

Real-Time Functional Magnetic Resonance Imaging¹

Mark S. Cohen²

UCLA Brain Mapping Division, 660 Charles Young Drive South, Los Angeles, California 90095

Magnetic resonance imaging (MRI) has been shown to be useful in the detection of brain activity via the relatively indirect coupling of neural activity to cerebral blood flow and subsequently to magnetic resonance signal intensity. Recent technical advances have made possible the continuous collection of successive images at a rate rapid compared with such signal changes and in the statistical processing of these image time series to produce tomographic maps of brain activity in real time, with updates of 10 frames/s or better. We describe here our preferred method of real-time functional MRI and some of the early results we have obtained with its use. © 2001 Elsevier Science

Traditionally, magnetic resonance imaging (MRI) has been a slow imaging modality. The physics of MRI require that increases in imaging speed result in signal losses. Specifically, the MR signal derives from the conversion of the sample (e.g., tissue) magnetization to a radio signal, and the magnetization recovers rather slowly. The relationship between MR signal (SI), the time between data collections (TR), and the longitudinal magnetization rate characteristic of the particular tissue (T_1) is described by

$$SI = k [1 - \exp(-TR/T_1)] \quad [1]$$

In Eq. [1], k includes a wide variety of additional factors including other tissue properties (T_2 , diffusion, flow, etc.), instrument parameters (magnetic field strength, radio coils, etc.), and geometrical issues. Although essentially all protons produce MR signal, the overwhelming majority of the signal in medical imaging derives from highly mobile aqueous and lipid proton nuclei. In

these tissues, the T_1 rates range from a few tenths of a second to several seconds (1). This dramatically limits the possibility of “real-time” MRI. For example, if data are collected in this fashion from a blood sample ($T_1 \approx 1.5$ s) every second, the signal intensity is reduced by more than 51% compared with a single data collection. This would not be a problem, but for the fact that the signal-to-noise ratio in the MR experiment is already quite small, and limiting, to normal applications. In conventional MRI, single images are formed from 32 to as many as 1024 repeated data sampling events, resulting in total imaging times of several minutes per picture.

Fortunately, over the past 15 years, technical advances in imaging have enabled substantial reductions in imaging time. The first and perhaps most clinically significant step was the development of low-flip angle of “FLASH” imaging (2), which reduced practical imaging times to several seconds. FLASH imaging makes much more efficient use of the available magnetization by taking advantage of the nonlinear conversion of magnetization to signal. Although not presently relevant to functional MRI, the technique known as RARE (3, 4), or Fast Spin Echo, resulted in a significant speed advantage to clinical imaging [for a review of these methods, see Cohen (5)]. The most dramatic speed advances, however, derive from the method known as echo-planar imaging (EPI). This technology was originally conceived as early as 1977 by Mansfield (6) and brought into laboratory practice only a few years later. However, the hardware available at the time was rather limited. Practical EPI became available a dozen years later, when high-field magnets became readily obtainable, when fast analog-to-digital convertors became cost-effective, and when advanced engineering made possible the production of the specialized very high speed, high-power, electromagnetic systems (gradients) needed for imaging applications (7). Mansfield’s group was the

¹ Portions of this material were presented previously at the International Society for Magnetic Resonance in Medicine.

² To whom correspondence should be addressed. Fax: (310) 794-7406. E-mail: mcohen@ucla.edu.

first to publish real-time MRI data by combining the EPI method with special-purpose imaging processing hardware, so that the data could be converted continuously to images (8). The fusion of EPI and low-flip-angle imaging made it possible to collect high-quality images (e.g., of the beating heart) at up to 16 frames per second (9). While EPI opened a large number of significant applications (10), the use of “real-time” MRI was constrained by the cost of special-purpose hardware and systems engineering and the limited range of identified applications.

Probably the most significant use of ultrafast imaging has been in the observation of blood flow and perfusion in the human brain. Rosen and colleagues realized that cerebral perfusion might be assessed with MRI by applying well-known tracer kinetic principles to study the signal changes that accompany the passage of intravascular contrast agents as they pass through the cerebrovasculature (11). On its own, this technique is exceedingly valuable in clinical assessment of primarily vascular problems, such as stroke (12, 13), as well as in conditions, such as cancer, that result in perfusion abnormalities (14).

Cerebral blood flow (and blood volume), however, has long been known to increase locally during periods of increased neural activity. In fact, the first observations of this association date back more than a century. That this coupling should occur was predicted from the earliest moments of the growth of physiological psychology. In one of his many insightful passages, James stated that: “We must suppose a very delicate adjustment whereby the circulation follows the needs of the cerebral activity. Blood very likely may rush to each region of the cortex according as it is most active, but of this we know nothing” (15). That same year, Roy and Sherrington (16) reported that the color of the brain reddened in response to local electrical stimulation, and inferred (correctly) that this was the result of increased blood flow. That vascular or metabolic changes might be used as markers for brain activity has been exploited in a variety of methods including autoradiography with 2-deoxyglucose (17), xenon-enhanced computed tomography (17a), positron emission tomography (18, 19), and, much more recently, direct optical methods (20).

Neurovascular coupling as an activity marker was exploited first with fast MRI in 1990 by Belliveau and colleagues (21). In these studies, the blood volume mapping methods developed by Rosen and co-workers were applied successively in conditions of “rest” and strong visual stimulation. By forming difference images, this group was able to show those regions whose blood volume was increased during visual stimulation. By comparison to PET, optical, autoradiographic, and xenon

methods, this new form of functional MRI was less invasive and had much better spatial resolution and was thus hailed as a revolutionary advance in the study of human brain function.

Concurrently, however, other investigators had noted that the MR signal has a strong dependence on the blood oxygen level. Ogawa *et al.* (22) reported that reducing the blood oxygen content, especially in high-field MRI, could be used to bring the venous microvasculature into high relief, while Turner *et al.* (23) showed that large global decreases in the signal from cat brain occurred when the animals breathed a high-nitrogen/low-oxygen atmosphere. Working at Harvard, Kwong showed that similar signal changes took place in the human brain during breath holding. Both Kwong’s and Ogawa’s groups realized quickly that this too could be used as a marker for neural activity and both showed in 1992 the first images from the technique that would become known as functional MRI (fMRI) (24, 25). Briefly, the method is understood to work as follows:

Increases in neural activity are metabolically costly and require an increase in substrate (presumably oxygen) delivery. Depletion of substrate in the capillaries reduces the transcapillary concentration gradients. To maintain the needed gradients, blood flow is increased. One consequence is an increase in venous oxygen (26). Although the physiological signaling that couples neural activity and blood flow is still not known precisely, it is likely to involve nitric oxide, vasoactive intestinal polypeptide, carbon dioxide, or a combination of these factors.

Mechanically, the method requires tracking the signal changes in the MR images that occur over time during cognitive, sensory, or other neural challenges. Although the brain electrical activity changes rapidly, over a millisecond time scale, the blood flow changes are relatively slow. Kwong *et al.*, for example, pointed out that during sustained visual stimulation, the MR signal continued to increase for about 7 s (24). Since that time, several groups have reported observing a predicted signal decrease within the first half-second of brain activation (27–29), although this remains an area of controversy. In any case this time scale is nearly ideal for study by echo-planar magnetic resonance; EPI is readily capable of imaging the entire brain every 1 to 2 s. At this sampling density, MRI can accurately follow the time course of brain activation. With this background understanding of fMRI, it became clear that the fusion of EPI, rapid computer processing and functional imaging would make a real-time implementation possible.

MOTIVATION

Clinical Applications of Functional MRI

During the initial years of its use, functional MRI attracted its greatest attention for the groundbreaking applications it allowed in the basic neurosciences (21, 30). Indeed, through fMRI we now have a much more sophisticated understanding of the neural architecture of higher-level vision [e.g., (31)], motor control (32), mental imagery (33), language processing (34, 35), and a wide variety of other tasks. From a public health perspective, however, there are other immediate gains to be realized. In this article we first explore the most important limits to the broader application of fMRI. In the latter sections, we show how we have developed a set of methods that directly address these key limitations.

Surgical Planning

The imaging tools available to guide neurosurgeons in both planning of the surgical approach and sparing of eloquent cortex have traditionally been limited. The method of electrocorticography, made famous by Penfield and Boldrey (36), is still used commonly today. In this approach, the surgeon studies the overt or subjective responses evoked in patients during or following direct galvanic stimulation of the exposed cortex. This may be done either during an awake surgery or through the chronic temporary implantation of surface electrodes. Aside from the obvious risks involved in such approaches, the technique has many technical limitations. In the vast majority of cases, for example, the electrodes can only stimulate effectively the outer few millimeters of the brain, whereas well over two-thirds of the cortical surface either is buried deep in sulci or is exposed only to the medial or inferior surfaces of the brain and is thus inaccessible. The complexity of the task that can reasonably be performed by the patient in surgery is clearly limited and the surgeon must, in any case, work very quickly to avoid secondary damage. The semipermanent surface electrodes have surprisingly ill-determined localizations; in some cases they may actually shift in position by several centimeters over a few weeks of implantation.

Localization of activity through surface EEG is another commonly used method in surgical planning. Though noninvasive, EEG offers only relatively crude localization of activity, on the order of a few centimeters (37), and its sensitivity drops rapidly with distance from the surface of the head. Even in the identification of epileptogenic foci, EEG is notoriously limited in its localizing power (38).

The continued popularity of the intracarotid amyntal,

or Wada, method (39) underscores just how limited the available imaging methods are. In the Wada technique, a small injection of sodium amyntal is administered through one side of the carotid system (a surgical procedure), and the patient is given a variety of either language or memory challenges. The time elapsed from injection to detectable decrease in function is then recorded and the procedure is repeated, after a suitable recovery period, with a contralateral anesthetic injection. Though something of a simplification, the primary end product of the study is the inference that if the patient's behavior collapses more rapidly with injection on one side, that half of the brain is implicated in dominant control of that behavior. On this basis, the surgeon must determine whether or not a critical function such as language is likely to be spared following cortical resection.

Functional MRI already offers significant promise in these applications (40). Binder *et al.* (41, 42), Benson *et al.* (43), and others, for example, each have demonstrated that fMRI methods may be more accurate than Wada examinations in assessing language dominance. Determinations of areas of hand and other motor activity are a routine component of the fMRI literature and are achieved with relative ease (24, 32, 44–46).

In fact, the basic science literature suggests that much more subtle definitions of eloquent cortex will be possible in the future. Interestingly, however, the surgeon's information requests tend to be quite coarse; generally the surgeon asks only, "can I operate safely in this brain region without loss of language/motor function?" In retrospect, the reasons for this are relatively clear. Before the functional examination, anticipating, for example, tumor resection, the surgeon will generally have available a structural study showing the location of a lesion. If that lesion lies in one of the obvious primary sensory or motor areas of the brain, the surgeon will see clear risk of damage and will want a functional assessment. In practice, the lesions often show mass effects that distort the cortical geography, and it is difficult to guess what functions are likely to be hampered. As only a limited number of functional assays may be performed with today's techniques (due to time constraints and other factors, see below), the surgeon must limit his or her questions to only the most severe risks.

Another important class of surgical planning applications will be in the preoperative assessment for pallidotomy or other stereotaxic lesions. In these cases, the surgeon should be able to directly visualize the brain regions associated with common disorders such as Parkinsonism, and use the activation studies to guide later electrode placement. These are technically complicated

studies, however, because the typical patient population is both elderly and impaired in ability to perform controlled motion. One consequence is a tremendous problem with motion artifacts (see below) that results in a large number of technical failures.

Neuropsychiatric Applications

It is by now well accepted that some classic neuropsychiatric disorders, such as schizophrenia, are accompanied by morphological changes in the brain (47–52), yet imaging is still used rarely in diagnostic workups. Our own work (53) and that of many others [e.g., Early *et al.* (54), Cleghorn *et al.* (55, 56), Buchsbaum (57), Buchsbaum *et al.* (58, 59), Friston (60), Friston *et al.* (61), Frith *et al.* (62), Liddle *et al.* (63), Gur and Pearlson (64), Silbersweig *et al.* (65, 66)] have suggested that there are distinct structural or functional abnormalities associated with psychiatric disorders that are visible in functional MRI or other imaging protocols.

We predict that functional MRI will have significant applications not only in the diagnosis of relatively common affective disorders such as schizophrenia and depression, but will be a major adjunct to the medical management of these diseases, guiding the internist in the effective prescription and dosing of pharmaceuticals. The majority of the published functional imaging studies in schizophrenia have been performed using PET imaging which, though obviously effective in demonstrating abnormalities, is not, and is not likely to become, as widely available as fMRI. The effective use of fMRI in this context is challenging, however, as the most sensitive PET measures have been based on FDG tracer studies. Since, at present, there is no clear MRI analog to such resting state descriptions of regional brain activity, we still need to develop effective functional protocols in the affective disorders. Some may be based on semiquantitative resting techniques such as spin-labeling (67, 68), but it is reasonable to expect that others will be based on the more dynamic confrontation protocols that are the traditional domain of fMRI. For example, Silbersweig and colleagues, using PET, (65, 66) and our own group, using fMRI, (53) have shown dynamic changes in regional blood flow that are correlated temporally with auditory hallucinations and bring us closer to the development of a meaningful clinical diagnostic protocol for the evaluation of central dysfunctions in schizophrenia. Others have shown (69, 70) that some of the characteristic behavioral abnormalities in the disease, such as eye-tracking changes, may be visible in fMRI studies as an alteration of the normal patterns of activation.

Spreading Disorders: Jacksonian Seizures, Migraine, Scotoma, and Stroke

Several relatively common neurological disorders are characterized by an apparent progressive spreading of the neural dysfunction that takes place over seconds or minutes. Classic Jacksonian epileptic seizures, for example, often start with a focal motor twitch in the extremities and progress centrally, ending, ultimately, in generalized seizures. Migraine headaches may have an analogous course, starting with focal signs such as visual scotomas, field cuts, or distortions, and progressing to generalized headaches. Woods *et al.* (71) have shown that the spread of migraine might be visible by blood flow imaging. The possibility of observing brain maps of such phenomena in real time may lead not only to a better understanding of such disorders, but also to the development of therapeutic interventions to arrest the symptom progression.

Why Must fMRI Methods Be Improved for Clinical Applications?

The majority of the published fMRI work has been fundamentally technological, with a smaller number of experiments addressed directly to questions in physiological or cognitive neuroscience. With the success of the technique has come both tremendous scientific and lay interest and many emerging clinical applications. fMRI is a relatively new method, whose interpretation relies on many assumptions; some of the important ones [in common with H₂¹⁵O PET (72)] are that: (1) MR signal change, presumably due to hemodynamic changes, bears a reliable relationship to neural activity; (2) other stimulus-correlated changes [e.g., head motion (73)] are small or readily corrected; and (3) “ceiling” and “floor” effects of signal changes can be neglected over the physiological range studied. Despite these theoretical limitations, the experimental results have been impressive and both confirm the results of more direct measures, such as local electrical recordings, and are predictive, for example, of losses following surgical extirpation of tissue. As fMRI matures and becomes increasingly the domain of the applications researcher, the underlying validity of the physiological assumptions and the stability of the experimental apparatus become increasingly important factors in data interpretation. In a recent editorial (74), we argued that, even now, many dubious interpretations of fMRI data are based on overambitious acceptance of the power of the current imaging and analysis technology, a regrettably large fraction of which are part of the current literature. Very clearly, there is much work to be done in the development of

accurate and reliable data collection and analysis strategies, particularly for applications in mental health research.

Avoidable Experimental Failures

As fMRI is a newly developed method, much of the early work has concentrated on results obtained from normal or supernormal (selected from the research staff) subjects, and has been collected over long, and often repeated, sessions. In these cases, the occasional experimental failure, though inconvenient, is only minimally limiting. In our hands, 20% or more of the data may be evaluated as uninterpretable due to technical failures. A casual review of the published fMRI literature suggests that this yield is not atypical. As the applications move to more complex populations [our own research concerns schizophrenic, elderly, and pediatric subjects (40, 53, 75, 76)] experimental failures become not only expensive, but in many cases fundamentally limiting. To work with these important subject groups, it is necessary to have a highly reliable data collection and analysis process that includes an immediate, objective, measure of the quality of the collected data. Only in these circumstances will it be possible to immediately repeat studies that fail or to continue imaging just until a satisfactory statistical threshold has been reached.

The causes of such technical failures are manifold. The most commonly reported are attributed to subject motion, a serious problem in fMRI (30, 73, 77). Generally speaking, motion artifacts cause two classes of outright failure. The first is obvious motion artifact, typically presented in “statistical images” as an embossed appearance with a dark edge on one surface and a corresponding light edge on the opposite (33, 73, 78). The second presents as excessive variance throughout the image, resulting in a failure to obtain a statistical

threshold (we prefer to avoid the term “statistical significance” in fMRI as the underlying assumptions in determining probability, such as sample independence, corrections for repeated measures, etc., are, in general, unverifiable, though significant progress has been made to develop more rigorous test, [e.g., (79)]).

Motion also can create a much more insidious set of errors. When it is correlated with the subject’s task, some motions can appear as cortical “pseudo-activations” (73), because of the large intensity gradients that exist at the cortical surface in MR images: small rotations or translations can appear as large signal changes in the first few millimeters of cortex, exactly where such activations might be expected if neurally based.

Though reported less frequently, scanner performance problems also plague the collection of reliable and repeatable functional images. Many researchers now favor echo-planar (Mansfield 1977) or spiral scan (80) MR acquisitions that take advantage of ultrahigh-performance gradient sets. To effectively cover the raw data space (k space) rapidly enough to collect functional images, it is ordinarily necessary to apply large voltages to the gradient coil sets. A variety of instrumentation approaches (7, 81, 82) have been developed to do so, some delivering thousands of volts to the gradient coils. Unfortunately, these high voltages result often in small corona discharges or arcs containing substantial spectral energy in the MR frequency range that introduce discrete high-amplitude “spikes” into the raw data. After Fourier transformation, such spikes appear as characteristic diagonal banding or Moiré patterns in the images that add considerable noise variance to the statistical functional images (Fig. 1). These image aberrations must be eliminated (hopefully prospectively, but often retrospectively) for reliable functional image calculation.




FIG. 1. Typical appearance of functional images contaminated by high-voltage spike artifacts during data collection. Image A is “clean,” whereas images B–D each were contaminated by one or more spikes in their respective raw data. As the data are Fourier transformed prior to reconstruction, the specific appearance of the spikes on the image can be quite variable. A single spike slightly to the left or right of the raw data origin would result in the appearance of image B. The artifacts in images C and D are likely the results of multiple spike events.

Processing Strategies and Magnitude Estimations

fMRI using the blood oxygenation level-dependent (BOLD) or T1-based approach is a method of differences; the raw data images provide little quantitative information. To localize regions of signal change, images acquired in different conditions are compared, voxel by voxel.

The simplest comparison is simply to subtract all images obtained in one condition from those obtained in another [as used in Kwong *et al.* (24)]. Subtraction yields images in arbitrary units of MR signal intensity, and is particularly prone to type I statistical errors, as high-intensity artifacts look like areas of functional signal change (73, 78). More commonly, at present, investigators use some manner of statistical method to normalize the signal change by the voxel variance (the z or t statistic, providing fundamentally similar measures), creating as output statistical parameter maps (SPMs). This greatly reduces the sensitivity of the functional maps to many sources of excess variance (83). The t statistic is used commonly in SPMs (84) because it gives a plausible picture of activation loci. Unfortunately, the MR signal changes are not well described by the difference in mean levels, as they contain important temporal characteristics. Figure 2, reprinted from one of the earliest published fMRI reports (24), emphasizes this point: Over the course of the 1-min “on” cycle, the signal takes nearly 10 s to reach its final intensity; the signal has an even longer return to baseline. Clearly, the t statistic will be overly conservative, in this case treating the signal rise and fall time as sample variance. It is also obvious that the signal change lags the stimulus and that this effect is not captured by a simple application of the t test. As in PET (most often analyzed using Friston’s SPM software), activation magnitude may be expressed as a z score, dividing the signal intensity change by the local variance. This statistic strongly attenuates the apparent magnitude of activation in regions with high local variance (e.g., blood vessels) and

has the advantage that bright areas on brain maps are likely to represent real, rather than artifactual, activation. Bandettini *et al.* have developed the use of correlation statistics to detect areas of activation (85). In this approach, the signal intensity time course in each voxel is cross-correlated with a function describing the task (such as a “box car” or “square wave”) or, when the activation blocks are short enough, with a sinusoid that mimics the observed vascular response. In the case of sinusoidal models, the response latency can also be controlled or determined by allowing a phase term in the model to be a free parameter.

Neither the z score, the t statistic, or the correlation maps, however, are useful in comparing degrees of activation across subjects, or even necessarily in the same subject across experimental runs. Each approach tends to scale the magnitude representation by an estimate of the variance which, because it is determined strongly by subject physiology, may vary across subjects or runs, and will depend strongly on experimental conditions such as time between images and scan time duration. Because the variance appears in the denominator on such statistics, when the signal-to-noise ratio (SNR) is large, even small fluctuations in the image variance can create very large changes in the magnitude of the derived statistic. As pioneered by Friston *et al.* (86), a separate t map may be generated to compare activations across runs on a single subject or, with suitable spatial normalization and image registration, across groups. Again, the magnitudes are transformed to dimensionless numbers that are difficult to relate to physiology. In the case of across subject comparison, they will also depend strongly on the variations in individual brain morphometry and task-related functional anatomy. In practice, the PET scanner does not yield images with sufficient signal-to-noise ratio to withstand a Bonferroni correction based on the number of voxels in the PET image. Friston *et al.* have developed methods to avoid this limitation, many of which are now used by practitioners in the field. In their 1991 paper (86) they proposed that, by quantifying its “smoothness,” one could develop an estimate of the number of effectively independent resolution elements in a PET image that allowed a much less conservative Bonferroni correction, and enabled PET images to reach statistical significance in reasonable circumstances. More recently, this group has advocated the use of relatively heavy Gaussian smoothing to ensure that the spatial noise distribution in the PET images is itself Gaussian. In so doing, this allows the application of the theory of Gaussian fields to the derivation of probability estimates (87, 88). Such aggressive smoothing also minimizes the effects of individual anatomical differences in pooled multisubject data. As interesting as these




FIG. 2. Signal intensity time course in primary visual cortex during presentation of 8 Hz patterned flash stimulation. Reproduced, with permission.

techniques are, we eschew their use in MRI for the following reasons: First, the blurring applied to the images can dramatically reduce the apparent effect magnitude and in some cases, entirely eliminate the detectability of otherwise reliable events (89). Second, the loss of spatial resolution implicit in these techniques obviates the advantages of MRI's intrinsic spatial resolution. Most importantly, transformation to smoothed t fields leaves no effective metric for the comparison of magnitude across subjects or tasks; it is useful only for hypothesis testing.

We note also that the sine-wave or Fourier analyses as originally proposed by Bandettini, though extremely powerful, place undue burdens on the experimental designs, essentially requiring that only two conditions be used, alternating at a fixed rate determined by hemodynamic considerations. Newer strategies, such as those of Buckner *et al.* (90) and our own as outlined below (91), relieve many of these constraints.

MR Data, and Data-Processing, Architecture

Typical MRI instruments are architected around clinical applications that involve reading the relative magnitude of the signal intensity in various regions on individual images. Few provisions are made for the sort of time series analysis that is essential for functional MRI. Until recently, the usual data acquisition times have been lengthy as well, ranging from tens of seconds to minutes for the collection of a single multislice image set. In such circumstances, image reconstruction times of a second or so per image were not limiting. With the advent of practical echo-planar imaging (6, 10, 92), however, good-quality scans can now be collected in less than 25 ms, so that the image reconstruction times can quickly become the determining factor in total scan time. Consider a typical functional imaging session on our unit, examining the signal changes during two experimental conditions A and B and a rest condition in a block design. In a typical block design, the A and B conditions are repeated twice each and are separated by rest periods (Fig. 3): 20 axial slice locations of 4-mm thickness cover the entire brain. At a TR of 2.5 s, one

image is acquired every 125 ms; 2160 images are acquired in the $4\frac{1}{2}$ -minute study. Without modification, our scanner processes 128×128 images at the rate of 2 per second, so that the reconstruction time is 16 min—almost four times the image acquisition time. It is seldom possible for “nonprofessional” subjects to tolerate MR examinations longer than 90 min. Given that practical sessions require a few minutes of setup time, and the acquisition of a variety of structural images, this means that no more than four “functional” runs are generally possible during a scan session.

With traditional data processing tools, the time burden is extended further in functional MRI, as the statistical image processing cannot begin until all images have been reconstructed. This is because most such computations process individually the vector describing the signal intensity time course at each voxel. All images must be acquired before any location can be completed. A second consequence is that for rapid processing, unrealistic memory requirements are placed on the host computer. As each time-intensity vector is formed from all images, the usual implementations require either that the entire data remain resident in RAM at once (2160 images * $128 * 128$ pixels/image * 4 bytes/pixel (floats) = 128 Mbyte)³ or that the data be swapped from disk to RAM multiple times for each pixel (Fig. 4). The data swapping results in significant increases in data processing time. One way to manage this is to defer all image reconstruction and data processing until all data acquisition has been completed on a single subject. While this allows a few more runs to be collected, it opens the overall experiment to additional failure mechanisms. If the subject moved, instrumental noise increased, the subject's breathing rate was abnormal, or the subject failed to perform the task (any of which increase the denominator in the statistical processing, thereby decreasing the power or sensitivity) the data acquisition cannot readily be repeated, for the subject will be long gone before the processing has been completed.

Block Designs and Predetermined Protocols

For many experimenter-driven protocols, it is acceptable that the timing of the behaviors or stimulation be fully determined in advance. In fact, until recently, nearly all published fMRI data used this sort of design. We are, however, aware, that this is severely limiting to applications and is a very poor model for real-world behaviors. In some cases it is interesting to determine the activation state of brain regions at a variety of time points in processing of a single task. In the study, for




FIG. 3. Schematic of a typical blocked protocol used at UCLA for functional imaging. Two conditions, A and B, are compared. Each is repeated twice with an intervening rest period.

³ At the time of this writing in 1999, such memory requirements are highly limiting.

example, of mental object rotation, (93–95), the actual task processing times can be quite long (up to 12 s), leaving open the possibility that fMRI can be used to study activity during components of the rotation task. Since the task processing time may be indeterminate (depending on the subjects ability) it is not realistic to use a blocked experimental design in this case. Similar examples abound in other cognitive applications.

In a blocked experimental design, the subject performs a single task repeatedly for the duration of a block and switches tasks in each new block. The most commonly used data analysis approaches (85) explicitly require blocked designs, or even periodic, blocked approaches. Most cognitive psychophysics experiments and studies of learning and memory do not fit well into this design, because the subject quickly learns that all trials in a task are the same; as a result the brain activation pattern may change during the block. In other experiments, the actual timing is not blocked, but may not be knowable in advance. For instance, there has been a great deal of interest in the use of fMRI potentially to isolate foci of epileptogenic activity based on the signal changes that might accompany interictal spike discharges (96–98). Obviously, the timing of such

activity is outside of the experimenter's control, and an appropriate analysis method is needed. Below we show an illustrative example.

Proprietary and/or Localized Image Analyses

There are already a few other tools, including particularly AFNI (written by Dr. Robert Cox and colleagues), SPM (written by Dr. Karl Friston *et al.*), and MedX (Sensor Systems), that offer reasonably convenient access to analyze images from a variety of scanners. Of these, at this writing only AFNI appears to have been extended to real-time use; the others are designed around off-line processing. Further, unix workstations, as needed for the available tools, are notoriously difficult for single users to maintain and typically require the help of individuals with tremendous computer sophistication. While many research centers have excellent access to high-end computing on expensive workstations, and have a budget appropriate to obtain site licenses for tools such as MatLab, the majority of clinical imaging centers do not. There is a remarkable lack of quality radiological review and analysis tools available for the sorts of small computers that might be more available to clinical centers, and there is very little attention in the marketplace to software packages that are easy enough for the computer nonspecialist to use.

Though many instrument vendors claim to support the ACR-NEMA format for image data interchange, in practice, it is usually necessary to use the manufacturer's proprietary data format to exchange and view images readily. Since there are few, if any, tools for small computers that already recognize such formats, the typical clinical user is bound to the scanner manufacturer or the expensive solutions suggested above, if she or he wishes to review and analyze images.

These are serious barriers to the dissemination of functional MRI, as too much is required of the user in learning and expense to get started readily on the applications of the method.




FIG. 4. Schematic view of the structure of the four-dimensional raw data describing a single functional imaging experiment. Each time point is represented by a 2D slab of about 320 kByte (for 20 image planes). To form the statistical images we must, in general, study the correlation between the time-intensity behavior of each voxel and a reference function describing the subject's behavior or task. In conventional strategies, where each voxel is studied individually (e.g., the gray voxels in the upper right of the figure), we must typically load the entire data set into main memory at once (placing enormous demands on memory), or, when memory is limited, we must repeatedly swap the data in and out of memory to sample all time points. As the data set becomes larger (more time points), the memory requirements scale proportionately. As currently configured, our scanner can acquire more than 6000 images, which is a realistic number for applications in cognition, but our computing systems simply cannot be extended to keep up with this processing need.

METHODS AND RESULTS

An MR Instrument for Applications in the Cognitive Neurosciences

Our group at the UCLA Brain Mapping Division has been active in the development of high-field MRI instruments designed specifically for functional and high-resolution imaging of the head and brain (99). By specifying an instrument for imaging of the head only, we

have been able to optimize many of the other MR performance factors. Briefly, limiting the homogeneous volume to 25 cm allowed us to use a very high field strength of 3 T without suffering undue cost. We were able to eliminate the need for a body RF coil. This allowed for a larger magnet clear bore, elimination of the body RF shield (and its associated eddy current losses), and the use of a solid-state RF power amplifier, as only 3000 W are needed (as opposed to the 30 kW that would be needed for body RF). The major advantages, however, will likely occur in the gradient technology. MRI speed is a strong function of gradient strength and rise time. These, in turn, are limited both by power requirements and by the possibility of current induction in the human body. Reducing the gradient volume to cover only the head mitigates the power problem dramatically, as the power requirements scale with the fifth power of the gradient diameter due to increases in efficiency. As dB/dt increases with distance from the gradient isocenter, it follows that smaller, head-only gradients, can switch at higher rates and remain below the threshold of sensation.

Figure 5 shows representative high-resolution images obtained from this unit. On the left is a 3-mm slice through the hippocampus, obtained in under 9 min as part of a multislice series. On the bottom right is a histological section from a similar brain slice (in a non-living subject). The high-field MR unit has sufficient

signal-to-noise ratio (SNR) and contrast that subtle details, such as the white matter of the perforant path, are readily apparent in the images. In many functional imaging applications, such high-resolution conventional images are used in certain circumstances in producing high-quality surface renderings of the image data as a base on which to superimpose and localize activation maps. The analysis techniques shown here take advantage of this performance, but the results are fully generalizable to the more conventional 1.5-T and small instruments.

The high intrinsic SNR of this unit, coupled with its unusual gradient performance, have allowed for dramatic EPI results as shown in Fig. 6. Figure 6 shows a multislice echo-planar series covering the entire human brain. Each 5-mm slice plane was scanned in 78 ms using a spin-echo sequence ("MBEST" (101) or "Instascan" (9, 102). Lipid suppression, immediately preceding each excitation (103), results in a maximum acquisition rate of about 10 such images/s. The relative absence of spatial distortions due to susceptibility effects (104) as compared with other high-field EPI data sets (105) is due largely to the high readout switching rate (1400 Hz in these images) and the image-based shimming (106) of the linear and second-order terms. We stress that Fig. 5 and 6 are of routine image quality. The echo-planar data sets are of particular relevance to the




FIG. 5. High-resolution MR image (A and magnified in B) obtained from the brain mapping research MR instrument, compared with a similar histological postmortem specimen (C). The MR images from this instrument achieve extremely high signal-to-noise ratio and contrast, and can thus demonstrate extremely subtle anatomical features, such as the perforant path from hippocampus to entorhinal cortex (arrow in C). MR parameters: TR, 4000; TE, 25 (effective); fast spin-echo sequence; ETL, 8; thickness, 3 mm; matrix 512×512 ; 2 NEX. Scan time: 8:38. (C) is reproduced, with permission, from DeArmond *et al.* (100).

functional imaging studies; images of this kind are the basis for the activation data shown later in this section.

Data Processing Rates

Image Reconstruction

As alluded to above, the practical requirements on traditional imaging systems have only entailed that a relatively small number of images (usually fewer than 100) be reconstructed in a time short compared with the typical scan times (5 to 10 min). Echo-planar-based functional imaging suffers from much more severe constraints. Working very closely with the scanner manufacturer, we developed modifications to the core image processing code and hardware on the Advanced NMR system that have increased the basic image processing speed from 2 to nearly 8 images/s for 128×128 MR images. The main modifications included the addition of a Sky array processor with four Intel i860 processors and 16 MB of dedicated memory with careful attention to vectorization and optimization of the image processing code, so that, for example, floating point operations and static variables were kept to a minimum. Most of these changes were implemented by Xiaole Hong, of Advanced NMR Systems, and the PI (107). Since such modifications are standard computing procedures, we will not detail them here. As detailed below, our initial implementation has also included modifications to the databasing code to allow us to save the images to a

remotely mounted disk. Commercial units are now available that reach this inconstructive speed readily.

Statistical Image Processing

We have discussed already how typical algorithms for calculation of the t statistic or correlation coefficient examine the time intensity behavior of each image voxel as a vector and require, therefore, that all data be acquired prior to beginning the data processing (108). This, of course, is anathema to the real-time data processing goals of the present project. Even after the data are collected, this approach is woefully inefficient, as it places extraordinary burdens on the random access memory (RAM) size of the processing computer or requires tremendous amounts of disk swapping (Fig. 4). Our approach is simple: rather than work on vectors of time data, we work on vectors of images. As an example, we consider the calculation of the correlation coefficient, r , using the standard formula of Pearson,

$$r = \frac{N\sum xy_i - \sum x \sum y_i}{\sqrt{\{N\sum x^2 - (\sum x)^2\}\{N\sum y_i^2 - (\sum y_i)^2\}}},$$

letting x represent the paradigm reference waveform (which is common to each voxel in the volume) and y_i represent the pixel intensity at volume location i . The calculation must therefore be repeated once for each voxel in the image. However, it is necessary only to maintain in memory three volume-sized vectors containing the current values of each of the sums in y , y^2 , and xy , and scalars containing the terms in x and x^2 to




FIG. 6. Echo planar imaging at 3.0 T. Shown are 20 4-mm axial sections through a normal brain, each acquired in 78 ms using a spin-echo EPI sequence. The field of view of 20×20 cm is covered with a matrix size of 128×128 pixels for an in-plane resolution of just over 1.5 mm. Other parameters are: tr, ∞ (EPI single shot); $t, 54$ ms; readout duration, 46 ms.

update the value of r for each voxel as each new image arrives. Figure 7 shows the three summary images after 75 images of processing. As we will see below, it is expedient in practice also to maintain vector descriptions of the x and x^2 terms for the convolution analysis. For the typical fMRI data set described above, assuming that we wish to perform this processing on all 20 slice locations continuously, our approach requires only about 6 Mbyte of RAM. Further, as the time series becomes longer, the RAM requirements are unchanged; only increasing the number of slice locations increases the memory need, and then only linearly. We can, in fact, process an unbounded number of time points using the same, limited, RAM store.

The advantages are immediately apparent. Table 1 shows the benchmark results for the traditional vector-based calculation and our modified algorithm in processing of typical data sets, using a DEC Alpha 3000/300 LX workstation. While the conventional processing dealt with each location sequentially, the modified algorithm handles all 20 locations simultaneously. In data sets of more than 200 time points, the conventional approach fails completely, even with 64 Mbyte of RAM and 500 Mbyte of swap space, whereas our modified algorithm maintains a constant processing speed of 39 images/s. This approach to image statistics can be applied just as easily to t tests though, as discussed above, we do not typically favor their use.

The MR Impulse Response

It is a basic result of linear systems analysis [see, for example, (109)] that for a linear, time-invariant system, the convolution of its impulse response with an arbitrary input will yield the system's response to that input. While the human brain a priori fails to meet the conditions of either linearity or time invariance, we

suggest that there exists an important regime over which such analysis is appropriate, as will be discussed below. Even with the limitations of these assumptions, our preliminary data suggest that modeling the fMRI responses based on convolution of its impulse response provides a reasonably accurate fit to the responses to real-world stimuli; these response models can then form the basis of correlation analysis to detect areas of activation.

For an estimate of the impulse response, R. Savoy of the Rowland Institute (Cambridge, MA) has provided us with data representing the averaged fMRI response to 10 repeated presentations of a brief (1-s) light flash. We have modeled these data by creating a three-parameter fit to a gamma variate function ($S = k t^{8.6} e^{-t/0.575}$, where k is simply a scaling function).

Because the system is causal, the convolution need only be applied to data following the behavior. One extremely important result is that the correlational analysis can still be performed in real time, as the convolution, in essence, is a weighting function to be applied only to future images (see below). In fact, for more than 400 ms following the stimulus, no response is seen.

We recognize many limiting assumptions in this analysis approach: The system is known to be nonlinear

TABLE 1

Data size	Conventional processing	Modified processing
20 locations, 25 time points	0:39 = 12 images/s	0:14 = 39 images/s
20 locations, 110 time points	2:27 = 14 images/s	0:56 = 39 images/s
20 locations, 200 time points	Could not complete	1:42 = 39 images/s




FIG. 7. Our real-time processing scheme takes advantage of the fact that both correlation and t statistics can be made from the ratio of summary values in y (A), y^2 (B), and xy (C). We maintain such summary images in memory during continuous statistical calculation and use the products and ratios of these summaries to derive our final statistical maps.

(see above), the system is surely time-variant (humans learn, accommodate, habituate, etc.), and the impulse response in the visual system is a questionable model for the impulse response in other areas. We, however, believe that each of these assumptions may be tested directly and, where they fall short, may be corrected for. We note also that similar approaches have been developed and tested by others [see, for example, Buckner, *et al.* (90)]. Deconvolution methods have promise in deriving the impulse response function from the raw data, given sufficient information about the stimulus and or behavior (110), and have, in fact, been used to develop event-related fMRI maps.

A Simple Example in the Motor System

Figure 8 shows a finger tapping protocol and its convolution with the impulse response above (the scaling constant, k , has been adjusted to give a unit amplitude response to extended stimuli). In this protocol, the subject alternates 20 to 40s blocks of rest with 20 s of finger tapping. Note that the predicted response lags the stimulus, as observed in early fMRI papers (24) and that several seconds are required for the response to reach its peak.

This experiment was performed on our 3-T imaging

system using the following parameters: Echo-planar sequence, TR = 2.5 s, TE = 42 ms (=T2*), slice thickness = 5 mm, 16 locations, FOV = 20 × 20, matrix = 64². The fMRI data, extracted from a region of interest in the precentral sulcus, show excellent concordance with the estimated fMRI response. In this region, the correlation coefficient exceeded 0.78. By comparison, in the same region, the correlation with a square wave (boxcar) describing the behavior in the same region was less than 0.61. In essence, the estimation procedure can be seen as reducing the data variance by incorporating the shape function. The image maps shown to the right further emphasize this point, showing maps in three slice locations made using either the t test or the correlation with the response estimate. The latter shows more focal and intense areas of activation and less extraneous “signal” outside of motor areas. Combined, these data are promising in suggesting increased sensitivity using this analysis procedure.

Linearity

In many cases, linearity need not be assumed: it is possible to suitably transform the input to model the nonlinearities of the neurovascular response. Figure 9

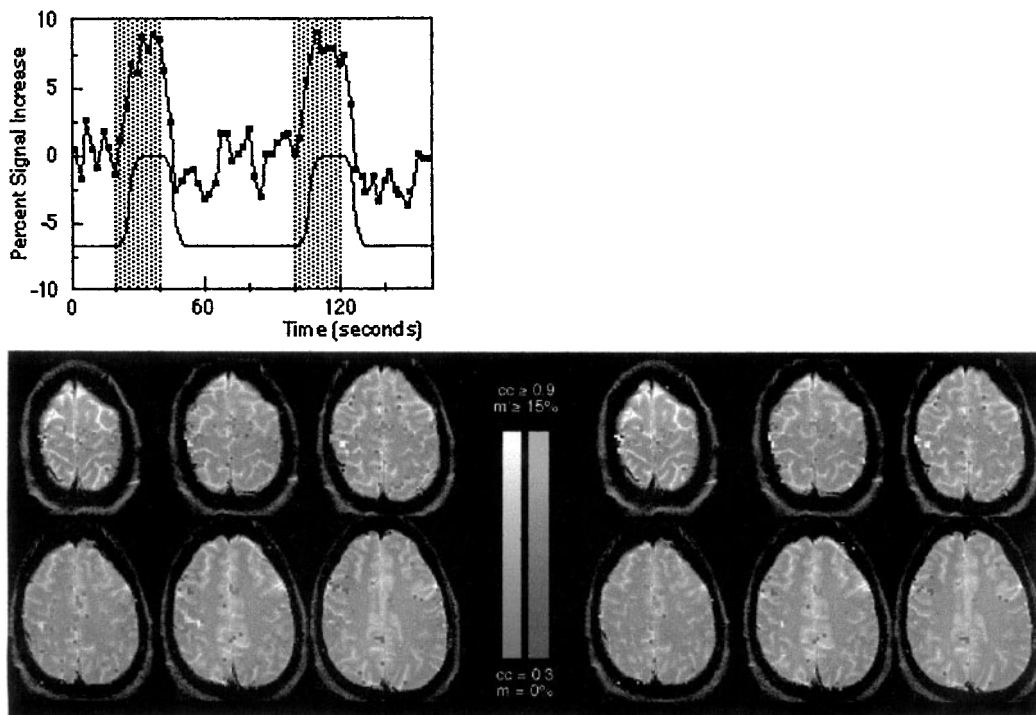


FIG. 8. Top: Finger tapping protocol and response estimates. Gray bars indicate periods of hand motion (sequential opposition of the thumb to each finger at a rate of 2 Hz). The raw time data are indicated by the filled symbols, and the response estimate based on the convolution method described is shown below. Bottom: “Activation” maps derived using the t statistic (top) or correlation with the response estimate (bottom) in three different slices. Arrows show activation clusters.

shows the response in paracentral gray matter as a function of finger tapping rate and the response in visual cortex as a function of light flash rate. Both patterns in the signal intensity (SI) are well-fitted by a monoexponential of the form

$$SI = k(1 - e^{-mB}),$$

where B models the behavioral intensity (finger tapping rate or visual contrast) and m and k are fit parameters. Because a ceiling effect must exist for stimuli over a certain level, a response of this general form is not unexpected.

Slope or Beta Maps

One advantage of the model fitting procedure described above is that it is a straightforward matter to transform these data into an estimate of the response amplitude by calculating the slope of the best fit between the observed signal change and the model. If the model parameters are scaled to unity at peak, then the slope can be expressed in percentage signal change per unit input. Unlike the t maps, z scores, or correlation maps, these maps provide a normalized measure of response amplitude that can be compared across subjects and trials. While maps of percentage signal change for a given stimulus share this virtue, it is far from clear whether such maps should express peak signal change, average signal change, or some other function. The information in the correlation and magnitude maps is complementary. The cc map gives an estimate of confidence, while the slope map estimates magnitude. In many real-world examples, the two are largely independent. This method and transformation are described in

more detail in (91). The most important result of the slope transformation, however, is its high degree of repeatability. While our data (and many published papers) show tremendous intertrial variability in the number of pixels above a correlation threshold, the slope parameter is highly stable (112).

Further, slope maps such as these allow an expression of the response magnitude for even complex tasks. Figure 10 is such an example, in which the slope map shows the percentage change in signal associated with changes in attention. For the tactile attention task the subject was to respond with a finger press using his dominant (right) hand each time he felt a touch on his left ring finger. In the verbal task, the subject was to monitor a semirandom stream of words for a target word (“Albania”). Both stimuli were presented simultaneously, but the subject alternated monitoring of the two sensory channels. The response estimates were treated similarly to those above. That is, they were convolved with the visual impulse response function. Our interpretation of these data is that directed attention can result in a decrease in signal in the unattended sensory channels.

A Basic Implementation for Real-Time Imaging

With the technologies for the rapid construction of quantitatively stable functional maps in hand, we have focused on the development of real-time imaging through the integration of the scanner, the stimulus presentation system, the response monitoring, and the image reconstruction. Our present implementation is shown schematically in Fig. 11.

In the General Electric/Advanced NMR hybrid system, the central scan control lies with the Signa (General Electric) hardware. This initiates the data collection for each image and specifies parameters such as




FIG. 9. Left: Mean fMRI signal paracentral gray matter as a function of finger tapping rate. Right: fMRI signal in V1 as a function of frequency of patterned flash presentation from 0.25 to 8 Hz (Note: The signal is known to decrease at higher flash frequencies (24, 111). Plotted are the mean and standard error of the fMRI signal intensity during the stimulus interval and a fit to the raw data based on a simple logarithmic response.




FIG. 10. Signal changes in precentral gyrus (hand sensory), inferior frontal gyrus (IFG), superior temporal gyrus (STG), and calcarine cortex during attention to touch or word stimuli delivered simultaneously.

spatial location, orientation, and slice thickness. The Advanced NMR components control the high-speed gradients necessary for echo-planar imaging, the data acquisition, and basic image processing. Our modified ANMR code presently processes the raw MR data to images at better than 10 images per second and allows us to save the data to a remotely mounted disk. All further image processing is performed on a Macintosh computer running scanSTAT, a program we have developed for this application. When the data for a complete volume (all slices) at a given time point are acquired, they are read into the Macintosh, together with the behavioral data, and a new statistical map is formed. On the 233-MHz G3 PowerPC, the new map for the entire volume is available in less than 1/20th of a second. Behavioral data are logged onto the Macintosh using the standard serial port, which is more than fast enough for this application. For example, the subject can depress an MRI-compatible keyboard to indicate his/her cognitive state or responses to questions. These data are convolved immediately with an impulse response estimate, and a file describing the convolved response is saved locally.

This implementation allows us to acquire real-time images with immediate feedback to the subject (107). With this system we are able to acquire and process functional images using behavioral protocols that are not predetermined, while still applying the full power of the nonlinear impulse response convolution approach.

Artifact Rejection

Since our initial work in functional MRI in 1991 (21), for better or worse, we have acquired an enormous

amount of experience in the typical artifacts associated with the technique. The overwhelmingly important concerns are with patient/subject motion and with scanner instabilities and spiking. Our generic approach to artifact detection and rejection has been to analyze each new image or image volume by eye, and to compare it to those acquired previously, flagging large deviations for later correction (e.g., reregistration) or rejection in the statistical analysis. Below, we present below data on an automated artifact detection and rejection approach.

To provide a quick and relatively simple detection for motion, we derive a noise threshold for the first image, and create a mask indicating which pixels are above and below noise. The algorithm we presently use to determine the noise threshold first creates a 256-point histogram of the pixel intensities (Fig. 12). It then determines the peak histogram value in the lower 15% of the full range (which should contain all of the noise voxels), and determines the local minimum above that value. While the formula is empirically based, it is fully determined by the data and correlates well with a visually defined threshold.

We then apply the same threshold to each subsequent




FIG. 12. Signal intensity histogram for a typical MR image of the brain. The arrow indicates the noise threshold that we calculate automatically.

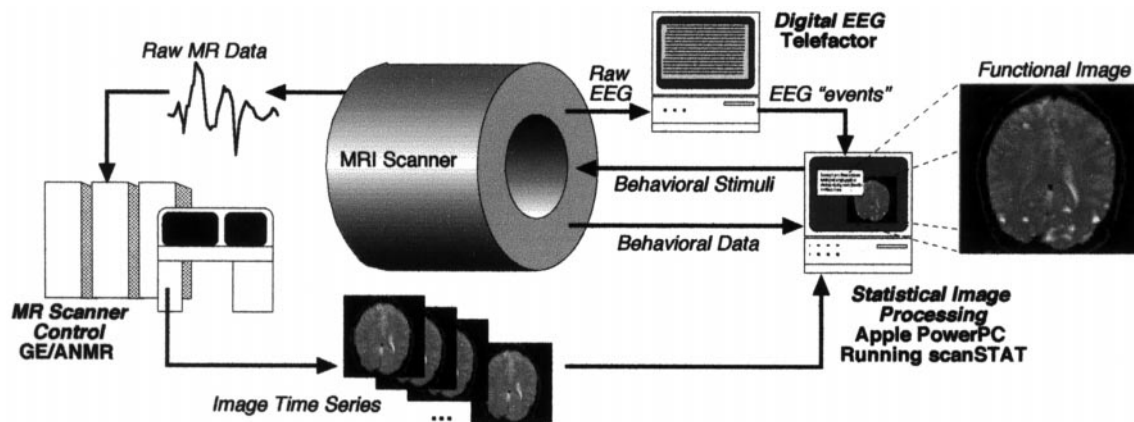


FIG. 11. Schematic of the real-time functional imaging system. In this implementation, the MR data are processed locally on the Sun system to form two-dimensional images at the rate of 10 images/s. Digital EEG data are recorded simultaneously on a specially developed Telefactor console. The actual statistical processing is performed on the Macintosh, integrating behavioral data acquired from the Macintosh computer. The Macintosh system, in addition, is used for stimulus display. Overall scan control is run from the Signa console. All of the computing is run asynchronously.

image and count the number of pixels whose identity changes from noise to signal (or vice versa) in that image as compared with the mask for the first image. We also compare each new image with a mask indicating the suprathreshold pixels in the image immediately prior. Since the mask indicates pixels that are not noise, it is a reasonable proxy for a representation of the edges of the head. Any changes in position are likely to cause such a change, but changes in intrinsic brain signal (e.g., from activation) are not. Further, since there are so many pixels that define this edge, this simple approach can detect motions of far less than one voxel.

Spike artifacts from factors such as static electricity discharges tend to create an increase in signal throughout the image, because the spikes generally contain energy at all frequencies and thus, after Fourier transformation, they add energy at all spatial locations superimposed onto the MR signal. When the signal magnitude is larger than the noise, depending on the signal phase such spikes can either increase or decrease the magnitude MR image intensities, accounting for the Moiré-like patterns that appear in Fig. 1. When the noise exceeds the signal, construction of the magnitude image effectively rectifies the noise, resulting in a Rayleigh noise distribution. The signal outside of the brain or sample is almost always increased by the presence of raw data spikes. We can therefore detect spikes by the large number of locations whose signal intensity has increased over threshold compared with prior images.

Motion artifacts will present as an approximately equal number of locations in the volume whose signals

increase and decrease, as the head translates in position. Unless the motions are quite large, the number of changed pixels is likely to be much smaller than for spike artifacts. In the event that the subject manages to move back to the original location, we can look at the cumulative, rather than the sequential, records of changed pixels to see the return to position. From the mask data it is a simple (and rapid) procedure to calculate the three-dimensional center of mass. It is therefore possible to discriminate among the leading sources of artifacts: spikes, transient and sustained shifts in position.

Based on this masking procedure, we calculate a variety of measures that are useful in artifact rejection. One output from the motion/noise detection software is a time series representing the percentage change in the number of voxels above and below threshold. Motion should result in an approximately equal number of transitions above and below threshold, while the presence of spike artifacts, which add energy to the image overall, should result in predominantly increased numbers of voxels above threshold. Figure 13 shows representative results in two different studies. The series on the left was contaminated by six raw data spikes. On the right is shown a series in which the subject moved during the acquisition. Note that the difference in the number of increasing and decreasing voxels is near zero (lower curve), as would be expected with motion, as the suprathreshold voxels are simply moved to a different location in the image volume. If we look, however, at

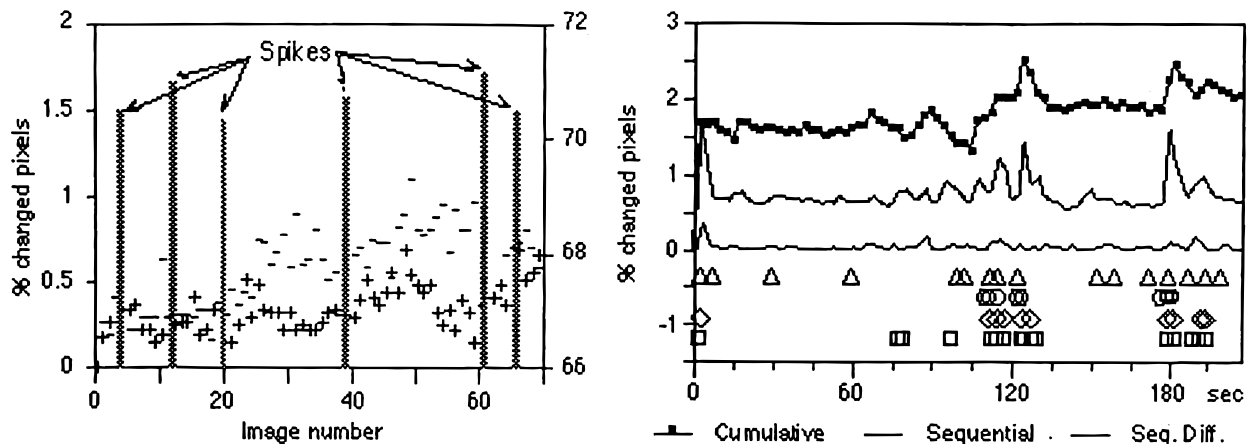


FIG. 13. Left: Automated spike (and motion detection) +, Percentage -, increase in number of voxels above threshold; percentage decrease as compared with the first image in the series. The gray bars represent (on a separate vertical axis) time points at which an individual image showed a large increase in intensity as compared with the prior image representing spikes in the MR raw data. Right: Motion detection results. The lines labeled Cumulative represent the fraction of changed voxels compared with the first image. The sequential curve is the total fraction of voxels that have changed from noise to image or vice versa compared with the image immediately prior, and the Seq. Diff. is the difference, in sequential images, in the number of voxels whose intensity has increased and decreased. Triangles, circles, diamonds, and squares are the results of visual identification of motion by four readers, blinded to the quantitative results.

the sum of the number of voxels increasing and decreasing, we can see the motion easily (middle curve). The top curve shows the cumulative fraction of voxels that have changed since the first image. In this study, the subject appears to have moved to a new location about 2 min into the study. To assess (somewhat crudely) the efficacy of the detection method, the image series were

evaluated visually by four readers blind to the computed results. While there was scatter (as might be expected) in the subjective impressions of motion, the concordance was good with the quantitative computed assays. It is clear that choosing an arbitrary threshold intensity (e.g., 0.8%) would effectively replicate the average observer results. Visual inspection, however, was


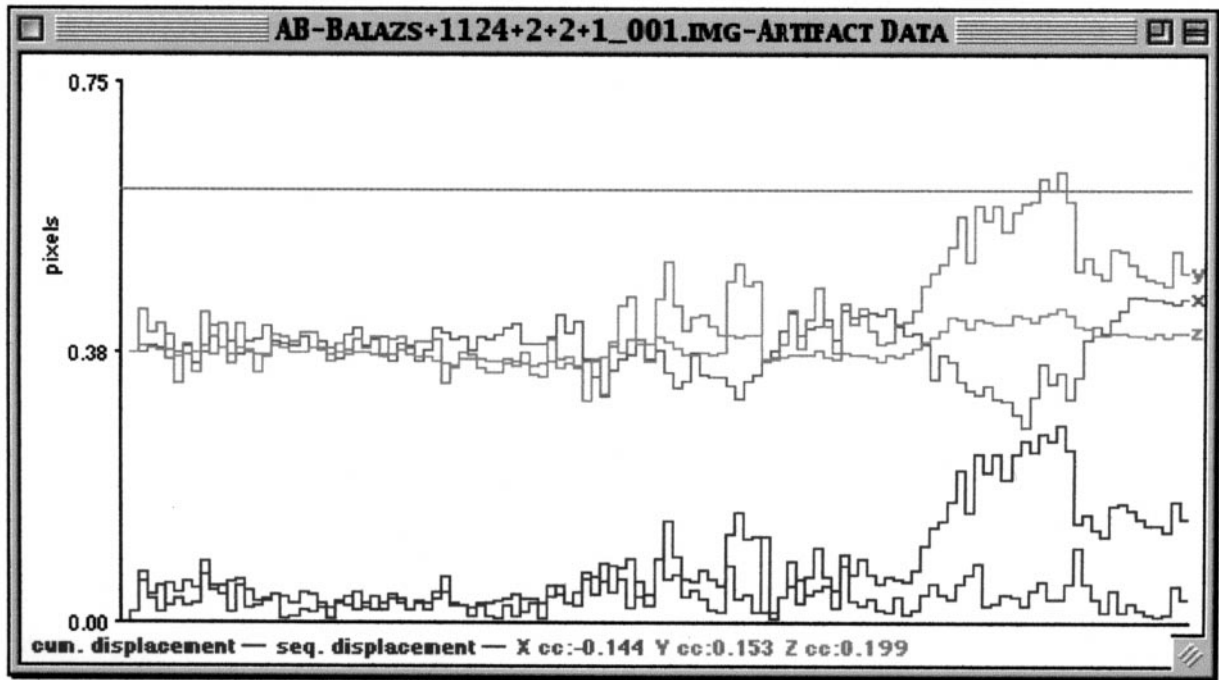


FIG. 14. Motion artifact data derived in real time. The center of mass in three dimensions as a function of time is plotted in pale blue, green, and dark yellow. The RMS displacement of the head is shown below in red (with respect to the first image) and blue (with respect to the prior image). The “bullseye display (right) shows the current head position (thick cross) compared with the position at the start of the scan. This display may be useful in providing feedback to the subject about head position.

not effective in detecting that the subject moved to a new position halfway through the examination. The center-of-mass data can provide a quantitative, though relatively crude measure of the head position, especially under assumption that the head moves as a rigid body. Representative data from the artifact detection tools, as represented by scanSTAT, appear in Fig. 14. This figure also shows a two-dimensional position indicator that is useful in tracking changes in head position during the scan. The architecture of the system allows this indicator to be presented directly to the subject, if desired, to offer feedback on the head position.

The corrective strategies for these different artifact sources also differ. In the event of an isolated spike event, usually instrument-related, the image must be omitted from further data analysis. If the subject moved, and returned to home position, the data can either be kept, if the statistical threshold needs can tolerate such motion, or disposed of. When the patient moves to a new position and stays, then we have additional alternatives: we can use an automated image registration scheme, such as Dr. Woods' AIR software (113), or the method of Pellizzari and co-workers (114, 115) to bring the images into register, or we can process the statistical data in separate blocks, possibly reregistering the functional maps at the end. Presently, our software performs the artifact detection with high reliability; scanSTAT allows the images to be discarded from data analysis (in real time) if they exceed user-defined thresholds for positional changes or other artifacts.

SOFTWARE DESIGN

The real-time fMRI package, scanSTAT, is currently written in ANSI C, with a few C++ components, and executes natively on a Macintosh computer. This implementation was selected because of the high performance/price ratio of such small computer systems, the familiar and easy-to-manage user interface on these computers, and the ready integration with other laboratory tools including stimulus devices and software. We have also made extensive use of a programming shell known as EasyApp written by Trudeau (116). Because the PowerPC processor used on the current generation of Macintosh computers is exceedingly fast, we have been able to perform all of the application programming in high-level languages, without, for example, resorting to the use of machine language for inner loops.

As attractive as this architecture is, we realize that it has certain limitations for the community at large. At present, we are rebuilding the software in Java, with

low-level functions written in ANSI C++. With this system design, we intend that it should be a relatively straightforward matter to port scanSTAT to a wide variety of platforms while retaining real-time performance. There will still be limitations at the scanner end, of course, as not all commercial scanners allow rapid access to the image or raw data early in the processing stream. Further, we intend to develop a "plug-in" feature that will allow users to readily extend the image processing with routines that can be added automatically to the processing pipeline. By the time this report is published, this Java-based version should be available at our web site: <http://www.brainmapping.org/scanSTAT/>.

SUMMARY

These studies indicate that real-time continuous image processing of functional MR data may be realized using conventional computing equipment. Further, they suggest that tools of linear systems analysis can be applied to fMRI data to yield improvements in sensitivity with practical importance. We have shown also, that based on the general approach of fitting fMRI data to careful estimates of the vascular response function, it is possible to derive fit parameters, particularly the slope of the signal change versus unit input, to form response magnitude maps that can be compared across subjects. The data presented on artifact rejection and detection indicate that we have a reliable method to identify the most common artifacts in functional MRI. Combined, these data suggest that we have the tools at hand to perform the needed computations for a practical and usable fMRI analysis system.

The use of fMRI is still growing explosively. It is highly attractive for its high sensitivity, spatial resolution, and noninvasiveness. Real-time processing, we believe, will have substantial value to the user community at large, not only in enabling a completely different class of experiments, but also in making the entire imaging process much more efficient and offering a much more interactive approach to image analysis than is available in the current generation of time-consuming offline processing tools.

ACKNOWLEDGMENTS

This work was supported in part by Grant R01-DA13054-01 from the National Institute on Drug Abuse, the Veteran's Administration VISN22 Mental Illness Research Education and Clinical Center. We acknowledge the generous support of the Brain Mapping Medical Research Foundation, The Ahmanson Foundation, the Pierson-

Lovelace Foundation, the Jennifer Jones Simon Foundation, the Tamkin Foundation, and the UCLA Council on Research.

REFERENCES

1. Bottomley, P., Foster, T., Argersinger, R., and Pfeiffer, L. (1984) *Med. Phys.* **11**, 425.
2. Frahm, J., Haase, A., and Matthaei, D. (1986) *Magn. Reson. Med.* **3**, 321–327.
3. Hennig, J., Nauwerth, A., and Friedburg, H. (1986) *Magn. Reson. Med.* **3**, 823–833.
4. Mulkern, R. V., Wong, S. T., Winalski, C., and Jolesz, F. A. (1990) *Magn. Reson. Imaging* **8**, 557–566.
5. Cohen, M. S. (1992) in *Radiology* (Taveras, J., and Ferruci, J., Eds.), Lippincott, New York.
6. Mansfield, P. (1977) *J. Phys. C* **10**, L55–L58.
7. Rzedzian, R. (1987) A method for instant whole-body MR imaging at 2.0 Tesla and system design considerations in its implementation [abstract], in *Society for Magnetic Resonance in Medicine Meeting*, 1987, p. 229.
8. Mansfield, P. (1984) *Br. Med. Bull.* **40**, 187–190.
9. Cohen, M. S., Weisskoff, R., and Rzedzian, R. (1989) Clinical methods for “single-shot” instant MR imaging of the heart [abstract], in *Radiological Society of North America Meeting*, Chicago, 1989, p. 359.
10. Cohen, M. S., and Weisskoff, R. M. (1991) *Magn. Reson. Imaging* **9**, 1–37.
11. Rosen, B., Belliveau, J., and Chien, D. (1989) *Magn. Res. Q.* **5**, 263–281.
12. Sorensen, A. G., Wray, S. H., Weisskoff, R. M., Boxerman, J. L., Davis, T. L., Caramia, F., et al. (1995) *Am. J. Neuroradiol.* **16**, 1753–1762.
13. Sorensen, A. G., Buonanno, F. S., Gonzalez, R. G., Schwamm, L. H., Lev, M. H., Huang-Hellinger, F. R., et al. (1996) *Radiology* **199**, 391–401.
14. Aronen, H., Goldberg, I., Pardo, F., Hochberg, F., Kennedy, D., Buchbinder, B., et al. (1992) Susceptibility contrast CBV imaging: Clinical experience in brain tumor patients [abstract], in *Society of Magnetic Resonance in Medicine Eleventh Annual Meeting*, Berlin, 1992; p. 739.
15. James, W. (1890) *The Principles of Psychology*, Henry Holt, New York.
16. Roy, C. S., and Sherrington, C. S. (1890) *J. Physiol. (London)* **11**, 85–108.
17. Sokoloff, L., Reivich, M., Kennedy, C., DesRosiers, M., Patlak, C., Pettigrew, K., et al. (1977) *J. Neurochem.* **28**, 897–916.
- 17a. Ingvar, D. H., and Philipson, L. (1977) *Ann. Neurol.* **2**, 230–237.
18. Phelps, M. E., Grubb, R. L., and Ter-Pogossian, M. M. (1973) *J. Appl. Physiol.* **35**, 274–280.
19. Phelps, M. E., Huang, S. C., Hoffman, E. J., Sokoloff, S. C., and Kuhl, D. E. (1979) *Ann. Neurol.* **6**, 371.
20. Frostig, R. D., Lieke, E. E., Tso, D. Y., and Grinvald, A. (1990) *Proc. Natl. Acad. Sci. USA* **87**, 6082–6086.
21. Belliveau, J. W., Kennedy, D. N., Jr., McKinstry, R. C., Buchbinder, B. R., Weisskoff, R. M., Cohen, M. S., et al. (1991) *Science* **254**, 716–719.
22. Ogawa, S., Lee, T. M., Kay, A. R., and Tank, D. W. (1990) *Proc. Natl. Acad. Sci. USA* **87**, 9868–9872.
23. Turner, R., Le Bihan, D., Moonen, C. T., Despres, D., and Frank, J. (1991) *Magn. Reson. Med.* **22**, 159–166.
24. Kwong, K. K., Belliveau, J. W., Chesler, D. A., Goldberg, I. E., Weisskoff, R. M., Poncelet, B. P., et al. (1992) *Proc. Natl. Acad. Sci. USA* **89**, 5675–5679.
25. Ogawa, S., Tank, D. W., Menon, R., Ellermann, J. M., Kim, S. G., Merkle, H., et al. (1992) *Proc. Natl. Acad. Sci. USA* **89**, 5951–5955.
26. Buxton, R. B., and Frank, L. R. (1997) *J. Cereb. Blood Flow Metab.* **17**, 64–72.
27. Hennig, J., Ernst, T., Speck, O., and Lauenberger, J. (1993) Functional spectroscopy: A new method for the observation of brain activation [abstract], in *Society for Magnetic Resonance in Medicine Twelfth Annual Meeting*, New York, 1993, p. 12.
28. Menon, R. S., Ogawa, S., Hu, X., Strupp, J. P., Anderson, P., and Ugurbil, K. (1995) *Magn. Reson. Med.* **33**, 453–459.
29. Hu, X., Le, T. H., and Ugurbil, K. (1997) *Magn. Reson. Med.* **37**, 877–884.
30. Cohen, M. S., and Bookheimer, S. Y. (1994) *Trends Neurosci.* **17**, 268–277.
31. Tootell, R. B., Reppas, J. B., Kwong, K. K., Malach, R., Born, R. T., Brady, T. J., et al. (1995) *J. Neurosci.* **15**, 3215–3230.
32. Kim, S. G., Ashe, J., Hendrich, K., Ellermann, J. M., Merkle, H., Ugurbil, K., et al. (1993) *Science* **261**, 615–617.
33. Cohen, M. S., Kosslyn, S. M., Breiter, H. C., DiGirolamo, G. J., Thompson, W. L., Bookheimer, S. Y., et al. (1996) *Brain* **119**, 89–100.
34. Hinke, R. M., Hu, X., Stillman, A. E., Kim, S. G., Merkle, H., Salmi, R., et al. (1993) *NeuroReport* **4**, 675–678.
35. Just, M. A., Carpenter, P. A., Keller, T. A., Eddy, W. F., and Thulborn, K. R. (1996) *Science* **274**, 114–116.
36. Penfield, W., and Boldrey, E. (1937) *Brain*; **60**: 389–443.
37. Gevins, A. S., and Schaffer, R. E. (1980) *CRC Crit. Rev. Bioeng.* **4**, 113–164.
38. Engel, J. J., Henry, T. R., Risinger, M. W., Mazziotta, J. C., Sutherling, W. W., Levesque, M. F., et al. (1990) *Neurology* **40**, 1670–1677.
39. Wada, J., and Rasmussen, T. (1960) *J. Neurosurg.* **17**, 266–282.
40. Bookheimer, S. Y., Cohen, M. S., Dapretto, M., Fried, I., Shewmon, A., Black, K. et al. (1995) Functional MRI in Surgical Planning [abstract] in *Society for Neuroscience Meeting*, San Diego, CA, 1995, p. 273.
41. Binder, J. R., Rao, S. M., Hammeke, T. A., Frost, J. A., Bandettini, P. A., Jesmanowicz, A., et al. (1995) *Arch. Neurol.* **52**, 593–601.
42. Binder, J. R., Swanson, S. J., Hammeke, T. A., Morris, G. L., Mueller, W. M., Fischer, M., et al. (1996) *Neurology* **46**, 978–984.
43. Benson, R., Kwong, K., Buchbinder, B., Jiang, H., Belliveau, J., Cohen, M. S., et al. (1994) Noninvasive evaluation of language dominance using functional MRI [abstract], in *Society for Magnetic Resonance Second Annual Meeting*, San Francisco, 1994; p. 684.
44. Bandettini, P. A., Wong, E. C., Hinks, R. S., Tikofsky, R. S., and Hyde, J. S. (1992) *Magn. Reson. Med.* **25**, 390–397.
45. Chapman, P. H., Buchbinder, B. R., Cosgrove, G. R., and Jiang, H. J. (1995) *Pediatr. Neurosurg.* **23**, 122–126.
46. Cosgrove, G. R., Buchbinder, B. R., and Jiang, H. (1996) *Neurosurg. Clin. North Am.* **7**, 313–322.
47. Andreasen, N. C. (1988) *Science* **239**, 1381–1388.
48. Andreasen, N. C., Flashman, L., Flaum, M., Arndt, S., Swayze, Vn., O’Leary, D. S., et al. (1994) *JAMA* **272**, 1763–1769.
49. Shenton, M. E., Kikinis, R., Jolesz, F. A., Pollak, S. D., LeMay, M., Wible, C. G., et al. (1992) *N. Engl. J. Med.* **327**, 604–612.
50. Howard, R., Mellers, J., Petty, R., Bonner, D., Menon, R., Almeida, O., et al. (1995) *Psychol. Med.* **25**, 495–503.

51. Menon, R. R., Barta, P. E., Aylward, E. H., Richards, S. S., Vaughn, D. D., Tien, A. Y., *et al.* (1995) *Schizophrenia Res.* **16**, 127–135.
52. Tien, A. Y., Eaton, W. W., Schlaepfer, T. E., McGilchrist, I. K., Menon, R., Powers, R., *et al.* (1996) *Schizophrenia Res.* **19**, 93–101.
53. Cohen, M. S., and Green, M. F. (1995) Where the voices come from: Imaging of schizophrenic auditory hallucinations [abstract], in Society for Neuroscience, San Diego, CA, 1995, p. 259.
54. Early, T. S., Posner, M. I., Reiman, E. M., and Raichle, M. E. (1989) *Psychiatr. Dev.* **7**, 109–121.
55. Cleghorn, J. M., Garnett, E. S., Nahmias, C., Brown, G. M., Kaplan, R. D., Szechtman, H., *et al.* (1990) *Br. J. Psychiatry* **157**, 562–570.
56. Cleghorn, J. M., Zipursky, R. B., and List, S. J. (1991) *J. Psychiatry Neurosci.* **16**, 53–74.
57. Buchsbaum, M. S. (1992) *Neuropsychopharmacology* **7**, 67–68; discussion 73–75.
58. Buchsbaum, M. S., Potkin, S. G., Marshall, J. F., Lottenberg, S., Teng, C., Heh, C. W., *et al.* (1992) *Neuropsychopharmacology* **6**, 155–163.
59. Buchsbaum, M. S., Potkin, S. G., Siegel, B. J., Lohr, J., Katz, M., Gottschalk, L. A., *et al.* (1992) *Arch. Gen. Psychiatry* **49**, 966–974.
60. Friston, K. J. (1992) *J. Neural Transmission Suppl.* **37**, 79–93.
61. Friston, K. J., Liddle, P. F., Frith, C. D., Hirsch, S. R., and Frackowiak, R. S. (1992) *Brain.*
62. Frith, C. D., Friston, K. J., Liddle, P. F., and Frackowiak, R. S. (1992) *J. R. Soc. Med.* **85**, 222–224.
63. Liddle, P. F., Friston, K. J., Frith, C. D., and Frackowiak, R. S. (1992) *J. R. Soc. Med.* **85**, 224–227.
64. Gur, R. E., and Pearlson, G. D. (1993) *Schizophrenia Bull.* **19**, 337–353.
65. Silbersweig, D. A., Stern, E., Frith, C., Cahill, C., Holmes, A., Grootoink, S., *et al.* (1995) *Nature* **378**, 176–179.
66. Silbersweig, D. A., Stern, E., Frith, C. D., Seward, J., Holmes, A., Schnorr, L., *et al.* (1995) Mapping the neural correlates of auditory hallucinations in schizophrenia: involuntary perceptions in the absence of external stimuli [abstract], in First International Conference on Functional Mapping of the Human Brain, Paris, France, 1995, p. 422.
67. Detre, J. A., Leigh, J. S., Williams, D. S., and Koretsky, A. P. (1992) *Magn. Reson. Med.* **23**, 37–45.
68. Detre, J. A., Zhang, W., Roberts, D. A., Silva, A. C., Williams, D. S., Grandis, D. J., *et al.* (1994) *NMR Biomed.* **7**, 75–82.
69. Thulborn, K., Davis, D., Erb, P., Strojwas, M., and Sweeney, J. (1996) *NeuroImage* **4**(3, Pt. 3), S101–S107.
70. Thulborn, K., Uttecht, S., Betancourt, C., Talagala, S., Boada, E., and Shen, G. (1997) *Int. J. Imaging Syst. Technol.*, **8**, 572–581.
71. Woods, R. P., Iacoboni, M., and Mazziotta, J. C. (1994) *N. Engl. J. Med.* **331**, 1689–1692.
72. Dettmers, C., Connelly, A., Stephan, K. M., Turner, R., Friston, K. J., Frackowiak, R. S., *et al.* (1996) *NeuroImage* **4**(3, Pt. 1), 201–209.
73. Hajnal, J. V., Myers, R., Oatridge, A., Schwieso, J. E., and Young, I. R. (1994) *Magn. Reson. Med.* **31**, 283.
74. Cohen, M. S. (1996) *J. Magn. Reson. Imag.* **6**, 273–274.
75. Bookheimer, S. Y., Dapretto, M. A., Cohen, M. S., and Wang, J. X. (1996) Functional MRI of the hippocampus during short-term memory tasks: Parametric response to task difficulty and stimulus novelty [abstract], in Second Annual Conference on Functional Mapping of the Human Brain, Boston, MA, 1996.
76. Dapretto, M. A., Bookheimer, S. Y., Cohen, M. S., and Wang, J. X. (1996) fMRI of language in dyslexic and normally developing children [abstract], in Second Annual Conference on Functional Mapping of the Human Brain, Boston, MA, 1996.
77. Cohen, M. S., Breiter, H., DiGirolamo, G., Thompson, W., Belliveau, J., Rosen, B., *et al.* (1995) Mental rotation studied by functional magnetic resonance imaging (fMRI) [abstract], in Brain Map '95, Paris, 1995.
78. Hajnal, J. V., Bydder, G. M., and Young, I. R. (1995) *NMR Biomed.* **8**, 97–100.
79. Forman, S. D., Cohen, J. D., Fitzgerald, M., Eddy, W. F., Mintun, M. A., and Noll, D. C. (1995) *Magn. Reson. Med.* **33**, 636–647.
80. Meyer, C., and Macovski, A. (1987) Square spiral fast imaging: Interleaving and off-resonance effects [abstract], in Society for Magnetic Resonance in Medicine Meeting, 1987, p. 230.
81. Macovski, A., and Meyer, C. (1986) A novel fast scanning system [abstract], in Society for Magnetic Resonance in Medicine Meeting, 1986, p. 156.
82. Mueller, O. M., McFarland, T. G., Park, J. N., Wirth, W. F., Vavrek, R. M., and Roemer, P. B. (1993) A new 'quasi-linear', high efficiency, non-resonant, high-power gradient system [abstract], in Society of Magnetic Resonance in Medicine Twelfth Annual Meeting, New York, 1993, p. 312.
83. Weisskoff, R. M., Baker, J. R., Belliveau, J. W., Davis, T. L., Kwong, K. K., Cohen, M. S., *et al.* (1993) Power spectrum analysis of functionally-weighted MR data: What's in the noise? [abstract], in Society of Magnetic Resonance in Medicine Meeting, New York, 1993, p. 7.
84. Friston, K. J., Dolan, R. J., and Frackowiak, R. S. J. (1991) Statistical Parametric Mapping. MRC Cyclotron Unit, Hammersmith Hospital, London.
85. Bandettini, P. A., Jesmanowicz, A., Wong, E. C., and Hyde, J. S. (1993) *Magn. Reson. Med.* **30**, 161–173.
86. Friston, K. J., Frith, C. D., Liddle, P. F., and Frackowiak, R. S. (1991) *J. Cereb. Blood Flow Metab.* **11**, 690–699.
87. Friston, K. J., Worsley, K. J., Frackowiak, R. S. J., Mazziotta, J. C., and Evans, A. C. (1994) *Hum. Brain Mapping* **1**, 210–220.
88. Friston, K. J. (1995) *J. Cereb. Blood Flow Metab.* **15**, 361–370.
89. Bookheimer, S., and Zeffiro, T. (1994) Localization of language operations with H2150 PET: individual differences in functional anatomy and their relation to image smoothing [abstract], in Society for Neuroscience 24th Annual Meeting, 1994, p. 531.11.
90. Buckner, R. L., Bandettini, P. A., O'Craven, K. M., Savoy, R. L., Petersen, S. E., Raichle, M. E., *et al.* (1996) *Proc. Nat. Acad. Sci. USA* **93**, 14878–14883.
91. Cohen, M. S. (1997) *NeuroImage* **6**, 93–103.
92. Rzedzian, R. (1988) Real time MRI at 2.0 Tesla [abstract], in Society for Magnetic Resonance in Medicine Meeting, 1988, p. 247.
93. Shepard, R., and Meltzer, J. (1971) *Science* **171**, 701–703.
94. Tagaris, G., Kim, S.-G., Menon, R., Strupp, J., Andersen, K., Ugurbil, K., *et al.* (1994) High field (4 Tesla) functional MRI of mental rotation [abstract], in Society for Neuroscience Meeting, Miami Beach, FL, 1994, p. 152.10.
95. Tagaris, G. A., Kim, S. G., Strupp, J. P., Andersen, P., Ugurbil, K., and Georgopoulos, A. P. (1996) *NeuroReport* **7**, 773–776.
96. Jackson, G., Connelly, A., Cross, J., Gordon, I., and Gadian, D. (1994) *Neurology* **44**, 850–856.
97. Bookheimer, S. Y. (1996) *Neuroimage* **4**(3, Pt. 3), S139–S146.
98. Warach, S., Ives, J. R., Schlaug, G., Patel, M. R., Darby, D. G., Thangaraj, V., *et al.* (1996) *Neurology* **47**, 89–93.
99. Cohen, M. S., Kelley, D. A., Rohan, M. L., and Roemer, P. A.

- (1996) An MR instrument optimized for intracranial neuroimaging [abstract] in *Human Brain Mapping 96*, Boston, MA, 1996, P1A1-007.
100. DeArmond, S. J., Fusco, M. M., and Dewey, M. M. (1976) *Structure of the Human Brain: A Photographic Atlas*, Oxford Univ. Press, New York.
101. Howseman, A. M., Stehling, M. K., Chapman, B., Coxon, R., Turner, R., Ordidge, R. J., *et al.* (1988) *Br. J. Radiol.* **61**, 822–828.
102. Rzedzian, R. (1987) High speed, high resolution, spin echo imaging by Mosaic scan and MESH [abstract], in *Society for Magnetic Resonance in Medicine Meeting*, 1987, p. 51.
103. Weisskoff, R., Cohen, M. S., and Rzedzian, R. (1989) Fat suppression techniques: a comparison of results in instant imaging [abstract], in *Society for Magnetic Resonance in Medicine Meeting*, 1989, p. 836.
104. Farzaneh, F., Riederer, S. J., and Pelc, N. J. (1990) *Magn. Reson. Med.*, **14**, 123–139.
105. Mansfield, P., Coxon, R., and Glover, P. (1994) *Comput. Assist. Tomogr.* **18**, 339–343.
106. Reese, T., Davis, T., and Weisskoff, R. (1995) *J. Magn. Reson. Imaging* **5**, 739–745.
107. Hong, X., Cohen, M., and Roemer, P. (1997) Functional EPI with Real Time Imaging Processing [abstract], in *Fifth Annual Meeting of the International Society for Magnetic Resonance in Medicine Vancouver, BC*, 1997, p. 321.
108. Press, W., Vetterling, W., Teukolsky, S., and Flannery, B. (1992) *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge Univ. Press, Cambridge.
109. Liu, C. L., and Liu, J. W. S. (1974) *Linear Systems Analysis*, McGraw-Hill, New York.
110. Josephs, O., Vandenberghe, R., Turner, R., and Friston, K. (1997) Event-Related FMRI: Word Reading [abstract], in *Society for Neuroscience, 17th Annual Meeting*, New Orleans, LA, 1997, p. 415.2.
111. Fox, P. T., and Raichle, M. E. (1984) *J. Neurophysiol.* **51**, 1109–1120.
112. Cohen, M. S., and DuBois, R. M. (1999) *J. Magn. Reson. Imaging* **10**, 33–40.
113. Woods, R. P., Mazziotta, J. C., and Cherry, S. R. (1993) *J. Comput. Assist. Tomogr.* **17**, 536–546.
114. Chen, G. T., and Pelizzari, C. A. (1989) *Comput. Med. Imaging Graphics* **13**, 235–240.
115. Pelizzari, C. A., Chen, G. T., Spelbring, D. R., Weichselbaum, R. R., and Chen, C. T. (1989) *J. Comput. Assist. Tomogr.* **13**, 20–26.
116. Trudeau, J. (1995) *Programming Starter Kit for Macintosh*, Hayden Books, Indianapolis, IN.