, MC 1451, Los Angeles, California 90089, USA.
e-mail: mel@inc.usc.edu

1. Lee, D. D. & Seung, H. S. *Nature* **401**, 788–791 (1999).
2. Bishop, C. M. *Neural Networks for Pattern Recognition* (Oxford Univ. Press, 1995).
3. Hinton, G. E. & Sejnowski, T. J. *Unsupervised Learning: Foundations of Neural Computation* (MIT Press, Cambridge, Massachusetts, 1999).
4. Penev, P. S. & Atick, J. J. *Comp. Neural Syst.* **7**, 477–500 (1996).
5. Biederman, I. *Psychol. Rev.* **94**, 115–147 (1987).

erratum Due to an editorial oversight, the picture accompanying the article "Neurobiology: Cognition by a mini brain" (*Nature* **400**, 718–719; 1999), did not carry the correct attribution. It should have been credited to Karlheinz Rein of the Lehrstuhl für Genetik, Biozentrum, Würzburg.

Daedalus

Setting and upsetting

Last week Daedalus presented his 'Wet Cement', a hygroscopic concrete which set but never dried out. A building made from it could creep rather than crack while settling into the soil. He is now taking the idea to extremes. His new, even more advanced cement does not merely creep under load; impact liquefies it completely. The result is the first thixotropic, re-usable concrete.

Anyone who has wiggled a foot in damp beach sand, stirring it into a mobile porridge, will understand the principle. A pile of sand can resist steady compression; but no such pile is fully compact. A few grains take all the load, leaving the rest essentially free and unloaded. Under sudden shock or stirring, the grains shift into a denser packing which occupies a smaller volume. The water around them, however, cannot shrink. The grains find themselves suspended in excess water under strong agitation. They slump into a strengthless fluidized mass.

So DREADCO's Thixocrete contains sand to carry the load, hygroscopic compounds and water-absorbing polymers to keep it wet, and a highly impure, mixed calciferous setting cement. As in ordinary cement, its crystals precipitate from the water to hold the sand grains in place. But the crystals, loaded with impurities and dislocations, are extremely brittle. They shatter to fragments under sudden shock or vibration, freeing the sand grains to fluidize in the water. Now small crystals preferentially dissolve in water whereas bigger ones crystallize from it (this is the 'ripening' action of a chemical precipitate). So when agitation ceases, new crystals slowly precipitate, and knit into a restraining matrix round the sand grains. The Thixocrete slowly sets again.

Thixocrete will transform the building trade. The new product will be made and transported not in slowly rotating mixers but in vigorously stirred tanks. Pumped into moulds or extruded along lines of bricks, it will set in minutes. Any mistake or design error will be easily rectified by a powerful vibrator or wrecker's ball, shocking the structure back to re-usable fluid Thixocrete. By the same token, a Thixocrete structure will be easily demolished. And an earthquake or bomb attack will not shatter it disastrously. Instead it will flow and slump. Falling pieces will be soft and plastic; escapers will wade to safety through them. And the deformed building will not be a hazardous, unstable ruin. It will slowly regain its strength.

David Jones

The fact that hibernation is probably not a part of the basking sharks' annual cycle may solve an irksome question — where and how do they hibernate? The meagre literature on basking sharks contains no observations of hibernation. Given that they are large, pelagic fish, some reaching lengths of over 10 m, surely they would have been observed if they hibernated close to the continental shelf. If, on the other hand, they hibernated in mid-waters in the open ocean (assuming this is possible at all), the sharks would be exposed to attack by predators or scavengers. At the very least they would have been colonized by symbionts and parasites, which attach themselves to immobile objects. Attacks of either kind, although probably not fatal owing to the size of these sharks, would result in tell-tale signs on their bodies. No signs of such depredations have been reported.

But if they don't hibernate, where do most full-sized sharks go during the winter? The lower energetic requirements for swimming, calculated by Sims³ using up-to-date information and theory, help to justify the assumption that basking sharks migrate to deep waters. However, feeding in deep waters is complicated by competition from more agile cephalopods and fish, and the lower feeding threshold may allow the sharks to use feeding grounds not frequented by potential competitors.

An important lesson from this story is that classical 'results' need to be checked and rethought in the light of current knowledge. This is especially relevant in studies of marine biology, because the open sea has not been extensively studied — it has been said that mankind knows more about the topography of the Moon than about the bottom of the ocean. However, both the hydrodynamic theory of how fish swim and our understanding of the biochemistry behind the ingestion and use of food have improved over the past

three decades. We can be confident in Sims' results because there are now well-established theoretical calculations for swimming thrust and energy requirements⁴, and these have been shown to fit a large range of fish sizes⁵, including several large species of shark.

Much still needs to be done. First, we need to search for the sharks' deep-water winter habitats, and to study their behaviour in these habitats. This can be facilitated using archival tags, which record parameters such as temperature, depth and swimming speed. Although such tags are still under development, some models can detach from the fish at predetermined times and transmit the data to satellites overhead. Second, in arriving at his results, Sims had to make several assumptions that need to be verified. These include catch efficiencies of the gaping mouth, the shape and intake of the ram filter feeding system, assimilation costs and metabolic efficiencies.

Here there is an interesting analogy to what used to be known as Gray's paradox. In the 1930s, Sir James Gray of Cambridge, one of the leading experimental biologists of this

century, estimated the efficiency of dolphin muscle. He did this by analogy to human muscle, and concluded that the swimming speeds achieved by dolphins cannot be accounted for by standard hydrodynamics. This resulted in a decades-long search for 'the dolphin's secret'. Many important discoveries and developments resulted from this search, including advances in hydrodynamic theory, and the study of flexible skin and polymer drag reduction. Finally, in the 1970s, when accurate data were used, the (by then improved) hydrodynamic theory was found⁴ to be sufficient to describe cetacean swimming. Studying sharks may yield similar dividends. □

Daniel Weis is at the Technion — Israel Institute of Technology, Haifa 32000, Israel.

e-mail: dweis@technion.ac.il

1. Matthews, L. H. *New Scientist* **280**, 756–759 (1962).
2. Parker, H. W. & Boeseman, M. *Proc. Zool. Soc. Lond.* **124**, 185–194 (1954).
3. Sims, D. W. *Proc. R. Soc. Lond. B* **266**, 1437–1443 (1999).
4. Webb, P. W. & Weihs, D. *Fish Biomechanics* (Praeger, New York, 1983).
5. Weihs, D. in *Principles of Animal Design* (eds Weibel, E. R., Taylor, C. R. & Bolis, L.) 121–128 (Cambridge Univ. Press, 1998).

Neurobiology

Cognition by a mini brain

Randolf Menzel and Martin Giurfa

We often expect little from small brains and cognitive miracles from big ones. But all brains were once small, in both evolution and development, reaching their respective levels of cognitive function only gradually. Elementary forms of cognition might, then, be present in small brains, and these could tell us what is really gained when brains become bigger. For example, the fruitfly *Drosophila melanogaster* (the geneticist's delight and the neurophysiologist's torment) has a particularly small brain. Very little is expected of this tiny brain, and, indeed, few impressive behavioural capacities have been found so far — a range of innate response patterns, some artistic sexual foreplay, and rather limited forms of learning^{1,2}. But new discoveries bring exciting prospects, using both molecular genetics and behavioural studies to analyse the fruitfly's elementary cognitive functions.

On page 753 of this issue, Liu *et al.*³ report their results on visual learning by *Drosophila* and its underlying neuronal substrate. They show that individual flies can do quite complex tasks. The authors first conditioned flies to associate visual patterns (the 'conditioned' stimulus) with the presence or absence of heat (the 'unconditioned' stimulus). The idea is that the animals should fly towards the appropriate patterns to avoid dangerous levels of heat. Their behaviour is 'operant', because the flight course that they choose determines delivery of the heat. The

experiments were done under particular illumination conditions, which form part of the general 'context' in which the associations are established.

The authors next showed that the flies could 'generalize' this trained response to several other, different environmental contexts. As contexts are ill-defined stimuli, comprising individual features from many modalities, the change in context that the authors introduced was a change in the illumination between training and test: between white and monochromatic broad-band light; between two monochromatic broad-band lights; and between constant white light and white light interspersed with 'dark flashes' (light is switched off for 200 ms). These changes did not affect the performance of the flies. These results show that context generalization — rather than context specificity — guides the insect's learning. But when the authors impaired the fly's normal brain function by eliminating the mushroom bodies (a central brain structure), they found that retention of the trained pattern was strictly bound to the context during learning, and that the flies did not generalize to other contexts.

The study by Liu *et al.*³ sheds light on two issues: the function of a central brain structure on an elementary form of cognition, and the role of context in learning. Mushroom bodies are prominent paired structures in the insect brain (Fig. 1), and are considered



Figure 1 Feeding time. The basking shark (*Cetorhinus maximus*) feeds on plankton, which it filters from sea water using bristle-like projections on the arches of its gills. In winter the concentration of plankton is reduced below a threshold level, and sharks were thought to escape starvation by hibernating or migrating. But a study by Sims³ indicates that the threshold of plankton needed by the sharks is much lower than previously estimated, and that the sharks do not need to hibernate or migrate in winter.




Figure 1 Mushroom bodies from the fruitfly *Drosophila melanogaster*. These paired structures are thought to be an information-processing centre, and Liu *et al.*³ now show that they are involved in context generalization during learning.

to be a high-order information-processing centre^{4,5}. In *Drosophila* they are related to olfactory processing and formation of olfactory memories¹, but they do not seem to be involved in visual processing and memory⁵.

The authors used three different techniques to remove the mushroom bodies (an *mbm* mutant lacking mushroom bodies; selective pharmacological ablation using the cell poison hydroxyurea during the first larval stage; and reverse-genetic ablation by enhancer-driven tetanus toxin), and showed that visual input affects the mushroom bodies. They suggest that the mushroom bodies do a dynamic analysis of the structure of the environmental stimuli, such that their temporal trace is compared to that of the unconditioned stimulus. (A similar function has been proposed for the honey-bee mushroom body, based on rather different data⁴.) Only the visual stimuli that change in a synchronous way with the heat stimulus can be useful predictors of dangerous heat. Other stimuli have a temporary structure, different from that of the unconditioned stimulus, so they are classified as 'context'. Flies that lack mushroom bodies cannot generalize their learned response to a different context because they are unable to dissociate the visual patterns from the illumination under which they were trained. In these flies, a compound between context and conditioned stimulus is formed. So, when the context changes, the compound changes too, impeding generalization to a new context.

The mushroom body seems to have three general functions. First, it separates stimuli according to their temporal structure. Second, it evaluates these stimuli with respect to the unconditioned stimulus (context or predictor). Finally, during the retention phase, it grades the stimuli according to predictive power. Such a proposal for a structure–function relationship is similar to a hypothesis by Solomon and Moore⁶ for the hippocampus in mammals — poorly associated conditioned stimuli are 'tuned out', and a hierarchy of stimulus predictiveness is reached. In this sense, the mushroom body may allow

selective attention, an attribute not yet experimentally accessible in insects.

Liu and colleagues' study also provides a fresh view of context dependence, an old issue in learning psychology. The textbooks tell us that animals can learn different meanings for the same stimulus when it is presented in different contexts. They maintain that contextual and conditioned cues appear to an animal as separate sensory entities, which may be associated in different ways⁷. *Drosophila*, on the other hand, seem to use a different system in which contextual cues and conditioned stimuli are filtered by the mushroom bodies. If this idea is correct, *Drosophila* may never have the flexibility to establish different, conditional relationships between contextual and conditioned stimuli.

But does the *Drosophila* mushroom body indeed act only on a basic sensory-processing level, as suggested by Liu *et al.*? Flies that lack the mushroom body perform tasks according to strict training conditions, indicating that sensory processing and association are done and stored outside the mushroom body, probably in the visual ganglia. So, the mushroom bodies might be related to the organization of multiple memories. It is possible that storage or retrieval initiates a search routine for additional conditional forms of memory traces — such as a trace indicating one stimulus as the contextual cue and the other as the conditioned stimulus. Such a trace would reside in the mushroom bodies, and would carry the information about whether the context or the conditioned stimuli were stable or variable. Flies deficient in the mushroom bodies may lack this trace, so would behave exactly as Liu *et al.* describe.

Although Liu and colleagues' data do not require such an elaborate model, their experiments have not yet tested it. The next steps may be to try and uncouple the function of mushroom bodies during acquisition, storage and retrieval, and to test combinations of contextual and conditioned cues in a balanced fashion. *Drosophila* will be an ideal organism in which to do these experiments, given the tools available for inducible knock-out mutants. In any case, the mini brain of this little fly is now well qualified as a model for basic cognitive functions. □

Randolf Menzel and Martin Giurfa are at the Freie Universität Berlin, Königin-Luise-Strasse 28/30, D-14195 Berlin-Dahlem, Germany.
e-mails: menzel@neurobiologie.fu-berlin.de
giurfa@neurobiologie.fu-berlin.de

1. Dubnau, J. & Tully, T. *Annu. Rev. Neurosci.* **21**, 407–444 (1998).
2. Greenspan, R. J. *Neuron* **15**, 747–750 (1995).
3. Liu, L., Wolf, R., Ernst, R. & Heisenberg, M. *Nature* **400**, 753–756 (1999).
4. Erber, J., Homberg, U. & Gronenberg, W. in *Arthropod Brain: Its Evolution, Development, Structure, and Functions* (ed. Gupta, A. P.) 485–511 (Wiley, New York, 1987).
5. Heisenberg, M. *Learning Mem.* **5**, 1–10 (1998).
6. Solomon, P. R. & Moore, J. W. *J. Comp. Physiol. Psychol.* **89**, 1192–1203 (1975).
7. Holland, P. C. & Bouton, M. E. *Curr. Opin. Neurobiol.* **9**, 195–202 (1999).

Daedalus

Superficial complexity

Metallic mercury has a high surface tension. But make it an electrode in an ionic solution, apply a voltage, and ions crowd its surface in surfactant fashion. Its surface tension drops dramatically. Non-ionic additives can reduce the tension even further. Daedalus is now taking this electro-capillary effect to extremes, using other metals.

He has set up a heated electrochemical cell with a molten metal electrode at the bottom, and a layer of molten salt above it as electrolyte. The vastly greater ion density of the melt will reduce the surface tension of the metal interface far more. Indeed, with a high enough voltage, a well-chosen salt, and some cunningly selected non-ionic additives, the surface tension of the metal could be driven negative.

A liquid with negative surface tension should increase its surface area vigorously. In the simplest case it might break into separate droplets; but not in Daedalus's system. Once detached from the bulk, a globule would cease to share its voltage. It would revert to positive surface tension, and reattach itself. Instead, the metal surface will fold into some nightmarishly complex fractal, dense with tiny corrugated corrugations of vast total surface area, but still all one surface.

Powerful applications beckon for this new metal treatment. Catalysis, heat exchange, electrode forming, battery and capacitor manufacture, all cry out for metals of the highest possible surface area. Daedalus is planning a 'continuous fractalization bath', in which a strip of base material, coated with the metal to be fractalized, runs at high speed through his molten-salt cell. The coating melts, acquires negative surface tension, blooms into vast surface area, and sets solid again in a cooler region of the bath downstream. Its adhering salt coating is washed off with water, and it is reeled up.

The pilot studies will use platinum, whose catalytic surface is crucial to so much chemistry. A possible salt for the job is potassium hydrogen sulphate; its melt cleans platinum so splendidly that it may be a natural surfactant for the metal. But Daedalus's fractal metal surfaces should find many other uses. In particular, they will be ideal for gluing and bonding. Unlike surfaces made by grinding or scoring, they will have re-entrant cavities and protrusions on all scales, and will keep an applied paint or resin immovably in place. Even PTFE-coated saucepans will at last retain their slippery coating.

David Jones