Review TRENDS in Cognitive Sciences Vol.11 No.11

Reading and controlling human brain activation using real-time functional magnetic resonance imaging

R. Christopher deCharms

Omneuron 3T MRI Research Center, Menlo Park, CA 94025, USA

Understanding how to control how the brain's functioning mediates mental experience and the brain's processing to alter cognition or disease are central projects of cognitive and neural science. The advent of real-time functional magnetic resonance imaging (rtfMRI) now makes it possible to observe the biology of one's own brain while thinking, feeling and acting. Recent evidence suggests that people can learn to control brain activation in localized regions, with corresponding changes in their mental operations, by observing information from their brain while inside an MRI scanner. For example, subjects can learn to deliberately control activation in brain regions involved in pain processing with corresponding changes in experienced pain. This may provide a novel, non-invasive means of observing and controlling brain function, potentially altering cognitive processes or disease.

Introduction

Although we experience the contents of our own awareness, we have little or no ability to observe the physiological processes taking place within our brains, which underlie conscious experiences and our behaviors. We also cannot directly observe brain processes in others in the way that we observe other aspects of their behaviour. People control brain activation all the time: every voluntary action engages the activation of specific brain mechanisms. However, the extent to which a person can learn to control the activation of a specific, anatomically localized region of their brain is still largely unknown. Real-time functional magnetic resonance imaging (rtfMRI; see Glossary) offers the intriguing opportunity for people to observe the brain processes that underlie their current thoughts and feelings, and to change them.

Bringing brain activation, which people are not normally aware of, into awareness may enable a person to learn greater explicit control over their own cognitive and neural activities by mimicking desired brain states, or mimicking the brain states of others [1]. It is unclear what the limits of this learned control may be. What brain areas is it possible to bring under conscious, volitional control? What cognitive or neural processes can be deliberately learned or modified? Is it possible to target neural plasticity to enhance function in particular brain systems through focused training strategies? Using rtfMRI to measure thousands of points in the brain simultaneously and in real-time, how much is it possible to know about what someone is thinking and experiencing or about their neural functioning, and about how they may learn to change it?

This review focuses on read-out and learned control over brain activation using real-time neuroimaging. This new research area is receiving increasing attention because of the theoretical interest in a new method of probing the relationship between brain and cognitive function in real-time, and the practical interest in potential applications of non-invasively controlling localized brain function. There are other important applications of rtfMRI, such as monitoring experiments, improved quality assurance during data acquisition and rapid mapping for surgical planning purposes, which are reviewed elsewhere [2,3].

Two stages are required for learning control over brain activation using this type of procedure: a method for reading out brain processes in real-time and a method for training a person to use this information to control those brain processes. In the read-out step, a person's current brain activation pattern can be compared with another

Glossary

BOLD: blood oxygen level dependent signal change. fMRI is measured as the fluctuation in signal intensity in a series of MR images. These fluctuations are caused by changes in blood volume, flow and oxygenation caused by changes in neuronal activation. fMRI measurements are inherently indirect, being based upon blood flow and oxygen utilization in response to neuronal energy consumption, rather than directly on neural activation. This limits temporal and spatial resolution to the scale of millimeters and seconds.

Pattern classifier: a mathematical algorithm that finds the similarity between different patterns of data (spatial pattern or patterns through time) or classifies a pattern according to its level of similarity to several pre-defined reference patterns.

Reference pattern: a pre-computed spatial pattern of brain activation that the current spatial pattern from a subject can be compared with. For example, the spatial pattern of the change in BOLD activation for each voxel when one imagines moving the left index finger. This could be computed using a contrast vs a rest condition, and may include voxels showing significant increases or decreases of different amplitude in signal.

ROI: a region of interest in the brain that is measured. This often corresponds to a single brain area. A ROI can be defined anatomically or using a localizer task that determines which brain voxels are associated with a particular task.

rtfMRI: real-time functional MRI. Normally this refers to fMRI where the pace of analysis display of relevant information keeps up with the pace of data acquisition, typically lagging by 1–4 s. A typical acquisition rate is one MR brain volume (16 slices of 64 × 64 voxels) per second.

Voxel: volume element. A 3D unit of MR imaging measurement, similar to a pixel, which is a 2D element.

Corresponding author: deCharms, R.C. (decharms@gmail.com).
Available online 7 November 2007.

 1364-6613/$ – see front matter doi:10.1016/j.tics.2007.08.014

person's, compared with that of a group of people or compared with an activation pattern hypothesized to correspond to a given function. The comparison can include a single brain area, multiple areas or spatial patterns spanning the brain. In the training step, subjects learn to select from, refine or create new cognitive strategies and neural processes to optimize a measure of their brain activation, for instance to increase or decrease the level of activation in a target region of interest (ROI).

Reading brain states in real-time using rtfMRI

rtfMRI methodology and development

The ability to read-out complex brain processes from fMRI in real-time has only recently become readily available [2–15]. Although other types of neurophysiology experiments, such as single-neuron recording or electroencephalogram (EEG) recording, have almost always actively monitored data as they were being collected, early neuroimaging studies were similar to a large cyclotron or space probe experiment: the subject was loaded into the scanner, gigabytes of data were collected from a large instrument, and researchers analyzed the results over days, weeks or months. Results can now be observed as the experiment unfolds.

Contemporary rtfMRI is a massively parallel measurement of brain physiology, measuring activation from $\sim 2^{16}$ spatial locations collected as a 3D stack of images every 1–2 s. The steps involved in rtfMRI analysis (Figure 1) are similar to offline fMRI [16,17]. Because analysis takes place during the experiment, results can be used to monitor and guide the subject or to guide the experimental procedure (Figure 1, Step 8).

The capability of analyzing fMRI data in real-time was originally developed by Cox *et al*. [4] not long after the advent of fMRI, but initially neither the applications of rtfMRI nor the possibility of performing rtfMRI on readily available computer hardware were generally recognized. Indeed, some early work focused on using supercomputers to accomplish rtfMRI [6]. rtfMRI is now performed using software available on the scanner and on a general purpose PC. The field has greatly benefited from technical progress in real-time anatomical MRI. For example, fast pulse sequences and image reconstruction can now provide non-invasive real-time MR visualization of the beating heart [18,19]. Many groups have contributed to rtfMRI becoming steadily more powerful and more accessible [2,4–14,20–27].

Limitations of the rtfMRI signal

Information derived from the brain using rtfMRI suffers from a number of important limitations in spatial and temporal resolution and noise. Neurons fire action potentials on a timescale of milliseconds, and precise timing may be an important feature of neuronal coding [28,29]. The fMRI signal is derived from slower changes in blood flow and oxygenation peaking $\sim$2–6 s after neural activation, with rtfMRI data processing adding another 1–4 s of delay [16,17,30–32]. The term 'real-time' in rtfMRI is used here to designate imaging where the full results from data analysis keep pace with data acquisition. Also, whereas neurons have micron spatial scale features, the fMRI signal is often measured from voxels on the order of $3 \times 3 \times 3$ mm.

The fMRI signal is also inherently noisy. MRI introduces many types of noise leading raw fMRI data images (Figure 1, top) to be of poor image quality compared with the anatomical MRI images on which fMRI data are normally overlaid (Figure 1, bottom). Raw fMRI images have a low signal to noise ratio, blurring, spatial distortion or signal loss in particular parts of the images (susceptibility induced 'drop out') and drifts in signal intensities measured over time [17,32]. Noise also arises from biological processes, including artifacts owing to respiration, heart rate, fluctuations in neuronal signals, fluctuations in blood flow and motion artifacts [17,32,33]. Given these many noise sources, it was not clear *a priori* whether the characteristics of an fMRI signal could be successfully used in real-time for interpreting continuously evolving brain information without averaging data over multiple trials or for guiding a person's behavior [3,27].

rtfMRI pattern classification for reading brain activation

Pattern classification algorithms may represent an important new method for reading brain states in real-time that benefit from the rich spatial detail of a massively parallel measurement such as rtfMRI. Pattern classification methods have been used increasingly to infer underlying brain states from spatial patterns observed in fMRI data offline [34–45], and more recently have been considered for or used in rtfMRI [22,45,46]. Pattern classification constitutes a broad class of mathematical methods that compare two or more spatial or temporal patterns of data to assess their similarity [47]. A 'reference pattern' can be created that corresponds to a particular brain state; for example, the spatial pattern of fMRI activation for each voxel from selected structures in the brain during a particular task compared with a rest period when the subject performs a simple unrelated task or no overt task, with each voxel having its own value of increased or decreased activation. The spatial pattern of activation at a given moment in time in a person's brain can then be compared quantitatively with this reference pattern, or with multiple reference patterns, by a pattern classification algorithm (Figure 2).

Pattern classification algorithms have the potential to exploit the inherent information richness of the large number of parallel spatial measurements collected using rtfMRI to extract greater information than can be derived from individual ROIs [48]. For example, whereas the overall average signal from a brain region may not be different during several different brain states, such as during perception of different stimuli, the detailed information contained within individual voxels within the brain region may make it possible to determine which state is taking place based upon the fMRI activation pattern. Pattern classifiers may be built automatically by providing many samples of actual data, which correspond to each brain state – the pattern classifier will seek to find the statistical similarities from data within each state. This process can be computationally intensive and may require substantial fMRI input data from many task repetitions. However, once a classifier has been built it

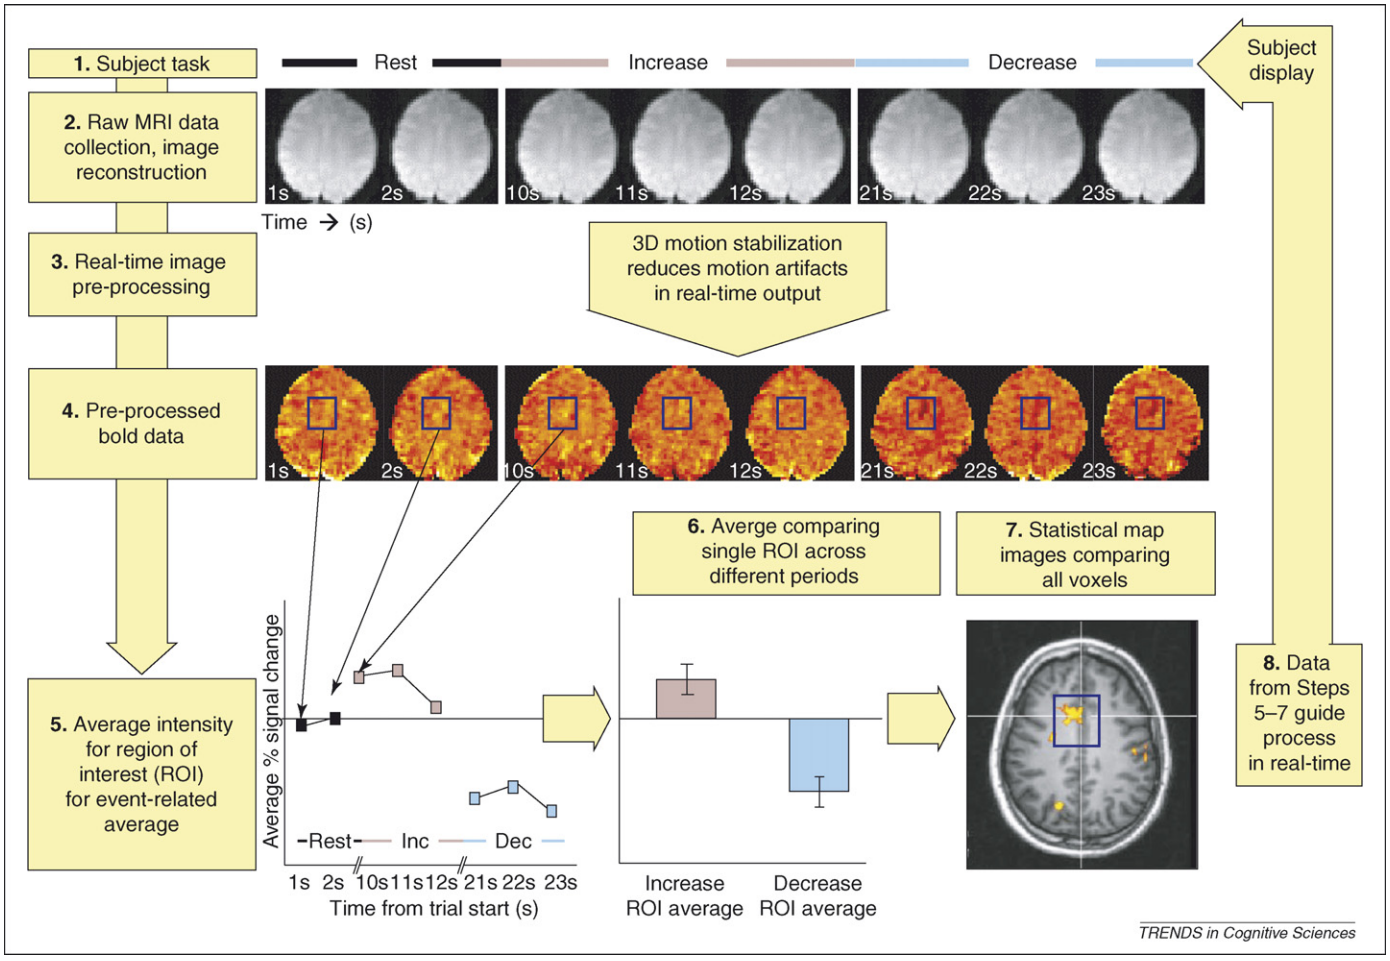

Figure 1. Process steps of an rtfMRI experiment. rtfMRI experiments include all of the procedures of conventional offline fMRI (Steps 1–7), with the new element that the results can be used to guide the experimental subject or experimental procedure (Step 8). **1.** Subjects engage in a task, while sensing or observing brain activation information derived from rtfMRI. For example, they may learn to increase or decrease activation in a target brain region while observing the level of activation in this region. **2.** Data collection and image reconstruction typically occurs within 1–2 s. Images may be collected using a variety of advanced imaging and reconstruction techniques, such as spiral imaging, multiple echo imaging, adaptive multiresolution echo-planar imaging and image distortion correction [27]. These advances represent the work of several MR physics groups working to improve the quality of the images used for measurement and the fMRI contrast measured. **3.** Pre-processing of fMRI data is performed, which may include motion correction, spatial and temporal filtering, and normalization to stereotactic space. 3D motion correction is particularly important for rtfMRI and may be performed either retrospectively by spatial transform of the images after they are collected [5,21], or on some MRI scanners it can be performed prospectively by adjusting the prescribed positions of each scan plane to remain fixed relative to the position of the head based upon the preceding image [82]. Signal intensity fluctuations caused by brain activation are typically small relative to the much larger changes in intensity caused by movement. For example, moving a $3 \times 3 \times 3$ mm voxel at the edge of the brain by 0.3 mm laterally to include more space outside the brain may lead to a 10% reduction in signal intensity, compared with $<1\%$ signal changes often associated with brain activation. **4.** Data are converted into a measure reflecting changes in image intensity, rather than absolute intensity. fMRI inherently measures differences in signal between different times, not the absolute level of activation at each time point. The conceptually simplest computation is therefore of the raw % signal change in each voxel at each time point (i.e. each voxel's difference from its own mean or baseline value), which can be presented to the subject with no further processing. The voxel's fluctuations can also be normalized into statistical units such as z-scores, correlated with the task to yield a correlation coefficient or fit to a general linear model. **5.** An average of the fluctuating % signal change values within an ROI is calculated for each time point. **6.** Signals are averaged over multiple trials to produce trial averages. Event-related averages may also be computed (not shown). **7.** Activation maps are created. Spatial analyses may include computing real-time fMRI brain activation maps, which evolve as further data are collected, fitting data to a general linear model, computing *t*-tests, correlation coefficients, or performing multiple regression or independent component analysis or pattern classification. Images can be displayed in 2D or 3D using cortical flattening or inflation visualization techniques. In addition, maps can be continuously computed on the entire dataset to determine when statistical significance criteria have been met so an experiment can be ended. Data can also be analyzed using a sliding window approach that produces maps of only the most recent portion of the experiment. **8.** The subject views displays based upon the results, and may use this information to guide their learning to control brain activation. Although extensive analysis is possible in real-time, the full results can be overwhelming to subjects during scanning. Therefore, typically only a single parameter is presented for each ROI for each time point; for example, presenting the subject with the mean % signal difference for an ROI at each time point versus the ROIs overall mean. This parameter can be presented in many ways, such as a scrolling line chart of the parameter, a graphical representation such as a flame that is proportional in size to the parameter, or an auditory stimulus whose pitch is proportional to the parameter.

can be applied to new data in real-time with only modest computational requirements.

LaConte and coworkers illustrated the power of pattern classification with rtfMRI data [46]. Using a support vector machine, a particular type of pattern classification algorithm, they demonstrated that a pattern classifier could discern when a subject was tapping the left versus the right finger or even thinking 'sad' versus 'happy' thoughts. They then used the output of this pattern classifier in real-time to train subjects to control their pattern of brain activation. This output could also be used as a brain computer interface (BCI) for communication or to control an external device. Further experiments are needed to compare what information real-time pattern classifiers can extract from brain data compared with traditional analyses. These results are important in that they lead the way toward training subjects to produce or control arbitrarily complex patterns of brain activation, or

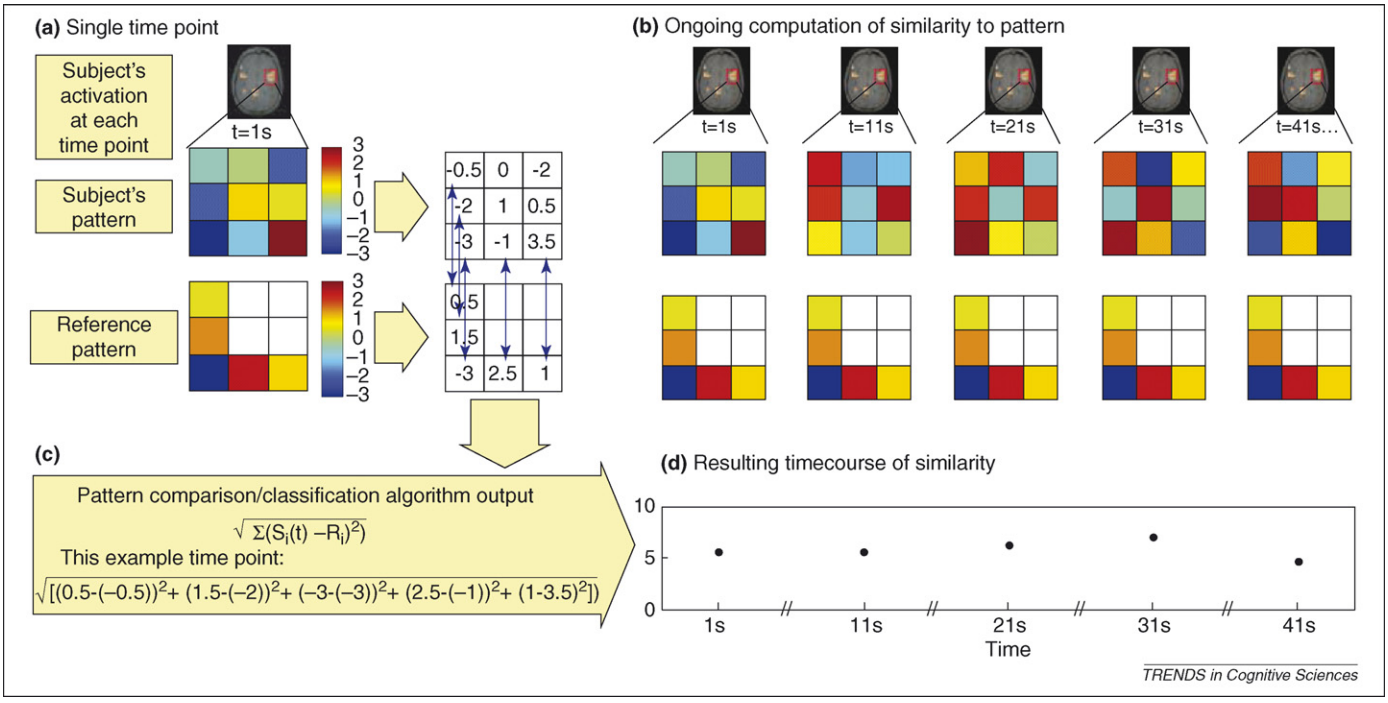

Figure 2. Real-time brain activation pattern comparison. The spatial pattern of a subject's brain activation can be compared quantitatively at a single time point (or using an average over time) to determine how similar the subject's activation is to one or more reference patterns corresponding to the performance of known tasks, or to the activation of selected brain areas. Many different methods for comparison are available, including traditional analytical methods, complex models such as neural networks, and others. The figure shows a single simple example computation of similarity, a distance measure. (**a**) Diagram of the comparison between a subject's pattern and a single reference pattern. From nine voxels shown, five voxels are selected to be included in the pattern comparison, and four are excluded from the comparison. Colors represent % BOLD signal intensities. All data in Figures 1 and 2 are created for illustration only, and magnitudes of changes in % BOLD signal are exaggerated to better illustrate the computation, and do not exactly match color scale. The voxels have % BOLD signal intensities between −3 and 0.5%. The algorithm computes the similarity at the first time point. (**b**) Shows this process in real-time to produce a time course of similarity of the subject's brain activation pattern with the reference pattern over time. As meaningful changes typically take several seconds to evolve, only times 1s, 11s, 21s, 31s, 41s are shown. If this were computed using only a single voxel, it would be similar to traditional ROI analysis, but computing distance from a target value (e.g. if the target value for a voxel is 1% BOLD, then both 0.5% and 1.5% would receive the same distance measure). (**c**) The time course can be plotted directly and used as one may use data from ROI analysis (e.g. in making event-related averages) or to statistically compare different periods. The comparison can also be presented to a subject during training or used to control an external device. The subject's brain activation may be compared with many different reference patterns to separately produce the time course of similarity to each pattern. The reference pattern that is most similar to the subject's activation, the closest match, can also be selected at each time point. The closest match output from a pattern classifier that was built using reference patterns corresponding to several different tasks can be compared at each time point with what task was actually being performed by a subject to assess the performance of the classifier. This comparison should be performed using data that the classifier was not built from to assess the classifiers prediction accuracy on new data, and avoid over-fitting the input data.

learning to produce desirable brain states through mimicking others.

Learned control over brain activation

Prior methods for physiological self-regulation

It has been known for several decades that subjects can learn volitional control over a variety of physiological functions not normally under conscious control by using feedback [49]. This has included autonomic measures, such as heart rate, skin conductance and muscle tone, and measures of brain activation including EEG rhythms (such as alpha or theta frequency power) [50–54], slow cortical potentials [51,55] and single motor neurons firing in prosthetic control [56–58]. If subjects can learn control over their heart rate, skin conductance, EEG rhythms or motor neurons, can they also learn volitional control over the highly specialized neurophysiological functions mediated by hundreds of individual brain areas?

Previous non-invasive techniques for measuring brain activation, principally EEG and Magnetoencephalography (MEG), were not able to anatomically localize activation as precisely as fMRI to serve as the basis for training. EEG and MEG have particular difficulty localizing activation sources in the brain from continuously acquired data; localization is normally performed using task paradigms that are time-locked to a brief stimulus and then averaged. Precise localization may be critical for learning control over a particular brain region and its highly specific functions, in contrast to controlling widespread neural events, such as changes in EEG power, which often correlate with much broader behavioral states, such as relaxation or attention [59,60]. However, EEG-based measurement has important advantages: it can be portable, it is inexpensive and it can provide precise temporal resolution. Current and foreseeable development efforts are also making MRI much smaller and less expensive, with some recently developed peripheral MRI scanners being as small as home appliances and suitable for a physician's office (e.g. Oni Medical, Wilmington, MA). Using precisely localized measurement techniques, such as single neuron recording and local field potential recording, subjects have also learned to control neural activation in the setting of prosthetic control [57,58,61–63], an approach requiring invasive surgical procedures for electrode implantation but indicating the potential power of neural plasticity and learning to shape brain function [58,64,65].

Although it is clear that there is a broad conceptual analogy between forms of feedback training based

upon different physiological measures, including EEG, Electromyography (EMG), single neurons and fMRI, this analogy may be superficial, with the stronger relationship being based upon the functions of the brain areas being trained and what is being learned, rather than whether feedback is used in training. For example, principles of learning precise control over motor cortex using fMRI may be much more analogous to principles of learning to play the piano (or other traditional motor skill learning) than to learning relaxation using EEG feedback or heart rate conditioning.

rtfMRI-based training methods

In rtfMRI training, one is performing a brain imaging experiment on oneself, trying to optimize a desired activation pattern by selecting cognitive and neural processes using feedback. A subject is often able to produce some activation in a target brain region at the outset (e.g. for a brain area involved in attention, by attending to the task). The goal is thus not just activation, but an enhancement in control over brain activation corresponding with an enhancement in control over the related cognitive process. The question often arises whether changes in brain activation derive from changes in cognitive processes or whether brain activation causes changes in cognitive processes. We presume brain activation and mental processes to be correlated manifestations of the same phenomena, which change together through learning.

In a typical protocol, subjects may be instructed to learn to increase or decrease brain activation in a target brain region during alternating time periods, with their success presented to them as a depiction of their brain activation data as it changes second by second. Subjects are typically instructed to avoid potential sources of signal artifact, such as movement, changes in respiration or overall changes in arousal and global brain activation, and they may also be presented with additional feedback measures regarding these artifacts. See Box 1 for further methodological considerations.

Box 1. Requirements for successful rtfMRI-based training

There are several core requirements for rtfMRI training to succeed in producing changes in cognitive processes and brain activation in a limited time:

Brain information correlated with a cognitive function to be trained
To alter a cognitive function, for example cognitive control over pain, the rtfMRI training signal must be functionally correlated with that cognitive function: that is, the investigator must be able to select a brain region(s) for training that is/are involved in mediating or controlling the function to be altered. The signal used could be the average level of activation in a single brain region, or the comparison of the current pattern of brain activation with a predefined spatial pattern spanning multiple regions, which may include areas with both increased and decreased activation. An important ongoing question is whether to select the brain region(s) for training based upon the trained individual's own brain data (which may capture that person's brain's unique anatomy and neurophysiology) or based on prior results pooled from many subjects (potentially more statistically reliable, potentially from subjects already highly skilled at a desired task), and whether to set anatomically or functionally defined boundaries for the signal.

Strategy instructions
Providing subjects with clear instructions may be critical. Before their first rtfMRI scan most experimenters and subjects tend to overestimate their ability to control brain activation, perhaps reasonably supposing that they will know how to control their own brain. fMRI activation is highly variable, difficult to control at the outset, and subjects may be wrong in their initial guesses as to which of countless mental strategies will produce a desired activation pattern. To avoid prolonged random search, subjects may be provided with a set of initial task strategies that reliably activate the brain region being trained, and provide rtfMRI signals that can serve as the basis for further selection and refinement.

Selecting a brain measure and function that can be controlled
rtfMRI in itself does not change the brain – it just provides information. rtfMRI-based training will only succeed with a cognitive or neural process that a person can learn to change, and in the context of appropriate training.

Training control over brain activation using real-time or near real-time fMRI feedback

Early training studies presented people with fMRI data from their brain, which were analyzed shortly after a single task trial to assess the feasibility of people using this information to learn to alter their cognitive strategies. Yoo and Jolesz [66] trained people to select a finger-tapping behavior by showing them fMRI data from their own somatomotor cortex collected on a single finger-tapping trial that had ended 20 s earlier and was then rapidly analyzed. People were able to interpret their fMRI data to select the best finger-tapping behaviors to increase their brain activation as measured using fMRI. Posse *et al.* [67] provided single trial fMRI feedback of amygdala activation to people as an adjunct to their trying to control their mood. After subjects attempted to self-induce sadness for 60 s, their data was immediately analyzed and feedback of amygdala activation was given to subjects verbally, leading to subsequent increased amygdala activation and increased ratings of sadness. These studies demonstrated that, in principle, the fMRI signal can be sufficiently statistically robust to serve as a meaningful basis for training. The approach of analyzing whole trials has the advantage that the signal may be averaged to generate more statistically robust information, with the disadvantages of introducing significant delay before the subject receives the information and removing temporal information contained in ongoing signal changes.

Nearly concurrently, early rtfMRI experiments demonstrated that subjects can learn control over localized brain activation through observing ongoing fluctuations in fMRI activation.

Weiskopf *et al.* [68] provided the first example of a single subject who learned control over activation in a localized brain region, the anterior cingulate cortex (ACC), using rtfMRI feedback. deCharms *et al.* demonstrated in a group of subjects that repeated training using rtfMRI can teach subjects to achieve greater control over activation in a specifically targeted brain region, the somatomotor cortex (regions engaged in tactile sensation, movement and motor mental imagery; Figure 3) [69]. The experimental subjects engaged in motor imagery and received continuous rtfMRI feedback of somatomotor cortex activation. Subjects did not move or tense muscles, as monitored by concurrent EMG, supporting the idea that they were learning mental

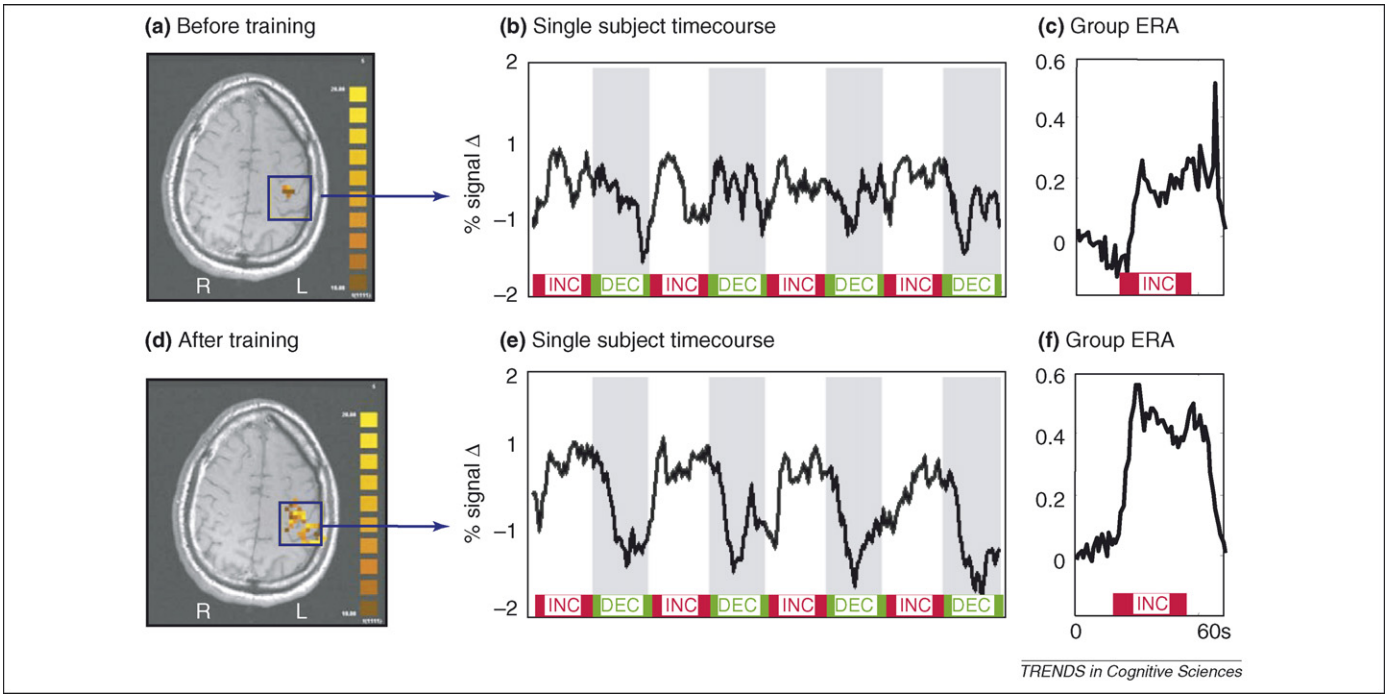

Figure 3. rtfMRI-based training leads to increased volitional control over activation in somatomotor cortex. Subjects were trained to increase the level of activation that they were able to achieve in the somatomotor cortex using rtfMRI information and motor imagery strategies. Compared with the activation seen at the outset of training (**a–c**), following training, subjects showed a statistically significant, spatially localized increase in brain activation in the trained ROI (**d–f**). The area of activation is centered on Brodmann areas 4 (primary motor cortex) and areas 1/2/3 (primary somatosensory cortex). For the ROI indicated by the blue box, the activation time course is shown before training for a single subject in (**b**) and after training in (**e**), demonstrating a clear, trial-by-trial improvement in modulation even in data presented raw, with no temporal filtering or spatial smoothing. (**c**) shows the event-related average (ERA) of fMRI activation from the ROI for a group of subjects before training, whereas (**f**) shows the ERA following training, indicating an increase in control over activation in the target ROI. The INC/DEC labels correspond to periods when subjects were instructed to increase or decrease the level of activation in the target ROI. Images represent *t*-maps computed in Brain Voyager (Brain Innovation, Maastricht, Netherlands). Figure adapted, with permission, from [69].

strategies as opposed to overt movements. A control group received otherwise similar training but saw sham fMRI feedback that did not correspond to their brain activation. The sham group subjects did not show improved control over brain activation, demonstrating that the learned control was due specifically to training using rtfMRI as opposed to other potential learning effects, such as repeated practice. Once they had learned volitional control over brain activation using rtfMRI, subjects could continue to exercise this control even if the rtfMRI information was later withdrawn.

rtfMRI data are temporally filtered, removing rapid fluctuations in the data due to noise and slow drifts. This leads to an implied temporal resolution of the signal on the order of ~1–5 s. However, the time required for a person to detect a change in an fMRI signal depends upon the magnitude of the change versus noise. A large change can be almost immediately apparent (e.g. a dramatic response to finger tapping in somatomotor cortex) whereas a small signal may never be detected by a subject (e.g. a subtle change in blood oxygen level dependent signal change (BOLD) requiring averaging over many trials to become statistically meaningful).

Brain areas subject to volitional control

rtfMRI experiments [12,14,21,70,71] have now demonstrated that subjects can learn to control activation in a number of different brain areas [3,68,69,71]. Initial targets for this approach have included brain areas that had been extensively studied, are straightforward to image, or that may be presumed to be readily controllable based upon their functions, such as the somatomotor cortex [21,66,69,71], the parahippocampal place area (PPA) [72], amygdala [67], auditory cortex [73], insular cortex [74] and ACC [68,71,75]. These brain areas are involved in a diverse array of functions, including volitional movement, tactile sensation, visual perception, hearing, emotion and pain. Bray *et al.* demonstrated that subjects can learn control over brain activation in an instrumental conditioning paradigm, a paradigm where subjects received a monetary reward for successfully learning control over brain activation [76]. This approach will become particularly interesting if it is shown in the future that subjects can learn control over brain activation in circumstances where pre-defined strategies cannot be made effective, such as in cases where the strategies could not readily be predicted, or in compromised patient groups. A rich area for further investigation is what other brain areas people can learn to control, what types of tasks and stimuli are required to teach control, and what behavioral consequences result from this control. We are not yet aware of brain areas that subjects have failed to learn to control when given adequate training, although many brain areas remain unexplored in this new form of brain cartography.

Future directions

rtfMRI training of multiple brain regions as a brain computer interface (BCI)

Studies have investigated the ability of subjects to control more than one brain region at once to generate a BCI to

communicate with subjects or control a prosthetic [61–63,72,77–79]. For example, Weiskopf *et al.* investigated whether subjects were able to differentially control activation in two cortical areas [72]. Subjects were provided with rtfMRI feedback of the difference between activation in the supplementary motor area (SMA) and PPA, and were encouraged to use strategies, such as visual imagery or motor imagery, known to activate these two structures. Following training, some subjects showed an increase in the differential signal between SMA and PPA. More recently, Yoo *et al.* [71] reported a near rtfMRI study where subjects navigated a computer cursor through a 2D maze by selecting from four cognitive tasks, one corresponding to each movement direction. Subjects used mental calculation to activate the medial superior frontal and anterior cingulate gyri, used left and right motor imagery to activate right and left somatosensory areas, and used mental speech generation to activate the posterior superior temporal gyrus. Subjects succeeded in learning to use their cognitive processes to move a cursor through a 2D maze. It will be interesting to compare the speed and fidelity of control achievable using approaches based on rtfMRI (with high spatial and low temporal resolution) with EEG-based approaches (with high temporal and low spatial resolution).

rtfMRI training leading to changes in cognition or behavior

Learned control over brain activation may lead to changes in cognition, behavior or disease processes. deCharms *et al.* demonstrated that when subjects learned control over activation in the rostral ACC (rACC) [75], a region associated with pain perception and pain control [80,81], there was a closely associated change in pain perception. In response to the same painful heat stimulus during periods when subjects increased activation in this brain area they experienced greater pain measured psychophysically inside the scanner compared with periods when they decreased activation in this brain area. This was not true for any of the four control groups of subjects who were trained using similar procedures but either with no rtfMRI, with rtfMRI from a different brain region, or with rtfMRI information measured previously from a different subject's brain. Subjects received detailed written instructions to use initial strategies to control brain activation, which were adapted from clinical practice for pain control, such as shifting attention, reappraising the painful percept or exerting cognitive control over the pain experience. Subjects were instructed to adapt these strategies and develop new ones to maximize control over the rtfMRI signal. Following training, subjects were not typically able to articulate what their cognitive strategies had been and often gave general anecdotal descriptions: 'I found a place where there was no pain', 'it was like that feeling of being very focused'. When a group of chronic pain patients were trained using similar rtfMRI procedures, they reported decreases in their ongoing chronic pain immediately following training using standard pain assessment questionnaires [75]. This decrease in pain was greater than for a control group of patients who received a similar period of biofeedback training using heart rate and respiratory information [75]. Following a similar design, Caria and coworkers [74] demonstrated that subjects could learn control over activation in the anterior insula, a region involved with emotional processing, potentially leading to changes in affective processing or interoception [74]. These studies suggest that learning to control localized brain activation may lead to corresponding changes in behavior and cognition in healthy subjects and potentially in patient groups.

The potential for clinical applications

This raises the possibility that rtfMRI may be developed into a clinically important tool for controlling localized brain activation using training, augmenting other forms of therapy, brain stimulation or pharmacology. A decade of basic research exploring fMRI has led to a new technology with the potential to measure and control brain activation that is non-invasive, non-pharmacologic, and potentially safe and reversible. As in other forms of rehabilitation through exercise, exercise of a brain system through this type of training could lead to long-term upregulation in brain function based upon use-dependent plasticity, and potentially to enduring changes in disease state [58,64,65]. Ongoing research is exploring a variety of clinical conditions potentially amenable to treatment through localized control over brain activation, including chronic pain, control over craving and addiction, depression, recovery from stroke, and as an adjunct to psychotherapy where both the patient and the psychotherapist have ongoing access to information about processes taking place in the patient's brain. Although it is not yet known what applications rtfMRI may come to have, the ability to monitor the brain's function in real-time may provide new avenues for diagnosis, to guide therapeutic interventions based upon this information or to assess the mechanism and efficacy of treatments.

Acknowledgements

This work was supported by National Institutes of Health grants R43MH067290, R44NS050642, N43DA-4–7748, R44DA021877–01A1 and N43DA-7–4408.

References

1 Kotchoubey, B. *et al.* (2002) Can humans perceive their brain states? *Conscious. Cogn.* 11, 98–113
2 Cohen, M.S. (2001) Real-time functional magnetic resonance imaging. *Methods* 25, 201–220
3 Weiskopf, N. *et al.* (2004) Self-regulation of local brain activity using real-time functional magnetic resonance imaging (fMRI). *J. Physiol. (Paris)* 98, 357–373
4 Cox, R.W. *et al.* (1995) Real-time functional magnetic resonance imaging. *Magn. Reson. Med.* 33, 230–236
5 Lee, C.C. *et al.* (1996) Real-time adaptive motion correction in functional MRI. *Magn. Reson. Med.* 36, 436–444
6 Goddard, N. *et al.* (1997) Online analysis of functional MRI datasets on parallel platforms. *J. Supercomputing* 11, 295–318
7 Hesser, J. *et al.* (1997) Real-time direct volume rendering in functional magnetic resonance imaging. *MAGMA* 5, 87–91
8 Busch, M. *et al.* (1998) Fast "real time" imaging with different k-space update strategies for interventional procedures. *J. Magn. Reson. Imaging* 8, 944–954
9 Gering, D.T. and Weber, D.M. (1998) Intraoperative, real-time, functional MRI. *J. Magn. Reson. Imaging* 8, 254–257
10 Lee, C.C. *et al.* (1998) Real-time reconstruction and high-speed processing in functional MR imaging. *AJNR Am. J. Neuroradiol.* 19, 1297–1300

11 Cox, R.W. and Jesmanowicz, A. (1999) Real-time 3D image registration for functional MRI. *Magn. Reson. Med.* 42, 1014–1018
12 Voyvodic, J.T. (1999) Real-time fMRI paradigm control, physiology, and behavior combined with near real-time statistical analysis. *Neuroimage* 10, 91–106
13 Yoo, S.S. *et al.* (1999) Real-time adaptive functional MRI. *Neuroimage* 10, 596–606
14 Gembris, D. *et al.* (2000) Functional magnetic resonance imaging in real time (FIRE): sliding-window correlation analysis and reference-vector optimization. *Magn. Reson. Med.* 43, 259–268
15 Weiskopf, N. *et al.* (2005) Single-shot compensation of image distortions and BOLD contrast optimization using multi-echo EPI for real-time fMRI. *Neuroimage* 24, 1068–1079
16 Moonen, C.T.W. *et al.* (1999) *Functional MRI. Medical Radiology,* Springer
17 Buxton, R.B. (2002) *Introduction to Functional Magnetic Resonance Imaging: Principles and Techniques,* Cambridge University Press
18 Nayak, K.S. *et al.* (2004) Real-time cardiac MRI at 3 tesla. *Magn. Reson. Med.* 51, 655–660
19 Santos, J.M. *et al.* (2006) Single breath-hold whole-heart MRA using variable-density spirals at 3T. *Magn. Reson. Med.* 55, 371–379
20 Mathiak, K. and Posse, S. (2001) Evaluation of motion and realignment for functional magnetic resonance imaging in real time. *Magn. Reson. Med.* 45, 167–171
21 Posse, S. *et al.* (2001) 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
22 Esposito, F. *et al.* (2003) Real-time independent component analysis of fMRI time-series. *Neuroimage* 20, 2209–2224
23 Bagarinao, E. *et al.* (2005) Enabling on-demand real-time functional MRI analysis using grid technology. *Methods Inf. Med.* 44, 665–673
24 Gasser, T. *et al.* (2005) Intraoperative functional MRI: implementation and preliminary experience. *Neuroimage* 26, 685–693
25 Stainsby, J.A. *et al.* (2005) Real-time magnetic resonance with physiologic monitoring for improved scan localization. *Magn. Reson. Med.* 53, 954–959
26 Nakai, T. *et al.* (2006) Dynamic monitoring of brain activation under visual stimulation using fMRI – the advantage of real-time fMRI with sliding window GLM analysis. *J. Neurosci. Methods* 157, 158–167
27 Weiskopf, N. *et al.* (2007) Real-time functional magnetic resonance imaging: methods and applications. *Magn. Reson. Imaging* 25, 989–1003
28 deCharms, R.C. and Merzenich, M.M. (1996) Primary cortical representation of sounds by the coordination of action-potential timing. *Nature* 381, 610–613
29 deCharms, R.C. and Zador, A. (2000) Neural representation and the cortical code. *Annu. Rev. Neurosci.* 23, 613–647
30 Jezzard, P. *et al.* (2001) *Functional MRI: An Introduction to Methods,* Oxford University Press
31 D'Esposito, M. (ed.) (2006) *Functional MRI: Applications in Clinical Neurology and Psychiatry*, Informa Healthcare
32 Bernstein, M.A. *et al.* (2004) *Handbook of MRI Pulse Sequences,* Academic Press
33 Kruger, G. and Glover, G.H. (2001) Physiological noise in oxygenation-sensitive magnetic resonance imaging. *Magn. Reson. Med.* 46, 631–637
34 Dehaene, S. *et al.* (1998) Inferring behavior from functional brain images. *Nat. Neurosci.* 1, 549–550
35 Cox, D.D. and Savoy, R.L. (2003) Functional magnetic resonance imaging (fMRI) "brain reading": detecting and classifying distributed patterns of fMRI activity in human visual cortex. *Neuroimage* 19, 261–270
36 Haynes, J.D. and Rees, G. (2005) Predicting the stream of consciousness from activity in human visual cortex. *Curr. Biol.* 15, 1301–1307
37 Haynes, J.D. and Rees, G. (2006) Decoding mental states from brain activity in humans. *Nat. Rev. Neurosci.* 7, 523–534
38 Kamitani, Y. and Tong, F. (2005) Decoding the visual and subjective contents of the human brain. *Nat. Neurosci.* 8, 679–685
39 O'Toole, A.J. *et al.* (2005) Partially distributed representations of objects and faces in ventral temporal cortex. *J. Cogn. Neurosci.* 17, 580–590
40 Carlson, T.A. *et al.* (2003) Patterns of activity in the categorical representations of objects. *J. Cogn. Neurosci.* 15, 704–717
41 LaConte, S. *et al.* (2005) Support vector machines for temporal classification of block design fMRI data. *Neuroimage* 26, 317–329
42 Haxby, J.V. *et al.* (2001) Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* 293, 2425–2430
43 Hanson, S.J. *et al.* (2004) Combinatorial codes in ventral temporal lobe for object recognition: Haxby (2001) revisited: is there a "face" area? *Neuroimage* 23, 156–166
44 Davatzikos, C. *et al.* (2005) Classifying spatial patterns of brain activity with machine learning methods: application to lie detection. *Neuroimage* 28, 663–668
45 Martinez-Ramon, M. *et al.* (2006) fMRI pattern classification using neuroanatomically constrained boosting. *Neuroimage* 31, 1129–1141
46 Laconte, S.M. *et al.* (2007) Real-time fMRI using brain-state classification. *Hum. Brain Mapp.* 28, 1033–1044
47 Duda, R.O. *et al.* (2001) *Pattern Classification,* (2nd edn), Wiley
48 Kriegeskorte, N. *et al.* (2006) Information-based functional brain mapping. *Proc. Natl. Acad. Sci. U. S. A.* 103, 3863–3868
49 Elbert, T. (1984) *Self-Regulation of the Brain and Behavior,* Springer-Verlag
50 Allen, J.J. *et al.* (2001) Manipulation of frontal EEG asymmetry through biofeedback alters self-reported emotional responses and facial EMG. *Psychophysiology* 38, 685–693
51 Rockstroh, B. *et al.* (1990) Biofeedback-produced hemispheric asymmetry of slow cortical potentials and its behavioural effects. *Int. J. Psychophysiol.* 9, 151–165
52 Manuck, S.B. *et al.* (1975) Role of feedback in voluntary control of heart rate. *Percept. Mot. Skills* 40, 747–752
53 Nowlis, D.P. and Kamiya, J. (1970) The control of electroencephalographic alpha rhythms through auditory feedback and the associated mental activity. *Psychophysiology* 6, 476–484
54 Lubar, J.F. and Deering, W.M. (1981) *Behavioral Approaches to Neurology (Behavioral Medicine Series),* Academic Press
55 Strehl, U. *et al.* (2006) Deactivation of brain areas during self-regulation of slow cortical potentials in seizure patients. *Appl. Psychophysiol. Biofeedback* 31, 85–94
56 Nicolelis, M.A. (2003) Brain-machine interfaces to restore motor function and probe neural circuits. *Nat. Rev. Neurosci.* 4, 417–422
57 Fetz, E.E. and Finocchio, D.V. (1971) Operant conditioning of specific patterns of neural and muscular activity. *Science* 174, 431–435
58 Jackson, A. *et al.* (2006) Long-term motor cortex plasticity induced by an electronic neural implant. *Nature* 444, 56–60
59 Friel, P.N. (2007) EEG biofeedback in the treatment of attention deficit hyperactivity disorder. *Altern. Med. Rev.* 12, 146–151
60 Loo, S.K. and Barkley, R.A. (2005) Clinical utility of EEG in attention deficit hyperactivity disorder. *Appl. Neuropsychol.* 12, 64–76
61 Hinterberger, T. *et al.* (2005) Neuronal mechanisms underlying control of a brain-computer interface. *Eur. J. Neurosci.* 21, 3169–3181
62 Lebedev, M.A. and Nicolelis, M.A. (2006) Brain-machine interfaces: past, present and future. *Trends Neurosci.* 29, 536–546
63 Mason, S.G. *et al.* (2007) A comprehensive survey of brain interface technology designs. *Ann. Biomed. Eng.* 35, 137–169
64 Merzenich, M.M. and deCharms, R.C. (1996) Experience, change, and plasticity. In *The Mind-Brain Continuum* (Llinas, R. and Churchland, P., eds), pp. 61–82, MIT Press
65 Buonomano, D.V. and Merzenich, M.M. (1998) Cortical plasticity: from synapses to maps. *Annu. Rev. Neurosci.* 21, 149–186
66 Yoo, S.S. and Jolesz, F.A. (2002) Functional MRI for neurofeedback: feasibility study on a hand motor task. *Neuroreport* 13, 1377–1381
67 Posse, S. *et al.* (2003) Real-time fMRI of temporolimbic regions detects amygdala activation during single-trial self-induced sadness. *Neuroimage* 18, 760–768
68 Weiskopf, N. *et al.* (2003) Physiological self-regulation of regional brain activity using real-time functional magnetic resonance imaging (fMRI): methodology and exemplary data. *Neuroimage* 19, 577–586
69 deCharms, R.C. *et al.* (2004) Learned regulation of spatially localized brain activation using real-time fMRI. *Neuroimage* 21, 436–443
70 Phan, K.L. *et al.* (2004) Real-time fMRI of cortico-limbic brain activity during emotional processing. *Neuroreport* 15, 527–532
71 Yoo, S.S. *et al.* (2004) Brain-computer interface using fMRI: spatial navigation by thoughts. *Neuroreport* 15, 1591–1595
72 Weiskopf, N. *et al.* (2004) Principles of a brain-computer interface (BCI) based on real-time functional magnetic resonance imaging (fMRI). *IEEE Trans. Biomed. Eng.* 51, 966–970

73 Yoo, S.S. *et al.* (2006) Increasing cortical activity in auditory areas through neurofeedback functional magnetic resonance imaging. *Neuroreport* 17, 1273–1278
74 Caria, A. *et al.* (2007) Regulation of anterior insular cortex activity using real-time fMRI. *Neuroimage* 35, 1238–1246
75 deCharms, R.C. *et al.* (2005) Control over brain activation and pain learned by using real-time functional MRI. *Proc. Natl. Acad. Sci. U. S. A.* 102, 18626–18631
76 Bray, S. *et al.* (2007) Direct instrumental conditioning of neural activity using functional magnetic resonance imaging-derived reward feedback. *J. Neurosci.* 27, 7498–7507
77 Birbaumer, N. and Cohen, L.G. (2007) Brain-computer interfaces: communication and restoration of movement in paralysis. *J. Physiol.* 579, 621–636
78 Wolpaw, J.R. *et al.* (2002) Brain-computer interfaces for communication and control. *Clin. Neurophysiol.* 113, 767–791
79 Birbaumer, N. (2006) Breaking the silence: brain-computer interfaces (BCI) for communication and motor control. *Psychophysiology* 43, 517–532
80 Mackey, S.C. and Maeda, F. (2004) Functional imaging and the neural systems of chronic pain. *Neurosurg. Clin. N. Am.* 15, 269–288
81 Rainville, P. (2002) Brain mechanisms of pain affect and pain modulation. *Curr. Opin. Neurobiol.* 12, 195–204
82 Thesen, S. *et al.* (2000) Prospective acquisition correction for head motion with image-based tracking for real-time fMRI. *Magn. Reson. Med.* 44, 457–465