

- Rideout, E.J., Billeter, J.C., and Goodwin, S.F. (2007). The sex-determination genes *fruitless* and *doublesex* specify a neural substrate required for courtship song. *Curr. Biol.* 17, 1473–1478.
- Lee, T., and Luo, L. (2001). Mosaic analysis with a repressible cell marker (MARCM) for *Drosophila* neural development. *Trends Neurosci.* 24, 251–254.
- Wong, A.M., Wang, J.W., and Axel, R. (2002). Spatial representation of the glomerular map in the *Drosophila* protocerebrum. *Cell* 109, 229–241.
- Datta, S.R., Vasconcelos, M.L., Ruta, V., Luo, S., Wong, A., Demir, E., Flores, J., Balonze, K., Dickson, B.J., and Axel, R. (2008). The *Drosophila* pheromone cVA activates a sexually dimorphic neural circuit. *Nature* 452, 473–477.
- Patterson, G.H., and Lippincott-Schwartz, J. (2002). A photoactivatable GFP for selective photolabelling of proteins and cells. *Science* 297, 1873–1877.
- Ferveur, J.F. (2005). Cuticular hydrocarbons: their evolution and roles in *Drosophila* pheromonal communication. *Behav. Genet.* 35, 279–295.
- von Schilcher, F. (1976). The role of auditory stimuli in the courtship of *Drosophila melanogaster*. *Anim. Behav.* 24, 18–26.
- Kurtovic, A., Widmer, A., and Dickson, B.J. (2007). A single class of olfactory neurons mediates behavioural responses to a *Drosophila* sex pheromone. *Nature* 446, 542–546.
- Kondoh, Y., Kaneshiro, K.Y., Kimura, K., and Yamamoto, D. (2003). Evolution of sexual dimorphism in the olfactory brain of Hawaiian *Drosophila*. *Proc. Biol. Sci.* 270, 1005–1013.
- Jefferis, G.S., Potter, C.J., Chan, A.M., Marin, E.C., Rohlffing, T., Maurer, C.R. Jr., and Luo, L. (2007). Comprehensive maps of *Drosophila* higher olfactory centers: spatially segregated fruit and pheromone representation. *Cell* 128, 1187–1203.
- Technau, G.M. (1984). Fiber number in the mushroom bodies of adult *Drosophila melanogaster* depends on age, sex and experience. *J. Neurogenet.* 1, 113–126.

University of Glasgow, IBLs-Division of Molecular Genetics, Anderson College, 56 Dumbarton Road, Glasgow G11 6NU, UK.
*E-mail: s.goodwin@bio.gla.ac.uk

DOI: 10.1016/j.cub.2008.03.035

Visual Categorization: When Categories Fall to Pieces

We cannot help but categorize the visual world into objects like cats and faces. An intriguing new study shows that observers automatically discover informative fragments of visual objects during category learning.

Quoc C. Vuong

We see the world in discrete categories in order to recognize and interact appropriately with objects in our environment [1]. How do we learn visual object categories? Our intuition suggests that, through experience, we acquire features found in members of one category but not in those from another category. For example, cats have whiskers; human faces, on the other hand, normally do not. There is empirical support for this intuitive view [2,3].

But a fundamental problem with this intuition is image variability. Familiar objects from the same category can have an enormous range of appearance; they are often occluded by other objects; how they appear to us can further be confounded by viewing conditions such as variable illumination; and so on [2]. These factors converge to make it extremely difficult to learn generic features that are reliable for visual categorization.

In work published recently in *Current Biology*, Hegd  et al. [4] offer a compelling solution to this problem, but one that highlights the need for us to re-think the pieces that make up objects and object categories. Armed with a set of novel visual categories [5] and a statistical means to select features [6,7], these authors have demonstrated that

observers automatically discover fragments — literally, bits and pieces of images — during category learning that are very effective for visual categorization. This provides a new and important link between visual category learning and visual categorization.

In this new study [4], observers classified a large number of unfamiliar objects into two categories. The

objects were synthesized from a novel virtual phylogenesis algorithm which simulated the evolution of biological forms [5], so that category members captured natural variations of categories we are more familiar with. The examples in Figure 1 show that this classification task is far from trivial, even with whole objects (see supplemental Figure S1 in the paper for more examples).

Two main sets of image fragments were extracted from trained objects using the same statistical procedure. Observers then classified all fragments, just as they had done with whole objects. This sounds like an even more daunting task. Amazingly though, observers were as accurate with one set of fragments as they were




Figure 1. Example objects synthesized by virtual phylogenesis. Observers were only trained on objects from two of the three categories A, B and C (from [4]).




Figure 2. Examples of informative and uninformative fragments.

The informative ones distinguished between two trained categories — say categories A and B from Figure 1 — whereas the uninformative ones distinguished between a trained (A) and untrained category (C).

with whole objects — nearly 100% correct! Surprisingly, the same observers struggled to perform above chance with the other set of fragments.

Figure 2 shows the remarkable lack of visual difference between informative and uninformative fragment sets. Clearly, observers do not acquire just any fragment during category learning. But what distinguishes informative fragments where performance is nearly at ceiling, from uninformative ones where performance is more or less at chance?

The answer lies in how the fragments are extracted from images. Following recent computational advances, Hegdé *et al.* [4] selected fragments which maximized their ability to distinguish categories, through a powerful statistic called mutual information [6,7]. This measure tells us how certain we can be about the presence of a category if a specific fragment is present in the image. For example, if a human eye is present in the image, then there is a good chance that a human face is also present in that image.

The informative fragments used by Hegdé *et al.* [4] distinguished between two trained categories (say categories A and B; see Figure 1). By comparison, the uninformative fragments distinguished between a trained and

untrained category (say categories A and C). Observers never saw untrained category members.

Thus, both informative and uninformative fragments were of comparable visual complexity and both contained diagnostic information to distinguish categories, but only the informative ones were relevant for the observers' task. In fact, observers could not classify informative fragments prior to any training, which underscores the importance of category learning to discover the right pieces for the task.

There are thousands of possible fragments of an image, but only a fraction of them will reliably indicate that a particular category is present. Feature selection based on mutual information is a powerful framework to extract those fragments [6,7]. These are typically of intermediate complexity [6], balancing how likely the fragments will occur in an image and how indicative they are of a particular category.

For example, a fragment containing the eyes and a bit of the nose probably indicates that a face is present in the image but it is very unlikely to find such a large fragment in many different images. Conversely, a much smaller fragment containing just the hair line (so it looks like an edge) is likely to occur in many different images which do contain faces but may accidentally

occur in images which do not contain faces.

This framework is very successful for familiar visual categories [7]. For example, the mutual information of familiar object fragments correlates with neural measures like visual evoked potentials [8] and haemodynamic brain responses [9]. So there is exciting new evidence that the brain may also extract fragments of intermediate complexity for everyday things.

One concern with using familiar objects is whether observers learn fragments out of necessity, as objects are often occluded, or whether fragment-based learning occurs automatically as a matter of course. Hegdé *et al.*'s [4] results clearly favour the latter, as observers learn novel whole objects. There was no need for them to discover fragments during training, but they did.

There is something to be said about Hegdé *et al.*'s [4] virtual phylogenesis algorithm for synthesizing objects. Like biological organisms, their objects evolve from a common ancestor. Objects from the same category inherit their common ancestor's shape characteristics but express individual shape variations. Indeed, this algorithm has a nice parallel to earlier work with an artificial taxonomy of 'caminalcules' used to study how taxonomists classify the evolutionary relationships between species [10].

Virtual phylogenesis gives rise to novel object categories with desirable properties: for example, objects have measurable natural within-class variations similar to biological organisms. It is also versatile: for example, objects can be structured into a hierarchy of categories, or other evolutionary mechanisms (such as sexual selection) can be incorporated into the algorithm. Importantly, it is a principled means to synthesize a large number of naturalistic objects without unknowingly pre-specifying the informative fragments studied.

The algorithm diverges from alternative methods of synthesizing novel objects, such as combining shape primitives [11,12] or clustering shapes on the basis of similarity [13]. Given its versatility, virtual phylogenesis is a significant addition to the repertoire of techniques for synthesizing objects that can be used for natural vision, machine learning, and even evolutionary taxonomy.

Hegd  et al.'s [4] findings provide strong support for a prominent computational model of object perception and categorization based on informative image fragments [6,7]. They also support observers' natural tendency to pick up statistical regularities in the visual input [14], which can develop as early as nine months [15]. Lastly, the results link visual category learning with visual categorization, in that informative fragments play a key role for both processes [4,6–9].

Category learning remains an important issue in visual cognition. There are ecological reasons for acquiring pieces of visual categories [1]; for example, to overcome very real problems like occlusions and image variability. The human visual system has evolved to automatically acquire informative fragments for visual categorization. Let's hope that we will likewise pick up the pieces.

References

1. Rosch, E., Mervis, C.B., Gray, W.D., Johnson, D.M., and Boyes-Braem, P. (1976). Basic objects in natural categories. *Cogn. Psychol.* 8, 382–439.
2. Tarr, M.J., and Vuong, Q.C. (2002). Visual object recognition. In *Steven's Handbook of Experimental Psychology (III edition)*, Volume 1: Sensation and Perception, H. Pashler and S. Yantis, eds. (New York: John Wiley and Sons, Inc.), pp. 287–314.
3. Palmeri, T.J., and Gauthier, I. (2004). Visual object understanding. *Nat. Rev. Neurosci.* 5, 291–303.
4. Hegd , J., Bart, E., and Kersten, D. (2008). Fragment-based learning of visual object categories. *Curr. Biol.* 18, 597–601.
5. Brady, M.J., and Kersten, D. (2003). Bootstrapped learning of novel objects. *J. Vision* 3, 413–422.
6. Ullman, S., Vidal-Naquet, M., and Sali, E. (2002). Visual features of intermediate complexity and their use in classification. *Nat. Neurosci.* 5, 682–687.
7. Ullman, S. (2007). Object recognition and segmentation by a fragment-based hierarchy. *Trends Cogn. Sci.* 11, 58–64.
8. Harel, A., Ullman, S., Epshtein, B., and Bentin, S. (2007). Mutual information of image fragments predicts categorization in humans: Electrophysiological and behavioural evidence. *Vision Res.* 47, 2010–2020.
9. Lerner, Y., Epshtein, B., Ullman, U., and Malach, R. (2008). Class information predicts activation by object fragments in human object areas. *J. Cogn. Neurosci.*, in press.
10. Sokal, R.R. (1983). A phylogenetic analysis of Caminalcules. I. The database. *Systematic Zool.* 32, 185–201.
11. Gauthier, I., Williams, P., Tarr, M.J., and Tanaka, J. (1998). Training "greeble" experts: A framework for studying expert object recognition processes. *Vision Res.* 38, 2401–2428.
12. Biederman, I., and Gerhardstein, P.C. (1993). Recognizing depth-rotated objects: evidence and conditions for three-dimensional viewpoint invariance. *J. Exp. Psychol. Hum. Percept. Perform.* 19, 1162–1182.
13. Edelman, S., and Duvdevani-Bar, S. (1997). A model of visual recognition and categorization. *Philos. Trans. R. Soc. Lond. B* 352, 1191–1202.
14. Kersten, D., Mamassian, P., and Yuille, A. (2004). Object perception as Bayesian inference. *Annu. Rev. Psychol.* 55, 271–304.
15. Fiser, J., and Aslin, R.N. (2002). Statistical learning of new visual feature combinations by infants. *Proc. Natl. Acad. Sci. USA* 99, 15822–15826.

Institute of Neuroscience, Newcastle University, Newcastle upon Tyne NE2 4HH, UK.
E-mail: q.c.vuong@ncl.ac.uk

DOI: 10.1016/j.cub.2008.03.036

Cellular Evolution: What's in a Mitochondrion?

Mitochondria and their relatives constitute a wide range of organelles, only some of which function in aerobic respiration. Mitochondrial remnants from different anaerobic lineages show a striking degree of functional convergence.

Christopher J. Howe

For many years, the view was widely held that mitochondria originated when a primitive eukaryotic cell acquired through endosymbiosis a prokaryote capable of oxidative phosphorylation. Some of the endosymbiont's genes were lost, some were transferred to the nucleus, and a stable relationship was established that has lasted very successfully for well over a billion years. The fact that anaerobic eukaryotic lineages exist today — such as the gut-dwelling pathogen *Giardia* — was attractively consistent with this view of mitochondrial origin. These anaerobic eukaryotes appeared to lack mitochondria and according to molecular phylogenetic trees seemed to have diverged from other eukaryotes very early — presumably before the acquisition of mitochondria. This group became known as the Archezoa [1]. However, a discovery that would

ultimately be crucial to the demise of the Archezoan concept had been made back in 1973 with the description of hydrogenosomes in anaerobic trichomonads [2]. Hydrogenosomes are now recognised as derived from mitochondria. They produce hydrogen and ATP and have been found in a range of anaerobic or almost anaerobic eukaryotes. Writing in *Current Biology*, Stechmann et al. [3] have now described another example of a mitochondria derived organelle that sheds light on their evolutionary fate.

A second development leading to the demise of the Archezoan concept was the recognition that the anaerobic, amitochondriate eukaryote *Entamoeba histolytica* contains nuclear genes for the mitochondrial proteins pyridine nucleotide transhydrogenase and the chaperonin cpn60 [4]. This discovery indicated that this supposedly amitochondriate organism had possessed mitochondria in the past

and might even have retained a remnant of the organelle. Although the placement of *Entamoeba* among the Archezoa was controversial, other members of the Archezoa were soon shown also to harbour genes for proteins of mitochondrial origin and remnant mitochondrial compartments [5]. We now recognize that all eukaryotes probably have mitochondria, or their remnants, and indeed it arguably was the acquisition of the mitochondrion that marked the birth of the eukaryotes [6]. Furthermore, the phylogenetic position of Archezoa as early-diverging eukaryotes is also questionable [7,8].

A Diversity of Mitochondrial Forms
Mitochondrial remnants are known as hydrogenosomes or mitosomes, depending on their function. In general, organelles derived from mitochondria can be ordered on a spectrum based on their structure and function (Figure 1). Classical mitochondria, with their cristae as well as their electron transfer chain and F_1F_0 ATPase for oxidative phosphorylation in aerobic conditions, represent one end of the spectrum. Close to these are the mitochondria of some anaerobic metazoa, such as those of parasitic worms, which lack some components of the electron transfer chain [9].