news and views

Trading faces

Anya Hurlbert

By demonstrating that adaptation to a particular face can bias the subsequent perception of familiar faces, a new study supports the prototype theory of face recognition. These results also suggest that the prototype face is formed by averaging and can change with experience.

How does the human brain identify faces with such exquisite accuracy? For example, few mothers would fail to distinguish their baby's face from another, even though on close inspection their features, such as large eyes, snub nose and toothless smile, might be similar. Even mothers of identical twins claim never to mistake one for the other, often relying on subtle facial differences. A normal adult can recognize almost all of his or her high school classmates as much as 50 years after graduation[1].

Theories of face recognition agree that the brain must store representations of familiar faces in memory. But they differ as to the form of the representation and on how the brain matches an incoming face to that representation. A new illusion (the 'face identity aftereffect'), demonstrated by Leopold and colleagues in this issue[2], supports the notion that the representation is centered on a prototype—a face never before seen, but formed by averaging together all the faces the brain has ever seen. The results suggest that incoming faces are compared to the prototype and identified according to how they differ from it. Fascinatingly, it seems that this prototype is not fixed for life. Rather, it remolds itself when necessary. For example, overexposure to a particular face changes the population average, pulling the prototype closer to that face. When the prototype changes, so too may the identity of familiar faces.

To demonstrate that the prototype may shift its shape when experience dictates, Leopold and colleagues used adaptation, a time-honored tool of psychophysics. The experimenters adapted people to certain faces by brief but repeated exposure, and then measured their ability to identify other faces. The technique of adaptation exploits the fact that our senses are optimally designed to signal relative changes against the status quo, not absolute values. As a classic example, consider the effects of staring 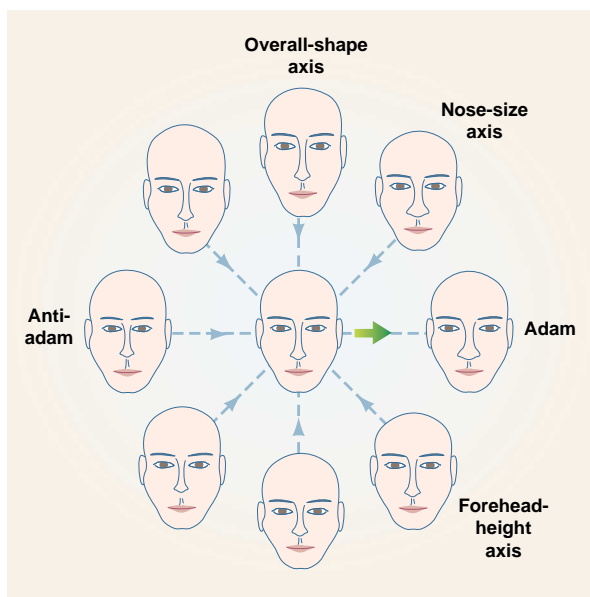fixedly at a red disk against a white background for a minute or so. When the red disk is abruptly removed, you will see its ghostly afterimage, a greenish-blue disk. This effect provides evidence for the existence of color-opponent mechanisms that pit red against green; desensitization to red shifts the neutral point, so that white now appears greenish. Adaptation experiments have also shown that many such color-opponent mechanisms may exist, corresponding to the multitude of hues that we can distinguish[3]. Similarly, the new face identity aftereffect suggests that faces may be encoded by 'face-opponent' mechanisms centered on a prototypical face—the 'white point' of face space.

At the heart of the experiment is a sophisticated morphing technique[4], which can seamlessly transform one face into any other, even creating its 'opponent face.' This technique does not simply align and blend two-dimensional images. Instead, it characterizes each face with a full description of its three-dimensional shape, color and texture. These physical parameters create a multi-dimensional 'face space' in which every face has its own unique location, corresponding to its own unique set of values (**Fig. 1**).

For the center of this morphable face space, Leopold and colleagues[2] created the average of a large population of normal adult faces (100 of each sex). As test faces, the experimenters choose four males: Adam, Jim, John and Henry. To create opponent faces for the suspects, they compared each of, say, Adam's values with the average, and created the face of 'anti-Adam' with equal but opposite features. For example, if Adam's only distinguishing features are his nose, which is 50 percent more bulbous than average, and his eyes, which are 50 per-

The author is Reader in Visual Neuroscience, Newcastle University Medical School, Framlington Place, Newcastle upon Tyne NE2 4HH, UK.
email: anya.hurlbert@ncl.ac.uk

Fig. 1. Cartoon of a multi-dimensional face space. Here the axes are defined by variations of certain facial structures. Faces with large noses are at the upper right end of one diagonal, with small-nosed faces at the lower left. The vertical axis captures the variation in overall face shape, with long, thin faces at the top, and short, wide faces at the bottom. In exemplar theory, incoming faces are located in face space according to the vector sum of their differences from each endpoint. The average face, therefore, sits at the origin. After adaptation to 'anti-Adam', sensitivity to 'anti-Adam' decreases. The vector sum (solid arrow) of responses to what was previously the average face now makes it appear more like Adam.

cent farther apart, then 'anti-Adam' has a nose only half as thick as the average, and his eyes are 50 percent closer together. Similarly, the experimenters created weak versions of Adam by starting with the average face and morphing toward him. At every step along the way, there is a perfectly plausible face that is intermediate in value between Adam's and the average for every parameter. Fractional versions of Adam have noses fractionally larger than average and eyes correspondingly a little farther apart.

The experimental subjects were taught to recognize and name Adam and the others in their weakened versions. After intensive training, subjects could perfectly identify any of the four faces at only 33 percent identity strength, in images flashed for only 200 milliseconds. Next came the adaptation phase. Subjects gazed at a particular anti-face for five seconds, then glimpsed the test face, which could be either the matching face or one of the other three, at small identity fractions. Now, subjects' abilities to identify the weakened test faces altered radically. After adapting to 'anti-Adam,' for example, subjects could perfectly identify Adam at much smaller fractions of 'Adam-ness'— around 10 percent identity strength. But they now failed to identify weak versions of Jim, John or Henry, even those versions that they had previously correctly identified on most trials. Furthermore, after adapting to 'anti-Adam,' when challenged with the average face (which possessed no prior identity whatsoever), subjects saw it as Adam on more than three-quarters of trials.

Like many other aftereffects, this one dies away in less than a second. Nonetheless, Leopold and colleagues argue that the identity aftereffect points to a mutable prototype at the center of face representation by the brain. Exposure to 'anti-Adam' shifts the prototype toward 'anti-Adam.' Weak versions of Adam now appear more like Adam than like the new prototype. Compared to the 'anti-Adam' prototype, weak versions of Jim also seem more like Adam and thus are harder to recognize as Jim. A 'face-opponent' neural mechanism centered on the prototype would account for these effects (**Fig. 2**).

But other explanations might work as well. Although face-opponent axes (each joining a face to its anti-face and all centered on the prototype) suffice to define 'face space,' they are not necessary. Theoretically, there are many other ways to parcel out face space, merely by aligning the axes along different dimensions. Indeed, neurophysiological data from the monkey suggests that neurons in face-specific brain areas collapse face space onto the smallest
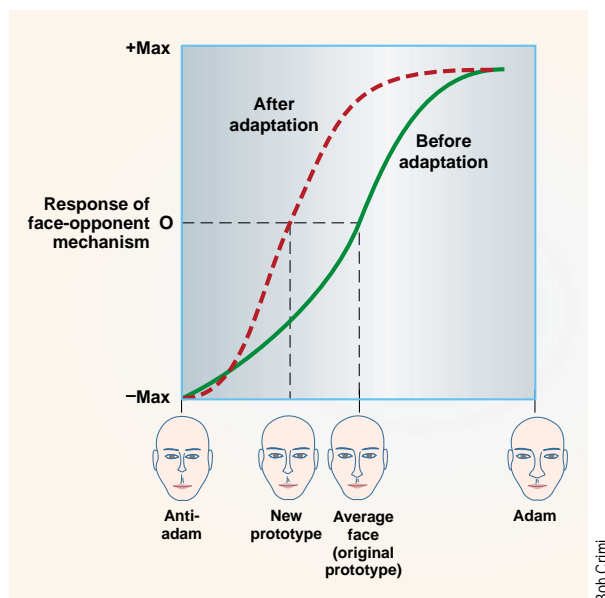

Fig. 2. Adaptation shifts subsequent face perception. The vertical axis represents the responses of a hypothetical 'face-opponent' neural mechanism, yielding maximum positive output to the face of Adam, intermediate positive outputs to faces morphed along the trajectory from the average face toward Adam, intermediate negative outputs to faces morphed from the average to 'anti-Adam' and maximum negative output to 'anti-Adam' itself. Before adaptation, the response to the average face is zero. After adaptation to 'anti-Adam', the response curve shifts toward 'anti-Adam'. Now the average face elicits a positive output, and therefore is signaled as Adam. A weak version of 'anti-Adam' now evokes the null response and acts as the new prototype.

number of dimensions possible, paying attention to only a few key dimensions along which physical features vary[5], such as face width or forehead height (**Fig. 1**).

It might also be that the axes themselves are not fixed features, but highly susceptible to change. Thus, exposure to the four test faces and their anti-faces might push the brain to re-align the dimensions of face space to those face-opponent axes. The observation that subjects improve markedly on the recognition task after extensive training, and the finding that the identity aftereffect is stronger after training than in a single one-hour session, might be the result of such realignment.

Other theories suggest that it might not be necessary for an explicit prototype to occupy center stage. Instead, face space might be peopled by a small number of 'exemplars' at key locations. Incoming faces would be identified by measuring their differences with each exemplar, then locating them according to the vector sum of the differences—a population code[6] (**Fig. 1**). The identity aftereffect could be explained by shifts in sensitivity to certain exemplars, which would ultimately shift the population code.

A further possibility is that the aftereffects are not due to changes in face space at all, but rather due to simpler changes in

sensitivity to isolated features—such as eyes, lips or just the outline of the face. Simple shape aftereffects do exist. For example, brief exposure to an isolated upward-curving 'mouth-like' shape causes a symmetrical 'diamond' shape to appear downward-curving[7]. In other words, the afterimage of a smile is a frown, under the right circumstances. But it is unlikely that a conjunction of simple shape aftereffects could explain the identity aftereffect.

To see why, consider another striking face aftereffect, the distortion aftereffect[8]. Here, subjects adapted to images of 'squashed' faces—faces whose features were artificially shrunk and compressed toward the center of the face. After five minutes of adaptation, subjects then saw normal images of the same face as unnaturally expanded: the eyes appeared widely separated and the mouth abnormally distended. This distortion aftereffect also worked in the opposite direction: adapting to expanded faces caused normal faces to appear squashed. But, crucially, adapting to normal faces had no effect on the appearance of distorted faces. If the distortion aftereffect were mediated by simple shape contrast, then normal face shapes should induce aftereffects just as effectively as distorted ones. That is, adapting to a normal-sized mouth should cause a distended

mouth to appear smaller. But it does not, suggesting that there is something special about the normal face: it is the center-point of the adaptation axis. Adapting to the center-point is utterly neutral, because it cannot push it further or closer from itself. Similarly, in the identity aftereffect, adapting to the average face had little effect on the appearance of weakened test faces.

These two illusions—the face identity and distortion aftereffects—share other characteristics, and together highlight the special nature of face recognition. Both aftereffects occur for upside-down faces. For example, when subjects adapted to an upside-down squashed face[8], they subsequently saw upside-down normal faces as expanded. But, after adapting to upside-down squashed faces, subjects saw upright normal faces as absolutely normal. Thus, it seems that there is separate circuitry for identifying upside-down faces, which is neither as streamlined nor as accurate as the well-honed machinery for upright faces. Other studies suggest that recognizing upside-down faces draws on a generic, feature-based system[9]. Both systems show effects of adaptation, but the effects do not interact. If so, one would predict that the identity aftereffect also should not transfer between upright and upside-down faces.

The identity aftereffect is also 'translation-invariant,' as the test face does not have to be presented at exactly the same location on the retina as the adapting face. Therefore, the affected neurons must also be indifferent to the location of the adapting stimulus, suggesting that the effect must occur at high levels in the visual system, where neurons have large, translation-invariant receptive fields. In macaque

monkeys, such neurons are found in inferotemporal cortex, and many are also face-specific[10]. This area corresponds to the site of the face-selective area found in human brain imaging studies[11], and to the site of damage in prosopagnosia, the inability to recognize faces[12].

The two aftereffects also point to two different roles for the prototype. The prototype at the center of the distortion aftereffect acts more like a generic prototype, establishing the norm for facial structure in general, as opposed to the structure of non-face objects. But the identity aftereffect demonstrates that the prototype also acts as the norm for individual identity. In other words, the prototype can act as standard-bearer for the whole class of faces, as well as for individuals. This finding is important, because the mechanisms for basic-level recognition (distinguishing between different classes of object) are often assumed to be different from those for subordinate-level recognition (distinguishing between instances within a class)[13]. But it has become increasingly clear that faces are special, and the routines for recognizing faces probably represent the acme of visual processing: streamlined, efficient and exquisitely sensitive.

Every day we see faces, and we look at some more than others. Are we influenced by the identity aftereffect? Perhaps the more we saw of Gore during the endless recounts, the easier it became to distinguish him from Bush—and vice versa. But the effect is fleeting, and subtle, so we probably did not notice. On the other hand, it might be that the effect is fleeting only when the exposure is fleeting, and we as yet know nothing about the effects of long-term exposure. It is entirely possible

that the prototype may shift more permanently in response to more permanent changes in the faces we see—if, for example, one moved from rural Iowa to an isolated village in mainland China. Prototypes might also build up during development: perhaps babies see all new faces as not-Mother because she is both the population and the prototype. Faces occupy a special space in our brains[14], and the more we explore it, the more dimensions we find.

1. Bahrick, H. P., Bahrick, P. O. & Wittlinger, R. P. *J. Exp. Psychol. Gen.* **104**, 54–75 (1975).

2. Leopold, D. A. *et al. Nat. Neurosci.* **4**, 89–94 (2001).

3. Webster, M. A. & Mollon, J. D. *Nature* **349**, 235–238 (1991).

4. Blanz, V. & Vetter, T. *ACM Computer Soc. Press SIGGRAPH '99 Proc.*, 187–194 (1999).

5. Young, M. P. & Yamane, S. *Science* **256**, 1327–1331 (1992).

6. Lewis, M. B. & Johnston, R. A. *Vis. Cognit.* **6**, 1–41 (1999).

7. Suzuki, S. & Cavanagh, P. *J. Exp. Psychol. Hum. Percept. Perform.* **24**, 1315–1341 (1998).

8. Webster, M. A. & MacLin, O. H. *Psychol. Bull. Rev.* **6**, 647–653 (1999).

9. Moscovitch, M., Wincour, G. & Behrmann, M. *J. Cogn. Neurosci* **9**, 555–604 (1997).

10. Perrett, D. I., Hietanen, J. K., Oram, M. W., & Benson, P. *J. Phil. Trans. R. Soc Lond. B* **335**, 23–20 (1992).

11. Kanwisher, N., McDermott, J. & Chun, M. M. *J. Neurosci.* **17**, 4302–4311 (1997).

12. DeRenzi E., in *Behavioural Neurology and Neuropsychology* (eds. Feinberg, T. E. & Farah, M. J.) 245–255 (McGraw-Hill, New York, 1997).

13. Damasio, A. R., Damasio, H. & Van Hoesen, G. W. *Neurology* **32**, 331–341 (1982).

14. Young, A. W. *Face and Mind* (Oxford Univ. Press, Oxford, 1998).

PI-3 kinase and IP3: partners in NT3-induced synaptic transmission

David R. Kaplan and Ellis Cooper

Neurotrophins are implicated in synaptic plasticity, but the signaling pathway was unknown. A new study shows that PI-3 kinase and IP3 are necessary and sufficient for NT3's actions.

David Kaplan is in the Brain Tumor Research Center, Montreal Neurological Institute, and Ellis Cooper is in the Department of Physiology, McGill University, Montreal, Quebec H3A 2B4, Canada.
e-mail: mcdv@musica.mcgill.ca

Although neurotrophins are best known for promoting neuronal survival, axon growth and differentiation, more recently they have also been implicated in synaptic development and plasticity. For example, brain-derived neurotrophic factor (BDNF) is required for specificity of developing connections in the visual system, dendritic remodeling and synaptic function[1,2], and neurotrophins modulate long-term potentiation (LTP) and possibly depression (LTD). How do the neurotrophins influence so many aspects of neuronal function? How can they concomitantly regulate survival, growth and activity in the same neuron? In this issue, Yang and colleagues[3] address these questions by identifying two signaling pathways that act together to regulate neurotrophin-induced synaptic potentiation and transmitter release.

In primary neurons, neurotrophins promote neuronal survival and growth through their Trk receptors, with NGF preferentially activating TrkA, BDNF and