

positive-intensity target in the middle frame. Again, this implies a sub-optimal strategy for detecting the target, as these early bar intensities are independent of the presence or absence of the target. Neri and Heeger interpret this to mean that these early high-contrast signals are engaging a separate ‘attentional’ mechanism that is used to detect the target. They ran a clever second experiment in which subjects had both to detect and to identify the polarity of the target, and the results demonstrated that the variance kernel accounted for the detection task and the mean kernel accounted for the identification task. The experiment thus lends support to the hypothesis that the visual system answers the questions of ‘what?’ and ‘where?’ using separate mechanisms [8].

Conclusion

The results of Neri and Heeger’s experiments are intriguing, and provide an elegant demonstration of the power of stochastic stimuli in characterizing visual mechanisms. It is worth considering the drawbacks of this approach, as well as possible generalizations. First, designing and executing this type of experiment is quite difficult, and relies on a number of decisions about how to instruct subjects, how much and what kind of training to allow, how strong a target signal to use, and whether to provide feedback. Second, the summary of their method in this brief review has been simplified to ignore the distinction between cases in which a target was present and those when there was no target, because most of the analysis presented in the paper by Neri and Heeger was done in this fashion (although they do present results for the ‘False Alarm’ case alone, which seem consistent with the simplified yes/no case). Analysis and interpretation of these sub-cases is more difficult, but can potentially offer further insights into the nature of the underlying visual mechanisms. It would also be interesting to extend the analysis to include interactions between stimulus bars (i.e. estimation of response-triggered *covariance*), as has been done in physiological settings [9–11]. This could provide a richer characterization of the underlying

mechanisms, at the expense of requiring more data for reliable estimation.

Finally, it would be interesting to see this technique applied to the detection, discrimination or identification of more complex stimulus features, such as those defined by orientation or motion (e.g. see [12]). Ultimately, refinement of these techniques could allow us to formulate a precise description of the mechanisms underlying all aspects of vision, from detection of complex features, to attentional and recognition processes.

References

- 1 deBoer, E. and Kuyper, P. (1968) Triggered correlation. *IEEE Transact. Biomed. Eng.* 15, 169–179
- 2 Marmarelis, P.Z. and Marmarelis, V.Z. (1978) *Analysis of Physiological Systems: The White Noise Approach*, Plenum Press
- 3 Ahumada, A.J. and Lovell, J. (1971) Stimulus features in signal detection. *J. Acoust. Soc. Am.* 49, 1751–1756
- 4 Ahumada, A.J. (1996) Perceptual classification images from vernier acuity masked by noise. *Perception* 26, 18
- 5 Eckstein, M.P. and Ahumada, A.J. (2002) Classification images: a tool to analyze visual strategies. *J. Vision* 2(1), Special Issue: *Classification Images*, DOI 10:1167/2.1.1x
- 6 Neri, P. and Heeger, D.J. (2002) Spatiotemporal mechanisms for detecting and identifying image features in human vision. *Nat. Neurosci.* 5, 812–816
- 7 Chichilnisky, E.J. (2001) A simple white noise analysis of neuronal light responses. *Netw. Comput. Neural Syst.* 12, 199–213
- 8 Sagi, D. and Julesz, B. (1985) ‘Where’ and ‘what’ in vision. *Science* 228, 1217–1219
- 9 de Ruyter van Steveninck, R. and Bialek, W. (1988) Coding and information transfer in short spike sequences. *Proc. R. Soc. Lond. B. Biol. Sci.* 234, 379–414
- 10 de Ruyter van Steveninck, R. and Bialek, W. (1988) Coding and information transfer in short spike sequences. *Proc. R. Soc. Lond. B. Biol. Sci.* 234, 379–414
- 11 Schwartz, O. *et al.* (2002) Characterizing neural gain control using spike-triggered covariance. In *Adv. Neural Information Processing Systems* (Vol. 14) (Dietterich, T.G. *et al.*, eds), pp. 269–276, MIT Press
- 12 Ringach, D.L. *et al.* (1997) A subspace reverse-correlation technique for the study of visual neurons. *Vision Res.* 37, 2455–2464

1364-6613/03/\$ - see front matter © 2003 Elsevier Science Ltd. All rights reserved.
doi:10.1016/S1364-6613(02)00043-8

Imagined movements that leak out

Margaret Wilson

Department of Psychology, University of California, Santa Cruz, CA 95064, USA

In a case study that fundamentally alters our understanding of motor imagery, Schwobel *et al.* report a patient who unintentionally carries out imagined movements. Furthermore, his ‘imagery’ movements are more accurate than his intended movements, which suggests that the inhibitory signal that normally prevents us from acting out our motor imagery can be selectively blocked. Removing this inhibition allows us

to observe motor imagery ‘in action’, and reveals that motor imagery and motor planning for execution are not identical.

In the last twenty years it has become accepted that ‘imagery’ of perceptual or motor events involves mental representations that, in some important sense, resemble the ‘real thing’. Visual imagery, for example, causes activation in visual processing areas of the brain [1], and motor imagery causes activation in motor areas [2]. But

Corresponding author: Margaret Wilson (mlwilson@cats.ucsc.edu).

how, then, does the system distinguish an imagined event from its overt counterpart? Why do visual images not become hallucinations, and why does motor imagery not result in muscular response? In the case of motor imagery, a kind of accepted wisdom has emerged: motor imagery is like motor production, with the addition of an inhibitory signal that prevents overt movement [3]. Now, in a single case study, comes a striking confirmation and partial disconfirmation of this account of motor imagery.

Unintended execution: a window into motor imagery

Schwoebel and colleagues [4] report the case of CW, 'the man who executed 'imagined' movements'. CW is a 67-year-old man with bilateral parietal damage due to two separate strokes. Two years after the second stroke, he shows some clumsiness with his hands, but largely preserved proprioception. Most startling, though, is the fact that when he imagines movements, he actually produces those movements without being aware that he is doing so. This behavior pattern, first observed informally, was confirmed in experimental settings. CW gives every indication of fully understanding the instructions to imagine certain hand movements, and reports no awareness that he overtly moves his hands when following the instructions.

In short, it appears that it is possible to selectively remove the inhibitory signal, so that motor imagery that is otherwise apparently normal can have the unintended consequence of following through to execution. This demonstration of the inhibitory signal by its absence is perhaps the most direct evidence to date that such an inhibitory signal exists.

The story does not end there, however. The fact that CW overtly executes his motor images conveniently allows us to observe the imagined movements. And for CW's left hand, those images are significantly more accurate than deliberate movements. When asked to touch his thumb to a particular finger, for example, CW makes significantly more errors under instructions to execute the movement than under instructions to imagine the movement.

How is this possible, if motor imagery is simply motor commands coupled with the presence (or in this case, the absence) of an inhibitory signal? The answer may lie in recent theories of the role of 'forward modeling' in motor control [5,6], an account which considerably complicates the view of motor imagery sketched above.

Explaining motor imagery with forward models

A forward-model is a simulation device. Using information about relevant properties of the system to be simulated, a forward-model runs in real time, paralleling and mimicking the activity taking place in the target system. The purpose of a forward model is to provide estimates of what is going on in the target system, and thus allow ongoing control and correction without waiting for feedback from the target system. Because the delays caused by neural transmission make it impossible for a pure feedback model to account for the fine-tuned temporal aspects of motor control [7], current theories incorporate forward modeling [5–7]. According to this view, forward models in the brain contain information about the biomechanics of the muscles, limbs and joints. Motor signals are sent

simultaneously to the pathways that result in overt movement and to the forward model, which runs a simulation of how the motor command is unfolding in the external world. By comparing the output of the forward model with the desired goal state, corrective commands can be sent in time to have their desired effect.

When the story of motor outflow is complicated in this fashion, we can see that there is now opportunity for motor imagery and motor execution to diverge. Rather than a single path of information flow from intention to execution, motor control involves multiple pathways and multiple representations. This makes it possible for dissociations and selective impairments to occur. One such dissociation appears to be occurring in CW.

Schwoebel *et al.* propose that the dissociation results from impairment of a mechanism that compares the output of the forward model with the actual proprioceptive and sensory feedback arriving from the external event. This comparator, which would be called upon in the case of executed but not imagined movement, might be producing faulty updating of the represented hand location. This would then result in inappropriate 'corrections' to the movement trajectory, but only in the case of an intended movement.

Alternative explanations for CW's performance

It is unclear whether this is the only, or even the best, explanation of CW's performance based on a forward-model account. As noted above, feedback from external movement is delayed, and it is unclear whether an impairment in integrating this feedback would occur in time to create the observed difficulties. Furthermore, other possible explanations suggest themselves when one considers how forward models work. For instance, there could be an impairment in the corrective motor commands derived from the forward model that are intended to fine-tune the movement in mid-trajectory. If such corrective commands are not sent in the case of motor imagery, the result could be a 'ballistic', uncorrected motor command that hits its target more often than motor commands that are modified by distorted corrections.

This is, of course, pure speculation, but it serves to illustrate how there could be other explanations for CW's performance. Further research should be directed at constraining the relationship between motor imagery, on the one hand, and the forward models that drive motor execution, on the other. Nevertheless, the case of CW both strengthens and fundamentally alters our understanding of how motor imagery occurs.

References

- 1 Farah, M.J. (1995) The neural bases of mental imagery. In *The Cognitive Neurosciences* (Gazzaniga, M.S., ed.), pp. 963–975, MIT Press
- 2 Jeannerod, M. and Frak, V. (1999) Mental imaging of motor activity in humans. *Curr. Opin. Neurobiol.* 9, 735–739
- 3 Lotze, M. *et al.* (1999) Activation of cortical and cerebellar motor areas during executed and imagined hand movements: an fMRI study. *J. Cogn. Neurosci.* 11, 491–501
- 4 Schwobebel, J. *et al.* (2002) The man who executed 'imagined'

- movements: evidence for dissociable components of the body schema. *Brain Cogn.* 50, 1–16
- 5 Wolpert, D.M. and Flanagan, J.R. (2001) Motor prediction. *Curr. Biol.* 11, R729–R732
- 6 Blakemore, S.-J. *et al.* (2002) Abnormalities in the awareness of action. *Trends Cogn. Sci.* 6, 237–242

- 7 Desmurget, M. and Grafton, S. (2000) Forward modeling allows feedback control for fast reaching movements. *Trends Cogn. Sci.* 4, 423–431

1364-6613/03/\$ - see front matter © 2003 Elsevier Science Ltd. All rights reserved.
doi:10.1016/S1364-6613(02)00041-4

Do we all look alike to computers?

Robert L. Goldstone

Psychology Department, Indiana University, Bloomington, IN 47405, USA

People are better able to recognize faces of their own race than those of other races. One explanation is that this ‘other-race effect’ is caused by lifelong experience in which faces from some races are more common than other races. A recent article by Furl *et al.* tests experience-based accounts of the other-race effect by giving computer algorithms of face recognition training on preponderantly Caucasian faces. The only algorithms that reliably produced other-race effects were those that created face representations distorted to emphasize features that individuated faces.

The chances are that you have had heard somebody, perhaps an uncle whom you suspect harbors racist attitudes, declare of some racial group, ‘*They all look alike to me.*’ Well, in fact, psychological experiments have confirmed this observation. In memory experiments, people are shown a set of faces, and are later shown a second set of faces that contains some of these old faces and additional new faces. People more accurately distinguish the old from new faces when the faces are from their own, typically more familiar, race than when they are from other races [1–3]. A recent study by Furl, Phillips and O’Toole [4] develops computational models of this ‘other-race effect’ Their modeling explores an account known as the ‘contact hypothesis’, which asserts that people develop specialized perceptual processes for representing familiar faces. An attractive aspect of the contact hypothesis is that it is consistent with a large corpus of empirical evidence showing robust perceptual learning for familiar objects that transfers to highly similar objects [5–7].

The only fly in the contact hypothesis’ ointment is that empirical evidence for life-long experience driving the other-race effect has been inconsistent. Some studies have shown that experience with other races modulates the other-race effect. For example, living among members of another race sometimes reduces the other-race effect [8,9]. However, other studies have failed to show an influence of contact with different races [1,10,11]. Given the empirical divergence of results, it is useful to step back and consider specific mechanisms for how perceptual systems might tune themselves to the properties of faces in their environment. Furl *et al.* pursue this strategy by testing

a set of computer algorithms for face recognition that tailor their representations to their environment. The algorithms were trained with a biased collection of faces containing a disproportionately large number of Caucasian faces. The researchers were interested in identifying the kinds of algorithms that would produce the same other-race effect observed in people using the same collection of faces. Is the creation of race-conscious computers really progress? It is if it uncovers principles that underlie our own skewed perception of races.

Training computers to recognize faces

In developing computers that exhibit an other-race effect, the authors started with 13 state-of-the-science computational algorithms for face recognition evaluated by the US government’s FERET (Face Recognition Technology) program. The 13 algorithms were tested on a database consisting of mostly Caucasian faces. Better recognition accuracy for Caucasian than Asian faces was taken as evidence for an other-race effect. These races were selected because the authors replicated the other-race effect with Caucasian and Asian participants using Caucasian and Asian faces from the FERET database. The 13 algorithms were divided into three categories based on how the face representations that they used are influenced by experience.

Three categories of face-recognition algorithms

The first category, non-contact hypothesis algorithms, represent faces in the same manner no matter what faces are used during training. The algorithms in this category could only account for an other-race effect if Caucasian faces were objectively more easily individuated than the Asian faces.

The second category consisted of algorithms that created perceptual representations based upon statistical properties present in the set of training faces. The eight algorithms in this set used Principal Component Analysis (PCA) to create dimensions that accounted for the most important sources of variation in the faces. For example, if broad foreheads, full lips, and round cheeks were all highly correlated, then PCA would create a single dimension that consolidated all three of these facial features. Every face can be represented in terms of the dimensions acquired by PCA during training, and the vector representation

Corresponding author: Robert L. Goldstone (rgoldsto@indiana.edu).