OPINION

Applications of real-time fMRI

R. Christopher deCharms

Abstract | For centuries people have aspired to understand and control the functions of the mind and brain. It has now become possible to image the functioning of the human brain in real time using functional MRI (fMRI), and thereby to access both sides of the mind–brain interface — subjective experience (that is, one's mind) and objective observations (that is, external, quantitative measurements of one's brain activity) — simultaneously. Developments in neuroimaging are now being translated into many new potential practical applications, including the reading of brain states, brain–computer interfaces, communicating with locked-in patients, lie detection, and learning control over brain activation to modulate cognition or even treat disease.

From the times of the ancient philosophers through to the modern pursuits of cognitive neuroscience and neurology, it has been a human passion to comprehend the physical basis of what we experience subjectively as the mind. The Western attempt to 'see' the mind in the brain has been ongoing since before Galileo looked at the planets through the telescope or Hooke looked at cells through the microscope. René Descartes famously used introspection and reason as tools to try to discover the basis of mental events[1]. In this historical context, our generation is the first to explore increasingly direct glimpses of the mind–body interface, through the science of neuroimaging. We can perceive the flow of our subjective experiences through introspection and, using neuroimaging technology, can simultaneously view a display of the physical processes that are taking place in our brain during those experiences ('introneuroimaging'). Imagine Descartes' wonder (FIG. 1) at our modern capabilities for mapping the physical substrates of our own minds.

Over the past few decades, neuroimaging has measured the patterns of brain activation that are associated with different cognitive processes, and has thereby illuminated previously hidden terrains of human brain function[2,3]. In attempts to create 'maps' of the functional roles of the many regions of the brain that are applicable to all people, neuroimaging measurements have typically been pooled and averaged across many individuals and across many repetitions of a task. Now, real-time functional MRI (rtfMRI) is exploring the possibility of watching one's own brain activation 'live'. The ability to observe one's own brain as the mind's processes unfold might allow us to become aware of and learn to control some of the most important aspects of human life: conscious experience, cognition, emotion, action, non-conscious functions, and even the breakdown of these processes in disease. On a technical level this possibility is brought about by recent advances in neuroimaging and computing[4–7]. Early experiments in this new field are just taking shape, as discussed in this Perspective: methods have been developed for reading patterns of brain activation in real time, for manipulating computerized devices using only the brain, for communicating with a patient who was thought to be in a vegetative state, for learning to control individual regions in one's own brain and thereby alter one's cognition, and potentially even for controlling disease symptoms such as chronic pain.

In the past twenty years there has been a revolution in our understanding of the human brain and the localization of processes that were largely outside the bounds of biological science a generation ago, such as executive function, mental imagery, emotion and conscious experience. As an acknowledged supporter of rtfMRI, in this Perspective article I describe where these developments — specifically the new technology of rtfMRI — have led, and I provide some conjectures regarding their possible applications in the foreseeable future. The more technical aspects of this field have been reviewed elsewhere[6–9].

rtfMRI methods and prior approaches

Compared with prior methods for measuring brain function, functional neuroimaging provides measurements of brain physiology that are highly distributed (sampling very large numbers of spatial locations, often spanning the brain) and highly parallel (providing an ongoing stream of information from each of the many measurement points). For example, MRI can currently sample from ~2^{16} spatial locations per second (FIG. 1), each location with a dimension on the order of 3x3x3mm[10–12]. The technical brilliance of MRI is that it provides a unique means by which to precisely 'address' each point in space on the basis of the physical properties of magnetic resonance, and thereby

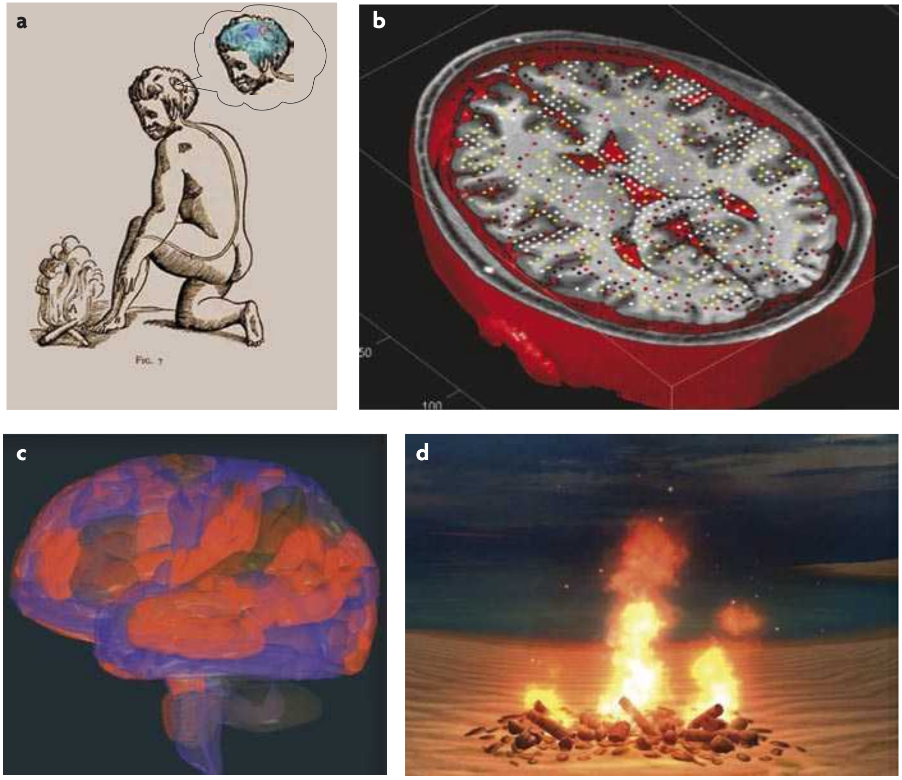

Figure 1 | **Imaging the human brain in real time.** **a** | Descartes used introspection as a way to perceive the mechanisms of the mind. This approach to observing the mind–brain interface is what people were limited to in the absence of technology. **b** | Today, using real-time functional MRI, it is possible to measure the level of activation from ~2^{16} brain locations per second. Here, measured activation levels are represented as colours that have been overlaid onto a three-dimensional rendered set of anatomical brain images. **c** | Information from individual spatial points can be segregated into multiple anatomically defined three-dimensional regions of interest. Here the activation levels (represented as colours) of three brain regions are rendered on a translucent 'glass brain' view. **d** | Activation in these regions can either be plotted second-by-second in real time or can be presented to subjects in more abstract forms, such as this virtual-reality video display of a beach bonfire, in which each of the three elements of the flickering fire corresponds to activation in a particular brain region. Brain activation can control arbitrarily complex elements of computer-generated scenarios.

to create two-dimensional (2D) images or three-dimensional (3D) volumes. This is possible because, in response to a radio-frequency energy pulse inside a high overall magnetic field, precise small differences in the magnetic field at each point in space will cause water at that point to emit a signal at a characteristic resonant frequency and phase that is related to the total magnetic field. Simplifying the actual process somewhat for illustration, magnetic fields that spatially vary (that is, they are lower at one end of an object being imaged, higher at the other end) are applied, so that each point in space receives a different magnetic field and therefore emits a magnetic-resonance signal with a different characteristic frequency and phase. The different signals can be measured simultaneously, and their spatial locations can then be sorted on the basis of their resonant frequency. Conceptually, this is like hearing many notes on a piano being played simultaneously, and then separating out the location of the key that produced each note on the basis of the note's pitch. Specialized antennas in the MRI machine that are positioned around the subject's head measure the signals that are emitted from many points all at once, and the spatial origin of each signal component is then separated mathematically from the total signal on the basis of each signal component's time, frequency and phase. For example, a 3D volume of the head might be collected in 0.5–2.0 seconds by collecting 16 individual 64 × 64 voxel images and then concatenating them to form a 3D volume. MRI can therefore capture something close to an instantaneous 3D snapshot of the internal features of an object (in this case, a head). Not all of the voxels will typically be useful. For example, depending on how a 3D volume that is being measured is oriented relative to a person's head, many of the voxels in the images will typically lie outside the brain or will lie in anatomical structures that are not expected to show meaningful signals, such as bone or cerebrospinal fluid.

To create a high-quality MRI image (that is, a static image of the internal structure of an object), a similar process is used but much more data are collected in order to produce images of higher spatial resolution and with less noise. This often requires 2–10 minutes rather than 0.5–2.0 seconds[12].

fMRI extends the MRI process by taking advantage of the speed with which 3D MRI volumes can be collected, although the quality of each individual fMRI image is lower. In fMRI the same volume is sampled repeatedly at short, regular intervals (for example, once per second) using an MRI measurement that is sensitive to changes in blood iron concentration and oxygenation. Fortunately, the magnetic properties of deoxygenated iron in blood haemoglobin are such that a voxel becomes slightly brighter or darker in response to changes in deoxygenated iron content in that voxel. This means that blood itself serves as a contrast agent, making the MRI signal sensitive to brain activation. Changes in activation can therefore be inferred from the voxels' changes in brightness at each time point. The fMRI procedure is thought to be quite safe, uses the subject's blood itself rather than an exogenous contrast agent, and requires only a strong magnetic field and radio-frequency energy. fMRI does not typically require injections, drugs, radioactive substances or X-rays. The spatial resolution of MRI is limited by the spatial-point spread function of the scanner, by the signal-to-noise ratio of the measurements and by other spatial artefacts[12]. Moreover, as fMRI measures changes in blood volume and flow in the complex vasculature that provides blood to neural tissue, the signal 'spreads' relative to the initial neural source.

The MRI process is inherently spatially addressed, and therefore it is conceptually quite different from the processes that are used in electroencephalography (EEG) and magnetoencephalography (MEG), in which measurements are made from locations near the surface of the head and mathematical models are used to infer the activation sources in the brain from the measured signals. Using a relatively small number of measurement points outside the brain (for example, 16–128), EEG and MEG methods essentially triangulate the source of a signal on the basis of its time of arrival, its phase and magnitude, and the coherence of the signals that reach the sensors. The 'inverse problem' of trying to derive the neural origin of signals that were measured outside the brain cannot be solved exactly, as an infinite number of activation patterns in the brain can produce a given pattern of signals at the external sensors. This manifests itself in uncertainty regarding where a signal originated from within the brain. Nevertheless, EEG and MEG have been applied intensively and with increasing spatial accuracy, and they can now be used to localize neural activation with millimeter-range accuracy if the same stimulus or task is repeated many times[13] to allow averaging. Moreover, these techniques have millisecond-range temporal resolutions, whereas the temporal resolution of fMRI is in the range of seconds. EEG is also both inexpensive and portable.

EEG and MEG can also be used for real-time measurements. Importantly, however, the spatial source localization that is achievable using EEG or MEG is far worse in real-time applications than in more typical applications in which the same trial is repeated many times. Without the opportunity to average over multiple identical trials, the mathematical models that are used to try to solve the inverse problem and localize signal sources in the brain produce broader spatial localization, often only localizing signals to large regions of the brain, such as the parietal lobe, rather than to individual brain structures, such as a particular cortical area. EEG signals have been used extensively as a basis for real-time feedback-based training, but here the emphasis has been primarily on the precise temporal character of the signals — for example, training subjects to control different frequency bands of oscillation — rather than on well-localized brain structures. fMRI can provide signals from well-localized individual brain regions in real time because it does not need to rely on event-related averaging to try to solve the inverse problem: it is based on the physical properties of magnetic resonance, which inherently provide spatial localization. An interesting compromise involves inexpensive and portable imaging techniques that are based on shallow-penetrating photons, such as near-infrared spectroscopy, which can produce relatively well-localized signals and provide additional information about blood oxygenation state but which currently only operates effectively for structures on or near the brain surface[14].

rtfMRI was originally developed by Cox[4]. The steps that are involved in real-time processing of fMRI images are conceptually similar to those that are used for offline fMRI analysis, but the process is accelerated by faster imaging sequences, computer processing and algorithms. What is new in rtfMRI is that the data are analysed as they are collected, so that the resulting information is immediately available and can thus be used to guide a person's cognitive processes, an experimenter's parameter selections or a clinician's interventions. Additional analysis of the images that have been obtained can take place offline.

Glossary

Biofeedback
A technique in which a continuous measure of some aspect of a person's biology is presented to that person for the purpose of training them to control the measure and, thereby, the corresponding biological function.

Contrast agent
A chemical agent that, when injected into a person, increases the measured contrast (the difference in image intensity) between different types of tissue. For example, a Gadolinium-based dye is sometimes used in MRI.

EEG
(Electroencephalography). A method for measuring the fast electrical activity in the brain that is associated with neuronal activation.

Locked-in syndrome
A medical state in which a patient has very limited or no ability to communicate with the world, often owing to extensive paralysis.

MEG
(Magnetoencephalography). A method for measuring the fast magnetic activity in the brain that is associated with neuronal activation.

Near-infrared spectroscopy
A method used for measuring brain blood flow and oxygenation near the cranial surface by shining near-infrared light through the skull and measuring the resulting emitted light spectrum, which is indicative of blood properties.

Pattern-classification algorithm
A computer-modelling method for classifying statistical patterns in complex multi-parameter data. For example, pattern-classification algorithms have been built that will classify spatial patterns of fMRI data (2D images or 3D volumes) by estimating what task a subject was undertaking when each particular fMRI pattern was measured.

Persistent vegetative state
A medical condition in which a patient shows sustained unresponsiveness and does not show evidence of awareness.

Region of interest (ROI) analysis
A method for measuring the time course of activation from a selected volume of the brain. This method can be used to infer the average activation in a region of a person's brain that has been caused by a stimulus or task, or conversely a level of ROI activation can be used to attempt to infer what task is being undertaken by a person.

Spatial-point spread function
The amount of spread, through space, of the measured signal that arises from an idealized single point in space. Spatial spread is caused by noise and imperfections in the measurement technique, for instance MRI.

Voxel
A 3D volume element of measurement (for example, a cube). Voxels are the 3D volume equivalent of a pixel in a 2D image.

Limitations of rtfMRI

It is important to define the limitations of a technique in order to understand its potential capabilities. The essential limits to rtfMRI are a consequence of its mechanism: fMRI measures changes in blood flow rather than neuronal activity. The technique is therefore inherently indirect and noisy, and does not allow the measurement of individual neuron's firing[10,11] or precise temporal coding of neuronal activity[15]. When neurons in the brain are activated, blood flow to their region increases with a delay of several seconds and with a spatial spread on the order of hundreds of micrometers to millimeters[16]. This spatial scale is large compared with that of individual neurons — one cubic millimeter of the brain can include more than a million individual cells.

Although the indirect mechanism, the artefacts, and the spatial and temporal limits of fMRI have all formed the basis of substantial and valid criticism of the technique, perhaps it is more interesting to think about what current fMRI methods can accomplish, rather than what they cannot. Spatially precise single-neuron measurements and temporally precise MEG and EEG have historically been windows on the brain from which it has been argued that neural signals depend on precise spatial or temporal coding[15,17–20]. Nevertheless, it is likely that many important brain processes can also be decoded using a large number of parallel measurements at low spatial and temporal resolution, which can be provided by fMRI. By analogy, in some cases it is preferable to have a measurement that shows the forest rather than the trees: for example, in assessing global patterns of vegetation using satellite imagery, very high-resolution data might be spatially averaged anyway. Even if it were possible to measure the billions of neurons of the brain simultaneously and at millisecond precision, the first step in data analysis for many applications would probably be spatial and temporal averaging to reduce data. Therefore, a clear goal is to find applications of fMRI that benefit both from covering large regions of the brain and from the spatial and temporal averaging that is inherent to fMRI. As we increase our understanding of the parallel data streams that fMRI measures from thousands of points in the brain, rtfMRI increases its potential to fundamentally alter our ability to 'read' mental states by decoding this information in real time. We do not yet know much about the optimal practical spatial or temporal scale for reading brain states, for example the trade-off between the amount of information that can be extracted using parallel, high-resolution spatial measurements (MRI) and the high temporal resolution that can be obtained using EEG and MEG, or how the two can best be combined. Which method extracts information more quickly (for example, at a higher bit rate) in a given task? Considering that many cognitive processes change slowly, over seconds or minutes (does someone fluctuate from being in pain to not being in pain and back again many times per second, or even per minute?), fMRI measurements might be well matched to these timescales.

Brain–computer interfaces

There have been several approaches to develop brain–machine interfaces[21–24]. rtfMRI provides a novel interface with the nervous system that offers the possibility of directly reading aspects of a person's brain-activation state. It has potential applications in many contexts, particularly when a patient cannot communicate verbally.

Extracting brain information from distributed patterns. fMRI data have increasingly been analysed using computer pattern-classification algorithms and related methods[25], with the goal of 'mind reading' (REFS 26–36). An essential finding of this emerging field is that it is possible to infer complex cognitive states from brain data collected from large numbers of fMRI voxels, even though the individual voxels have quite poor spatial and temporal resolution relative to neurons. In fact, even fMRI patterns obtained from single trials can in some cases provide enough information to allow a good prediction of what a person is experiencing or doing[34,37,38]. Examples of fMRI-pattern classification include: determining whether individuals were viewing faces or houses on the basis of the spatial pattern of activation in their temporal lobe[26,27,39], establishing the orientation[40,41] or retinal position[42] of images that someone is viewing on the basis of the spatial pattern of activation in their visual cortex, and uncovering what someone is remembering during memory retrieval[43]. Recently it has even become possible to predict the particular image that a person is seeing by using detailed models to decode fMRI measures of activation in the visual cortex[37], a step that is tantalizingly close to directly reading a person's ongoing mental images.

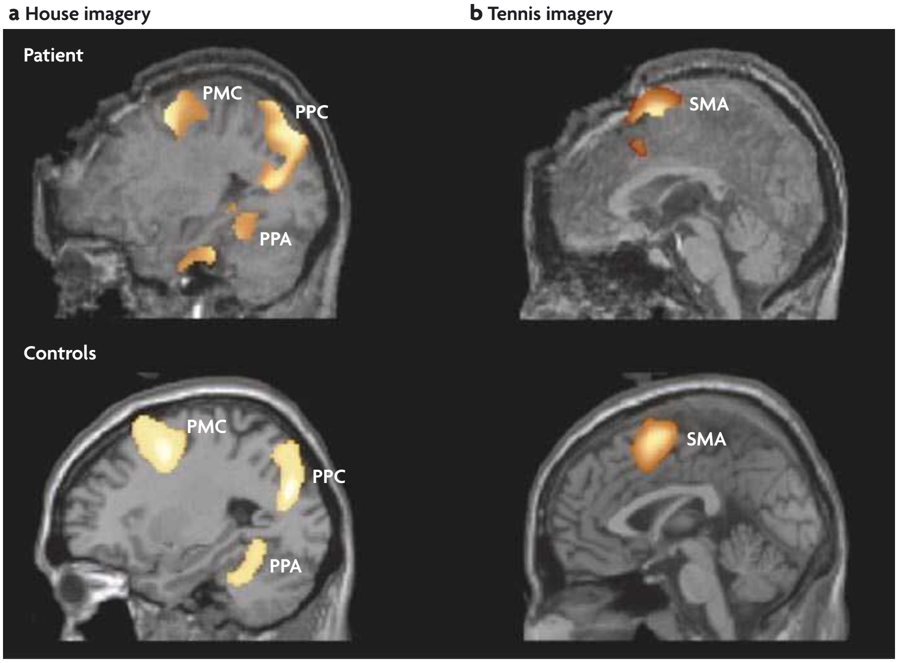

Figure 2 | **Communication with a patient in a vegetative state using functional MRI.** Functional neuroimaging has been used to detect 'awareness' in a patient who had been diagnosed as being in a vegetative state. Following instructions to imagine moving around a house (**a**), activation was observed in the parahippocampal gyrus (PPA), the posterior parietal lobe (PPC) and the lateral premotor cortex (PMC) in both the patient and in a group of 12 healthy volunteers (controls). Following instructions to imagine playing a game of tennis (**b**), activation was observed in the supplementary motor area (SMA) in both the patient and the volunteer subjects. Figure reproduced, with permission, from REF. 46 © (2006) American Association for the Advancement of Science.

Using pattern-classification methods, rtfMRI can potentially read complex brain states in real time. To initially 'build' an fMRI pattern-classification algorithm — a computer model, such as a neural-network model that uses many parameters to 'fit' the input fMRI data — typically requires significant computing time and data, in order to adjust the model's parameters so that it will be able to correlate cognitive states with fMRI input data. However, once a classifier model has been built in this way, it can be run in real time with only modest computational requirements. LeConte *et al.*[44,45] were the first to demonstrate that it is possible to measure patterns of fMRI activation in real time using a support-vector-machine approach: they built an algorithm that assessed rtfMRI patterns that corresponded to movement of the left or the right hand, or to happy or sad emotions, and presented the model's outputs to participants during scanning in an attempt to train them to control complex brain states (see below). If advances in pattern classification provide increasingly precise information about brain and cognitive processes, then this information can potentially increase the power of most rtfMRI applications. Although the approach is conceptually appealing, the key question is whether information that is derived using pattern-classification methods truly provides meaningful practical advantages over much simpler traditional approaches, such as region-of-interest (ROI) analysis.

Communication with locked-in patients and patients in a vegetative state. rtfMRI could potentially be used to communicate with communication-impaired people, such as locked-in patients, or to determine the state of consciousness of a patient in an apparently persistent vegetative state. Owen *et al.*[46] investigated whether it was possible to explore the state of consciousness in a 23-year-old woman who had been diagnosed as being in a persistent vegetative state after an automobile accident (FIG. 2). The woman had been unresponsive for 5 months. The experimenters verbally instructed her to imagine either

Task		rtfMRI		Translation
Mental calculation (A)	→	Brain-activation pattern A	→	Move up
Mental speech generation (B)	→	Brain-activation pattern B	→	Move down
Motor imagery (right hand) (C)	→	Brain-activation pattern C	→	Move right
Motor imagery (left hand) (D)	→	Brain-activation pattern D	→	Move left

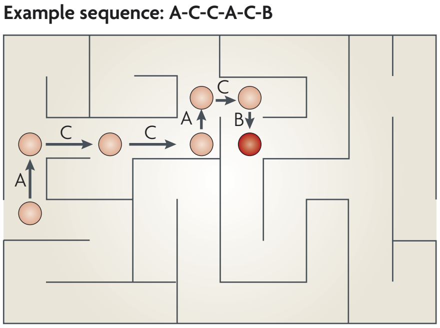
Example sequence: A-C-C-A-C-B

Figure 3 | **Direct mental control over navigation through a maze by monitoring distributed brain activation with fMRI.** Functional MRI (fMRI)–detected levels of activation from four brain regions were used to control the movement of a cursor through a maze[65]. In an initial set-up phase, the subject's brain activity as they engaged in four distinct cognitive processes ('tasks') was measured by fMRI and then arbitrarily assigned to one of the four cardinal directions. The subject was subsequently asked to complete a series of cognitive tasks, and their brain activity as recorded by fMRI during these tasks was then translated by an algorithm into a list of directions that was used to move a cursor around a maze.

playing tennis or walking through her house, although it was unclear whether she could hear or understand them. These instructions resulted in activation in the patient's supplementary motor area (SMA) and parahippocampal place area (PPA), which the authors suggested indicated that she had performed the task. This study highlights two important and currently debated issues, namely whether data from a single individual on a single trial are sufficiently robust to provide information that is reliably correlated to that individual's mental state, and how to determine the mental state of a non-responsive person for the basis of this correlation. A series of comments and further studies related to these issues have subsequently been published (see REFS 47–51). rtfMRI could potentially enable ongoing back-and-forth communication with patients who are not able to communicate by conventional means. Less technology-intensive approaches, such as EEG, could also be used, again with the caveat that there are significant problems in source localization in real-time EEG[23,24]. To the patient or their family, even a single communication session could be of immense value, and could potentially lead to life-or-death decisions. Further work in this area has been reviewed in REFS 50,52,53.

Lie detection. There has been considerable interest, and controversy, regarding the possibility that fMRI will become useful in reading the brain states of people who are unwilling to provide truthful replies to questions[30,54,55]. Studies to date have suggested that in carefully controlled experimental settings it might be possible to distinguish patterns of activation that correspond to deliberate deceit from patterns that correspond to truthfulness. For example, Spence *et al.* demonstrated greater activation in the ventrolateral prefrontal cortex when people were lying than when they were telling the truth[56]. This finding has been replicated, and a number of groups have contributed to this research area[55,57–60].

The ability to reliably detect lies would have dramatic societal implications. If a lie-detection method becomes available, it will have tremendous potential in law enforcement and in the legal system generally[61,62]. Theoretically, such a system could put criminals in prison and keep innocent people out on the basis of their accounts of what happened. In national security it could be used for screening out foreign countries' agents during security clearance, and for counter-terrorism[63]. Similar to polygraph testing, fMRI-based lie detection also has dramatic potential for causing harm if its results are inaccurate or over-interpreted. It also has the potential to be abused, such as by inappropriately screening employees or insurance applicants. Indeed, current polygraph testing remains controversial and is generally inadmissible in court owing to limitations in its accuracy. Neuroimaging-based lie detection methods might follow a similar course[64]. There is a great difference between demonstrating differences across a group of study participants in a carefully controlled fMRI experiment and detecting lies reliably in a single individual in real-world settings in which people might have a strong incentive not to cooperate. Further, the brain is often not an accurate and faithful recorder of events, and the very nature of what it means for a disputed statement to be 'true' is often far from straightforward. Finally, there is a significant challenge regarding how to attain veridical reference brain states corresponding to 'true' or 'false' statements in an individual that could then be compared with the brain states that arise when the individual makes contested statements. Reference states could be derived from the individual in question during pre-testing (leading to problems regarding single-subject reliability), or they could be derived from a group (leading to problems regarding the question of whether the brain state from a given individual would match that of a group of other people).

Nevertheless, rtfMRI could potentially be used during repeated questioning to assess the truthfulness of answers, with the flow of questions being guided by the information that is being gathered, and with the questions being repeated until the results are statistically adequate. Two private companies that aim to commercialize this new application have already been created; this has raised the valid concern that the technology could be interpreted as being authoritative before the data support this conclusion. Although an accurate lie detector could be a truly transformative technology, it remains to be seen whether the technique will be validated across large numbers of individuals in realistic settings. In the meantime, there is significant need for caution in interpretation.

Controlling external devices with the brain. rtfMRI can allow an individual to control an external object using only their brain activation[9,65]. For example, early work has demonstrated that people can play an analogue of the computer game Pong using their brain-activation patterns[66]. In addition, Yoo *et al.*[65] have shown that a participant can control the movement of a cursor in a maze on a computer screen using only mental processes that were read by rtfMRI (FIG. 3). Measurements were made from four brain regions that were activated by separate tasks, and the algorithm was programmed so that activation of each region would lead to movement of the computer cursor in one of the four cardinal directions. Participants engaged in a series of mental tasks or strategies, and the sequential pattern of activation in the brain of each person was used to determine the trajectory of the cursor through the maze after they completed the sequence of tasks.

As new technologies for brain measurement evolve and the equipment that is involved becomes smaller and less expensive, the knowledge that is gained from today's early experiments and technologies might lead to future applications for reading brain signals, such as for communication or for prosthetic control. Various technologies are being explored as potential brain–computer interfaces, including EEG and implanted multi-unit and single-unit recording electrodes[21–24,67]. In addition, rtfMRI might provide important insights into how to extend the current limits of what can be extracted from the brain using fMRI, and might lead to a better understanding of how people can learn to control brain activation in order to manipulate devices[6,7,9].

Learned control over brain activation

Technology now allows us to investigate the functioning of our own brains to the extent that we might be able to answer questions that were previously impossible even to ask. What brain regions and processes are accessible to consciousness and subject to conscious control (BOX 1)? If one can learn to control the activation of a particular brain system, can one also control its cognitive processes or disease processes? Earlier studies conducted over several decades have demonstrated that people can learn to control various physiological functions using biofeedback[67–69]. These functions were often closely linked to autonomic processes, such as heart rate and skin conductance, and to EEG measures[70–73]. I think that it is not particularly helpful to look for a close analogy between rtfMRI-based training and methods of biofeedback, and that it is more insightful to look for similarities among training paradigms on the basis of what the subject learns. For example, in my group's experience, what a subject does internally while learning highly specific cognitive strategies for controlling pain using rtfMRI seems to be closely related to what the same subject might do while using other methods for learning cognitive control over pain, such as in cognitive behavioural therapy. This seems to be quite unrelated to the cognitive processes that would be involved in learning to use heart-rate information to decrease autonomic tone, or EEG information to reduce anxiety. We can expect this new technique to yield different training effects that are as unrelated to each other as are the various functions of the many regions of the brain. Real-time neuroimaging provides us with a high-resolution portal into our own brain, allowing us to investigate which regions in the vast landscape of brain processes our awareness can be trained to reach, and which of the brain's myriad cognitive functions can be consciously controlled.

Learning brain states through mimicry. Much of what we learn is learnt through various forms of mimicry: copying the movements of a parent, imitating the actions of a violin instructor or tennis coach, seeking to emulate humanity's greatest paintings or literature, even learning the skills of mathematics. It has not previously been possible to mimic another person's brain states, because we could not see anyone else's brain states, or even our own. As it becomes possible to measure complex brain states in real time, perhaps it will also become possible to learn to control them through mimicry. Imagine copying a brain-activation pattern that corresponds to enhanced focus or performance. Imagine emulating the brain pattern of someone exercising great compassion. Imagine a man learning to mimic the pattern of brain activation used by a woman to solve a problem, or a woman mimicking the brain activation of a man's cognitive strategy. Imagine if an old person could try to reproduce patterns of activation from when they were young. Imagine mimicking the pattern of brain activation of a person experiencing profound happiness, perhaps involving the controlled release of endogenous modulators of mood. We do not yet know how to mimic such patterns of activation. Indeed, we do not yet know which of these ideas are fanciful and which could one day come to pass; however, we do know that mimicry is a powerful means by which to drive learning and shape one's cognition and behaviour. Perhaps mimicry can now be applied more directly to brain states.

Control over localized brain activation. Several teams have set out to determine the extent to which someone can learn to control activation in individual brain regions, initially focusing on the somatomotor cortex as this region seemed likely to be particularly easy to control[74–76]. Indeed, prior to any training people can produce modest activation of specific parts of their somatomotor cortex by imagining moving the part of their body that corresponds to the area of the somatomotor cortex that is being monitored[74] (FIG. 4). After repeated training using the rtfMRI signal from a brain region to learn to guide the strategies they use to alter the activity of that brain region, most people succeed in learning increased control over activation in the targeted brain region. Thus, people can learn to increase and decrease activation in the somatomotor cortex, even without making overt movements (as monitored by electromyography)[74]. Control participants who received sham fMRI information that did not correspond to their own brain signals did not learn to control activation in this brain region. The effect was quite spatially specific: people learned to target changes in activation specifically to this brain region. This new technology

Box 1 | Explicit versus implicit control over brain activation

Implicit control over brain activation is learned as a skill or through normal development; a person is unaware of exerting it, and it is the more typical form of control that underpins most behaviour and cognition. Individuals exhibit unwitting, implicit control over brain activation all the time: every voluntary action engages the activation of specific brain mechanisms, as does every perception. For example, when learning to play the piano one develops considerable implicit control over one's motor cortex, because it in turn produces motor performance. However, one normally has no awareness of what is going on inside one's brain, no ability to make decisions on the basis of it, and therefore no ability to shape it.

Explicit control over brain activation is learned deliberately; a person is conscious of it and it allows them to take volitional charge over their brain activation through deliberate cognitive choices. For example, one can learn to control activation in a brain region by exerting exactly the type of cognition, emotion, intention, action, perception or mental imagery that will maximize or minimize its activation.

The distinction between implicit and explicit control of brain activation is analogous to the distinction between implicit and explicit memory. The interesting question is the extent to which people can learn to gain greater explicit control over their brain activation through specific training, and what practical consequences result. Real-time functional MRI has the potential to bring normally non-conscious brain processes into conscious awareness, and to transform implicit control over brain activation into explicit control.

allows us to see for the first time that it is possible to select a brain system and teach someone to control it, much like one can isolate a muscle group for training. This type of finding has been reproduced in studies that involved many brain regions, including the somatomotor cortex[38,65,74,76], the PPA[8], the amygdala[77], the auditory cortex[78], the insular cortex[79] and the anterior cingulate cortex (ACC)[65,75,80]. These brain areas have a wide range of functions, including movement, tactile sensation, visual perception, hearing, emotion and pain.

As brain-activation patterns and behaviour are assumed to be closely linked, it seems likely that learned control over brain activation would lead to changes in cognition and behaviour[81]. Neuroimaging techniques can be used to test this hypothesis by training subjects to control their brain activation and then measuring what changes in behaviour or cognition result. Brain activation can be studied as an 'independent variable' when exploring the role of a brain region, reversing the typical methodology of neuroimaging experiments and providing a new means for investigating function.

Control over pain. Our research group proposed that people might be able to learn to control their pain by learning to control the brain regions that mediate pain perception[80]. Chronic pain represents an important medical issue given the large population of patients with pain that is refractory to existing treatment options. The pain modulatory system is comparatively well understood, has clear anatomical targets and has well-defined psychophysical measures. To investigate this possibility in a model system, eight healthy volunteers were trained in an MRI scanner to control activation in a ROI in the rostral ACC (rACC)[82]. Continuous rtfMRI activation information from the rACC was depicted both as a scrolling line graph and as a changing image of a virtual fire. Participants were instructed to alternately increase and then decrease the activation level of the rACC, which they saw continuously, in alternating blocks of 60-second duration. In order to probe the participants' experience of pain, a psychophysically controlled painful thermal stimulus was applied during each period. The stimulus given was always the same for each participant. Participants learned to control the rtfMRI signal in the rACC, and again the effect was spatially specific: the rACC showed a greater increase in activation than any other forebrain region (FIG. 5). The participants' control over rACC activation increased during three training runs of 13 minutes each. In addition, training to modulate rACC activity also engaged other elements of the pain modulatory system, including the secondary somatosensory cortex, the insula, the supplementary motor cortex, the superior cerebellum and the superior temporal gyrus. All of these brain regions have shown activation in tasks that involve pain processing, pain control or the placebo effect[83–88].

Importantly, as participants learned to modulate activation of their rACC, there was a corresponding change in their perception of pain. The same painful stimulus was rated as more painful during periods in which the participants increased their rACC activation than during periods in which they decreased it. To determine whether this effect was specifically due to rtfMRI-induced learning, rather than to other learning effects or to nonspecific effects, participants in the experimental group were compared with four control groups of participants who received either: extended practice without rtfMRI information, twice the duration of training to focus attention away from pain, training using rtfMRI data taken from a different brain area, or training using sham rtfMRI data taken from a different participant. The learned control over pain perception demonstrated by the experimental (rtfMRI) group was significantly larger than for any of the four independent control groups[80], none of which showed significant effects. The fourth control group is particularly important, in that each control subject saw exactly the same displays that an experimental subject had seen during successful training — that is, the information taken from the experimental subject's brain. Therefore, any positive expectation of learned control over pain that might have been created in the experimental subject as they watched their success in learning control over their brain activation would have been highly similar in the control subject, but the information would have been of no use to the control subject because it did not reflect their own cognitive or neural processes.

Modes of training. It is not yet clear what types of learning paradigms will benefit most from rtfMRI information. An important area for future research is the development of new types of task paradigms. A particularly important question is whether it is possible to train participants without them having to engage in a deliberate cognitive process. Bray *et al.* provided the first example of this approach: they trained participants to control activation in the somatomotor cortex, but rather than showing them rtfMRI feedback information they provided them with rewards when they completed the task correctly[89]. It is also possible that rtfMRI-based training might

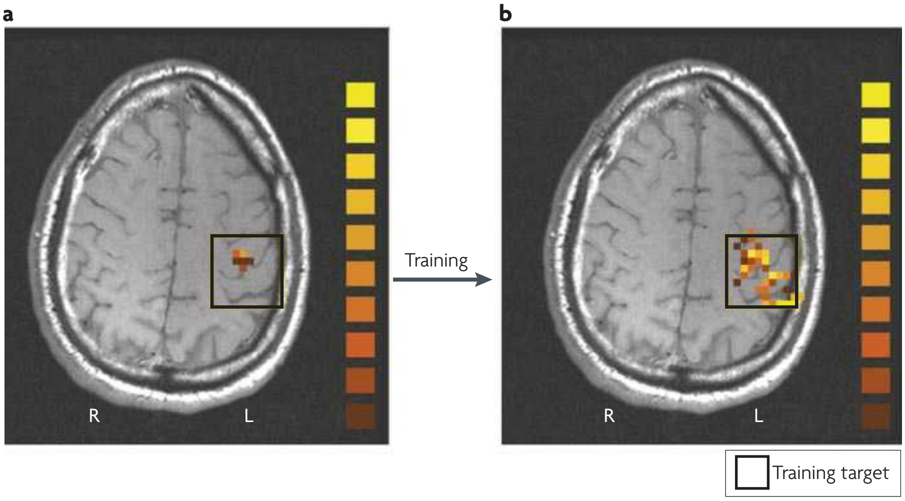

Figure 4 | **Learned explicit control over activation in a targeted brain region.** Functional MRI blood-oxygen-level-dependent (BOLD) activation was measured before and after training during a task that involved subjects imagining moving their right hand. The images show BOLD activation of a representative participant. **a** | Before training the subject was able to produce a modest level of activation by imagining moving their right hand. **b** | Following three blocks of twenty minutes during which the subject watched a live signal from this brain region and attempted to increase their control over its activation, there was a significant increase in brain activation that was specific to the target region. The activation patterns are shown here superimposed on a T1-weighted anatomical image. The black square designates the selected region of interest that was used to generate the training signal. Figure reproduced, with permission, from REF. 74 © (2004) Academic Press.

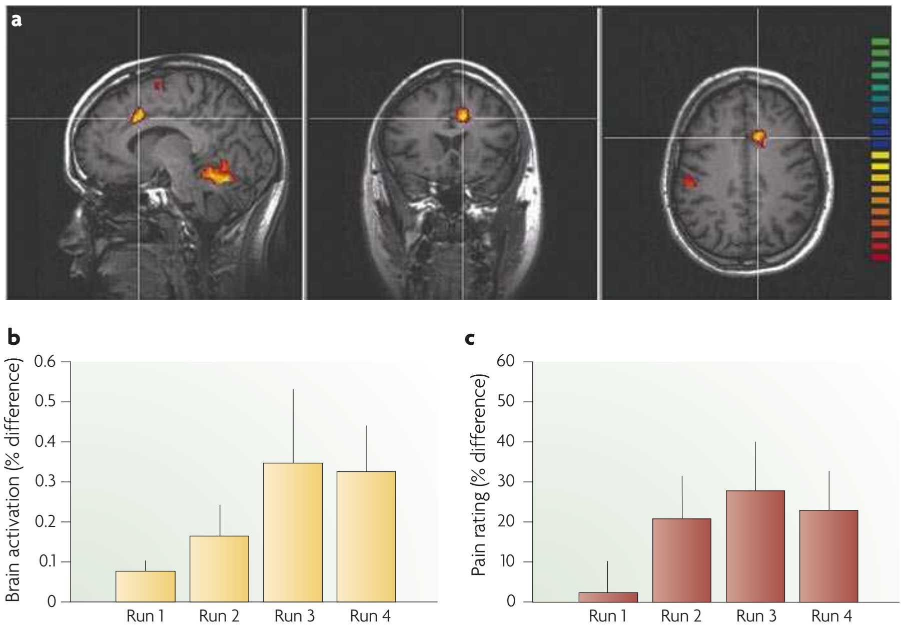

Figure 5 | **Training leads to spatially specific explicit control over brain activation, with corresponding changes in pain perception. a** | Subjects were able to learn (through training sessions) to control activation in a region of their rostral anterior cingulate cortex (rACC) (indicated by the cross-hairs), which is thought to be involved in pain perception and pain control. The images show the change in activation that occurred from the first training session to the last training session. **b** | The graph shows the increase in the average difference in functional MRI (fMRI) signals from the rACC between periods in which the participants had to increase rACC activity and periods in which they had to decrease rACC activity (for each of four runs). Each run consisted of five alternating increase–decrease cycles, demonstrating that with practice subjects showed increasing control over rACC activation. **c** | Over the course of training, subjects also showed a parallel and statistically significant increase in control over their subjective experience of pain induced by a psychophysically controlled painful heat stimulus that was identical on all trials. Subjects underwent three training sessions (Runs 1–3) and were then tested in a final run (Run 4). Four control groups who received similar training but without real-time fMRI feedback did not show this effect (data not shown). Figure modified, with permission, from REF. 80 © (2005) National Academy of Sciences.

have performance-enhancement applications, for example it might be of use for boosting memory by increasing activation in memory-related brain areas.

Therapeutic applications

Potential for a novel treatment method for targeting CNS disease. Current therapeutic approaches to spatially targeting and controlling brain activity in disease are limited: pharmacology targets systems of neurons on the basis of their selective expression of drug receptors, 'talk therapy' targets cognitive processes and mental constructs, and invasive approaches target structures surgically. rtfMRI-based training might provide a novel means for targeting the source of disease anatomically. If rtfMRI can be used to regulate brain activation in precisely selected neural systems, this would open up the possibility of therapeutic intervention that is non-invasive and non-pharmacologic and that consequently might have few side effects. As brain systems can be enduringly changed through use[90,91], rtfMRI might produce long-term benefit for patients, a topic that is under active investigation. There is already a long, occasionally controversial, history of studies that investigated EEG-based training for the treatment of disease, particularly for attention-deficit disorder[72,92–94].

The potential for therapeutic applications of rtfMRI was investigated in a study in which eight patients with chronic pain were taught to control the activity of the brain's pain system in order to assess whether this would lead to a reduction in pain[80]. Patients reported a significant decrease in pain immediately following a single session of rtfMRI training, as measured using the McGill Pain Questionnaire and a visual analogue scale. This effect was significantly greater than that in a group of patients who received a similar period of autonomic heart-rate and respiration-rate biofeedback. Although encouraging, these results are from a small initial study and should not be over-interpreted; it will be important to determine whether they are replicable and whether they are truly due to control over brain activation on the basis of rtfMRI rather than due to expectation or placebo effects. Indeed, ongoing clinical trials are attempting to determine whether changes of this type are enduring and will produce lasting benefit for patients.

As rtfMRI requires expensive equipment, cost is an important factor in considering whether rtfMRI-based treatments are feasible for any particular condition. The focus should be on severe conditions for which less-expensive forms of therapy are not available. For example, if it is efficacious, rtfMRI might be less costly and more successful than other, invasive approaches (particularly surgical interventions) that are already in widespread use for the treatment of chronic refractory pain. In addition, as the patents surrounding MRI technology continue to expire, manufacturers are developing new technologies to make MRI scanners smaller, cheaper and increasingly suitable for the physician's office. Several manufacturers of MRI scanners now offer rtfMRI software on their scanners, developed largely for use in pre-surgical planning.

In the future, rtfMRI-based approaches might be applied to other clinical conditions that are thought to be associated with specific brain structures. For example, they might be used in stroke rehabilitation (to target remaining viable tissue) and in the treatment of epilepsy (to control seizure foci), depression (to control neuromodulatory structures or their targets) and addiction (to control systems involved in reward, craving and addiction). The fundamentals of neurology and neuroscience suggest that if the structures that underpin disease can be controlled, then disease itself can be modified. The essential question is: if patients can see brain activation from a target structure using rtfMRI, can they also learn to control it?

Neuroimaging therapy. Another potential therapeutic application for rtfMRI is in augmenting psychotherapy and in measuring the effects of psychotherapy on the brain. A central goal of psychotherapy is that the patient and the therapist better understand and ultimately change the patient's cognitive (and neurophysiological) processes. Patients and clinicians might benefit from watching the brain as it functions during this process. Experiments have shown that participants can be trained to control their

Box 2 | Ethical considerations

The potential power to read information from a person's brain also leads to a new frontier of personal confidentiality: mind privacy. The foreseeable ability to read the state of a person's brain and thereby to know aspects of their private mental experience, as well as the potential to incorrectly interpret brain signals and draw spurious conclusions, raises important new questions for the nascent field of neuroethics. One can imagine the scenario of someone having their private thoughts read against their will, or having the contents of their mind used in ways of which they do not approve. Many of these concerns also apply to non-real-time applications.

There are also important questions about the use of real-time functional MRI (fMRI) training in controlling brain activation, such as potential health risks, particularly during the early investigative stages of this approach. Although concern is warranted, it is important that concerns be based in the realities of the technology. There are inherent limits in the technology that mitigate against the possibility of reading a person's mental states against their will: participants in fMRI experiments currently have to be willing and active participants. For example, they must typically remain extremely still and perform prolonged or repeated mental processes that are inherently voluntary because they are internal. However, approaches are being developed to allow fMRI to function even during movement and to overcome physiological noise.

Who should control a person's fMRI data and the circumstances under which someone might be compelled to submit to an fMRI exam are potentially important future legal/policy issues. As has been true in many fields of life science and biotechnology when a potentially powerful new technology has been developed, dialogue with the public, caution and careful consideration of ethics will be important, particularly among investigators and technology developers, whereas alarmist rhetoric will probably not be productive. As research in this area continues, the potential benefits of each application will need to be weighed against potential harm as data regarding utility accumulate.

emotion while viewing real-time activation from areas that are involved in emotional processing, including the insular cortex[79] and the amygdala[77,95]. Our group performed a pilot study using real-time brain imaging during a psychiatric treatment session for an anxiety disorder[96], demonstrating the use of rtfMRI during cognitive behavioural therapy as a potential means for measuring and augmenting the technique. In this study, patients with obsessive–compulsive disorder received a specially adapted form of cognitive therapy while lying inside a scanner and watching the activation in brain regions that are associated with anxiety disorder. The patients were primarily asked to monitor and learn to control their own brain activation while they also interacted with their therapist through a two-way audio system.

Therapy during brain imaging has obvious challenges: the patient and therapist have far less direct interaction than is typically desired and the setting is confining, foreign and possibly frightening. The pilot work that has been performed to date demonstrates that, even in patients with moderate-to-severe anxiety disorders, it is possible to perform extended neuroimaging therapy sessions that the patients themselves describe in positive terms. The therapist and patient watch in real time as cognitive processes unfold and brain activation changes, and so observe the direct impact of the cognitive intervention. rtfMRI data can potentially be used as a quantitative endpoint to objectively measure the impact on the brain of therapeutic interventions. However, this is only meaningful to the extent that rtfMRI data can provide relevant measures (surrogate endpoints) that truly correlate with successful treatment. It is also possible that successful traditional therapy changes cognitive or neural processes that rtfMRI does not measure.

Conclusions and future directions

Real-time functional brain imaging enables us for the first time to look inside our own brains and view the biological underpinnings of our own unique conscious experiences and behaviours, and potentially the causes of psychiatric or neurological diseases. Although the future outcomes of this approach cannot be clearly seen today, it is possible that, just as neuroimaging has been transformative to basic science and to our understanding of the brain over the past few decades, rtfMRI could have equally transformative new practical applications in the coming decades, potentially in fields as widespread as neurology, psychiatry and law (BOX 2). In the initial phase of what will probably be decades of exploration, neuroimaging has arguably already been used to detect consciousness in patients in what was diagnosed as a vegetative state, to detect lying, to control brain–computer interfaces, and to train people to control their own brain activation as a means of modulating cognitive function and disease. No one is yet in a position to accurately predict the ultimate success or failure of any of these applications as they grow from compelling ideas and initial pilot experiments into robust, widespread practical applications. Although the limitations of early attempts to use a new technology should not be considered failings, neither is it helpful to predict world-changing positive outcomes of a technology that is truly still in its infancy. Instead, hard work, focus on ethical considerations (BOX 2), and creative efforts to improve areas of current weakness will continue to stimulate exploration in this new technology that allows us to see inside our own minds and brains like never before.

R. Christopher deCharms is at the Omneuron 3T MRI Research Center, 99 El Camino Real, Menlo Park, California 94025, USA.
e-mail: decharms@gmail.com

doi:10.1038/nrn2414

1. Descartes, R., Haldane, E. S. & Ross, G. R. T. *The Philosophical Works of Descartes* Vol. 2 (Cambridge Univ. Press, 1981).
2. Cabeza, R. & Kingstone, A. *Handbook of Functional Neuroimaging of Cognition* Vol. 3 (MIT Press, Cambridge, Massachusetts, 2001).
3. Gazzaniga, M. S. *The Cognitive Neurosciences* 3rd edn (MIT Press, Cambridge, Massachusetts, 2004).
4. Cox, R. W., Jesmanowicz, A. & Hyde, J. S. Real-time functional magnetic resonance imaging. *Magn. Reson. Med.* **33**, 230–236 (1995).
5. Cohen, M. S. Real-time functional magnetic resonance imaging. *Methods* **25**, 201–220 (2001).
6. deCharms, R. C. Reading and controlling human brain activation using real-time functional magnetic resonance imaging. *Trends Cogn. Sci.* **11**, 473–481 (2007).
7. Weiskopf, N. *et al.* Real-time functional magnetic resonance imaging: methods and applications. *Magn. Reson. Imaging* **25**, 989–1003 (2007).
8. Weiskopf, N. *et al.* Self-regulation of local brain activity using real-time functional magnetic resonance imaging (fMRI). *J. Physiol. (Paris)* **98**, 357–373 (2004).
9. Weiskopf, N. *et al.* Principles of a brain-computer interface (BCI) based on real-time functional magnetic resonance imaging (fMRI). *IEEE Trans. Biomed. Eng.* **51**, 966–970 (2004).
10. Buxton, R. B. *Introduction to Functional Magnetic Resonance Imaging: Principles and Techniques* (Cambridge Univ. Press, Cambridge, UK, 2002).
11. Bandettini, P. A., Aguirre, G. K. & Moonen, C. T. W. *Functional MRI. Medical Radiology* (Springer, Berlin; New York, 1999).
12. Bernstein, M. A., King, K. F. & Zhou, Z. J. *Handbook of MRI Pulse Sequences* (Academic, Amsterdam; Boston, 2004).
13. Lopes da Silva, F. Functional localization of brain sources using EEG and/or MEG data: volume conductor and source models. *Magn. Reson. Imaging* **22**, 1533–1538 (2004).
14. Irani, F. *et al.* Functional near infrared spectroscopy (fNIRS): an emerging neuroimaging technology with important applications for the study of brain disorders. *Clin. Neuropsychol.* **21**, 9–37 (2007).
15. deCharms, R. C. & Zador, A. Neural representation and the cortical code. *Annu. Rev. Neurosci.* **23**, 613–647 (2000).
16. Zarahn, E. Spatial localization and resolution of BOLD fMRI. *Curr. Opin. Neurobiol.* **11**, 209–212 (2001).
17. deCharms, R. C. & Merzenich, M. M. Primary cortical representation of sounds by the coordination of action-potential timing. *Nature* **381**, 610–613 (1996).
18. Averbeck, B. B. & Lee, D. Coding and transmission of information by neural ensembles. *Trends Neurosci.* **27**, 225–230 (2004).
19. Gawne, T. J. Temporal coding as a means of information transfer in the primate visual system. *Crit. Rev. Neurobiol.* **13**, 83–101 (1999).

20. Reich, D. S. *et al.* Interspike intervals, receptive fields, and information encoding in primary visual cortex. *J. Neurosci.* **20**, 1964–1974 (2000).
21. Lebedev, M. A. & Nicolelis, M. A. Brain-machine interfaces: past, present and future. *Trends Neurosci.* **29**, 536–546 (2006).
22. Mason, S. G. *et al.* A comprehensive survey of brain interface technology designs. *Ann. Biomed. Eng.* **35**, 137–169 (2007).
23. Birbaumer, N. Breaking the silence: brain-computer interfaces (BCI) for communication and motor control. *Psychophysiology* **43**, 517–532 (2006).
24. Birbaumer, N. & Cohen, L. G. Brain-computer interfaces: communication and restoration of movement in paralysis. *J. Physiol.* **579**, 621–636 (2007).
25. Duda, R. O., Hart, P. E. & Stork, D. G. *Pattern Classification* 2nd edn (Wiley, New York, 2001).
26. Cox, D. D. & Savoy, R. L. Functional magnetic resonance imaging (fMRI) “brain reading”: detecting and classifying distributed patterns of fMRI activity in human visual cortex. *Neuroimage* **19**, 261–270 (2003).
27. Haxby, J. V. *et al.* Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* **293**, 2425–2430 (2001).
28. Norman, K. A. *et al.* Beyond mind-reading: multi-voxel pattern analysis of fMRI data. *Trends Cogn. Sci.* **10**, 424–430 (2006).
29. Carlson, T. A., Schrater, P. & He, S. Patterns of activity in the categorical representations of objects. *J. Cogn. Neurosci.* **15**, 704–717 (2003).
30. Davatzikos, C. *et al.* Classifying spatial patterns of brain activity with machine learning methods: application to lie detection. *Neuroimage* **28**, 663–668 (2005).
31. Hanson, S. J., Matsuka, T. & Haxby, J. V. Combinatorial codes in ventral temporal lobe for object recognition: Haxby (2001) revisited: is there a “face” area? *Neuroimage* **23**, 156–166 (2004).
32. Martinez-Ramon, M. *et al.* fMRI pattern classification using neuroanatomically constrained boosting. *Neuroimage* **31**, 1129–1141 (2006).
33. O’Toole, A. J. *et al.* Partially distributed representations of objects and faces in ventral temporal cortex. *J. Cogn. Neurosci.* **17**, 580–590 (2005).
34. Dehaene, S. *et al.* Inferring behavior from functional brain images. *Nature Neurosci.* **1**, 549–550 (1998).
35. LaConte, S. *et al.* The evaluation of preprocessing choices in single-subject BOLD fMRI using NPAIRS performance metrics. *Neuroimage* **18**, 10–27 (2003).
36. Mitchell, T. M. *Machine Learning*. *McGraw-Hill Series in Computer Science* (McGraw-Hill, New York, 1997).
37. Kay, K. N. *et al.* Identifying natural images from human brain activity. *Nature*, **452**, 352–355 (2008).
38. Posse, S. *et al.* A new approach to measure single-event related brain activity using real-time fMRI: feasibility of sensory, motor, and higher cognitive tasks. *Hum. Brain Mapp.* **12**, 25–41 (2001).
39. Haynes, J. D. & Rees, G. Decoding mental states from brain activity in humans. *Nature Rev. Neurosci.* **7**, 523–534 (2006).
40. Haynes, J. D. & Rees, G. Predicting the stream of consciousness from activity in human visual cortex. *Curr. Biol.* **15**, 1301–1307 (2005).
41. Kamitani, Y. & Tong, F. Decoding the visual and subjective contents of the human brain. *Nature Neurosci.* **8**, 679–685 (2005).
42. Thirion, B. *et al.* Inverse retinotopy: inferring the visual content of images from brain activation patterns. *Neuroimage* **33**, 1104–1116 (2006).
43. Polyn, S. M. *et al.* Category-specific cortical activity precedes retrieval during memory search. *Science*, **310**, 1963–1966 (2005).
44. LaConte, S. *et al.* Support vector machines for temporal classification of block design fMRI data. *Neuroimage*, **26**, 317–329 (2005).
45. Laconte, S. M., Peltier, S. J. & Hu, X. P. Real-time fMRI using brain-state classification. *Hum. Brain Mapp.* **28**, 1033–1044 (2006).
46. Owen, A. M. *et al.* Detecting awareness in the vegetative state. *Science* **313**, 1402 (2006).
47. Nachev, P. & Husain, M. Comment on “Detecting awareness in the vegetative state”. *Science* **315**, 1221; author reply 1221 (2007).
48. Greenberg, D. L. Comment on “Detecting awareness in the vegetative state”. *Science*, **315**, 1221; author reply 1221 (2007).
49. Coleman, M. R. *et al.* Do vegetative patients retain aspects of language comprehension? Evidence from fMRI. *Brain* **130**, 2494–2507 (2007).
50. Owen, A. M. *et al.* Using functional magnetic resonance imaging to detect covert awareness in the vegetative state. *Arch. Neurol.* **64**, 1098–1102 (2007).
51. Di, H. B. *et al.* Cerebral response to patient’s own name in the vegetative and minimally conscious states. *Neurology* **68**, 895–899 (2007).
52. Bernat, J. L. & Rottenberg, D. A. Conscious awareness in PVS and MCS: the borderlands of neurology. *Neurology* **68**, 885–886 (2007).
53. Owen, A. M. & Coleman, M. R. Functional neuroimaging of the vegetative state. *Nature Rev. Neurosci.* **9**, 235–243 (2008).
54. Spence, S. A. *et al.* A cognitive neurobiological account of deception: evidence from functional neuroimaging. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **359**, 1755–1762 (2004).
55. Langleben, D. D. *et al.* Telling truth from lie in individual subjects with fast event-related fMRI. *Hum. Brain Mapp*. **26**, 262–272 (2005).
56. Spence, S. A. *et al.* Behavioural and functional anatomical correlates of deception in humans. *Neuroreport* **12**, 2849–2853 (2001).
57. Langleben, D. D. *et al.* Brain activity during simulated deception: an event-related functional magnetic resonance study. *Neuroimage* **15**, 727–732 (2002).
58. Ganis, G. *et al.* Neural correlates of different types of deception: an fMRI investigation. *Cereb. Cortex* **13**, 830–836 (2003).
59. Kozel, F. A. *et al.* Detecting deception using functional magnetic resonance imaging. *Biol. Psychiatry* **58**, 605–613 (2005).
60. Nunez, J. M. *et al.* Intentional false responding shares neural substrates with response conflict and cognitive control. *Neuroimage* **25**, 267–277 (2005).
61. Thompson, S. K. The legality of the use of psychiatric neuroimaging in intelligence interrogation. *Cornell Law Rev.* **90**, 1601–1637 (2005).
62. Appelbaum, P. S. Law & psychiatry: the new lie detectors: neuroscience, deception, and the courts. *Psychiatr. Serv.* **58**, 460–462 (2007).
63. Wild, J. Brain imaging ready to detect terrorists, say neuroscientists. *Nature* **437**, 457 (2005).
64. Meegan, D. V. Neuroimaging techniques for memory detection: scientific, ethical, and legal issues. *Am. J. Bioeth.* **8**, 9–20 (2008).
65. Yoo, S. S. *et al.* Brain-computer interface using fMRI: spatial navigation by thoughts. *Neuroreport* **15**, 1591–1595 (2004).
66. Peplow, M. Mental ping-pong could aid paraplegics. *Nature News* 23 Aug 2004 (doi: 10.1038/news040823-18).
67. Fetz, E. E. & Finocchio, D. V. Operant conditioning of specific patterns of neural and muscular activity. *Science* **174**, 431–435 (1971).
68. Nowlis, D. P. & Kamiya, J. The control of electroencephalographic alpha rhythms through auditory feedback and the associated mental activity. *Psychophysiology* **6**, 476–484 (1970).
69. Manuck, S. B. *et al.* Role of feedback in voluntary control of heart rate. *Percept. Mot. Skills* **40**, 747–752 (1975).
70. Allen, J. J., Harmon-Jones, E. & Cavender, J. H. Manipulation of frontal EEG asymmetry through biofeedback alters self-reported emotional responses and facial *EMG. Psychophysiology* **38**, 685–693 (2001).
71. Elbert, T. *Self-Regulation of the Brain and Behavior* (Springer, Berlin; New York, 1984).
72. Friel, P. N. EEG biofeedback in the treatment of attention deficit hyperactivity disorder. *Altern. Med. Rev.* **12**, 146–151 (2007).
73. Lubar, J. F. & Deering, W. M. *Behavioral Approaches to Neurology. Behavioral Medicine Series* (Academic, New York, 1981).
74. deCharms, R. C. *et al.* Learned regulation of spatially localized brain activation using real-time fMRI. *Neuroimage* **21**, 436–443 (2004).
75. Weiskopf, N. *et al.* Physiological self-regulation of regional brain activity using real-time functional magnetic resonance imaging (fMRI): methodology and exemplary data. *Neuroimage* **19**, 577–586 (2003).
76. Yoo, S. S. & Jolesz, F. A. Functional MRI for neurofeedback: feasibility study on a hand motor task. *Neuroreport* **13**, 1377–1381 (2002).
77. Posse, S. *et al.* Real-time fMRI of temporolimbic regions detects amygdala activation during single-trial self-induced sadness. *Neuroimage* **18**, 760–768 (2003).
78. Yoo, S. S. *et al.* Increasing cortical activity in auditory areas through neurofeedback functional magnetic resonance imaging. *Neuroreport* **17**, 1273–1278 (2006).
79. Caria, A. *et al.* Regulation of anterior insular cortex activity using real-time fMRI. *Neuroimage* **35**, 1238–1246 (2007).
80. deCharms, R. C. *et al.* Control over brain activation and pain learned by using real-time functional MRI. *Proc. Natl Acad. Sci. USA* **102**, 18626–18631 (2005).
81. Fetz, E. E. Volitional control of neural activity: implications for brain-computer interfaces. *J. Physiol.* **579**, 571–579 (2007).
82. Rainville, P. Brain mechanisms of pain affect and pain modulation. *Curr. Opin. Neurobiol.* **12**, 195–204 (2002).
83. Apkarian, A. V. Functional magnetic resonance imaging of pain consciousness: cortical networks of pain critically depend on what is implied by “pain”. *Curr. Rev. Pain* **3**, 308–315 (1999).
84. Moisset, X. & Bouhassira, D. Brain imaging of neuropathic pain. *Neuroimage* **37** (Suppl. 1), S80–S88 (2007).
85. Petrovic, P. & Ingvar, M. Imaging cognitive modulation of pain processing. *Pain* **95**, 1–5 (2002).
86. Peyron, R., Laurent, B. & Garcia-Larrea, L. Functional imaging of brain responses to pain. A review and meta-analysis. *Neurophysiol. Clin.* **30**, 263–288 (2000).
87. Rainville, P. Brain mechanisms of pain affect and pain modulation. *Curr. Opin. Neurobiol.* **12**, 195–204 (2002).
88. Tracey, I. Imaging pain. *Br. J. Anaesth.* **101**, 32–39 (2008).
89. Bray, S., Shimojo, S. & O’Doherty, J. P. Direct instrumental conditioning of neural activity using functional magnetic resonance imaging-derived reward feedback. *J. Neurosci.* **27**, 7498–7507 (2007).
90. Buonomano, D. V. & Merzenich, M. M. Cortical plasticity: from synapses to maps. *Annu. Rev. Neurosci.* **21**, 149–186 (1998).
91. Merzenich, M. M. & deCharms, R. C. in *The Mind-Brain Continuum* (eds Llinas, R. & Churchland, P.) (MIT Press, Boston, 1996).
92. Moore, N. C. A review of EEG biofeedback treatment of anxiety disorders. *Clin. Electroencephalogr.* **31**, 1–6 (2000).
93. Sterman, M. B. & Egner, T. Foundation and practice of neurofeedback for the treatment of epilepsy. *Appl. Psychophysiol. Biofeedback* **31**, 21–35 (2006).
94. Gruzelier, J. & Egner, T. Critical validation studies of neurofeedback. *Child Adolesc. Psychiatr. Clin. N. Am.* **14**, 83–104, vi (2005).
95. Phan, K. L. *et al.* Real-time fMRI of cortico-limbic brain activity during emotional processing. *Neuroreport* **15**, 527–532 (2004).
96. Adcock, R. A. *et al*. Real time fMRI during the psychotherapy session: toward a methodology to augment therapeutic benefit. *11th Ann. Meeting Organ. Hum. Brain Mapp*. (Toronto, Ontario, Canada, 2005).

Acknowledgements

This work was funded by US National Institutes of Health Grants and Contracts MH067290, NS050642, NS049673, N43DA-4-7748, DA021877, and N43DA-7-4408.

Competing interests statement

The author declares competing financial interests: see web version for details.

FURTHER INFORMATION

Christopher deCharms’ homepage:
http://www.omneuron.com/decharms.html

ALL LINKS ARE ACTIVE IN THE ONLINE PDF