OPINION

Neural worlds and real worlds

Patricia S. Churchland and Paul M. Churchland

States of the brain represent states of the world. But at least some of the mind–brain's internal representations, such as a sensation of heat or a sensation of red, do not resemble the external realities that they represent: mean kinetic energy (temperature) or electromagnetic reflectance (colour). The historical response has been to distinguish between objectively real properties, such as shape and motion, and subjective properties, such as heat and colour. However, this approach leads to trouble. A challenge for cognitive neurobiology is to characterize, in general terms, the relationship between brain models and the world. We propose that brains develop high-dimensional maps, the internal distance relationships of which correspond to the similarity relationships that constitute the categorical structure of the world.

Consider the following simplified theory: the brain constructs models, of the body, of the world external to the body, and of some activities of the brain itself. By 'model', we loosely mean an organized representational scheme. The body is represented in its somatosensory and motor aspects, as well as in its inner milieu, including its drives, CO_2 levels and blood pressure. Representation of the external world is probably anchored in a computational platform that is organized to predict the edibility, congeniality, hostility and so on of objects in space–time. Motor plans are run through emulators that simulate the body in its environment to predict the consequences of movements before their execution[1–6].

Some neural activities represent other neural activities as feelings, such as pain, fear and fatigue; some represent sensory states as resulting from an encounter between an external object (for example, a bee) and oneself ('that stung me'). According to Damasio[7], this fundamental representation of causal interactions between the external world and the body anchors conscious self-representation; that is, the representation of self-as-inner versus world-as-outer. How much integration normally exists across various inner models and how integration, within and across models, is achieved are open questions[8]. From the point of view of the user, of course, that the brain constructs its distinctions between inner-me and outer-world is anything but self-evident. The brain does not have an introspective representation of its world-modelling activities in the way that it has of its pains, needs and emotions.

To a first approximation, this theory 'cartoon' is the conceptual framework that is assumed, implicitly or explicitly, by many cognitive scientists and neuroscientists. The point of outlining the theory is to make the main load-bearing pieces easily visible, as visibility aids exploration of what the framework entails, whether it constitutes a coherent package, what parts are radically incomplete and whether some parts of the story might be entirely misconceived.

These are not questions to be answered by a single, crucial experiment. Some might not be answerable at all at this early stage in the development of the brain sciences. Nevertheless, they are worth asking because background assumptions can turn out to be problematic. In addition, such assumptions, despite living a quiet life mainly in the background, do motivate experimental research. They inevitably have a significant role in constructing and testing hypotheses. Finally, unless they are brought to the foreground occasionally, background assumptions tend to pass as obvious, God-given, beyond-question truths. This is not a good thing. Dogma is undesirable in any science, but especially so in a young field, such as neuroscience, in which the fundamental principles that govern brain function and organization are only beginning to be understood.

Primary and secondary qualities

One matter that deserves attention concerns how we understand the relationship between the brain's models of the world and the world itself. In particular, what can we learn about the fidelity of the brain's representations, relative to the things that are represented? A nest of interconnected problems resides in this domain, and it will be useful to extract the most troublesome for further dissection.

According to conventional wisdom, some properties that are represented by the brain as in the world are not genuinely in the world at all, but are mere products of brain activity. The standard examples are colours, smells and sounds. So, it will be argued, a peppermint smell is what some brains create in response to depolarization of receptor cells in the olfactory epithelium by particular molecules, whereas other brains are indifferent to those molecules. Although the molecules are in the external world, the smell itself is not. By contrast, the argument continues, certain properties, such as mass, motion and spatiotemporal relationships, are really in the world and our representations do resemble them. These are 'primary qualities', whereas smells and colours are 'secondary qualities' that are caused in nervous systems by primary qualities.

Galileo (1564–1642) was perhaps the first to postulate this distinction[9]. He aimed to explain the fact that the way the world appears in experience might not be the way the world really is. Galileo's hypothesis was occasioned by the phenomenon of heat. He reasoned that, in reality, heat was probably the motion of tiny 'corpuscles': atoms, as conceived by Democritus (460–370 BC). Galileo was struck by the fact that heat is not experienced in this way. The feeling of heat seems to have nothing to do with the perceived motion of anything. To explain why the appearance does not resemble the reality, he suggested that particle motion kicks off a causal process in the body that results in something quite different: heat as we experience it. Given this and other differences, discovered by science, between appearance and reality, thinkers such as John Locke (1632–1704) saw wisdom in making a general distinction between real-world properties (primary qualities) and brain-constructed properties (secondary qualities)[10].

Something akin to the primary–secondary distinction is still commonly wheeled out in discussion of the conceptual framework outlined above. Despite its *prima facie* plausibility, the distinction carries the seeds of its own destruction.

Figure 1 | **The slippery slope to Idealism.** The slippery slope begins with the seemingly plausible distinction between primary and secondary qualities, which was first conceived by Galileo when puzzling about the differences between heat as a subjective experience and heat as an objective property in the world. John Locke was moved by similar considerations and explored the distinction further, glimpsing but not worrying too much about the slope in front of him. Immanuel Kant struggled to keep from sliding down, Bishop Berkeley was convinced that sliding was the only logical recourse, and Georg Hegel was a thoroughgoing Idealist. Image courtesy of M. Churchland.

Notice that the distinction privileges primary qualities as uniquely real, by running very fast by the fact that there is no brain-independent or representation-free access to reality. If colour and smell representations are the brain's causal response to certain external stimuli, then so are spatial representations and motion representations. The brain cannot directly compare its representation of the external world with the external world itself, as we might compare the on-stage Wizard of Oz with the man behind the curtain.

To be sure, instruments can provide the brain with further data, as Galileo's clever thermal-expansion thermometer provided him with data concerning heat. But instrumental data require human observation and theory-backed interpretation, both of which involve filtering through the lens of representational models. Hence, 'objective' instruments are not the general solution to making a principled distinction between primary and secondary qualities.

Here, then, is the dilemma: if the distinctions between inner-me and outer-world are made within the brain's representational models, then how does the distinction between brain-constructed properties and real-world properties gain objective grounding? And why should we believe that primary qualities accurately characterize reality whereas secondary qualities are merely representational fabrications that are incident on the brain's interactions with reality?

Into Idealism and out again

The philosopher Bishop Berkeley (1685–1753) realized that the arguments that support the secondary qualities as mind creations start us down a very slippery slope (FIG. 1). Once on the slope, we find that we have slipped to the next stage: primary qualities, just like secondary qualities, are nothing but mind-created responses to a real world, the true nature of which we can never know. At the bottom of the slippery slope is the proposition that the so-called external world is, after all, nothing more than my idea of an external world. Ideas are the only things there really are. Furthermore, my apparently physical brain also must be nothing but an idea. In this case, only non-physical minds — constellations of ideas — genuinely exist. Classically, this view is known as 'Idealism', and it is what awaits us at the bottom of the slope.

Berkeley gleefully took the trip to the bottom of the slope, and was an uncompromising Idealist to the end[11]. In the *Critique of Pure Reason*, Immanuel Kant struggled mightily to stop two-thirds of the way down the slope, but hit bottom even so[12]. Georg Hegel (1770–1831)

Figure 2 | **The refraction illusion.** The pencil appears to have a kink where it enters the water. This appearance is explained by the refractive properties of the water medium. Image courtesy of D. Stack.

and other German Idealists in the nineteenth century found their intellectual home at the bottom[13]. Their efforts went into explaining the apparently 'physical' world in terms of the allegedly more basic world of mental ideas. Because some contemporary neuroscientists consider the primary–secondary distinction to be unavoidable and the slippery slope to Idealism to be inevitable, we shall briefly discuss the disadvantages of Idealism before returning to the question of the fidelity of brain models to the world modelled.

An obvious and fundamental objection to Idealism is that it cannot account for the coherence and regularity of the (idea of the) external world, or even of one's own mental life. Even simple regularities — objects thrown in the air regularly fall towards the Earth, dry wood regularly burns, water reliably quenches fire — are inexplicable. Berkeley's solution appealed to the supernatural: God keeps the flux of ideas coherent. But what happens in deep sleep, or in coma, when ideas vanish? What of the existence of the Universe before there were minds? Luckily, Berkeley's God saves the day by having all of that as ideas in His ample mind. The supernatural solution is transparently *ad hoc* — a pixie-dust, magic-wand solution — and Berkeley's contemporaries tore it to shreds. However, unfortunately for Idealism, no one ever contrived a solution that was both intelligible and less *ad hoc* than Berkeley's.

Idealists have also been unable to account for the progress of science in falsifying hypotheses (for example, malaria is caused by bad air, the Earth is flat, the Earth is motionless). For Idealism, so-called falsification has nothing special to do with getting closer to the nature of objective reality. There are just ideas, and then some more ideas. As an explanation for scientific progress, Hegel, essentially a scientific naïf, ground out famously obscure rigmarole about a dialectical process whereby the Universal Mind is inexorably coming to know Itself.

More generally, Idealism cannot begin to account for the spectacular developments in the sciences that falsified intuition-friendly ideas (for example, that heat is a kind of fluid) and replaced them with explanatorily powerful but unobvious theories, such as the atomic theory, statistical mechanics and evolutionary biology. Ironically, the same developments that provoked Galileo and Locke to seek epistemological ballast for the appearance–reality distinction are developments that stubbornly resist the Idealist's approach. Finally, and most awkward for the Idealist, it turns out that the brain sciences are making progress in explaining why things sometimes seem to be different from how they probably are; for example, why a straight pencil in water looks bent (FIG. 2), why two identical patches of grey appear to have different luminances (FIG. 3), and why a full moon looks much larger when seen on the horizon than when seen overhead[14]. In other words, the evidence increasingly supports the view that science can explain, in great and systematic detail, mental properties in terms of physical things (nervous systems), whereas the Idealists' programme is left clawing at the air[15–17].

Like any hypothesis, large-scale or small-scale, Idealism's value has to be measured in terms of its distinct explanatory and predictive results. On this criterion, Idealism scores in the hopeless range. It does no explanatory or predictive work in science or in ordinary life. Indeed, as Berkeley more or less admitted, even if one believes Idealism to be true, one has no choice but to act as though it were false.

Brain models and real-world survival

The epistemological question now before us is how to address productively the relationship between representational models and the world modelled. One response is to circumvent the whole mess with a pre-Galilean resolve to believe that we do, in fact, perceive reality precisely as it is. In other words, dodge the primary–secondary distinction and boldly claim that there is a perfect match between appearance and reality. Known as 'Naive Realism', this view is undone by its specious innocence about the developments of modern physics, astronomy, biology and chemistry. In addition, it has to ignore the gathering evidence from cognitive science and neuroscience concerning the brain's constructive processes.

Consider, for example, ambiguous figures, which can be seen as different objects even though the stimulus is unchanged (FIG. 4). The brain constructs contours where no luminance differences exist in the stimulus (FIG. 5); it constructs the perception of colour in regions that are actually white. Three-dimensional depth is regularly constructed from two-dimensional retinal arrays[18,19] (FIG. 6).

Contrary to Naive Realism, the conclusion that there is a perfect match between reality and representation is untenable. Nevertheless, a reasonable solution can be found by exploiting contemporary science to articulate a more likely relationship, such as an informational correspondence between brains and the world.

Perhaps representational models in nervous systems are roughly like a map, in the sense that their internal, abstract relationships map onto external relationships between the various categories in the world[14,20]. The rough and low-dimensional analogy is the road map of a city, in which the real spatial relationships between roads are represented in the relationships between road-lines on the paper map. Just as road maps come in varying degrees of fidelity and detail, so brain models of the external world map the categorical and causal structure of the world with varying degrees of

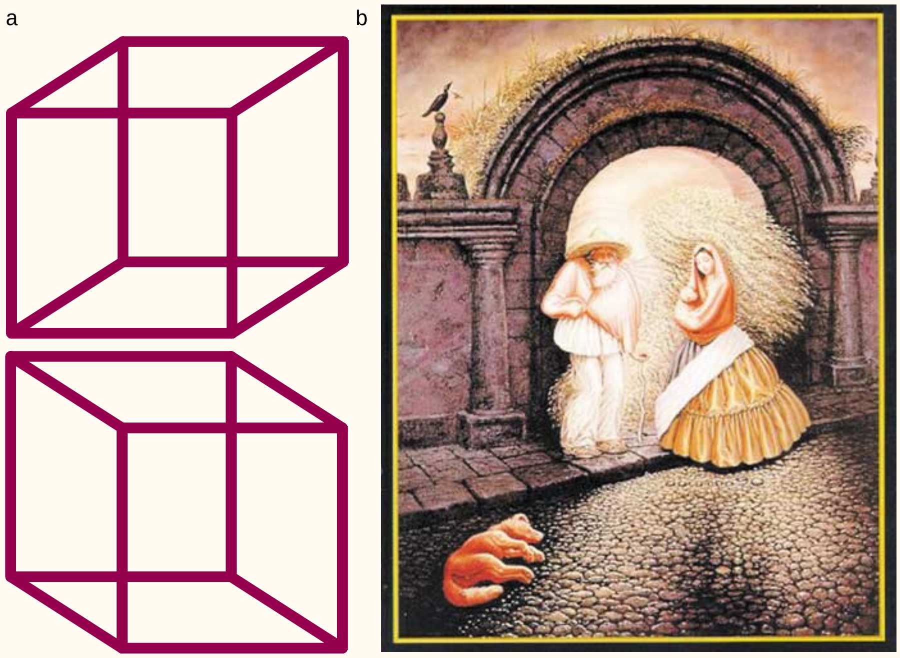

Figure 4 | **Ambiguous figures. a** | The Necker cube is the classic ambiguous figure that is entirely neutral between the two interpretations of its orientation. **b** | In this image, there are nine regions that can be seen either as a face or as part of the background scene. So, the old man's ear can also be seen as the upper body of a woman, and his nose is the right shoulder and arm of another man. Unlike the Necker cube, this picture is not perfectly neutral between competing interpretations, but the compelling perception of the large profile of the old man illustrates the powerful effect of top-down processing in defeating texture and stereoscopic depth information. Image courtesy of R. L. Gregory, Department of Experimental Psychology, University of Bristol, UK.

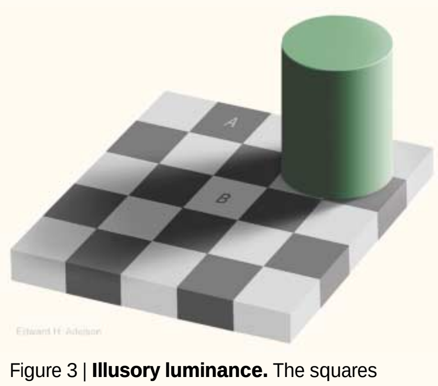

Figure 3 | **Illusory luminance.** The squares marked A and B have identical grey levels and measure the same luminance on a photometer. B looks much lighter than A because the brain deploys the assumption that the light source is behind and to the right of the green pillar, B is in its shadow, and the alternating squares on the checkerboard have constant grey levels. The brain therefore constructs the visual experience so that if B were out of the shadow, it would form a consistent part of the checkerboard pattern. Image courtesy of E. H. Adelson, Brain and Cognitive Sciences, Massachusetts Institute of Technology.

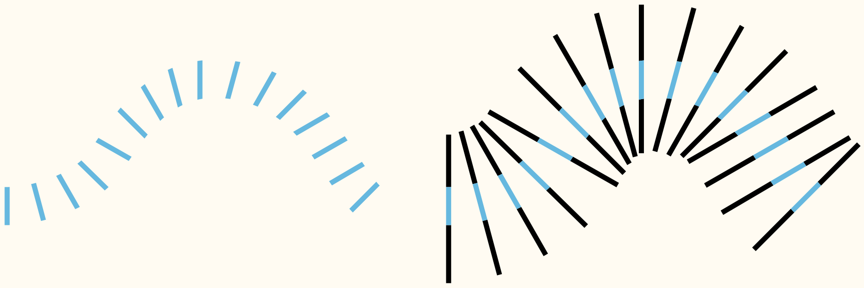

Figure 5 | **Neon colour spreading.** On the left is a sequence of light blue lines on a plain white background. The visual system groups them together to form a single curved form. The figure on the right is identical, except for the fact that black extensions have been added to the ends of each of the blue lines. Now, the spaces between the lines appear to be blue, although the background is really uniform, and one sees a semi-transparent blue worm with clear subjective boundaries. The subjective appearance of the colour to the intervals between the lines is known as neon colour spreading. Adapted, with permission, from REF. 1 © 1998 W. W. Norton & Co.

fidelity and detail. A frog's brain maps less of the categorical structure of the world than a raven's brain; an infant's brain maps less of the categorical structure of the world than an adult's brain; pre-scientific human brains map less of the categorical structure of the world than scientifically trained brains. Note also that, just as area maps focus on specific features of interest, so the features mapped by brains are generally those that matter to the organism and to how it makes its living: 'me-relevant' features. And unlike two-dimensional paper maps, the representational models in nervous systems will be multi-dimensional, probably very high-dimensional maps[5,21].

The coherence and predictive power that representational models enjoy is explained not by Berkeley's God, but by biological evolution and by empirical learning. Animals are movers, and nervous systems earn their keep by servicing movement. Other things being equal (and there are a lot of other things), the better and faster the brain's predictive capacities relative to the animal's *modus vivendi*, the better the organism's behavioural portfolio in the cut-throat competition to survive and reproduce.

In the broadest terms, the solution found by evolution to the problem of prediction is to modify motor programmes by sensory information. The value of the sensory impact is greater if it can signal me-relevant features and causal regularities between events. To achieve this, the system needs neural populations that are interposed between sensory receptors and motor neurons to find and embody higher-order regularities. The richer the interposed neuronal resources, the more sophisticated the statistical capacities and the greater the isomorphisms achievable between the brain's categorical/causal maps and the world's categorical/causal structures. Importantly, much of the brain's input will be consequent on the organism's own movements, exploratory and otherwise. In a given time interval, this dynamical loop extracts vastly more information about the causal properties of the external world than could a purely passive system. In nervous systems generally, testing expectations and having them met or surprised is the key to the falsification and revision of representational models, and the plasticity of our predictive capacities allows the long-term use of the results of trial-and-error learning[22,23].

Brains that represent certain higher-order regularities as allocentric objects and structures enduring in objective space–time have a powerful representational tool for exploring and finding out yet more about the categorical and causal structure of the world. Nervous systems that can use external tools, such as microscopes and telescopes, extend their predictive capacities and expand the range of causal structures within their ken. Nervous systems that can invent and deploy theoretical tools, such as the notion of 'valence' or 'gravity' or 'gene', extend that range even further.

Organisms whose brains map knowledge sources in social groups can economize on what a single brain must know by exploiting the distribution of knowledge across individuals[24]. Organisms whose nervous systems enable them to deposit new knowledge into offspring, artefacts and social institutions can revise and improve their internal maps across many generations[25].

Science, in its broadest sense, thus provides the general answer to the Galilean question: what epistemological grip can we get on the appearance–reality distinction? To provide a more detailed answer, we have to go through examples of scientific discovery. We can begin with simple examples, such as the calculations showing that the moon is not the size of a barn and as far away as a high cloud, or why neon colour spreading is illusory, or why we think malaria is caused by *Plasmodium falciparum*, transmitted by female mosquitoes as they penetrate the skin for a blood meal. We could move on to Lavoisier's experiments showing that, in the process of burning, so-called 'dephlogisticated air' (oxygen) rapidly combines with material in the wood, instead of sucking something (phlogiston) out of it. We could move next to the predictive and explanatory revolution wrought by Newton's laws of motion, then on to Einstein's revolutionary updating in the special theory of relativity, and to contemporary developments in the brain sciences. Learning a lot of science and a lot of history of science allows us to get a feel for the principles by which science moves forward in mapping the categorical and causal structure of the world.

There is no algorithm for making scientific progress, just as there is no algorithm for being rational. There are, however, instructive prototypes and useful rules of thumb:

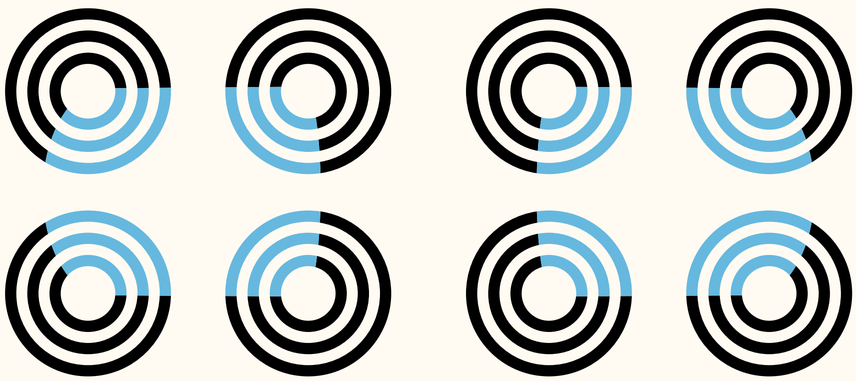

Figure 6 | **Subjective effects in stereopsis.** Only the lines on the segments of the circles are actually blue. Nevertheless, we perceive a single piece of blue film with a clearly visible shape lying on the black circles. This illustrates both neon colour spreading and the construction of subjective contours. The images on the right and left are stereoptic pairs and can be fused to form a single image in depth (cross-fuse). Now, the subjective contours of blue film curve out in space towards the viewer. Adapted, with permission, from REF. 1 © 1998 W. W. Norton & Co.

observe, think, test, don't be dogmatic but don't change your mind too easily, don't get in a rut but don't give up too soon, take advantage of statistical tools but don't suppose that good theories will simply waft up from statistical analyses, and so on[26–29].

Ultimately, we want to understand in detail the neural mechanisms by which common sense and science develop coherent representational models, and what 'coherence' means in neural terms. Meanwhile, we can begin to see what the metaphor of 'isomorphisms between models and world' might amount to in non-metaphorical terms.

Assume that a representational model can be characterized in terms of a parameter space, the dimensions of which are those neurons that participate in the model. Distinct representations (for example, the face of Clint Eastwood and the face of Sophia Loren) will involve distinct patterns of activation across the participating neurons and will thus individuate distinct points in the neuronal parameter space. In a neuronal parameter space for faces, for example, the activation pattern for Queen Elizabeth's face will be closer to the pattern for Sophia Loren's face than to that for either Clint Eastwood or George Bush. In the neuronal taste space, the activation patterns for apricots, raspberries and honey cluster together, and are more distant from the cluster of patterns for bitter things, such as choke cherries, quinine and urea, which in turn are separated from the cluster of patterns for salty things.

Whence isomorphism? The various distance relationships between the learned clusters (prototype points), within the activation space of a given population of neurons, are collectively and literally isomorphic with the similarity relationships that objectively exist between the various categories in the external world. The same holds for causal regularities in the world and prototype trajectories in neuronal-activation space. The greater the degree of isomorphism, the greater the fidelity of the model to the world. As we cannot directly compare model and world modelled, predictive success is the measure of fidelity and the guide to the need for model revision. Somehow, the various mechanisms of neuronal plasticity, including dendritic growth, the emergence of new synapses, changes in the probability of vesicle release and changes in transmitter released per spike, are orchestrated to enhance the fidelity of the basic model.

We call this hypothesis 'domain-portrayal semantics', because it proposes that the primary representational relationship holds between the high-dimensional map as a whole, and the categorical/causal domain as a whole. Traditional semantics, by contrast, assumes the primary representational relationship to hold between our internal concepts taken one by one, and external features taken one by one. According to the domain-portrayal hypothesis, single concepts derive their representational significance entirely from the larger neural model in which they are embedded. Intuitively, of course, it may seem otherwise, but 'folk semantics' is undoubtedly as misconceived as were folk physics and folk cosmology.

The proposal to distinguish between 'primary' and 'secondary' properties runs amok because it attempts to explain successful representation in terms of a resemblance relationship between inner concepts and outer properties taken one by one. The slippery slope to Idealism is avoided not by working ever more feverishly to establish the resemblance between individual primary qualities and individual properties in the real world, but by looking to the general function of representational domain models in nervous systems. This ushers in the recognition that representational utility depends on the higher-order, multidimensional 'resemblances' that nervous systems produce; that is, on the relative richness of the isomorphism between the representational model and the world modelled, as indexed by the model's predictive profile.

Domain-portrayal semantics is, so far, described only in the most general terms. Nevertheless, reorienting semantics away from the one-by-one paradigm and towards a model-to-domain paradigm that is more consilient with current neuroscience motivates continued exploration of specific neuronal populations to discover what relationship-preserving mapping constitutes the particular representational success of particular neuronal populations. Contrary to both Locke and the Idealists, the mind–brain does model the real world, including that part of the world that is the mind–brain. The reality–appearance distinction ultimately rests on comparisons between the predictive merits of distinct representational models, and the best explanation for why one theory out-predicts another is that one theory is closer to the truth than the other.

Patricia S. Churchland and Paul M. Churchland are at the Philosophy Department, University of California San Diego, La Jolla, California 92093, USA.
Correspondence to P.S.C.
e-mail: pschurchland@ucsd.edu

doi:10.1038/nrn958

1. Hoffman, D. D. *Visual Intelligence* (Norton, New York, 1998).
2. Wolpert, D. M., Ghahramani, Z. & Jordan, M. I. An internal model for sensorimotor integration. *Science* **269**, 180–182 (1995).
3. Grush, R. Self, world and space: on the meaning and mechanisms of egocentric and allocentric spatial representation. *Brain Mind* **1**, 59–92 (2001).
4. Churchland, P. S. *Brain-Wise: Studies in Neurophilosophy* (MIT Press, Cambridge, Massachusetts, 2002).
5. Churchland, P. M. *The Engine of Reason, The Seat of the Soul* (MIT Press, Cambridge, Massachusetts, 1996).
6. Suri, R. E. Anticipatory responses of dopamine neurons and cortical neurons reproduced by internal model. *Exp. Brain Res.* **140**, 234–240 (2001).
7. Damasio, A. R. *The Feeling of What Happens* (Harcourt Brace, New York, 1999).
8. Gazzaniga, M. S. Cerebral specialization and interhemispheric communication. *Brain* **123**, 1293–1326 (2000).
9. Galileo *The Assayer* (1623) in *The Scientific Background to Modern Philosophy: Selected Readings* (ed. Matthews, M. R.) 56–60 (Hackett, New York, 1989).
10. Locke, J. *An Essay Concerning Human Understanding* (1690). Modern edition (ed. Nidditch, P. H.) 544–547 (Oxford Univ. Press, Oxford, UK, 1975).
11. Berkeley, G. in *Berkeley's Philosophical Writings* (ed. Armstrong, D.) (Macmillan, New York, 1965).
12. Kant, I. *Critique of Pure Reason* (1781). Modern edition (eds Guyer, P. & Wood, A.) (Cambridge Univ. Press, Cambridge, UK, 1998).
13. Hegel, G. *Phenomenology of Spirit* (1807). Translated by Miller, A. V. (Oxford Univ. Press, Oxford, UK, 1977).
14. Kaufman, L. & Rock, I. The moon illusion: I. *Science* **136**, 953–961 (1962).
15. Palmer, S. E. *Vision Science* (MIT Press, Cambridge, Massachusetts, 1999).
16. Purves, D., Lotto, R. B., Williams, S. M., Nundy, S. & Yang, Z. Why we see things the way we do: evidence for a wholly empirical strategy of vision. *Phil. Trans. R. Soc. Lond. B* **356**, 285–297 (2001).
17. Ramachandran, V. S. Filling gaps in perception: part II. Scotomas and phantom limbs. *Curr. Dir. Psychol. Sci*. **2**, 56–65 (1993).
18. von Helmholtz, H. L. F. *Treatise on Physiological Optics* (1866). Translated by Southall, J. P. C. (Dover, New York, 1925).
19. Gregory, R. L. *The Intelligent Eye* (McGraw–Hill, New York, 1970).
20. Shepard, R. N. & Chipman, S. Second-order isomorphism of internal representations: shapes of states. *Cogn. Psychol.* **1**, 1–7 (1970).
21. Glymour, C. *The Mind's Arrows: Bayes Nets and Graphical Causal Models in Psychology* (MIT Press, Cambridge, Massachusetts, 2001).
22. O'Brien, G. & Opie, J. in *Representations in Mind: New Approaches to Mental Representation* (eds Clapin, H., Staines, P. & Slezak, P.) (Greenwood, Westport, Connecticut, in the press).
23. Sutton, R. S. & Barto, A. G. *Reinforcement Learning: an Introduction* (MIT Press, Cambridge, Massachusetts, 1998).
24. Hutchins, E. *Cognition in the Wild* (MIT Press, Cambridge, Massachusetts, 1995).
25. Tomasello, M. *The Cultural Origins of Human Cognition* (Harvard Univ. Press, Cambridge, Massachusetts, 1999).
26. Bacon, F. *Novum Organum* (1660). Modern edition (eds Jardine, L. & Silverthorne, M.) (Cambridge Univ. Press, Cambridge, UK, 2000).
27. Medawar, P. B. *Advice to a Young Scientist* (Basic Books, New York, 1981).
28. Crick, F. H. C. *What Mad Pursuit: a Personal View of Scientific Discovery* (Basic Books, New York, 1988).
29. Spirtes, P., Glymour, C. & Scheines, R. *Causation, Prediction, and Search* (Springer, New York, 1993).

Acknowledgements
We are grateful for advice and assistance from F. Crick, A. Damasio, L. Goble, E. McAmis and S. Rickless.

Online links

FURTHER INFORMATION
Al Seckel's laboratory: http://neuro.caltech.edu/~seckel/
Dale Purves's laboratory: http://www.purveslab.net/
Donald D. Hoffman's laboratory:
http://aris.ss.uci.edu/cogsci/personnel/hoffman/
Edward H. Adelson's laboratory:
http://www-bcs.mit.edu/people/adelson/
Grand Illusions: http://www.grand-illusions.com/
Illusionworks: http://psylux.psych.tu-dresden.de/i1/kaw/diverses%20Material/www.illusionworks.com/
MIT Encyclopedia of Cognitive Sciences:
http://cognet.mit.edu/MITECS/
Kant, Immanuel
Richard L. Gregory online: http://www.richardgregory.org/
Access to this interactive links box is free online.