STATE OF THE ART IN MAGNETIC RESONANCE IMAGING

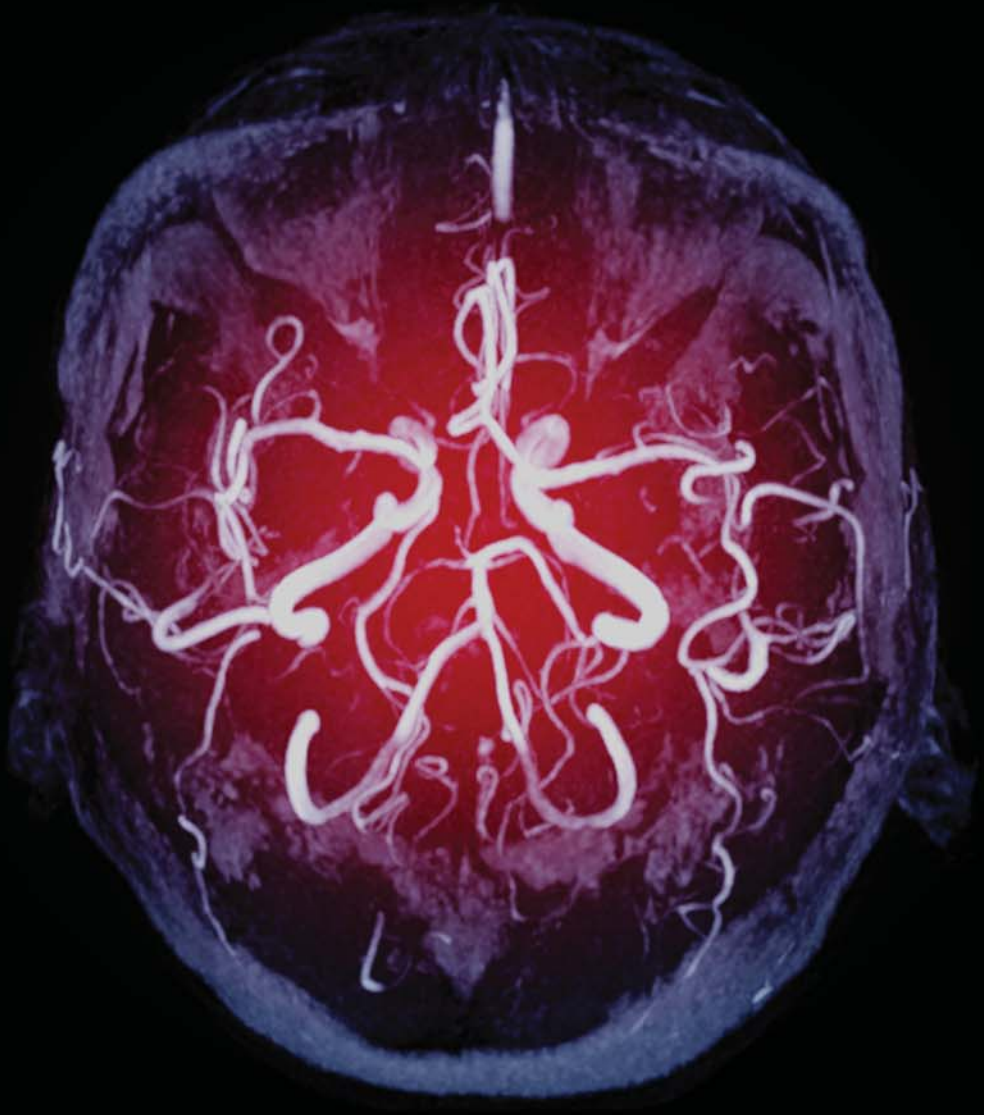

ISTOCK.COM/MR.SUPHACHAI PRASERDUMRONGCHAI

David Jordan

As a clinical technology, MRI offers unsurpassed flexibility to look inside the human body.

Dave Jordan is chief medical physicist at University Hospitals and associate professor of radiology at Case Western Reserve University, both in Cleveland, Ohio.

In 1977, inspired by the observation that cancerous and healthy tissues produced different nuclear magnetic resonance signals, Raymond Damadian, Michael Goldsmith, and Lawrence Minkoff performed the first MRI scan of a live human body.[1] In the early days of clinical MRI, scans took hours and provided low spatial resolution, but they have become essential for distinguishing between healthy and diseased tissues. By its 40th anniversary, MRI was a must-have tool in hospitals and clinics of all sizes. And it has found applications in image-guided interventions and surgeries, radiation therapy, and focused ultrasound. Advances in technology, meanwhile, have pushed the envelope of scanner performance with improvements to speed and spatial resolution.

At the frontiers of MRI development, work is focused on fast, quantitative imaging. Clinical needs increasingly demand functional information—on heart-muscle contractions,[2] brain activity,[3] chemical concentrations in tumors,[4] and blood flow in and out of tissue[5]—in addition to anatomical structures. New approaches must also maintain a patient's comfort and safety; MRI is well-known for sparing patients any exposure to ionizing radiation, yet it is not without hazards.[6]

How MRI works

When biological tissue is placed in a magnetic field, nuclei with magnetic moments become magnetized. RF pulses are then applied that match the resonance, or Larmor, frequency of the nuclei, causing them to tip out of alignment with the external magnetic field and precess about it. The precessing nuclei, in turn, induce oscillating magnetic fields at the Larmor frequency; those oscillations are detected via Faraday induction of an electromotive force in a nearby coil of wire.[7]

In practice, many nuclei must precess in phase with each other to produce a detectable signal. The loss of phase coherence among precessing nuclei over time is called T2 relaxation. And the orientations of the nuclei eventually return to their equilibrium orientation in the external magnetic field—a process called T1 relaxation. Early NMR experiments revealed that various tissues have distinct T1 and T2 relaxation times. For certain diseases, including cancer, changes in either time can distinguish between diseased and healthy tissue. That feature is useful in the case of lesions whose absorption of x rays is similar to that of surrounding healthy tissue, which makes them difficult to detect using radiography or x-ray computed tomography.

In NMR measurements, the timing of the applied RF pulse and of the RF readout signal from the tissue can be chosen so that the strongest signal is produced by tissue with the shortest T1 relaxation time. A measurement whose timing is chosen that way is called a T1-weighted measurement. Alternatively, the sequence timing can be chosen so that the strongest signal comes from tissue with the longest T2 relaxation time, a T2-weighted result.

In tissue, the hydrogen nucleus is the most abundant magnetizable nucleus. Its gyromagnetic ratio is 42.56 MHz/T, which results in operating at Larmor frequencies of roughly 64 MHz and 128 MHz for 1.5 T and 3 T MRI scanners, respectively. Magnets with a range between 0.2 T and 7 T are used for clinical scanning, and human scanning up to 10.5 T is currently available in research settings.

Spatial encoding

Using the NMR signals from tissue for clinical diagnosis requires that they be localized in three dimensions to form images. Three sets of electromagnetic gradient coils in the MRI scanner accomplish that task. Each produces a linearly varying magnetic field along one of three orthogonal axes. And each gradient can be switched on and off to produce different strengths depending on the current applied; the gradient fields are superimposed on the main magnetic field—usually 1.5 T or 3 T—to create a spatially dependent variation in Larmor frequency. If a gradient is applied for some time and then turned off, all signals have the

same frequency, but their relative phase shifts, accumulated while the gradient was on, vary according to position along the gradient axis.

Frequency encoding—collecting the NMR signal while a gradient field is applied—produces a spatially dependent variation in resonant frequency along one spatial dimension. Phase encoding is applying a gradient field along one axis for some time and then removing it before collecting the NMR signal. In MRI, frequency encoding and phase encoding are performed along perpendicular axes to localize the signal in two dimensions.

Gradients also allow tomographic slices of tissue to be imaged selectively. By applying an orthogonal gradient through the slice plane, the resonant frequencies of nuclei can be shifted so that only those in the slice of interest are matched to the frequency of the incident RF magnetic pulse used to excite the tissue. Much clinical MRI uses that method. When the in-plane frequency and phase encoding and data readout are subsequently performed, only those nuclei at the desired location and within the selected slice thickness produce signals.

In early MRI, various techniques were used to arrange signals spatially to produce images. But the spatial-encoding approaches developed by Paul Lauterbur and Peter Mansfield, for which they shared the Nobel Prize in Physiology or Medicine in 2003 (see PHYSICS TODAY, December 2003, page 24), have become the standard for clinical MRI.[8–10] Magnetic resonance scanning is usually an iterative procedure. It involves slice-selective excitation followed by phase encoding, frequency encoding, and signal detection repeated a few hundred times to produce an image of one slice of tissue. Each time, a clinician uses a different phase encoding gradient strength, collects the RF waveform emitted by the tissue during each cycle, and then stores it in one row of an array called *k* space—so named for the wavenumber or spatial frequency.

The gradient spatial encoding has the effect of decomposing signal variations from the tissue into a Fourier series of spatial frequencies. The collected RF waveforms represent the spatial frequencies of the signal intensities, and a 2D inverse Fourier transform of the *k*-space data produces a two-dimensional image of the tissue. Using fast-Fourier-transform techniques, MRI image reconstruction can be performed quickly and efficiently, and dozens or even hundreds of images appear onscreen virtually instantaneously upon completion of a scan.

In an MRI scan, the number and relative timings of activations of the RF transmitter, receiver, and gradient coils make up the pulse sequence. That sequence is the recipe that determines the amount of signal that's collected and the type of NMR information, such as contrast weighting, encoded in the image.

Qualitative evaluation of weighted images by a trained radiologist is an effective diagnostic tool (see figure 1 for an example), but interest is growing in quantitative measurements of tissue parameters to improve diagnosis. Water diffusivity and chemical concentrations can be measured quantitatively. The signal intensities in T1- and T2-weighted images are usually relative, however, and do not inherently represent absolute quantitative measurements of T1 or T2 relaxation times in the imaged tissue. Such measurements are possible, and efficient techniques to do so are an active area of research. Figure 2 shows an example of differences in quantitative T1 measurements for normal and diseased lung tissue in cystic fibrosis. Lung tissue is difficult to image using qualitative MRI because the signal level is inherently low, but quantitative T1 maps provide a clear visualization of disease.

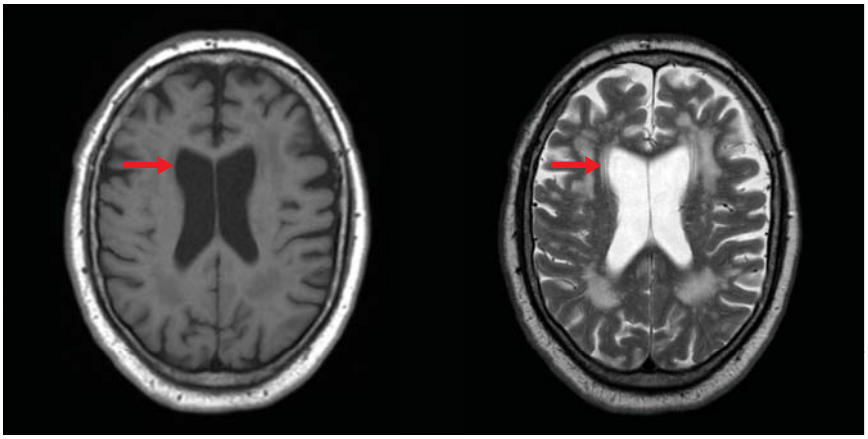

FIGURE 1. MRI OF A BRAIN with T1 (left) and T2 (right) weighting. T1 refers to the time required for precessing nuclei in tissue to relax to their equilibrium orientation in a magnetic field, and T2 refers to the time it takes the nuclei to lose their phase coherence. The timing of RF pulses can be chosen so that the MRI signal comes from tissue with the shortest T1 times or the longest T2 times. Images taken with those different timings demonstrate how MRI can be tuned to produce different signal intensities in the same tissue. Fluid-filled ventricles (red arrow) are darker than the brain tissue on the T1-weighted image and brighter on the T2-weighted image. (Images by David Jordan, University Hospitals Cleveland Medical Center.)

Contrast enhancement

In x-ray imaging, radiologists can use contrast agents such as barium and iodine to fill blood vessels, bowel loops, and other structures that are otherwise difficult to see. Contrast dye flowing in the bloodstream strongly attenuates radiation and makes it much easier to image the vessel itself and determine where in the tissue blood is being delivered.

In MRI, the workhorse contrast agent is gadolinium, chelated to various molecules to prevent toxic interactions with the human body; the free Gd^{3+} ion competes with Ca^{2+} metabolism.[11] The paramagnetic Gd nucleus induces a strong local magnetic field and substantially shortens the T1 relaxation time of nearby hydrogen nuclei. As a result, Gd dye in a fluid or tissue increases the local brightness on T1-weighted MRI, as illustrated in figure 3.

Several commercial, Gd-based contrast agents are commonly used in MRI exams of all body parts and diseases. The two main safety concerns are the stability of the chelate and retention of the agent in tissue. Most chelates eventually break down and release free Gd^{3+} ions over a long time period; that process is not a safety concern if it happens on a time scale much longer than the time it takes dye to be cleared from the body. However, toxicity does become a concern if dye is retained in tissue long enough for the chelate to break down.

Recent research indicates that some tissues may retain Gd-based contrast agents much longer than previously suspected.

Fast imaging

In MRI, it takes time to acquire images, typically two to five minutes for each. Patients must lie still, often in unnatural poses, while the scan data are collected. Some scans are influenced by involuntary (cardiac or peristaltic) motion; to capture clear images of those moving structures, fast scanning is essential. Patients can be coached to control their breathing and swallowing, but there are limits to how long they can reliably hold their breath or lie still.

Many image-quality parameters can be sacrificed to reduce scan time, but doing so may limit the physician's ability to make a confident diagnosis. Reducing scan time can also reduce spatial resolution and produce grainy, speckled images. So-called parallel-imaging techniques increase scan speed by using a phased array of RF receiver coils positioned at different locations around the anatomy of interest. To reduce time, fewer data points are collected than would normally be needed to scan the field of view at the desired spatial resolution.[12]

In one family of parallel MRI techniques, a radiologist maintains the spatial resolution but undersamples the data to save time, a process that produces a smaller field of view. He may combine a collection of small fields of view to produce a larger one, in which case each coil in the array needs calibration. Those sensitivity encoding (SENSE) methods provide a new feature to spatial encoding—the relative signal that tissue produces in multiple receiver coils—and supplement the frequency- and phase-encoding information. In another family of parallel techniques, some *k*-space data are omitted during the scan to save time and are synthesized prior to reconstruction.

Parallel imaging speeds up the process of collecting all the samples needed to achieve good spatial resolution. Compressed sensing can further accelerate scanning. In that technique, a portion of the sampling is omitted entirely to save scan time.[13] The data are processed with discrete cosine and wavelet transforms, which are used extensively for image compression in digital photography and video. The transforms generate MR images with the required spatial resolution but from a fraction of the measured data that would normally be required. In essence, the smaller data set contains a compressed version of the MR image, which is decompressed using those transforms during image reconstruction.

By combining parallel imaging with compressed sensing, high-quality scans of the liver or lungs are becoming practical without requiring the patient to hold their breath. Images of the beating heart can now be scanned faster and with higher quality than ever.

MRI mammography

Recently, x-ray mammography for breast-cancer screening has become controversial as government and medical professional bodies have diverged on recommendations for screening various populations of women. MRI has long been used to refine a breast-cancer diagnosis, and in high-risk populations—women with dense breasts or a family history of breast cancer—it has been a successful screening strategy.

MRI can typically detect cancers at an earlier, more treatable stage and those that start small but are aggressive, fast growing, and more likely to metastasize. X-ray mammography is best for imaging tumors that are denser than surrounding tissue, contain calcifications, or distort the surrounding tissue. Those gross structural changes are less likely to be evident in early cancers, and some evidence suggests that aggressive cancers may not even develop the changes in tissue before spreading. With injected Gd contrast dye, MRI can highlight microscopic tissue changes in early cancer, such as changes in T2, water diffusivity, or blood flow to certain locations.

The high sensitivity of breast MRI would be a boon to all patients who need screening, but widespread tests have never been feasible because of cost and time constraints. Clinics that perform breast MRI screening typically use protocols borrowed from their diagnostic problem-solving work. That means spending 20–40 minutes scanning a patient to get images with several different contrast weights and images after contrast dye has been injected into the tissue to best determine the location, size, type, and prognosis of the lesion.

Those scans produce a large and complex data set. From each image series, the radiologist identifies a given lesion as malignant or benign and estimates its aggressiveness. In an hour, breast radiologists can typically interpret dozens of mammograms but only a handful of MRI exams.

Researchers have been working to overcome those limitations and bring breast MRI screening to a wider population. The breakthrough in feasibility for large-scale screening came from redesigning the imaging protocol to focus on detection sensitivity and speed while eschewing the additional scans usually needed to characterize a tumor.

In fast-screening breast MRI, a single scan is performed on both breasts simultaneously; Gd dye is injected and the same scan is then repeated. The first image is subtracted from the second, which ideally removes all the tissue from view except for the locations containing dye

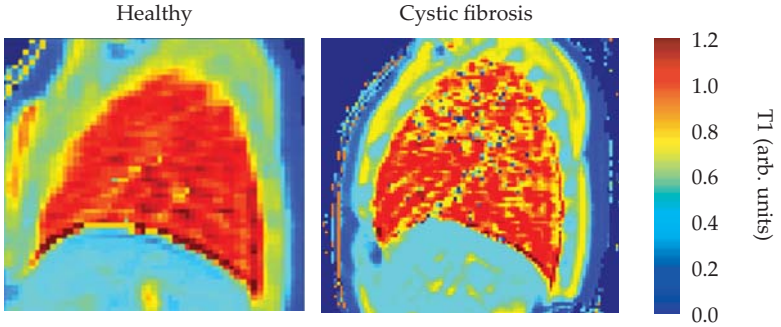

FIGURE 2. QUANTITATIVE T1 MRI MAPPING of the lungs of healthy volunteers (left) and cystic fibrosis (CF) patients (right). Low T1 values (blue, green, and yellow regions) in the lungs of CF patients indicate reduced blood flow to scarred lung tissue. Researchers hope to use those measurements of T1 to monitor the progression of disease and the effectiveness of treatment. (Courtesy of Chris Flask, Case Western Reserve University.)

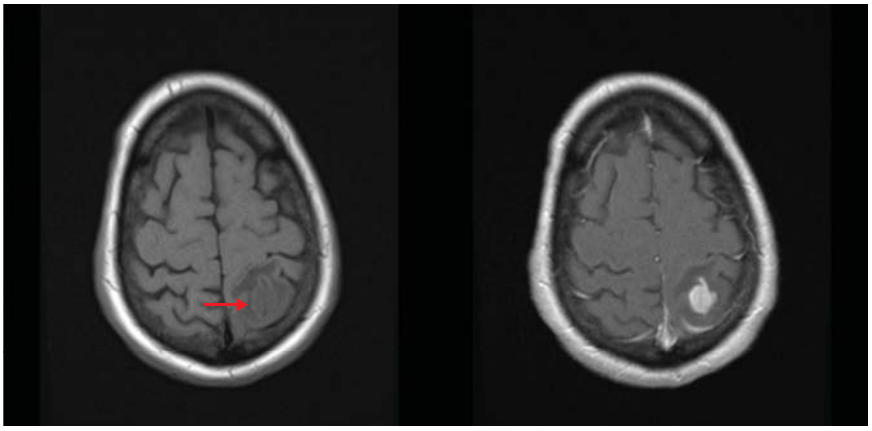

FIGURE 3. T1-WEIGHTED IMAGES OF A BRAIN without (left) and with (right) gadolinium contrast dye. The tumor (red arrow) is only subtly visible without the contrast; new blood vessel formation in the tumor results in high uptake of the dye, which makes the tumor much more visible in the contrast-enhanced image. (Images by David Jordan, University Hospitals Cleveland Medical Center.)

in the second scan. Individual slices of imaged tissue are fused together into a volume using the so-called maximum intensity projection technique. The fusion allows the radiologist to view all the dye-flow information in a pseudo-3D slab without having to scroll through images of individual slices (see figure 4 and the image on the title page of this article). She can thus easily detect tumors that recruit new blood flow in a matter of seconds following a three-minute MRI scan.[14]

Fingerprinting

Magnetic resonance fingerprinting (MRF) is a recent advance that takes a fundamentally different approach to quickly acquiring and processing signal data.[15] Conventionally, MRI scans are often used only to detect and localize possible disease, whereas the ultimate diagnostic judgment follows a biopsy and histopathological analysis. With more robust tools to characterize tissue, radiologists could make definitive diagnoses directly from MRI scans, and that could reduce the need for invasive biopsy. The practice would save time, cost, pain, and potential complications.

The MRF technique is conceptually similar to fingerprinting techniques used by law enforcement agencies to identify people. The pattern of ridges on someone's fingertips does not contain particularly interesting or useful details about the person. But it is unique, and if it can be matched to the person in a database, the match provides access to a much richer set of identifying details.

Here's how it works: A pulse sequence is synthesized with a pseudo-random variation in such parameters as the repetition time between pulses and the RF power applied to tissue. (In conventional MRI, it is critical to keep those parameters constant while generating and reading out spatially encoded signals.) A computer models the signal that a theoretical tissue would produce in response to the pulse sequence, given the T1 and T2 relaxation times of the tissue. The responses are then calculated using the NMR Bloch equations, which describe the magnetization of the tissue over time.[16]

A library of responses is created from the signal models for a wide range and numerous combinations of T1 and T2 values. The responses may not represent real physical tissues, but each mathematically represents the physical response of a unique combination of T1 and T2 values. The MRF scan uses the synthesized MRF pulse sequence to scan a patient (with spatial encoding to produce images), and the real MRF signals are recorded and compared with the library responses. For each pixel, the best match is determined between the recorded MRF signal and a response in the library, and the library match's T1 and T2 values are assigned to the pixel.

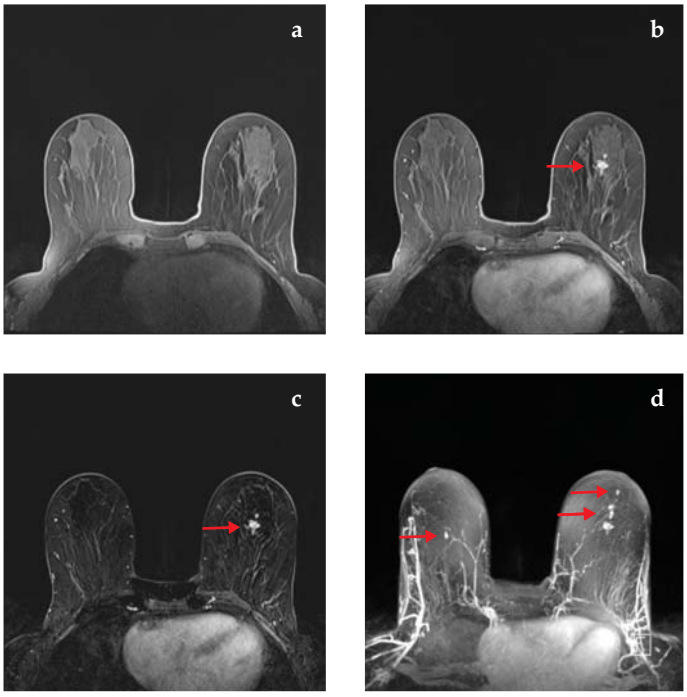

FIGURE 4. FAST-SCREENING BREAST MRI. (a) In this T1-weighted image without contrast dye, no abnormalities are apparent. **(b)** With contrast dye injected into the tissue, this T1-weighted image shows a large lesion (red arrow) that's easily seen. **(c)** This subtracted image, panel b minus panel a, improves the tumor's visibility and the brightness of dense glandular tissue. **(d)** A so-called maximum intensity projection of *all* subtracted slices (such as panel c) reveals other small lesions (red arrows), including one in the contralateral breast. A radiologist can much more rapidly detect all suspicious lesions by viewing a single image of this type than by individually reviewing dozens of image slices. (Images by David Jordan, University Hospitals Cleveland Medical Center.)

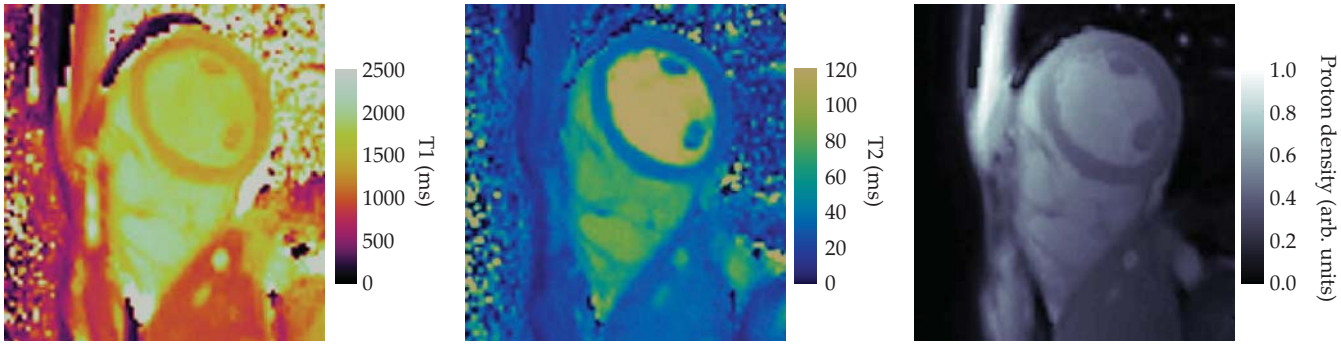

FIGURE 5. MAGNETIC RESONANCE FINGERPRINTING. In this image of a human heart, the technique provides quantitative maps of T1, T2, and proton density. All three images come from a single scan, whereas conventional MRI would require at least three separate scans to collect the same information. Clinical MRF research is underway to measure the effect of diseases and injuries on the three parameters. (Courtesy of Nicole Seiberlich, University of Michigan.)

That approach allows clinicians to determine the absolute T1 and T2 values from a single scan. Matching the tissue response to a library profile is relatively insensitive to noise; therefore, very fast MRF acquisitions can reliably provide robust quantitative results without the need for longer scans to achieve high signal-to-noise ratios.

Once the T1 and T2 values for each voxel have been extracted from the MRF library, the values can be displayed as quantitative maps. (Figure 5 shows an example.) In addition, the response of each voxel to traditional MRI pulse sequences can be calculated, so the MRF data can be used to synthesize traditional T1- or T2-weighted images without having to spend additional scan time to acquire them.

Using quantitative images of tissues, clinicians can potentially supplement their qualitative evaluations of the appearance of anatomy and pathology on conventional weighted images. With additional research, it is likely that they will use MRF to definitively identify specific disease signatures directly. That process would simplify the task of identifying a disease or judging its severity on the basis of its qualitative appearance and quantitative measurements.

Radiation therapy

Image-guided radiation therapy relies increasingly on MRI in treating cancer and other disorders. Two key roles are treatment planning and online image guidance during treatment.[17] Treatment planning uses images scanned before treatment to determine the placement of radiation beams to deliver a targeted dose and destroy a lesion. In online guidance, modern treatment machines use images taken during treatment to modify the targeting and account for the patient's motion and changes in tumor shape and position.

To deliver a large, precise radiation dose to a tumor while sparing healthy tissue, radiation therapy requires detailed information about the tissue environment. Computed tomography (CT) is useful in that regard since its image is a map of photon attenuation. A clinician measures a scanner's response by scanning calibration phantoms—reference objects with known size and electron density—and uses the patient's CT image to model where a treatment will deposit energy. The treatment can then be simulated and optimized to destroy the tumor and spare damage to the surrounding tissue. (See, for example, the article by Paul Moran, Jerome Nickles, and James Zagzebski, PHYSICS TODAY, July 1983, page 36.)

CT images are a good model for determining radiation-dose deposition in the treatment area, but they often fail to show a clear view of the tumor itself. Yet the tumor's visibility is crucial when the radiation oncologist defines how to target the treatment. MRI overcomes that difficulty: With many contrast-weighting options available to control how bright or dark the tissue appears, there is almost always a way to distinguish the tumor from surrounding soft tissue.

Unfortunately, no straightforward way exists to interpret MRI images as maps of radiation attenuation. NMR signals arise from the nucleus's magnetic behavior, which correlates poorly with the orbital electron behavior that governs x-ray absorption and scattering. Thus it is extremely difficult to accurately model the radiation-dose distribution using MRI images that best show the target lesion.

Radiation oncologists seek the best of both worlds by scanning patients with both CT and MRI. Although image-fusion techniques align multiple images using rigid structures such as bones, small differences in patient positioning or breathing motions between two scans can produce errors when the target, defined on an MRI image, is projected onto the CT image set for radiation planning. In radiation therapy, millimeters of misalignment can severely injure healthy tissue or leave part of an aggressive tumor untreated.

To resolve the problem, physicians and scientists are turning to machine-learning techniques to extract information from multiple MRI image sets and to determine the radiation attenuation of tissue without x-ray images. Many of those algorithms reconstruct "pseudo CT" images—maps of photon attenuation—from the MRI data. That approach provides both tumor visibility and the radiation-dose distribution from a single scan, without the risk of errors introduced by repositioning the patient on another scanner.

Tumors that are difficult to see using x-ray imaging during treatment planning are also difficult to monitor during live treatments. The trend in radiation therapy is to deliver higher, more focused doses to tumors in each treatment session. Doing so reduces the spatial margin for error and places greater importance on techniques for managing moving structures. Imaging and tracking the target in real time during treatment is a powerful way to do that, and MRI offers distinct imaging advantages.

Performing MRI on a patient during a radiation-treatment session also introduces fundamental challenges. A linac is the common tool for modern radiation therapy, but magnetic fields

interfere with its operation; indeed, strong magnetic fields can cut off the electron beam entirely. For the first commercial MRI-guided treatment machine, engineers turned to an earlier mainstay of radiation treatment—the cobalt-60 teletherapy machine. The presence of a magnetic field doesn't affect gamma-ray production in ^{60}Co, and instrument designers were able to focus on other issues, such as integrating the imaging and treatment devices and customizing how mechanical elements functioned inside the magnet.

Putting the system into clinics and hospitals allowed radiation oncologists to treat patients using MRI scans acquired in the treatment room, while medical physicists worked on the problems of measuring and calibrating therapy beams in the strong magnetic field and understanding changes in dose distribution. In particular, when the treatment beam is absorbed or scattered in tissue, energetic electrons are released; they are usually absorbed in nearby tissue, a consequence accounted for in the treatment plan. In a magnetic field, the paths of those electrons are deflected, and the deflection must be modeled in the treatment plan to determine accurate radiation doses.

MRI-guided ^{60}Co therapy opened the door to MRI-guided treatment, but many treatments performed using linacs can't be done on ^{60}Co systems. The quest to develop integrated MR-guided linacs has led to innovative systems in research centers around the world, and commercial systems are now coming to market.

Physicists have managed to produce a strong, uniform magnetic field for real-time MRI in the treatment room while eliminating magnetic interference with the attached linac.[18] There is still much to learn about optimizing MRI scanning during radiation treatments to give the best therapeutic results, but the technology has broad applications for radiotherapy and radiosurgery of all body regions and types of disease.

REFERENCES

1. R. Damadian, M. Goldsmith, L. Minkoff, *Physiol. Chem. Phys.* **9**, 97 (1977).
2. E. B. Schelbert, D. R. Messroghli, *Radiology* **278**, 658 (2016).
3. C. D. Barras et al., *Aust. Fam. Physician* **45**, 798 (2016).
4. Y. L. Dai, A. D. King, *Clin. Radiol.* **73**, 45 (2018).
5. M. Viallon et al., *Neuroradiology* **57**, 441 (2015).
6. L. P. Panych, B. Madore, *J. Magn. Reson. Imaging* **47**, 28 (2018).
7. D. B. Plewes, W. Kucharczyk, *J. Magn. Reson. Imaging* **35**, 1038 (2012).
8. P. Mansfield, A. A. Maudsley, *Br. J. Radiol.* **591**, 188 (1977).
9. P. C. Lauterbur, *Nature* **242**, 190 (1973).
10. D. Hoult, P. C. Lauterbur, *J. Magn. Reson.* **34**, 425 (1979).
11. M. Rogosnitzky, S. Branch, *BioMetals* **29**, 365 (2016).
12. A. Deshmane et al., *J. Magn. Reson. Imaging* **36**, 55 (2012).
13. M. Lustig, D. Donoho, J. M. Pauly, *Magn. Reson. Med.* **58**, 1182 (2007).
14. C. K. Kuhl et al., *J. Clin. Oncol.* **32**, 2304 (2014).
15. D. Ma et al., *Nature* **495**, 187 (2013).
16. F. Bloch, *Phys. Rev.* **70**, 460 (1946).
17. J. J. W. Lagendijk et al., *Phys. Med. Biol.* **59**, R349 (2014).
18. J. J. W. Lagendijk, B. W. Raaymakers, M. van Vulpen, *Semin. Radiat. Oncol.* **24**, 207 (2014).

PT