4 Principles of CT and Hybrid Imaging

Christoph J. Trauernicht

4.1 Introduction

Computed tomography (CT) entered its sixth decade of clinical use and has proved an exceptionally valuable and useful imaging tool. The first head CT scanner was introduced in 1972. Each pair of slices took over 4 min of scan time and over 1 min of reconstruction time. While this would be considered quite terrible by today's standards, it was considered revolutionary at the time and earned Godfrey Hounsfield and Allan Cormack the Nobel Prize in Medicine in 1979.

4.2 Principles of CT

4.2.1 The X-ray Tube

At the heart of a CT scanner is an X-ray tube (Fig. 4.1). A tungsten filament is heated and emits electrons by a process known as thermionic emission. The emitted electrons, having a negative charge, are accelerated across a potential difference towards a copper anode, which sits at a positive potential. All this happens inside an evacuated glass housing. The vacuum is required to prevent electrons from interacting with any materials inside the housing, other than the tungsten target.

When energetic electrons come into the vicinity of an atomic nucleus, the positively charged nucleus attracts the negatively charged electrons, and these are decelerated in the process. The electric field of the nucleus exerts a force on the incoming electron and forces it to change its velocity (i.e., energy) and direction.

The energy difference of the initial electron energy and the deflected electron energy shows up as an X-ray photon. This process is known as bremsstrahlung (Fig. 4.2), which is the German word for "braking radiation". Bremsstrahlung production depends on the square of the number of protons (Z) in the target nucleus; therefore, the target should consist of a material with a high atomic number, like tungsten ($Z = 74$). Bremsstrahlung production is very inefficient, only around 0.9% for 100 keV electrons, the other 99% of energy is lost through other interactions that do not produce X-rays, but that do result in heat. Therefore, there must be a cooling mechanism for the X-ray target, and the target material must have a high melting point.

An electron that is travelling close to a nucleus will experience a larger force of attraction to the nucleus than one that is passing the nucleus at a larger distance. The electron will experience a larger energy loss, resulting in a higher energy X-ray photon. If the potential dif-

C. J. Trauernicht (✉)
Division of Medical Physics, Tygerberg Hospital and Stellenbosch University, Cape Town, South Africa
e-mail: cjt@sun.ac.za

J. H. Bortz et al. (eds.), *CT Colonography for Radiographers*,
https://doi.org/10.1007/978-3-031-30866-6_4

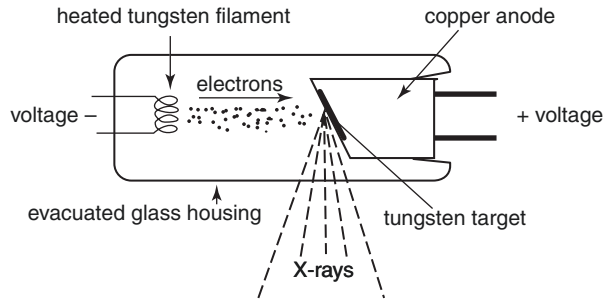

Fig. 4.1 Simplified diagram of an X-ray tube

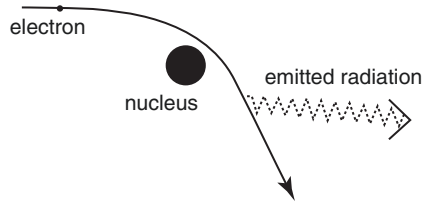

Fig. 4.2 Bremsstrahlung production

ference between the tungsten filament and the target is, for example, 120.000 V (or 120 kVp), then the resultant X-rays will have an energy range from near zero up to 120 kVp. When the operator selects a kVp before the acquisition of an image series, this adjusts the potential difference inside the X-ray tube through the use of a transformer. If the kVp is increased, the energy of the X-rays is increased. This is often referred to as beam quality. In addition, bremsstrahlung production becomes more efficient at higher energies, resulting in more X-rays. This is referred to as the "quantity". Dose to the patient increases with the square of the kVp. It should be noted that the average X-ray energy is only about 1/3–1/2 of the peak kilovoltage. The very low energy X-rays do not contribute to useful image information and are filtered out before the beam enters the patient. Additionally, superimposed on the continuous bremsstrahlung X-ray spectrum are discrete peaks (characteristic radiation) that happen when electrons transition from one energy shell to another one inside the target material after they were knocked out of their shell by the incoming energetic electrons. The manipulation of this spectrum is important in dual-energy CT (DECT). Principles of DECT are presented in Chap. 26.

A direct current (DC) must be used to accelerate the electrons in the tube housing; an alternating current (AC) would result in electrons being accelerated back and forth in the tube, without any useful output. There are additional circuits in an X-ray tube (not indicated in Fig. 4.1), for example, to provide power to focusing cups for the electrons before they are accelerated, so that they travel in a more focused beam towards the anode. The tungsten filament also has its own circuit, and the number of electrons that are emitted can be increased by increasing the filament current, which results in more heat in the filament, and thus more emitted electrons. When an operator adjusts the mA on the control console, this refers to the number of electrons that are accelerated in the X-ray tube, i.e. the beam current. Increasing the mA on the console increases the current to the filament, which means more electrons are emitted and accelerated. Increasing the mA does not affect the beam energy (or beam quality) but does increase the number of X-rays (quantity). Dose is linearly proportional to the beam current or the current-time product (mAs).

CT scanner X-ray tubes have substantially higher requirements than ordinary X-ray tubes because they do not just take a single image, but a series of many projection images. One way in which the additional heat that is generated is dissipated, is by situating the tungsten target inside the X-ray tube on a rotating anode disk, which rotates at a few thousand revolutions per minute. If a rotating anode has a focal track radius of 5 cm and a 1 mm track width, then the annular area that the electrons hit is 314 times ($2\pi r$) larger than that of a fixed anode with a focal spot of 1 mm × 1 mm, resulting in substantially improved heat loading.

Of course, the X-ray tube must be shielded. X-rays are uncharged photons, and it is thus not possible to steer or focus these. The only way to stop them from going where they shouldn't is to either shield or collimate them. An X-ray tube has an exit window that allows X-rays to escape in a particular direction: the rest of the tube is shielded. For a more detailed description of X-ray tubes and computed tomography, the reader is referred to [1–3].

4.2.2 What Are We Imaging?

A planar X-ray of a patient is, in essence, a representation of how the X-ray beam was attenuated through the body of a patient. This is shown visually for a single line of response, also known as a ray, in Fig. 4.3. The initial beam intensity I_0 gets reduced as the beam penetrates through the patient. An attenuated beam intensity I_x arrives at the detector. In mathematical terms, this is given by $I_x = I_0 \times e^{-\mu \times x}$, where x is the thickness of the patient and μ (Greek letter "mu") represents the linear attenuation coefficient.

If we now superimpose a matrix over a patient (Fig. 4.4), then the total attenuation along the line of response will be the sum of each pixel attenuation contribution along the same line of response. This is given by:

$$I_x = I_0 \times e^{-(\mu 1+\mu 2+\mu 3+\cdots+\mu n)\times d},$$

where d is the pixel dimension and μ_1 to μ_n is the linear attenuation coefficient of each pixel in that line of response.

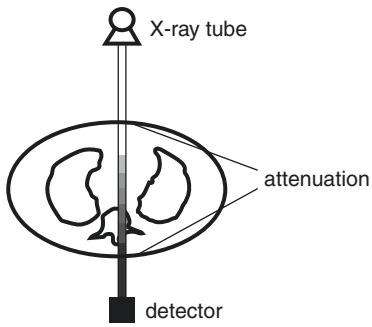

Fig. 4.3 Attenuation

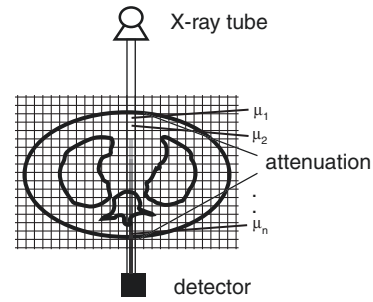

Fig. 4.4 Pixel map

For a CT image a third dimension, the slice thickness, is included as well. Therefore, each pixel on a CT image represents the average attenuation properties of the tissue in this volume element (voxel). A series of rays that pass through the patient at the same orientation is called a view or a projection. A single axial CT image may involve about 800 rays taken at 1000 different projection angles. A typical CT slice has dimensions of 512 × 512 pixels.

One of the biggest advantages of CT imaging is the reconstruction of many projections into cross-sectional images. Interestingly, the maths for this was developed in 1917 already, when Johann Radon showed that the image of a three-dimensional object can be constructed from an infinite number of two-dimensional images of the object.

4.3 Tomographic Reconstruction: Backprojection

In the previous section, the relationship between the initial and transmitted beam intensity for each ray was given by $I_x = I_0 \times e^{-\mu \times x}$. The initial intensity can be determined by doing a blank scan (as is done every morning), and I_x is measured with the patient in the beam. Therefore, the equation can be rearranged to solve for μ, the linear attenuation coefficient, the only unknown in the equation: $\mu = \frac{1}{x} \times \ln (I_0/I_x)$.

This can ultimately be done for each pixel along each line of response by tomographic reconstruction, the simplest of which is the back-

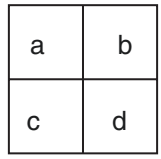

Fig. 4.5 Simple backprojection explained

projection algorithm. The following example will explain this algorithm for a 2 × 2 matrix and four projection angles. In this recipe, each pixel has a unique value (*a*, *b*, *c*, or *d*), which represents the linear attenuation coefficient for that pixel as shown in Fig. 4.5. The first part of the recipe requires the acquisition of the four projections.

4.3.1 Backprojection: An Example

The first ray of the first projection angle is shown in step 1.

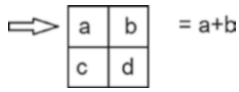

Step 1
The second ray of the first projection angle is shown in step 2.

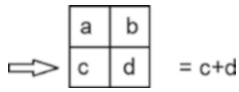

Step 2
This is the first projection angle. The second angle, after a 45° rotation, is shown next in step 3.

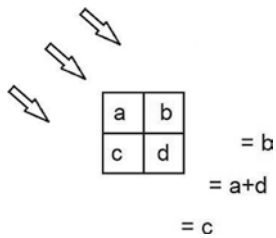

Step 3
Another 45° rotation gives us the third projection angle as shown in step 4.

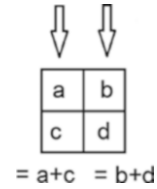

Step 4
Finally, the last projection angle for the backprojection example is shown in step 5.

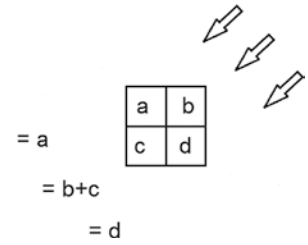

Step 5
At this point, all projections have been acquired, now the reconstruction begins, the second component of the recipe. This is done as follows: since the attenuation coefficients in each pixel are not known, they are initially all set to zero. Then each line of response must be backprojected (step 6) into the empty matrix, starting with the first line of response from step 1 above.

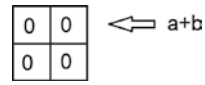

Step 6
Since the values of *a* and *b* are not known at this point and only their sum is known for this particular ray, that value is backprojected into each pixel along that line of response to give step 7—first line of response backprojected.

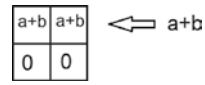

Step 7
The same is done for the next ray (second ray) from step 2, as shown in step 8.

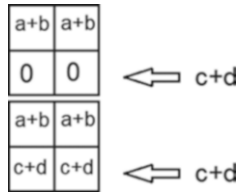

Step 8
Now the second projection from step 3 is backprojected into the result from step 8 for step 9.

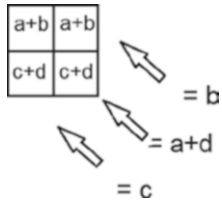

Step 9
At this point, the resultant matrix starts filling up. Step 10 is after the second projection is backprojected.

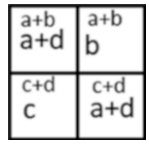

Step 10
In step 11, the third projection from step 4 is now done (backprojected).

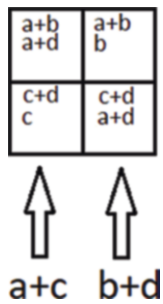

Step 11
Step 12 gives the following after three projections are backprojected.

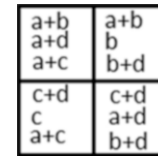

Step 12
Finally, the last projection angle is included as well in step 13 (final projection angle).

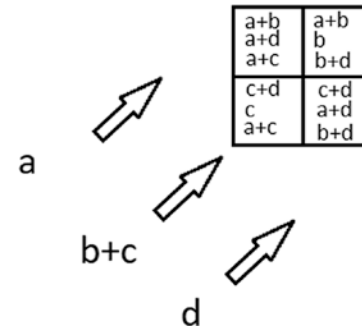

Step 13
After all projections are included in the backprojection algorithm, the following is obtained in step 14: all projections are backprojected.

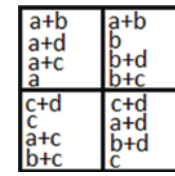

Step 14
The data in the matrix of step 14 can now be rearranged to result in step 15.

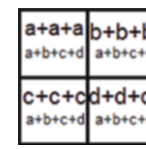

Step 15
Each pixel contains a common $a + b + c + d$. This number is known as it is the sum of all pixels in each projection and can thus be subtracted to give the below in step 16.

3a	3b
3c	3d

Step 16
Now each pixel value can be divided by 3 to give the answer in final step 17.

a	b
c	d

Step 17
This algorithm explains how to do simple backprojection for a 2 × 2 matrix with four projection angles. The reader is encouraged to follow this recipe on a piece of paper with actual numbers substituted for *a*, *b*, *c*, and *d* (say 1, 2, 3, and 4). Obviously, when a 512 × 512 matrix gets reconstructed from a thousand projections, the maths is substantially more challenging.

A simplified solution as the one presented also does not consider things like scatter. Scatter adds counts in a pixel where there should be none, thus altering the linear attenuation coefficient ever so slightly. In a backprojection algorithm, the slight alterations are also backprojected into the image matrix, resulting in noisy images. Projections can be filtered before backprojection to reduce the image noise and to give images a smoother appearance. This is known as filtered backprojection. Unfortunately, smoothing an image also means that fine detail is potentially lost, and therefore the correct filter has to be found for the task at hand. For example, a bony filter used during reconstruction will generally give a higher spatial resolution, but the images also tend to be noisier than images that are reconstructed with a soft tissue filter.

A similar argument can be made for metal artefacts, where the very high CT numbers are also backprojected along the line of response, resulting in streaks across an image (see Figs. 7.4a(i) and 12.12a(ii)). These are not easily filtered out and their removal often requires advanced software solutions.

4.4 Tomographic Reconstruction: Iterative Reconstruction

A more advanced family of reconstruction algorithms are the iterative reconstruction algorithms. The principle of iterative reconstruction is to find a solution by successively estimating an image that gave the original projections. Like in the backprojection algorithm, the first step consists of acquiring the projections. However, unlike backprojection, where the projections are backprojected to form the tomographic image, in iterative reconstruction an estimate is made of what the tomographic image could look like, and projections of this image are made at the same angles of the original acquisition. Then the projections of the original image are compared to the projections of the estimate. The result of this comparison is used to modify the current estimate, thereby creating a new estimate. Iterative reconstruction algorithms differ in the way the measured and estimated projections are compared and what kind of correction is applied to the current estimate. The recipe for iterative reconstruction is as follows.

1. Make a first estimate of the slice—this can just be a homogeneous image.
2. Forward-project the estimated slice into projections analogous to those measured by the CT scanner. In this step, physical corrections can be introduced.
3. Compare the projections of the estimate with the measured projections and apply the correction factor.
4. Decide whether to stop or continue. If the difference between iterations is small, or the maximum number of iterations is reached, then stop. Else,

5. Apply the correction to the estimate. A simple correction factor would be to add or subtract half the difference of each pixel value to the first estimate, before forward projecting again.
6. Repeat from step 2.

A simple example analogous to the filtered backprojection example may help to explain this. Consider the same 2 × 2 matrix, but only two projection angles will be used in this example. For the sake of simplicity, only the top left pixel containing "*a*" will be considered in the example.

The first ray will be the same as in the backprojection example above.

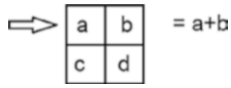

The second ray is shown next:

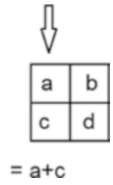

The first estimate of the solution can consist of a uniform image, denoted by "*x*". Let *x* be the average pixel count $(a + b + c + d)/4$.

The analogous projections of the first estimate are shown below.

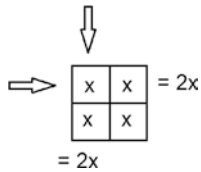

Now $2x$ is compared to $a + b$ for the one ray and to $a + c$ for the other ray. The first correction factor that must be applied to the top left pixel of the estimate can be based on the differences between the original and estimated projections. Since there are two pixels in each row and column, the differences will be halved to prevent additional counts being added into the image. Correction factor = $(a + b - 2x)/2 + (a + c - 2x)/2$.

Therefore, the top left pixel of the updated estimate now becomes x + correction factor = $x + (a + b - 2x)/2 + (a + c - 2x)/2$. A similar argument must be followed for all other pixels, before proceeding to the next iteration. Imagine the original matrix has the following values.

According to this recipe for iterative reconstruction, the following is found for the top left pixel.

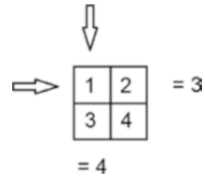

Let us assume the first estimate is $(a + b + c + d)/4 = 2.5$ in each pixel. Then the forward projection of the first estimate along the same rays gives the below.

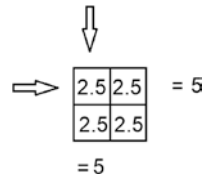

According to the recipe, the first correction factor is then given by: $(a + b - 2x)/2 + (a + c - 2x)/2 = (3–5)/2 + (4–5)/2 = -1 - 0.5 = -1.5$. The top left pixel now becomes $2.5 - 1.5 = 1$ for the next estimate.

Top left pixel done

If one follows this recipe for the other four pixels, then, for this particular example, the solu-

tion converges after a single iteration. Even from this example, it is clear that iterative reconstruction techniques are computationally intensive. However, they offer more versatility than back-projection and allow reconstruction of noisier projections, which means that substantial dose savings are possible.

Since many slices, and not just a single one, are reconstructed, a three-dimensional (3D) data set is obtained. This data set is usually viewed in transverse, sagittal or coronal viewing planes. However, it is also possible to cut the dataset along any viewing plane through software manipulation.

4.5 CT Numbers

CT images that are used clinically are typically displayed in a grayscale, and each pixel's grayscale is a representation of the CT number (or Hounsfield unit) of that voxel. CT numbers are rescaled linear attenuation coefficients to give a CT number of zero for water.

$$\mathrm{CT}(x,y) = 1000 \times \left(\frac{\mu(x,y) - \mu(\mathrm{water})}{\mu(\mathrm{water})} \right).$$

Since the attenuation of X-rays in air is almost zero (μ(air) $\approx$ 0), the CT number for air is very close to −1000.

4.6 Multi-Slice CT

Multiple detector arrays allow for simultaneous data acquisition along several slices, and slice thickness is now determined by the detector size and no longer by the beam collimation. Detectors can be electronically binned to give different slice thicknesses as desired by the operator. Figure 4.6a is an example of a single detector and single slice CT. Figure 4.6b shows a CT with multiple detectors and slices. Multi-slice acquisitions mean that much larger regions can be scanned in a single gantry rotation. The largest number of true slices by a commercial scanner is 320 slices.

The pitch is a parameter that comes into play when helical scan protocols are used. It is defined as the table increment per gantry rotation, divided by the beam collimation width (collimator pitch), or divided by the detector width (detector pitch). A pitch of less than one implies overlapping slices and thus a higher patient dose.

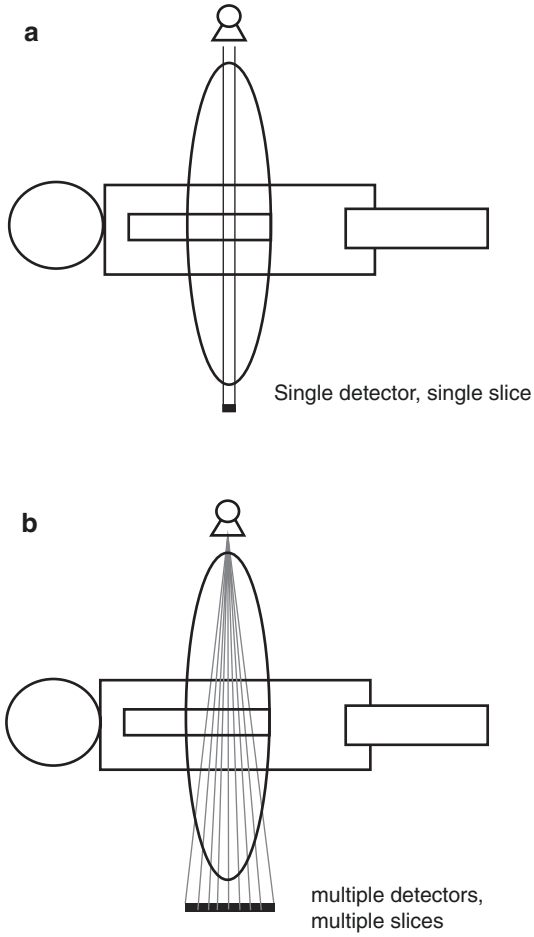

Fig. 4.6 (**a**) Single CT detector and single slice. (**b**) Multiple CT detectors and slices

4.7 Other Considerations

Modern CT scanners typically use a fan-beam geometry, rather than a parallel beam geometry, and reconstruction algorithms are amended accordingly. The introduction of slip-ring technology allowed for continuous gantry rotation without wires getting tangled up around the gantry. This in turn allowed for data acquisition while the CT table was moving, resulting in helical acquisitions patterns.

Surface rendering is a software tool that identifies voxels on the edge of a structure, and then displays these voxels. This approach is useful to examine tubular structures, like the colon.

Dual-energy CT and photon counting CT are covered in Chap. 26.

4.8 Hybrid Imaging

Hybrid imaging refers to the fusion of two or more imaging modalities. The idea behind this is that the second modality will add additional useful diagnostic information, thus making hybrid imaging more powerful than each imaging modality on its own. Existing hybrid imaging modalities include the combination of positron emission tomography (PET) with CT or magnetic resonance imaging (MRI), single photon emission computed tomography (SPECT) with CT or MRI, ultrasound with CT or MRI, or CT and MRI. Of course, not all combinations are useful in all settings.

Most modern PET scanners are integrated with a CT scanner, or, to a lesser degree, with an MRI scanner. These systems can acquire PET images, along with spatially registered CT or MRI images within a short time frame. The nuclear medicine imaging modalities (PET or

SPECT) image radiotracers that are designed for functional imaging, often showing increased (or decreased) metabolic activity, but not necessarily showing the anatomical position of the uptake very well. Additionally, the spatial resolution of nuclear medicine imaging equipment is poor compared to CT or MRI scanners (spatial resolution approaching 1 cm vs sub-millimetre imaging capabilities). However, the correlation of functional images with anatomical images offers a powerful diagnostic tool. The integration of two imaging systems allows for the acquisition of the two scans in quick succession, increasing the probability of the data to be aligned spatially and temporally. Particularly for PET-CT and SPECT-CT, the CT dataset can also be used to do the attenuation correction for the PET or SPECT component of the study. In addition, a single imaging session is logistically more efficient for the patient [4].

PET-CT colonography using fluorine-18-fluoro-deoxyglucose (F-18 FDG) is useful in the accurate anatomic correlation of both malignant and premalignant lesions [5, 6]. Additionally, PET-CTC offers accurate whole-body tumour staging, with integrated morphological and functional images, but at the cost of a higher radiation dose [7].

The integration of PET or SPECT with MRI has also been done. MRI offers better soft tissue contrast than CT, making it more useful in some clinical situations. Although commercial solutions are available, the very strong magnetic field remains a challenge in the integration process with the nuclear medicine imaging modalities. In addition, MRI imaging takes more time than a CT scan (minutes or tens of minutes vs a few seconds), and the systems are very expensive [8]. However, there is a growing field of clinical indications for PET-MRI.

Ultrasound can also be fused with CT data. Stang et al. [9] found that this technique may be useful in staging patients with colorectal cancer, but no literature was found that specifically described ultrasound and CT colonography image fusion.

The value of fusing MRI images was CT data in the diagnosis and staging of colon cancer has been described [10] although this did not specifically refer to CTC.

Generally, the baseline dataset is the higher resolution and anatomically correct CT or MRI dataset. The fusion of two imaging modalities requires the co-registration of the datasets. For a combined imaging modality, like for example a PET-CT scanner, the distance between the imaging planes between the two modalities is known, and fusion can occur by just moving the one dataset that known distance on top of the other one. If the patient did not move and there were no internal organ motions, then co-registration is usually acceptable. Differences in pixel size must be considered.

The co-registration can also be done manually by the operator, who matches points or landmarks that are visible on both datasets, to place the image sets on top of each other. This could include the use of fiducial markers that are visible on both datasets [11, 12]. There are also several automatic or semi-automatic algorithms that can do this using various techniques like edge-detection, gradient methods, voxel similarity, surface matching, intensity differences, standard deviation of voxel differences, or similar [12]. Mutual information methods compare the amount of information that one image contains about the other and when this number if a maximum, then the images are correctly aligned [13]. These techniques work well for “rigid-body” transformations.

Challenges for image fusion include patient movement, tissue deformation, or patient positioning. Deformable image registration allows the re-calibration and modification of images (or parts of images) to allow for better fusion of two datasets where the alignment is not ideal [11]. Deformable algorithms vary significantly in complexity. Affine transformations preserve straight lines and distances between them. Non-affine transformations (image warping) are based on theoretical models that describe tissue properties, like flexibility or rigidity [12].

However, warping techniques can easily hide small features (like for example small polyps) and thus only offer a limited use in medical applications.

Electromagnetic navigation can be used in ultrasound image matching, where sensor coils are tracked in real-time for the orientation and position of the ultrasound probe, and the signal is re-formatted to match previously acquired CT, MRI, or PET-CT data.

Techniques such as a sliding window technique or checkerboard images can be used to view a set of fused images as a single image. The fused image should have more complete information than either of the two input images and help in clinical diagnosis.

Key Messages

- Computed tomography utilises a set of two-dimensional projections taken at many angles around a patient to reconstruct a three-dimensional dataset.
- Tomographic reconstruction is usually done through (filtered) backprojection or iterative reconstruction.
- A CT scan is in essence an attenuation map of the human body. The CT numbers are rescaled attenuation coefficients, relative to water.
- Advances in CT, such as multi-slice CT, helical data acquisition, and many software improvements, have expanded the number of clinical indications for which CT is useful, including CT colonography.
- Hybrid imaging refers to the synergistic fusion of two imaging modalities, with each imaging modality adding unique data for improved clinical decision making.

4.9 Summary

Computed tomography is a powerful imaging modality, allowing the reconstruction of a three-dimensional dataset from two-dimensional image data. This solves the issues of overlying or underlying structures that are present in planar radiography. One of the fields where CT advances have enabled the expansion of clinical indications is CT colonography, which allows the visualisation of the colon in a minimally non-invasive manner.

References

1. Bushberg J, Seibert A, Leidholdt E, Boone J. The essential physics of medical imaging. 2nd ed. Philadelphia: Lippincott Williams & Wilkins; 2002.
2. Mahesh M. MDCT physics: the basics: technology, image quality and radiation dose. Philadelphia: Lippincott Williams & Wilkins; 2009.
3. Dance D, Christofides S, Maidment A, McLean I, Ng K. Diagnostic radiology physics: a handbook for teachers and students. Vienna: International Atomic Energy Agency; 2014.
4. Cherry S, Sorensen J, Phelps M. Physics in nuclear medicine. 4th ed. Philadelphia: Saunders; 2012.
5. Gollub M, Akhurst T, Markowitz A, Weiser M, Guillem J, Smith L, Larson S, Margulis A. Combined CT colonography and 18F-FDG PET for colon polyps: potential technique for selective detection of cancer and precancerous lesions. Am J Roentgenol. 2007;188(1):130–8. https://doi.org/10.2214/AJR.05.1458.
6. Quon A, Napel S, Beaulieu C, Gambhir S. "Flying through" and "flying around" a PET/CT scan: pilot study and development of 3D integrated 18F-FDG PET/CT for virtual bronchoscopy and colonoscopy. J Nucl Med. 2006;47:1081–7.
7. Sun L, Wu H, Guan Y. Colonography by CT, MRI and PET/CT combined with conventional colonoscopy in colorectal cancer screening and staging. World J Gastroenterol. 2008;14(6):853–63.
8. Cal-Gonzalez J, Rausch I, Sunda L, Lassen M, Muzi O, Moser E, Papp L, Beyer T. Hybrid imaging: instrumentation and data processing. Front Phys. 2018;6:47. https://doi.org/10.3389/fphy.2018.00047.
9. Stang A, Keles H, Hentschke S, Seydewitz C, Keuchel M, Pohland C, Dahlke J, Weilert H, Wessling J, Malzfeld E. Real-time ultrasonography-computed tomography fusion imaging for staging of hepatic metastatic involvement in patients with colorectal cancer. Investig Radiol. 2010;45(8):491–501. https://doi.org/10.1097/RLI.0b013e3181ddd3da.
10. Zhao Z, Shou Y, Jiang M, Dang L. Application value of MRI combined with MSCT in diagnosis and staging of colon carcinoma. Comput Math Methods Med. 2022;2022:2593844. https://doi.org/10.1155/2022/2593844.

11. Singh B, Cazacu I, Deza C, Rigaud B, Saftiou A, Gruionu G, Guionu L, Brock K, Koay E, Herman J, Bhutani M. Image fusion involving real-time transabdominal or endoscopic ultrasound for gastrointestinal malignancies: review of current and future applications. Diagnostics. 2022;12:3218. https://doi.org/10.3390/diagnostics12123218.
12. Ewertsen C, Saftiou A, Gruionu L, Karstrup S, Nielsen M. Real-time image fusion involving diagnostic ultrasound. Am J Roentgenol. 2013;200(3):W249–55. https://doi.org/10.2214/AJR.12.8904.
13. Maes F, Loeckx D, Vandermeulen D, Suetens P. Image registration using mutual information. In: Paragios N, Duncan J, Ayache N, editors. Handbook of biomedical imaging. Boston: Springer; 2015. https://doi.org/10.1007/978-0-387-09749-7_16.