Review of Medical Imaging

Topic 1 – Ultrasound. We first discussed medical ultrasound imaging in some detail. Ultrasound's popularity stems from the fact that it yields high-resolution images of soft tissue and internal organs and can be used to monitor blood flow. In addition, ultrasound is portable, fast, safe, and relatively inexpensive.

Imaging

The concepts underlying ultrasound-based imaging are relatively simple. Ultrasonic waves with frequencies between ~2 and 15 MHz are sent into the body, where they undergo reflection/scattering and transmission when incident on an interface, similar to light. However, in the case of ultrasound, reflection and transmission occur at interfaces between tissues with different acoustic impedances, whereas, in the case of light, reflection and transmission occur at interfaces between media with different refractive indices. Reflected and scattered ultrasound creates a signal (an echo) that is characterized by two important parameters: echo amplitude and arrival time. These two parameters are determined by the strength of any detected specular reflection, diffuse reflection, and/or scattering and by the depth of the reflecting tissue, as described in more detail below.

Ultrasonic waves are generated and detected by piezoelectric crystals, which are components of an ultrasound transducer. Piezoelectric crystals convert electrical energy into mechanical energy by oscillating when deformed by an applied voltage. Similarly, piezoelectric materials convert mechanical energy into electrical energy by producing a voltage when set into vibration. Most transducers used in medicine contain many crystals, making it possible to steer and focus the beam and sweep it to generate a rectangular or sector-like field of view, as shown at right.

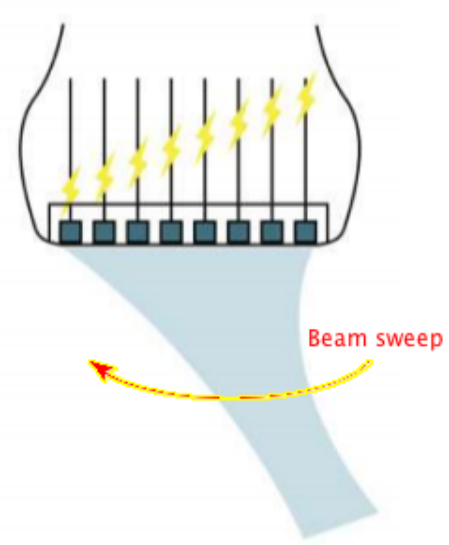

We discussed one- and two-dimensional ultrasound imaging; two-dimensional imaging is the most common diagnostic mode in obstetrics. One-dimensional imaging was developed first and now is largely outmoded except in cardiology, where it is commonly used to monitor valve motion. One-dimensional imaging can

be divided into two categories – A (amplitude) and M (motion) mode. In A mode, echoes are collected along one line within the body, and echo arrival time and amplitude are recorded. The arrival time, t, can be used to determine the depth of the reflecting structure in A mode (and other modalities) using the relationship

$$depth = \frac{v_{sound} t}{2}$$

where v_{sound} is the speed of sound in the body (~1500 m/s in soft tissue). The echo amplitude determines echo height in one-dimensional images and brightness (contrast) in two-dimensional images. A schematic and a real A mode image are shown below; these reveal the depth and dimensions of a blood vessel and of an eye, respectively.

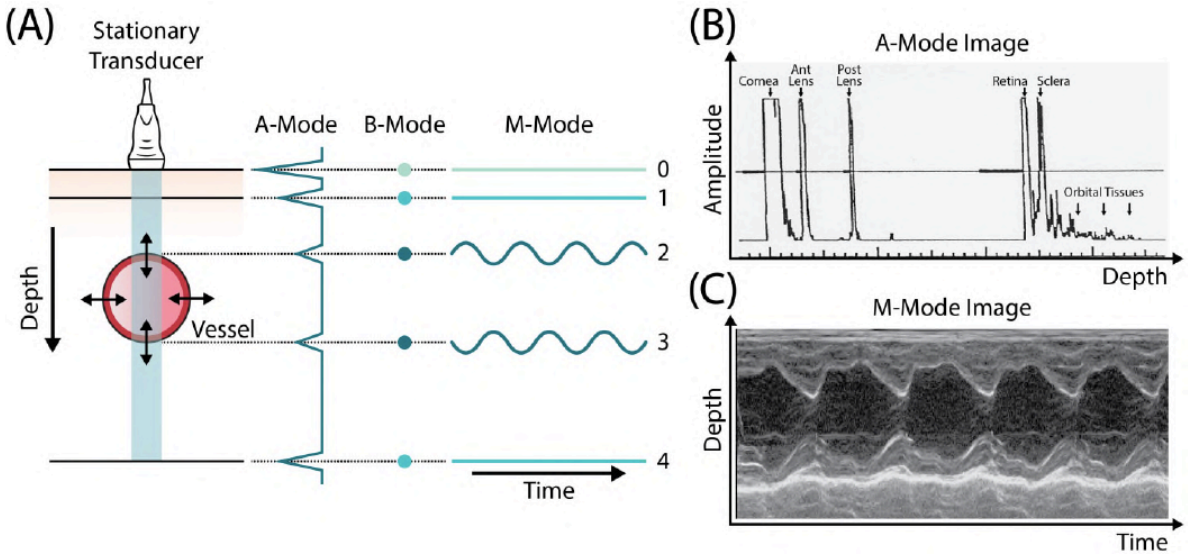

In M mode, motion is added to A mode images. Depth is plotted versus time using dots whose brightness is determined by echo amplitude. Thus, changes in dot depth with time reveal motion, as shown in the figure above. M mode is used in cardiology to study valve motion.

Two-dimensional B (brightness) mode generates conventional-looking images that are similar to cross-sectional slices. In B mode, many line scans are collected. For each line scan, time and amplitude are recorded, similar to an A scan, and then the depth of the reflecting interface is plotted along the vertical using dots whose brightness reflects echo amplitude, similar to M mode. However, in B mode, the horizontal axis is a distance, whereas in M mode it is time. Thus, B mode images are two-dimensional plots of tissue depth versus transverse location with contrast generated by echo amplitude. Like M mode, B mode can be used to monitor motion in real time because two-dimensional images are acquired rapidly. A B mode image of the heart is shown below.

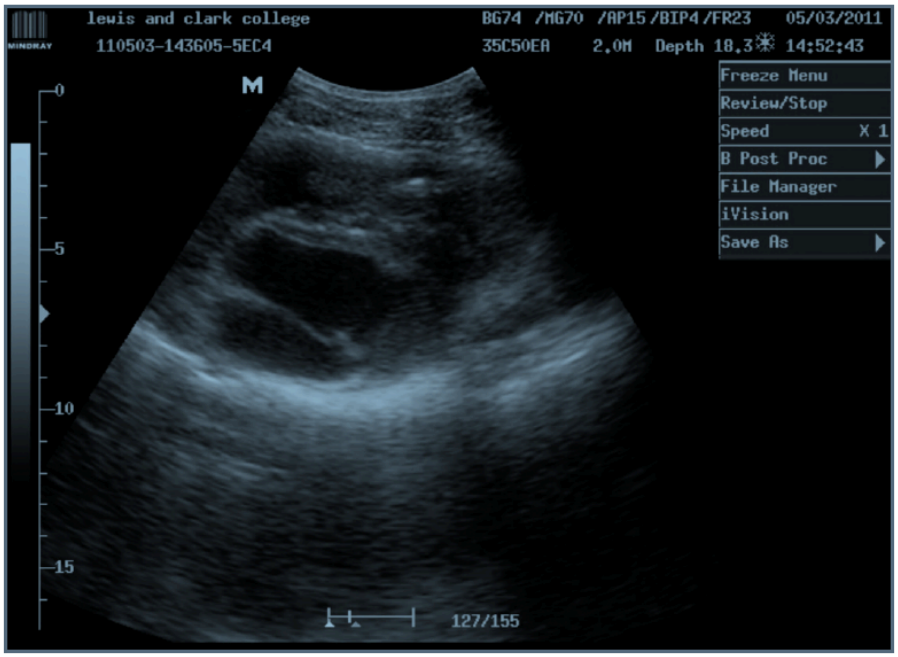

After learning the basics of ultrasound image generation, we discussed some attributes of ultrasonic waves that affect image quality. These include beam spread, reflection strength, intensity attenuation, and frequency.

Reflection strength

Reflection strength is a very important determinant of image quality. A structure must reflect (or scatter) some of the incident ultrasound wave to be detected, but excessive reflection is not desirable. Transmission also is useful because transmitted waves generate additional echoes, facilitating visualization of the interior of reflecting structures as well as the visualization of structures at different depths. In addition, transmission of ultrasound into the body is essential.

Most soft tissues in the body have relatively similar acoustic impedances, and thus interfaces between soft tissues only reflect $\lesssim$1% of the incident wave. Fortunately, this percent reflection produces an echo with adequate signal strength, and it leaves a transmitted beam that is strong enough to generate additional echoes. In contrast, air and bone have acoustic impedances that differ markedly from those of soft tissue, and thus interfaces between air or bone and soft tissue do not transmit ultrasound well. As a consequence, ultrasound will not penetrate into the body unless an impedance-matching gel fills the space between the transducer and the skin, and bone imaging is not a strength of ultrasound.

Intensity attenuation

Despite the fact that most reflections are weak, ultrasonic waves lose intensity quite rapidly with depth. Intensity fall-off is caused by absorption and scattering and is strongly dependent on beam frequency. Unfortunately, attenuation increases and resolution improves with increasing frequency, and therefore frequency choice in ultrasound imaging of deeper structures involves a compromise between penetration and resolution. Higher-frequency ultrasound is used to image structures that are close to the surface and thin tissues, whereas lower-frequency ultrasound is used to image deep structures at lower resolution.

It is well established that ultrasound echoes attenuate exponentially with distance and thus ultrasound instruments can correct for intensity fall-off mathematically. This correction is known as time-gain compensation.

Resolution

We discussed two kinds of resolution – lateral and axial. The first refers to the ability to discriminate objects that are side-by-side in a plane perpendicular to the beam, and the second refers to the ability to discriminate nearby objects along the direction parallel to the path of the beam, as shown in the figure below. Objects are laterally resolvable if their separation exceeds the beam width, and therefore lateral resolution degrades with depth as the beam spreads. Objects are axially resolvable if the pulse is sufficiently narrow (i.e., the spatial pulse length is less than twice the axial separation). Since a pulse typically is a few wavelengths in width, narrow pulses and improved axial resolution are obtained as frequency increases.

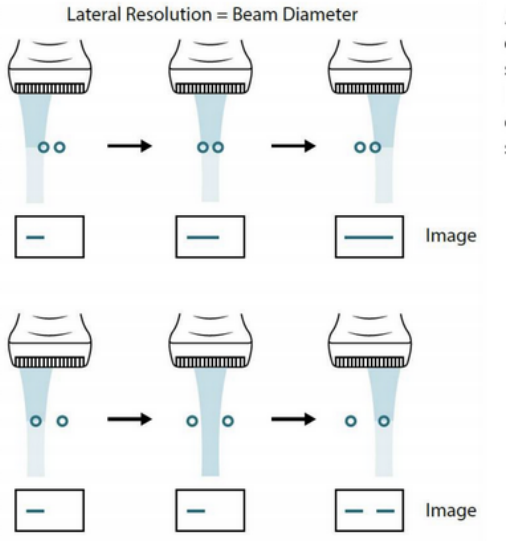

Monitoring blood flow

Ultrasound also is used routinely to measure and monitor blood flow. Blood flow is measured by detecting Doppler shifts in frequency, which are generated when ultrasound from the transducer impinges on flowing blood. Two Doppler shifts are generated by blood flow; the first shift is produced because the blood acts as a moving receiver of a wave sent out by a stationary source (the transducer), and the second shift is produced because the blood acts as a moving source of (Doppler-shifted) ultrasound that is scattered back to, and detected by, a stationary transducer. The transducer monitors both the sign and the magnitude of the change in frequency; positive changes (increases in) frequency arise when the blood is flowing towards the transducer, and negative changes arise when the blood is flowing away from the transducer. The magnitude of the frequency change, $\Delta\nu$, is used to compute flow speed, u_{blood}, using the equation from lecture.

Doppler flow data also can be superimposed on B mode images to create color flow Doppler images. In these images, flow direction relative to the transducer is encoded in color, and flow magnitude is encoded in brightness. A common mapping format is BART (blue away red towards the transducer).

Topic 2 – Conventional radiography and computed tomography.

Our second topic was X-ray imaging. In this section, we discussed conventional radiography and X-ray computed tomography (CT). Conventional radiographs are familiar to most people because they are so commonly used in dentistry and to generate images of the skeleton. CT is a less familiar, but much more powerful, X-ray imaging technique that can be used to image the skeleton and soft tissues in two and three-dimensions. A primary weakness of X-ray based imaging is the use of ionizing radiation; another is an inability to monitor function.

Conventional radiographs are just shadows, or projections, created when X-rays are absorbed by radiopaque structures in the body, most notably bone. Because radiographs are projections, they yield no information about the depth of structures within the body. In addition, structures can be partially or completely missing from radiographs if they are obscured by other structures. Some of the differences between conventional radiographs and CT images are exemplified in the figure at right.

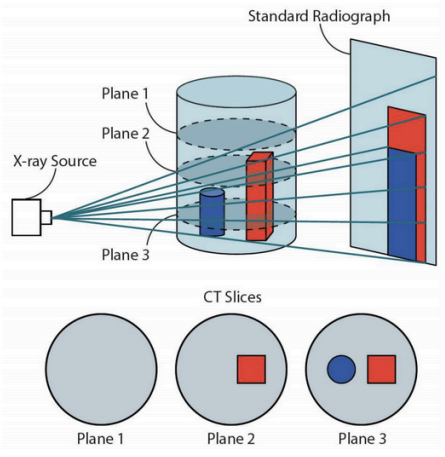

Conventional radiography

The ideas underlying image formation in conventional radiography also are relatively simple. In the ideal case, X-rays from a source travel along a straight line through the body (i.e., there is no scattering). Transmission, and the detected signal, are high in regions of the body that lack X-ray absorbing structures, whereas they are low in regions containing significant X-ray absorbing material. In radiographs, radiolucent areas appear black, and radiopaque areas appear white, as shown at right. In reality, X-ray imaging is far from ideal; for example, scattering is a major problem.

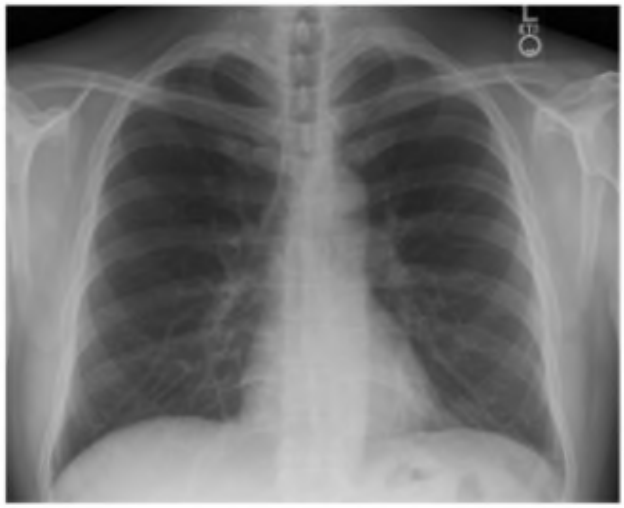

X-rays used in diagnostic imaging have energies that lie in the range from ~20-150 keV. In this energy regime, X-rays have three fates; they can be absorbed via the Photoelectric (PE) effect, scattered via the Compton effect, or transmitted. Absorption and transmission yield useful diagnostic information by creating contrast that differentiates radiopaque and radiolucent structures in the body. Scattering, in contrast, creates an image-degrading fog by causing rays to travel along nonlinear paths.

As its name suggests, in the PE effect a photon is absorbed, and the energy is used to free a bound electron. Electrons that are weakly bound in metals can be freed via the PE effect using relatively low energy (violet and ultraviolet photons). In contrast, longer-wavelength photons cannot free electrons from metals because their energy is not as large as the binding energy of electrons.

X-ray photons are very energetic, and thus absorption of X-rays via the PE effect leads to the release of tightly bound, inner shell electrons from atoms like carbon, oxygen, hydrogen, and nitrogen. In X-ray imaging, it often is useful to enhance the incidence of the PE effect and thereby increase absorption. This can be accomplished by introducing atoms with high atomic number, Z, into the body. The PE effect also can be made more prevalent by decreasing photon energy and/or by matching photon energy to an “absorption edge.” At an absorption edge, photon energy slightly exceeds the energy required to eject an electron, and thus the probability of the PE effect increases markedly.

In the Compton effect, photons free weakly bound electrons by giving up part of their energy. Importantly, lower-energy, scattered photons remain as products of the Compton effect, and these photons degrade X-ray images. Absorption edges and Z do not influence the probability of occurrence of the Compton effect, whereas density and photon energy do.

Often, only a few percent of X-rays are transmitted through the body and reach the detector. Despite their minority status, transmitted rays are key to image formation because some rays must reach the detector to produce an image. In view of its importance, we analyzed transmission quantitatively for a beam that is monoenergetic. In this case, the transmitted intensity decays exponentially with penetration distance, x, according to the formula:

$$I_{trans} = I_0 e^{-\mu x}$$

where μ is the tissue- and energy-dependent attenuation coefficient. Not surprisingly, μ is larger for bone than for softer tissues like fat and muscle, leading to greater attenuation by bone.

Several attributes of μ are important to image formation. First, μ decreases with increasing energy. Thus, images of thick objects are generated using more energetic (better penetrating) X-rays. Second, disparities in μ enhance contrast. Tissues that have significantly different μ's at a particular X-ray energy will generate good contrast, whereas tissues with similar μ's tend to generate poor contrast unless their thicknesses differ markedly. To circumvent this problem and enhance contrast, radiopaque materials (contrast media) containing high Z elements such as barium and iodine can be introduced into soft tissue in the body. Contrast media can improve the local contrast of a vessel in soft tissue by an order of magnitude.

Contrast also can be improved by using lower-energy X-rays, if a tissue is not too thick, and by reducing the amount of scattering and/or the detection of scattered rays. To reduce scattering, physicians can reduce the irradiated area. Radiolucent materials surrounded by lead strips also can be placed edge on between the patient and the detector. These reduce detection of scattered rays because scattered rays preferentially hit the lead strips and are absorbed, whereas transmitted rays preferentially pass through the radiolucent material and reach the detector.

Computed tomography (CT)

The limitations of conventional X-ray ray imaging were largely overcome in the 1970s with the introduction of X-ray computed tomography. Tomography means picture of a plane, and, consistently, CT converts a series of projections collected at different viewing angles into a cross-sectional view that contains the depth and structural information that are missing from projections. In addition, CT, like ultrasound and MRI, can produce three-dimensional images.

The basic idea underlying CT is that projections/views collected at different angles reveal complementary information about structure that can be used to reconstruct two- and three-dimensional images. A typical CT scanner will collect 300-1000 views spanning 360° to generate one cross-sectional slice. This takes a few seconds to minutes using a CT scanner that rotates the source and detector together around a stationary patient. Modern (spiral) CT scanners also translate the patient while slice data are collected. Spiral CT, and the use of multi-slice CT (MSCT) scanners, markedly accelerate data collection. Thus, spiral CT is very useful for severely traumatized patients that need rapid treatment.

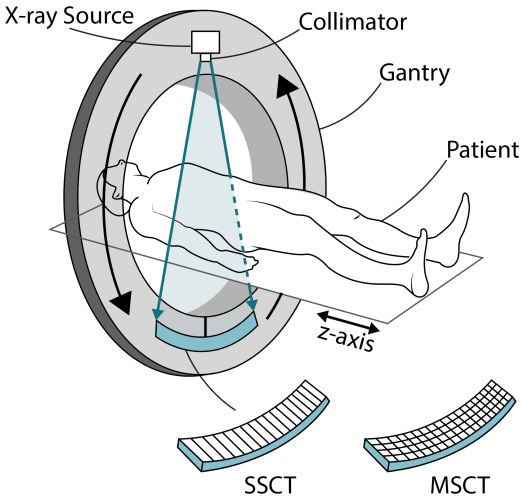

After many projections are collected, two-dimensional slices are reconstructed. Reconstruction of slices from projections is a difficult mathematical problem that was solved by the mathematician Johann Radon in 1917. Unfortunately, the medical community remained unaware of Radon's work for a long time, despite its relevance to CT.

We discussed one popular method of reconstruction, filtered back projection (FBP). The essence of back projection is not too difficult to grasp. Many views revealing attenuation along lines are collected; for the sake of simplicity two views are shown in the figure below. Projection data are then smeared back uniformly

along the line of collection and data from different views are added. This process is simple back projection, and, as shown in the (unnormalized) figure below, it yields a reconstruction that is blurred but reveals the basic attributes of the object.

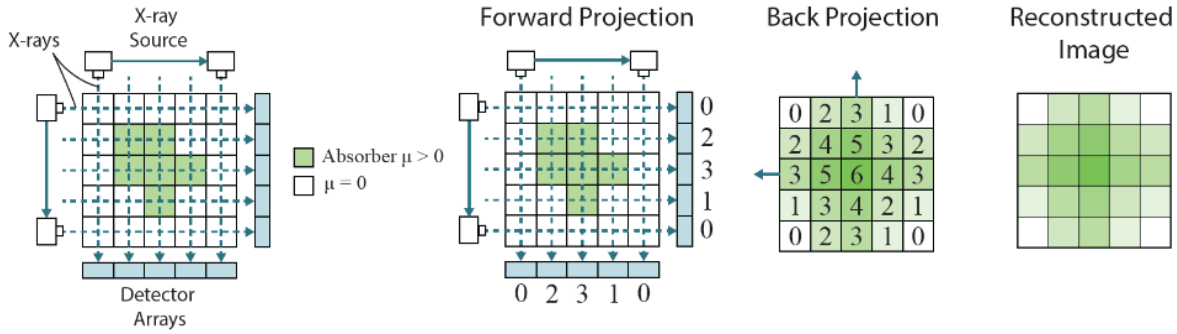

FBP removes the blur generated by simple back projection. The value of filtering can be rationalized by noting that blur suggests a need to suppress low-frequency information in the image. Filtering can be used to remove undesirable low-frequency information and thereby improve image quality. FBP will yield an exact reconstruction of the object if an infinite number of views are collected with an infinite number of points per view.

To close our section on X-ray imaging, we watched a movie showing three-dimensional images obtained from a CT scan of the coronavirus-infected lungs of a man who caught one of the earlier, very deadly variants.

Topic 3 – Radionuclide imaging.

Our third topic was radionuclide imaging, which is a functionally-oriented analog of X-ray imaging. High-energy photons again create the image, but in nuclear imaging these are gamma rays that are generated by unstable nuclei inside the patient. Nuclear images generally have much poorer resolution than X-ray images, but the contrast of nuclear images typically is quite good.

To generate a radionuclide image, the patient is exposed to a radiopharmaceutical, which consists of a radioactive material appended to a pharmaceutical that localizes to a particular site in the body. The labeled site becomes a source of high-energy, gamma-ray photons that are generated by nuclear decay events. The gamma rays escape the body, and the spatially varying “count profile” that they produce on a detector is used to generate images of the distribution of the radiopharmaceutical at the site of accumulation.

Planar nuclear imaging, like projection radiography, generates images that lack depth information and that suffer from structural superposition. In contrast, nuclear

emission tomography, like CT, yields computer-reconstructed cross-sectional and three-dimensional images of radionuclide distribution.

Gamma emitters used in medical imaging

Radionuclide imaging is based on the use of unstable nuclei that decay and emit energy (directly or indirectly) in the form of high-energy gamma-ray photons. We focused on two types of gamma emitters – direct emitters, including technetium 99-m (Tc-99m), and indirect emitters, including fluorine-18 (^{18}F).

Tc-99m is a metastable excited nuclear state that decays into a lower-energy state with half-life of ~6 hours. The decay event leads to the generation of a single, 140-keV gamma ray photon that can participate in image formation if it escapes the patient. Tc-99m is commonly used in planar imaging and in single-photon emission computed tomography (SPECT).

In contrast, ^{18}F is an unstable nucleus with a shortage of neutrons. To correct the deficit, ^{18}F emits a positron, which converts a proton in a neutron. A key feature of positron emission in a patient is downstream production of two 511-keV photons, which travel in anti-parallel directions, as shown at right. ^{18}F is commonly used in positron emission tomography (PET).

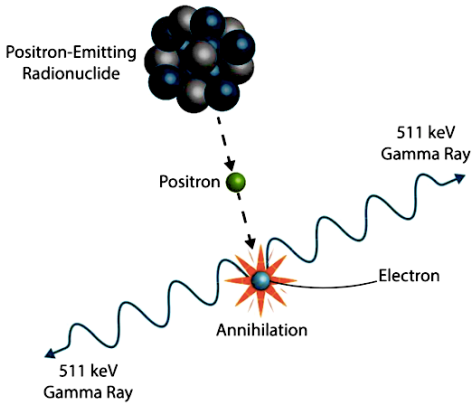

Planar imaging

Projection radiographs ideally are a map of total X-ray attenuation along a line between the source and elements on the detector. Similarly, planar nuclear images ideally are a map of total gamma ray counts arising from a well-defined line of response (LOR) in the patient, as shown at right. In X-ray imaging, scattering is an important non-ideality that degrades quality. Isotropic emission, attenuation, and scattering have a similarly negative effect on planar nuclear imaging.

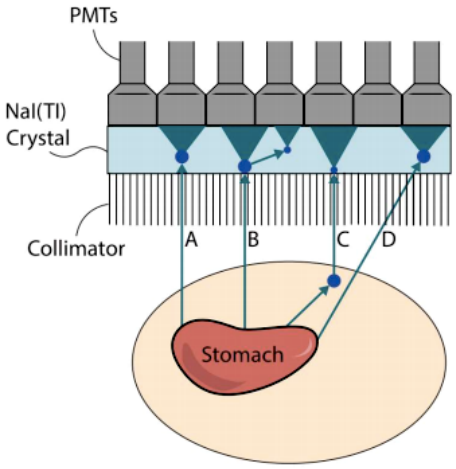

To create a meaningful projection image, gamma ray counts that impact the detector must arise from a unique LOR in the patient. To meet this requirement, gamma rays, which are emitted isotropically, must be collimated, as shown in the schematic above. We discussed collimators in some detail because they are a major determinant of resolution in planar nuclear imaging and in SPECT. Collimators also drastically reduce count rates by rejecting gamma rays that do not travel along a normal to the detector. Notably, PET does not require collimators, and this is a major advantage of the approach. Signal strength in PET is two to three of orders of magnitude higher, and, as a consequence, resolution also is improved because images require less smoothing.

Gamma rays are scattered and attenuated in the patient, similar to X-rays. Scattering degrades contrast; it can partially be eliminated via energy discrimination because scattered rays have lower energy.

Attenuation artificially reduces detected counts from source points deeper in the patient. One simple method used to reduce attenuation is to collect a view (e.g., posterior) in which the structure of interest (e.g., a kidney) is as close as possible to the detector. In tomographic imaging, attenuation typically is corrected using dual-mode nuclear emission/X-ray transmission imaging systems, such as combined SPECT/CT and PET/CT systems. CT is used to map attenuation coefficients, and these are incorporated into reconstruction algorithms to facilitate attenuation correction. Dual-mode imaging also is a powerful method of obtaining images that combines functional and anatomical data, as shown at right.

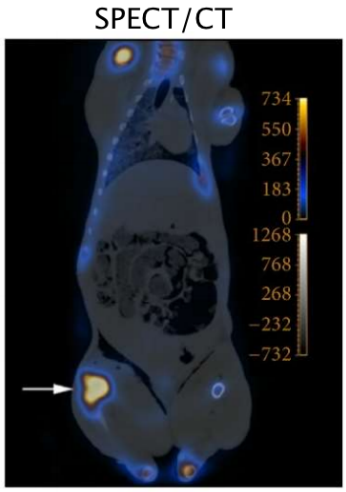

Emission computed tomography (ECT)

ECT usually is sub-divided into two techniques – SPECT and PET. The goal in both techniques is to collect a set of projection images that permit reconstruction of the distribution of activity in the patient. SPECT is less expensive than PET, but also generates images with lower resolution. Resolution in SPECT typically ranges from 10 – 20 mm, whereas resolution in (brain) PET can be less than 5 mm.

As the name suggests, SPECT uses radiopharmaceuticals, commonly labeled with Tc-99m, that generate a single photon per decay event. Similar to X-ray CT, image acquisition in SPECT involves the collection of many (e.g., 64 – 128) views (projections) of radionuclide distribution. Each view consists of a set of data points representing the sum of activities along LORs. Different views are generated by

rotating the camera(s) around the patient, as shown in the schematic below. Again, similar to CT, cross-sectional and three-dimensional SPECT images are reconstructed from multiple views using FBP or iterative approaches.

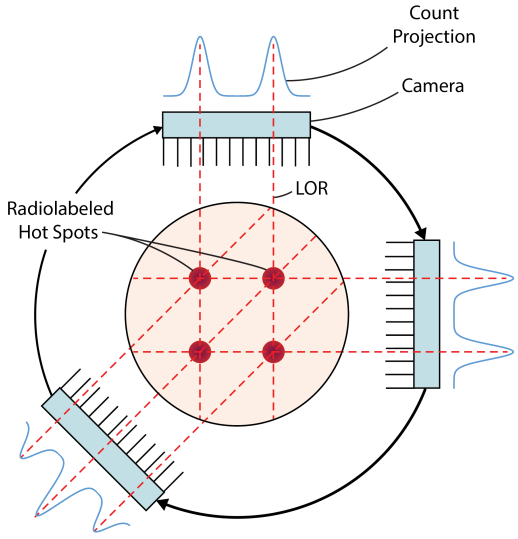

Unlike SPECT, PET uses radiopharmaceuticals, commonly labeled with ^{18}F, that decay and generate a positron. Individual positrons travel a short distance in the body (~0.13 mm) and then are annihilated via interaction with an electron. Each annihilation yields two 511-keV photons that ideally move in anti-parallel directions and arrive, nearly simultaneously, at two detectors. The associated detection circuitry registers a coincidence event if certain energy and timing conditions are met.

Unlike SPECT, collimator grids are not mandatory in PET because their function, which is to define an LOR, is fulfilled by coincidence detection. Specifically, coincidence detection assigns the LOR to the line connecting the two detectors that register the events, as shown in the schematic at right.

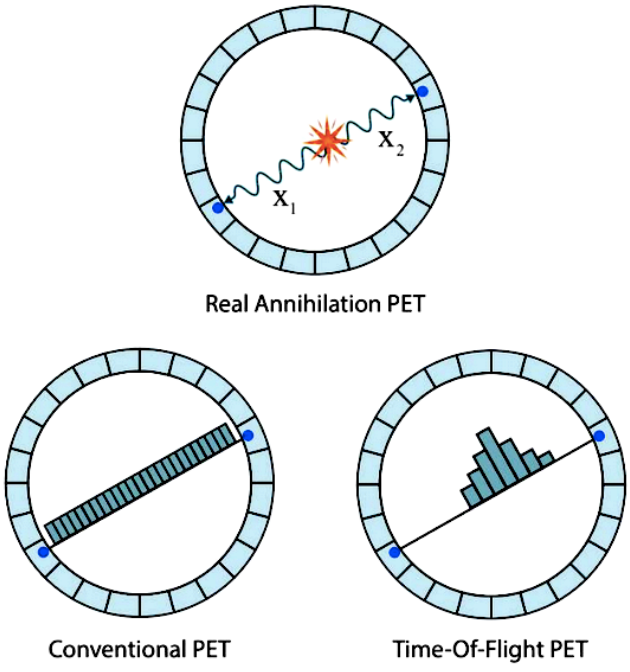

Another distinct attribute of PET is the possibility of implementing "time-of-flight" (TOF) detection, which uses differences in arrival times to refine the location of the decay event along the LOR. Identifying decay positions is the goal of reconstruction, and thus highly accurate TOF detection is an alternative approach that could supplant reconstruction. Unfortunately, current temporal resolution is insufficient to eliminate the need for reconstruction, but TOF data can be used to a develop a probability distribution for the decay event along the LOR, as shown above. This information is enough to produce significant improvements in image quality. As a consequence, TOF-PET/CT systems are increasing in number.

The reconstruction problem in ECT is similar to that in X-ray CT, and so too are the image reconstruction approaches, which include FBP and iterative techniques. Iterative reconstruction is increasingly preferred in nuclear imaging because it is less prone to artifacts and noise and is better suited to attenuation correction, and so we emphasized this approach.

Iterative reconstruction techniques attempt to find a solution via successive estimates, as outlined in the schematic at right. The basic idea is to use the current "estimate" of activity to compute estimated projection data. The estimated and measured projection data are compared, and their difference is used to update the activity estimate. Iteration is terminated when agreement between calculated and measured projection data is sufficiently good.

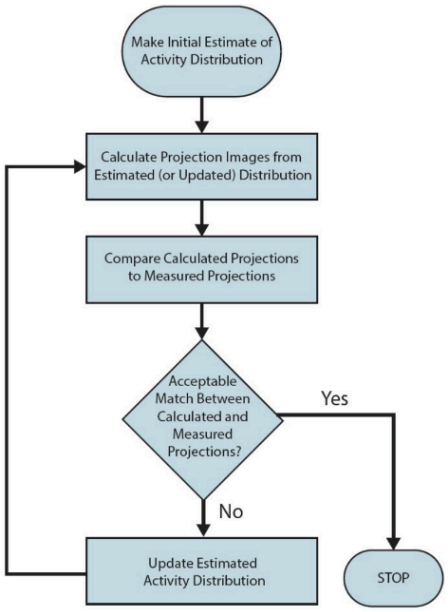

Topic 4 – Magnetic resonance imaging (MRI).

Our last topic was magnetic resonance imaging (MRI), which is the premier method of imaging soft tissue, especially in the brain, for patients who do not have metal implants or a pacemaker.

Most descriptions of MRI use a mix of concepts from classical and quantum mechanics. No matter which view is adopted, the starting idea is that atomic nuclei can possess a magnetic moment $\vec{\mu}$ and that, in the presence of an external

magnetic field (B_0) along z, it is energetically favorable for $\vec{\mu}$ to align with its z component parallel to the field. Even in strong magnetic fields, the bias in alignment is small, ~0.001%, because the energy difference between high and low energy states is small relative to thermal energies. However, for hydrogen, which is abundant in most body tissues and which has a large $\vec{\mu}$, the small bias leads to a net magnetic moment/magnetization along z that can be manipulated to generate a detectable signal in MRI.

Signal generation

When signal generation is discussed, the classical and quantum mechanical descriptions begin to diverge substantially. Signal generation involves perturbing the static net magnetization generated by the static external field and then letting the magnetization return to its equilibrium value. This perturbation process leads to the time varying signal that is monitored in MRI.

The classical view of signal generation is that, before perturbation, the net equilibrium magnetization is oriented along the direction of the external magnetic field (i.e, the z axis). This static situation is perturbed by subjecting the nuclei to a brief magnetic pulse at the Larmor frequency, $\nu_L = \bar{\gamma}B$, where $\bar{\gamma}$ is the gyromagnetic ratio. This pulse is on resonance with the nuclei and tips the net magnetization by 90° into the xy plane, where it initially is large because the nuclei oscillate and emit in synchrony. With time the nuclei lose synchrony, and the transverse magnetization decays exponentially to zero with a time constant T_2, generating the MRI signal. A schematic of signal generation in MRI is shown below.

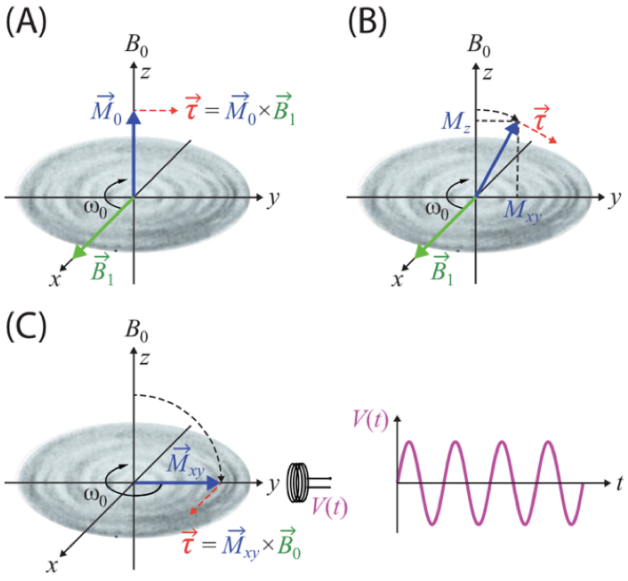

The quantum view of signal generation is that nuclei can absorb radiation from a pulse and transition into a higher-energy state if E_{photon} matches the energy separation between low- and high-energy states. In class, we analyzed nuclear transitions from a quantum perspective and found that the frequency that will induce transitions is the Larmor frequency. We also showed that this frequency lies in the radiofrequency range of the electromagnetic spectrum. For example, for a proton in a field with a magnitude of 1 Tesla, the transition frequency/Larmor frequency is 42.58 MHz. In summary, from a quantum perspective, a signal is generated in MRI using a pulse that causes nuclei to transition into the higher energy state and that thereby reduces the bias in alignment. After the pulse, the nuclei relax back to the lower energy state and re-establish the biased distribution in alignment and its associated magnetization.

As the transverse magnetization M_{xy} decays (exponentially) to zero, the longitudinal magnetization M_z must necessarily rise, and it does so exponentially with a time constant $T_1 \geq T_2$. Importantly, both T_2 and T_1 vary markedly with tissue type, whereas density does not, and thus contrast in MRI images typically is better if the signal amplitude is manipulated to reflect variations in relaxation times rather than variations in density, ρ.

We discussed two remaining aspects of MRI that add to the complexity of image formation. First, generating signals that are weighted towards relaxation times is not straightforward; it involves the use of pulse sequences. Second, MRI signals must be made to vary with position; this involves the use of field gradients.

Pulse sequences

One reason for the use of pulse sequences is that they facilitate generating images with contrast that is weighted by T_2, T_1, or ρ. To illustrate the idea, we analyzed a commonly used pulse sequence in MRI – the two-dimensional spin echo pulse sequence.

The sequence uses RF pulses and two repetition times – TE and TR. TE is the time at which data are collected. TR is the time when the experiment is repeated. If TR is less than a few T_1, the longitudinal magnetization will have recovered incompletely at TR, i.e., $M_z(T_R) = M_0\left(1 - e^{-TR/T_1}\right)$, will be less than M_0. The second 90° RF pulse tips this reduced longitudinal magnetization into the x/y plane. As a consequence, the signal at the second echo time TE will be reduced relative to the first echo.

We discussed two ways to decide how to achieve different types of weighting in a two-dimensional spin echo pulse sequence. One is based simply on consulting a summary weighting table, which is shown below.

Parameter	T_1 Contrast	Proton Density Contrast	T_2 Contrast
TR (ms)	400-600	2,000-4,000	2,000-4,000
TE (ms)	5-30	5-30	60-150

A more illuminating approach is to calculate the signal amplitude for different values of TE and TR and for different types of tissues, like cerebrospinal fluid (CSF) and fat. The expression for the signal amplitude is:

$$Signal\ Amplitude \propto M_{xy}(T_E) \propto \rho(1 - e^{-TR/T_1})e^{-TE/T_2}$$

We did a few examples in class. One demonstrated that CSF typically appears bright on T_2-weighted images and dark on T_1-weighted images. Results for fat were the opposite. The examples also revealed which of the two exponential terms was creating contrast differences and thus whether the contrast was weighted by T_1 or T_2.

The important take-home message is that MRI signals can be made to reflect effects of T_1, T_2, or ρ with judicious choices of TE and TR. T_1-, T_2-, and ρ-weighted images of brain, are shown below.

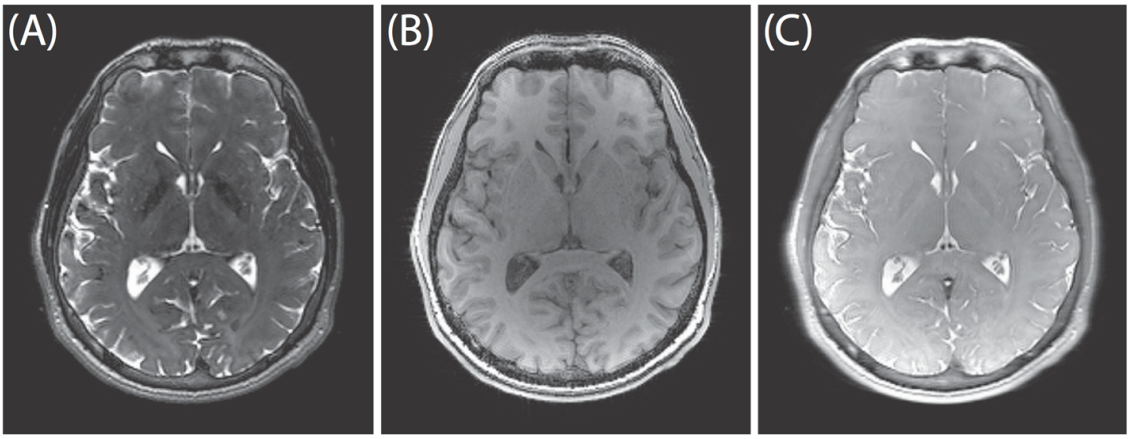

Field gradients and signal encoding

We also discussed how signals in MRI are made to depend on position. The key idea is that the resonance frequency is directly proportional to B and thus if B varies

with position so does the resonance frequency. Thus, MRI uses three field gradients to localize signals first along z, and then along y and x.

The first step in signal mapping is to apply a gradient along z and simultaneously to send in an RF pulse that will be absorbed only by protons that localize to a z plane where the Larmor frequencies match the frequencies of the pulse, as shown in the figure below. This slice selection process defines the image plane.

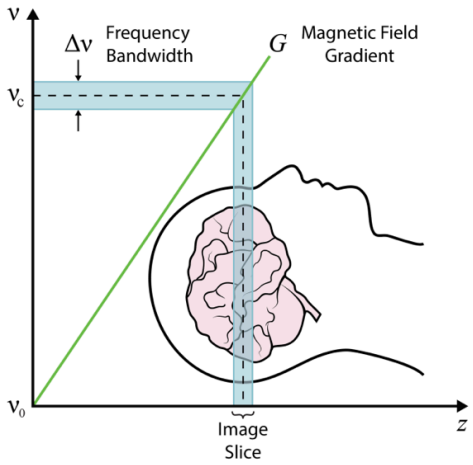

After a slice is selected, two more gradients are applied to make the MRI signal vary with x and y in the image plane. First, a y gradient is applied transiently and then turned off before signal readout. Frequencies do not vary with y at readout because the y gradient is off; however, phases do vary with y because, during the gradient, protons at different y locations "rotate" at different rates and therefore acquire different phase shifts. Thus, along y, signals are distinguished based on phase differences. Second, an x gradient is applied, and the signal simultaneously is read out, allowing x locations to be distinguished based on frequency differences, as shown in the figure below.

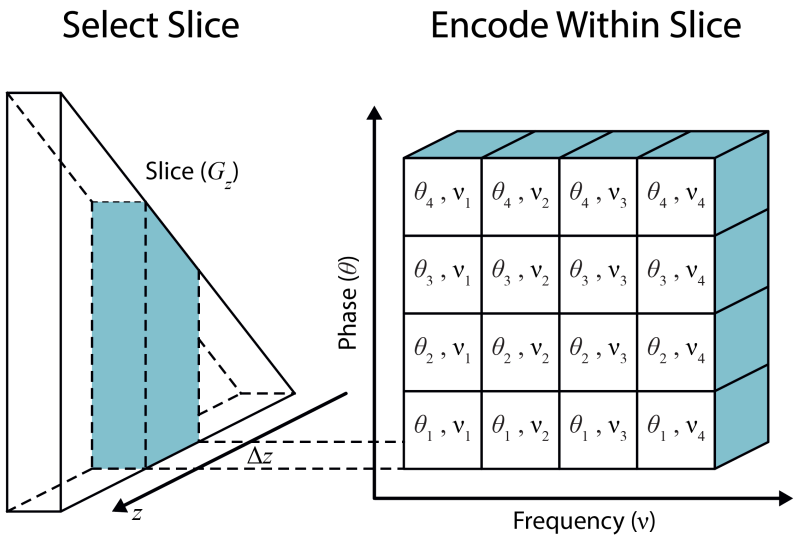

In summary, the two-dimensional spin echo pulse sequence involves the use of RF pulses to create signal and the use of field gradients to achieve slice selection, followed by phase and frequency encoding. The entire process is summarized in the pulse sequence diagram shown below.

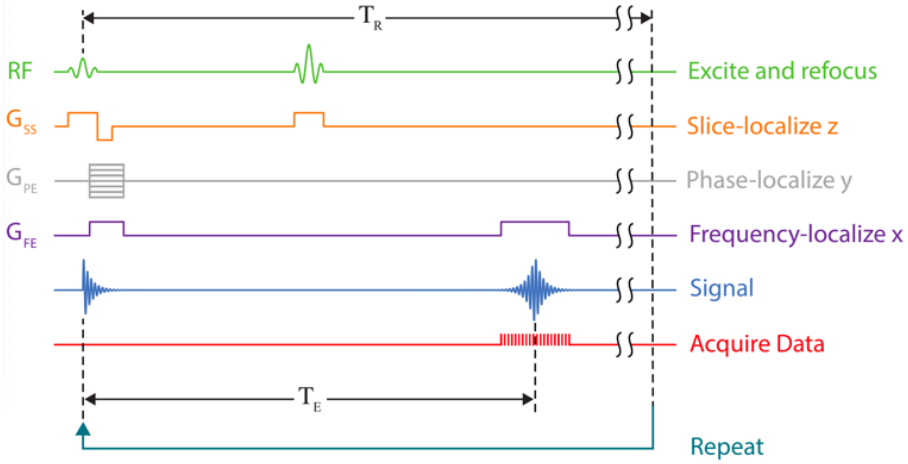

Comparison of ultrasound, radiography, radionuclide imaging, and MRI

We close with a comparison of the four main imaging techniques discussed in class. Ultrasound yields high quality images of soft tissue at relatively low cost, using instruments that are portable, and without causing any apparent harm. Ultrasound also provides useful information about blood flow. Conventional radiography and CT, in contrast, yield high quality images of bone. CT also yields excellent images of soft tissue because it is much more sensitive than conventional radiography, and spiral CT produces three-dimensional images very rapidly and at low cost relative to MRI. Unfortunately, conventional radiography and CT carry with them the dangers of ionizing radiation. Moreover, contrast in X-ray imaging is generated by variations in a parameter, the attenuation coefficient, that is not a sensitive measure of anatomy and that is a not measure function. Radionuclide imaging generates high-contrast, low-resolution images that tend to be directed at assaying function. For example, radionuclide images can reveal sites of enhanced or reduced glucose metabolism, which can reveal the presence of cancer or reduced brain function, respectively. MRI yields exceptionally high quality two- and three-dimensional images of soft tissue, including the brain, without any apparent harm. MRI also can map function, particularly activity of the brain, via a signal that reveals the extent of blood flow to different areas of the brain. Unfortunately, not everyone can have an MRI. Moreover, image acquisition is relatively slow, and MRI instruments are expensive.