The Mathematics of Medical Imaging

by

Charles L. Epstein

Chapter 1

Measurements and modeling

A *quantitative* model of a physical system is expressed in the language of mathematics. A qualitative model often precedes a quantitative model. For many years clinicians used medical X-rays without employing a precise quantitative model. X-rays were thought of as high frequency 'light' with three very useful properties:

(1). If X-rays are incident on a human body, some fraction of the incident radiation is absorbed, though a sizable fraction is transmitted. The fraction absorbed is proportional to the total 'density' of the material encountered.

(2). A 'beam' of X-ray light travels in a straight line.

(3). X-rays darken photographic film. Taken together, these properties mean that using X-rays one could "see through" a human body to obtain a shadow or projection of the internal anatomy on a sheet of film see figure 1.1[1].

The model was adequate given the available technology. In their time, X-rays led to a revolution in the practice of medicine because they opened the door to non-destructive examination of internal anatomy. They are still useful for locating bone fractures, dental caries and foreign objects but their ability to visualize soft tissues and more detailed anatomic structure is very limited. There are several reasons for this. The X-ray image is a two dimensional projection of a three dimensional object which renders it impossible to deduce the spatial ordering in the missing third dimension. Photographic film is not very sensitive to X-rays. To get a usable image, a light emitting phosphor is sandwiched with the film. This increases the sensitivity of the overall 'detector,' but even so, large changes in the intensity of the incident X-rays still produce small differences in the density of film. This means that the contrast between different soft tissues is poor. Because of these limitations a qualitative theory was adequate for the interpretation of X-ray images.

[1]Chest x-ray provided courtesy of Dr. David S. Feigen, ENS Sherri Rudinsky and Dr. James G. Smirniotopoulos of the Uniformed Services University of the Health Sciences, Dept. of Radiology, Bethesda, MD

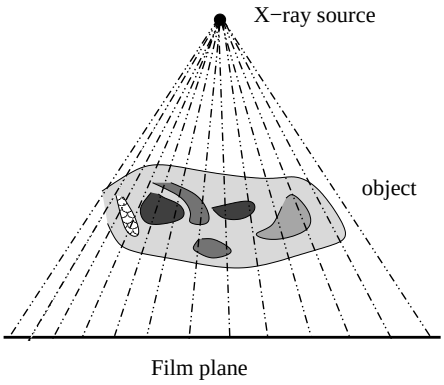

(a) Depth information is lost in a projection.

(b) A old fashioned chest X-ray.

Figure 1.1: The world of old fashioned X-rays.

A desire to improve upon this situation led Alan Cormack and Godrey Hounsfield, to independently develop X-ray *tomography* or slice imaging, see [30] and [10]. The first step in their work was to use a quantitative theory for the absorption of X-rays. Such a theory already existed and is little more than a quantitative restatement of (1) and (2). It was not needed for old fashioned X-rays because they are read "by eye," no further processing is done after the film is developed, Both Cormack and Hounsfield realized that mathematics could be used to infer 3-dimensional anatomic structure from a large collection of *different* two dimensional projections. The possibility for making this idea work relied on two technological advances: 1. The availability of scintillation crystals to use as detectors. 2. Powerful, digital computers to process the tens of thousands of measurements needed to form a usable image. A detector using a scintillation crystal is about a hundred times more sensitive than film which makes possible much finer distinctions. As millions of arithmetic operations are needed for each image, fast computers are a necessity for reconstructing an image from the available measurements. It is an interesting historical note that the mathematics underlying X-ray tomography was done in 1917 by Johan Radon, see [59]. It had been largely forgotten and both Hounsfield and Cormack worked out solutions to the problem of reconstructing an image from its projections. Indeed, this problem had arisen and been solved in contexts as diverse as radio astronomy and statistics.

This book is a detailed exploration of the mathematics which underpins the reconstruction of images in X-ray tomography. The list of mathematical topics covered is dictated by their importance and utility in medical imaging. While our emphasis is on understanding these mathematical foundations, we constantly return to the practicalities of X-ray tomography and explore the relationship of the mathematical formulation of a problem and its solution, to the realities of computation and physical measurement. There are many different imaging *modalities* in common use today, X-ray computed tomography (CT), magnetic resonance imaging (MRI), positron emission tomography (PET), ultrasound, optical imaging, impedence imaging, etc. Because each relies on a different physical principle, each provides different information. In every case the mathematics needed to process and inter-

pret the data has a large overlap with that used in X-ray CT. We concentrate on X-ray CT because of the simplicity and clarity of the physical principles underlying the measurement process. Detailed descriptions of the other modalities can be found in [39] or [4].

1.1 Mathematical modeling

Mathematics is the language in which any quantitative theory or model is eventually expressed. In this introductory chapter we consider a variety of examples of physical systems, measurement processes and the mathematical models used to describe them. These models illustrate different aspects of more complicated models used in medical imaging. The chapter concludes with a consideration of linear models.

Mathematics is used to model physical systems from the formation of the universe to the structure of the atomic nucleus, from the function of the kidney to the opinions of voters. The first step in giving a mathematical description of a "system" is to isolate that system from the universe in which it sits. While it is no doubt true that a butterfly flapping its wings in Siberia in mid-summer will effect the amount of rainfall in the Amazon rain forest a decade hence, it is surely a tiny effect, impossible to accurately quantify. To obtain a practical model such effects are ignored, though they may come back to haunt the model, as measurement error and noise. After delineating a system, we need to find a collection of numerical parameters which describe its state. In this generality these parameters are called *state variables*. In the idealized world of an isolated system the exact measurement of the state parameters would uniquely determine the state of the system. In general the natural state parameters are not directly measurable. The model then describes relations between the state variables which suggest feasible measurements with which one might determine the state of the system.

1.1.1 Finitely many degrees of freedom

See: A.1, B.5, B.6, B.7.

If the state of a system is described by a finite collection of real numbers $\mathbf{x} = (x_1, \ldots, x_n)$ then the system has finite many *degrees of freedom*. Most of the systems encountered in elementary physics and electrical engineering have this property. The mathematical model is then expressed as relations that these variables satisfy, often taking the form of functional relations,

$$\begin{array}{rcl} f_1(x_1, \ldots, x_n) & = & 0 \\ \vdots & & \vdots \\ f_m(x_1, \ldots, x_n) & = & 0 \end{array} \tag{1.1}$$

If there are more state variables than relations, that is $m < n$ then, heuristically $n - m$ measurements are needed to determine the state of the system. Of course this counting parameters arguments assumes that the relations in (1.1) are functionally independent. For linear relations the number of functionally independent equations does not depend on the state. If the relations are non-linear then counting the number of independent conditions

can be quite involved and the result generally depends on $\mathbf{x}$. For measurements to be useful they must be expressible as functions of the state variables.

Example 1.1.1. Suppose the system is a ball on a rod. The state of the system is described by (x, y), the coordinates of the ball. If the rod is of length r and one end of it is fixed at the point $(0, 0)$, then the state variables satisfy the relation

$$x^2 + y^2 = r^2. \tag{1.2}$$

Imagine now that one dimensional creatures, living on the x-axis $\{y = 0\}$ can observe a shadow of the ball, cast by very distant light sources so that the rays of light are perpendicular to the x-axis. The line creatures want to predict whether or not the ball is about to collide with their world. Locating the shadow determines the x-coordinate of the ball, using equation (1.2) gives

$$y = \pm\sqrt{r^2 - x^2}.$$

To determine the sign of the y-coordinate requires additional information not available in the model. On the other hand this information is adequate if one only wants to predict if the ball is about to collide with the x-axis. If the x-axis is illuminated by red light from above and blue light from below, then a ball approaching from below would cast of red shadow while a ball approaching from above would cast a blue shadow. With this additional data, the location of the ball is completely determined.

Example 1.1.2. We would like to find the height of a mountain without climbing it. To that end, the distance l between the point P and the top of the mountain, as well as the angle θ are measured. If l and θ are measured exactly then by a trigonometric identity, the height h of the mountain is given by $l \tan\theta$. Measurements are never exact, let us use the model

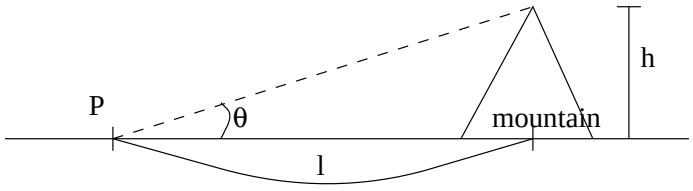

Figure 1.2: Using to trigonometry to find the height of a mountain.

to relate the error in measuring θ to the computed value of h. This requires a basic tool from calculus, the Taylor expansion. If $f(x)$ is a smooth function of the variable x then we can approximate the behavior of f near to x_0 by using the Taylor expansion

$$f(x) = f(x_0) + f'(x_0)(x - x_0) + \frac{1}{2}f''(x_0)(x - x_0)^2 + \cdots + \frac{1}{n!}f^{[n]}(x_0)(x - x_0)^n + R_n(x - x_0).$$

Here R_n is called the n^{th}-remainder, it is a term that goes to zero faster than $|x - x_0|^n$. The complicated function $\tan\theta$ can be replaced by a much simpler function, though at the expense of replacing an exact (but not very useful) formula with an approximate (but more usable) formula.

Suppose that we measure $\theta + \Delta\theta$ where θ is the exact angle. The error in h as a function of $\Delta\theta$ is given approximately by

$$\begin{aligned} \tan(\theta + \Delta\theta) &= \tan(\theta) + \partial_\theta \tan(\theta)\Delta\theta + O(\Delta\theta^2) \\ &= \tan(\theta) + \sec^2(\theta)\Delta\theta + O(\Delta\theta^2). \end{aligned} \tag{1.3}$$

The notation $O(\Delta\theta^2)$ refers to an error term which is bounded by a constant times $\Delta\theta^2$, as $\Delta\theta$ goes to zero. The height predicted from the measurement of the angle is

$$h_m = l\tan(\theta+\Delta\theta) = l(\tan\theta + \frac{\Delta\theta}{\cos^2\theta} + O(\Delta\theta^2)).$$

Disregarding the quadratic error terms $O(\Delta\theta^2)$, the *absolute error* is

$$h_m - h \approx l\frac{\Delta\theta}{\cos^2\theta}.$$

The absolute error is a number with the same units as h; in general it is not a very interesting quantity. If, for example the true measurement is $10,000$m then an error of size 1m would not be too significant whereas if the true measurement is 2m then it would. To avoid this obvious pitfall one normally considers the *relative error*. In this problem the relative error is

$$\frac{h_m - h}{h} = \frac{\Delta\theta}{\cos^2\theta\tan\theta} = \frac{\Delta\theta}{\sin\theta\cos\theta}.$$

Generally the relative error is the absolute error divided by the correct value. It is a dimensionless quantity that gives a quantitative assessment of the accuracy of the measurement. If the angle θ is measured from a point too near to or too far from the mountain, i.e. θ is very close to 0 or $\pi/2$ then small measurement errors result in a substantial loss of accuracy. A useful feature of a precise mathematical model is the possibility of estimating how errors in measurement affect the accuracy of the parameters we wish to determine.

Example 1.1.3. In a real situation we cannot measure the distance l either. Suppose that we measure the angle from two different positions i.e. θ_1 and θ_2 as in the figure below. Then we have $\tan\theta_1 = h/l_1, \tan\theta_2 = h/(l_1+l_2)$, for the same h. Using trigonometry we

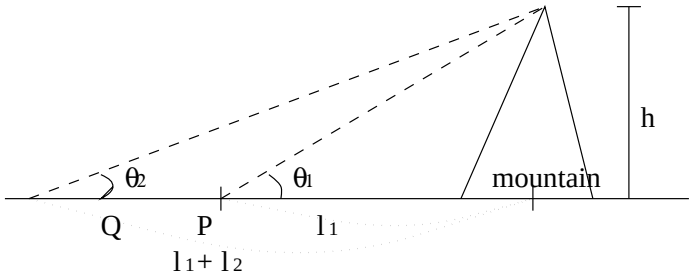

Figure 1.3: A more realistic measurement.

deduce

$$\begin{aligned} l_1 &= \frac{l_2}{\tan\theta_1/\tan\theta_2 - 1}, \\ h &= (l_1+l_2)\tan\theta_2 = \left(\frac{l_2}{\tan\theta_1/\tan\theta_2 - 1} + l_2\right)\tan\theta_2 = l_2\frac{\sin\theta_1\sin\theta_2}{\sin(\theta_1-\theta_2)}. \end{aligned}$$

Assuming that l_2, the distance between P and Q can also be measured, then h can be determined from θ_1 and θ_2. We complete the discussion of this example by listing different ways that this model may fail to capture important features of the physical situation.

- If the shape of a mountain looks like that in figure 1.4and we measure the distance and angle at the point P, we are certainly not finding the real height of the mountain. Some *a priori* information is always incorporated in a mathematical model.

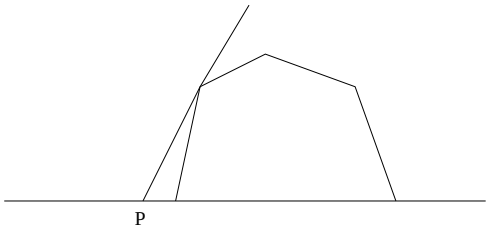

Figure 1.4: Not exactly what we predicted!

- The curvature of the earth is ignored. A more sophisticated geometric model is needed to correct for such errors. This becomes a significant problem as soon as the distances, l, l_1, l_2 are large compared to the distance to the horizon (about 25km for a 2 meter tall person). The approximations used in the model must be adapted to the actual physical conditions of the measurements.

- The geometry of the underlying measurements could be very different from the simple Euclidean geometry used in the model. To measure the angles θ_1, θ_2 one would normally use a transit to sight the peak of the mountain. If the mountain is far away then the light travels on a path from the mountain to the transit which passes through air of varying density. The light is refracted by the air and therefore the ray path is not the straight line assumed in the model. To include this effect would vastly complicate the model. This is an important consideration in the very similar problem of creating a map of the sky from earth based observations of stars.

  Analogous problems arise in medical imaging. If the wavelength of the energy used to probe the human anatomy is very small compared to the size of the structures that are present then it is reasonable to assume that the waves are not refracted by the medium through which they pass, i.e. X-rays can be assumed to travel along straight lines. However for energies with wavelengths comparable to the size of structures present in the human anatomy, this assumption is simply wrong. The waves are then bent and diffracted by the medium and the difficulty of modeling the ray paths is considerable and, in fact largely unsolved! This is an important issue in ultrasound imaging.

Example 1.1.4. Refraction provides another example of a simple physical system. Suppose that we have two fluids in a tank as shown in the figure and would like to determine the height of the interface between them. Suppose first of all that the refractive indices of the fluids are known. Let n_1 be the refractive index of the upper fluid and n_2 the refractive index of the lower one, Snell's law states that

$$\frac{\sin(\theta_1)}{\sin(\theta_2)} = \frac{n_2}{n_1}.$$

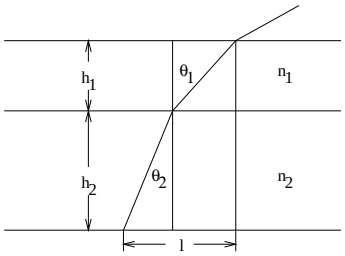

Figure 1.5: Using refraction to determine the height of an interface.

Let h denote the total height of the fluid, then

$$h_1 + h_2 = h.$$

The measurement we make is the total displacement l, of the light ray as it passes through the fluids. It satisfies the relationship

$$h_1 \tan(\theta_1) + h_2 \tan(\theta_2) = l.$$

Using these three relations h_1 and h_2 are easily determined. The assumption that we know n_1 implies, by Snell's law that we can determine θ_1 from a measurement of the angle of the light ray above the fluid. If n_2 is also known, then using these observations we can determine θ_2 as well:

$$\sin(\theta_2) = \frac{n_1}{n_2} \sin(\theta_1).$$

The pair (h_1, h_2) satisfies the 2×2-linear system

$$\begin{pmatrix} 1 & 1 \\ \tan(\theta_1) & \tan(\theta_2) \end{pmatrix} \begin{pmatrix} h_1 \\ h_2 \end{pmatrix} = \begin{pmatrix} h \\ l \end{pmatrix}. \tag{1.4}$$

In example 1.3.2 we consider a slightly more realistic situation where the refractive index of the lower fluid in not known. By using more measurements n_2 can also be determined, though the equation determining n_2 is no longer linear.

Exercise 1.1.1. Suppose that in example 1.1.1 light sources are located at $(0, \pm R)$. What is the relationship between the x-coordinate and the shadow?

Exercise 1.1.2. In example 1.1.3 work out how the absolute and relative errors depend on θ_1, θ_2 and l_2.

1.1.2 Infinitely many degrees of freedom

See: A.3, A.6.

In the previous section we examined some simple physical systems described by a finite collection of numbers. Such systems are said to have finitely many degrees of freedom. In

these examples, the problem of determining the state of the system from feasible measurements reduces to solving systems of finitely many equations in finitely many unknowns. In imaging applications the state of a system is usually described by a function or functions of continuous variables. These are systems with infinitely many degrees of freedom. In this section we consider several examples of this type.

Example 1.1.5. Suppose that we would like to determine the shape of a planar object, D that cannot be seen. The object is lying inside a disk and we can fire particles at the object which bounce off. Assume that this scattering process is very simple: each particle strikes the object once and is then scattered along a straight line off to infinity. The outline of the object can be determined by knowing the correspondence between incoming lines, l_{in} and outgoing lines, l_{out}. Each intersection point $l_{\text{in}} \cap l_{\text{out}}$ lies on the boundary of the object. Measuring $\{l^j_{\text{out}}\}$ for finitely many incoming directions $\{l^j_{\text{in}}\}$ determines finitely many points $\{l^j_{\text{in}} \cap l^j_{\text{out}}\}$ on the boundary of D. In order to use this finite collection of points to make any assertions about the rest of the boundary of D, more information is required. If we know that D consists of a single piece or component then these points would lie on a single closed curve, though it might be difficult to decide in what order they should appear on the curve.

Convex obstacles is a class which satisfy these simple assumptions and for which a finite number of points on the boundary might carry a lot of useful information. A region D in the plane is convex if it has the following property: for each pair of points p and q lying in D the line segment $\overline{pq}$ is also contained in D. An equivalent condition is the following: for each point p on the boundary of D there is a line l_p which passes through p but is otherwise disjoint from D. This line is called a *support line* through p. If the boundary is smooth at p then the tangent line to the boundary is the unique support line.

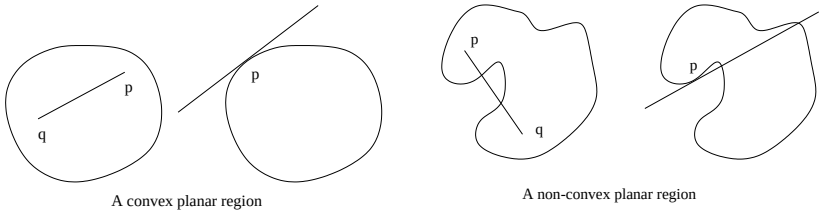

(b)

Figure 1.6: Convex and non-convex regions.

If the object is convex and more is known about the scattering process, for example if the angle of incidence is equal to the angle of reflection, then from a finite number of incoming and outgoing pairs, $\{(l^i_{\text{in}}, l^i_{\text{out}}) \quad : i = 1, \ldots, N\}$ we can determine an approximation to D with an estimate for the error. The intersection points, $\{l^i_{\text{in}} \cap l^i_{\text{out}}\}$ lie on the boundary of the convex region, D. If we use these points as the vertices of a polygon, P^{in}_N then the first convexity condition implies that P^{in}_N is completely contained within D. On the other hand, as the angle of incidence equals the angle of reflection we can also determine the tangent lines to the boundary of D at the points of intersection. A line divides the plane into two half planes, since D is convex, it lies entirely in one of the half plane determined by each of

its tangent lines. By intersecting the half planes defined by the tangent lines through the points $\{(l^i_{\text{in}}, l^i_{\text{out}})\}$ we obtain another convex polygon, P^{out}_N which contains D. Thus with these N-measurements we obtain the both an *inner* and *outer* approximation to D :

$$P^{\text{in}}_N \subset D \subset P^{\text{out}}_N.$$

It is clear that the boundary of a convex region has *infinitely many degrees of freedom* as it is conveniently described as the image of a map $s \mapsto (x(s), y(s))$, where s lies in an interval $[0, 1]$. On the other hand the images of such maps can be approximated by polygons. Once the number of sides is fixed, then we are again considering a system with finitely many degrees of freedom. In all practical problems, a system with infinitely many degrees of freedom must eventually be approximated by a system with finitely many degrees of freedom.

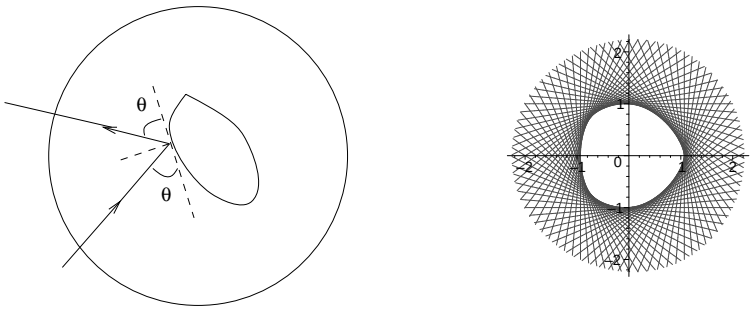

(a) The angle of incidence equals the angle of reflection.

(b) The outer approximation as an intersection of half spaces.

Figure 1.7: Using particle scattering to determine the boundary of a convex region.

Remark 1.1.1. For a non-convex body the above method does not work as the correspondence between incoming and outgoing lines can be quite complicated: some incoming lines may undergo multiple reflections before escaping, in fact some lines might become permanently trapped.

Example 1.1.6. Suppose that the *surface* of a sea is mapped by coordinates (x, y) belonging to a region $D \subset \mathbb{R}^2$. The depth of the bottom of the sea is described by a function $h(x, y)$. One way to determine h would be to drop a weighted string until it hits the bottom. There are problems with this method: 1. It is difficult to tell when the weight hits the bottom. 2. Unknown, underwater currents may carry the string so that it does not go straight down. A somewhat less direct approach would be to use sonar to measure the distance to the bottom. The physical principle underlying the measurement is that the speed of sound is determined by the density and temperature of the water which are in turn determined by the depth. Let $c(z)$ denote the *known* speed of sound, as a function of the depth. A speaker underneath the boat emits a loud, short pulse of sound and the time it takes for the sound to return is measured. Here we assume that the sound travels in a straight line

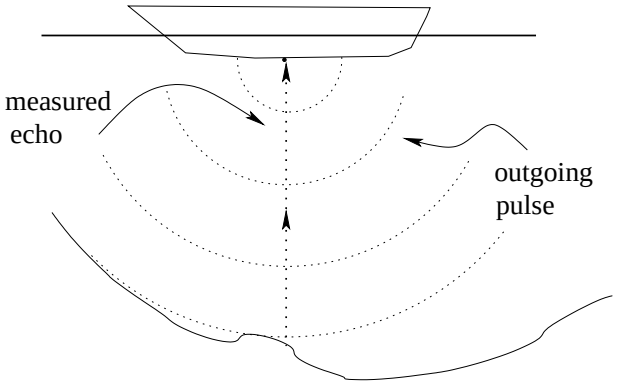

Figure 1.8: Using sound to measure depth.

to the bottom and the microphone only detects the direct reflection, traveling back along the straight line. Using $c(z)$ the transit time can be related to the depth.

A simple model, valid for shallow seas, is that the speed of sound is a constant, c. The measurement, T is the time it takes for the sound pulse to go down and back,

$$2h = cT.$$

This assumes that the boat is stationary from the time the pulse is emitted until the return is received. With such a measurement, $T(x, y)$ for each position $(x, y) \in D$, the depth is determined, everywhere by

$$h(x, y) = \frac{cT(x, y)}{2}.$$

In reality such continuous measurements are not possible. Instead the boat is placed at a finite set of locations $P = \{(x_j, y_j) \, : \, j = 1, \ldots, N\}$ and $T(x_j, y_j)$ is measured. The finite set of values

$$h(x_j, y_j) = \frac{cT(x_j, y_j)}{2}$$

are then determined.

Again, what use is a finite set of values? Without qualitative, *a priori* information about the nature of the function h, this finite data set is indeed useless! On the other hand it is reasonable to assume that h is a *continuous* function of (x, y). With this assumption, values of h for points **not** in P can be *interpolated* from the measured values. The minimum necessary separation between the points in P is determined by a quantitative assessment of how continuous h is expected to be. Suppose it is known that there is a constant M so that

$$|h(x, y) - h(x', y')| \leq M\sqrt{(x - x')^2 + (y - y')^2}.$$

If every point (x, y) is within d of a point (x_j, y_j) in P then we have the estimate

$$|h(x, y) - h(x_j, y_j)| \leq Md.$$

This then gives an estimate for the accuracy of the interpolated values. A small value of M indicates that the depth is varying slowly, while a large value indicates rapid variations. In the former case a larger value of d provides acceptable results, while in the latter case a smaller value of d is needed to get an accurate picture of the bottom.

Example 1.1.7. Now assume that the sea, in the previous example is one dimensional, but that the sound speed is not constant. To use the measurements described above to determine the depth $h(x)$ requires more mathematical apparatus. Let $z(t)$ denote the depth of the sound pulse at a time t after it is emitted. Using calculus we can express the assertion that the 'speed of sound at depth z is $c(z)$' as a differential equation

$$\frac{dz}{dt}(t) = c(z(t)). \tag{1.5}$$

Formally this is equivalent to

$$\frac{dz}{c(z)} = dt.$$

The transit time T is a function of the depth, h integrating this equation gives

$$G(h) \stackrel{d}{=} \int_0^h \frac{dz}{c(z)} = \frac{T(h)}{2}. \tag{1.6}$$

The function G is monotonely increasing and therefore its inverse is well defined. Using G^{-1} we can determine the depth, h from the available measurement, T

$$h = G^{-1}\left(\frac{T}{2}\right).$$

To use this model, the function G^{-1} needs to be explicitly determined. If $c(z)$ is simple enough then an analytic formula for G might be available. Otherwise the integral defining G is computed for a finite collection of depths $\{h_1, \ldots, h_m\}$, with $t_i = G(h_i)$. From this table of values, the inverse function is also known for a finite collection of times

$$h_i = G^{-1}(t_i).$$

If $c(z)$ is a differentiable function, then a linear approximation of the form

$$c(z) \approx c + az$$

is valid for small values of z. Integrating gives

$$G(h) \approx \log(1 + \frac{ah}{c}),$$

solving for $G^{-1}(T)$ we find

$$h(T) \approx \frac{c}{a}(e^{a\frac{T}{2}} - 1).$$

Using Taylor's formula for e^x gives

$$h(T) \approx c\frac{T}{2} + \frac{caT^2}{8} + O(T^3).$$

Here as usual $O(T^3)$ is an error term which goes to zero, as T goes to zero, at the same rate as T^3. This agrees, to leading order with the previous computation.

Example 1.1.8. The one dimensional model in the previous example can be used to solve the two dimensional problem. Suppose that the area we are interested in mapping corresponds to the rectangle $[-1,1]\times[-1,1]$ in the (x,y)-map coordinates. For each y define the function of one variable

$$h_y(x) \stackrel{d}{=} h(x,y).$$

Knowing the collection of functions $\{h_y(x) : y \in [-1,1]\}$ for $x \in [-1,1]$ is evidently exactly the same thing as a knowing $h(x,y)$, for $(x,y) \in [-1,1]\times[-1,1]$. Because the measuring apparatus only observes the sound returning on the straight line from the boat to the bottom of the sea, the analysis in the previous example applies to allow the determination of $h_y(x)$ from measurements of $T_y(x)$,

$$h(x,y) = h_y(x) = G^{-1}\left(\frac{T_y(x)}{2}\right).$$

In this way a two dimensional problem is sliced into simpler one dimensional problems. In real applications, only finitely many measurements are made. A typical strategy if to pick an equally spaced set of y-values,

$$y_k = \frac{k}{N}, \quad k = -N,\ldots,N$$

and determine $h_{y_k}(x_j)$ at finitely many, equally spaced x-values

$$x_j = \frac{j}{N} \quad j = -N,\ldots,N.$$

These examples capture many of the features that we will encounter in X-ray tomography: by using a mathematical model for the measurements, an inaccessible, physical quantity can be determined using feasible measurements. The model is itself an approximation, but is subject to improvements.

Exercise 1.1.3. Describe parameters to describe the set of polygons with n-vertices in the plane. For the case of triangles, find the relations satisfied by your parameters. Find a condition, in terms of your parameters implying that the polygon is convex.

Exercise 1.1.4. Find an example of a planar region such that at least one particle trajectory is trapped forever.

Exercise 1.1.5. Why is G a monotonely increasing function?

Exercise 1.1.6. Suppose that $c(z)$ is piecewise constant, so that

$$c(z)\begin{cases} c_1 & \text{if } 0 \le z \le z_1, \\ c_2 & \text{if } z_1 < z. \end{cases}$$

Find G and G^{-1}.

Exercise 1.1.7. Why is it reasonable to model $c(z)$ as a linear function under the assumption that it is a differentiable function? Suggest a method for determining a.

Exercise 1.1.8. In the examples above it is assumed that all returns not arriving on the straight line path from the bottom of the ocean are ignored. Analyze the problems that result if return signals are accepted from all directions. What impact would this have on using the slicing method to reduce the dimensionality of the problem?

Exercise 1.1.9. Repeat the analysis in example 1.1.7 assuming that the boat is traveling at constant velocity v. Continue assuming that only returns meeting the bottom of the boat at right angles are detected.

1.2 A simple model problem for image reconstruction

The problem of image reconstruction in X-ray tomography is sometimes described as reconstructing an object from its "projections." Of course these are projections under the illumination of X-ray "light." In this section we consider the analogous, but simpler problem, of determining the outline of an object from its shadows. As is also the case in medical applications, we consider a two dimensional problem. Let D be the convex region in the plane. Imagine that a light source is placed very far away from the body. Since the light source is very far away, the rays of light are all traveling in essentially the same direction. We can think of the rays of light as a collection of parallel lines. We want to measure the shadow that D casts for each position of the light source. To describe the measurements imagine that a screen is placed on the "other side" of D perpendicular to the direction of the light rays, see the figure below. The screen is the detector, in a real apparatus sensors would be placed on the screen, allowing us to determine where the shadow begins and ends.

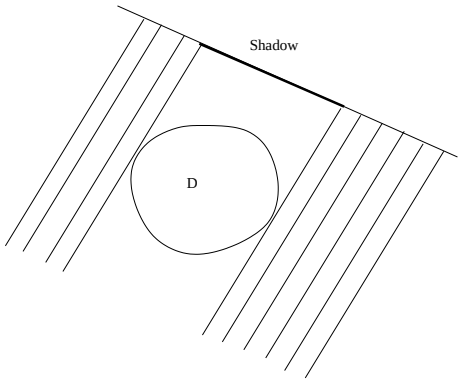

Figure 1.9: The shadow of a convex region

The region, D blocks a certain collection of light rays and allows the rest to pass. Measuring the shadow is therefore determining the "first" and "last" lines in this family of parallel lines to intersect D. To completely describe the object we need to rotate the source and detector through 180°, measuring, at each angle, where the shadow begins and ends.

The first and last lines to intersect a region just meet it along its boundary. These lines are therefore tangent to the boundary of D. The problem of reconstructing a region from its shadows is mathematically the same as the problem of reconstructing a region from a knowledge of the tangent lines to its boundary. As a first step in this direction we need a

good way to organize our measurements. To that end we give a description for the *space of all lines in the plane.*

1.2.1 The space of lines in the plane

A line in the plane is the set of points which satisfy an equation of the form

$$ax + by = c$$

where $a^2 + b^2 \neq 0$. We get the same set of points if we replace this equation by

$$\frac{a}{\sqrt{a^2+b^2}}x + \frac{b}{\sqrt{a^2+b^2}}y = \frac{c}{\sqrt{a^2+b^2}}.$$

The coefficients, $(\frac{a}{\sqrt{a^2+b^2}}, \frac{b}{\sqrt{a^2+b^2}})$ define a point ω on the unit circle and the constant $\frac{c}{\sqrt{a^2+b^2}}$ can be any number. The lines in the plane are parametrized by a pair consisting of a unit vector, ω and a real number t. The line $l_{t,\omega}$ is the set of points satisfying the equation

$$(x, y) \cdot \omega = t.$$

Here $(x, y) \cdot \omega$ is the *dot*-product

$$(x, y) \cdot \omega = x\omega_1 + y\omega_2$$

where $\omega = (\omega_1, \omega_2)$. As the set of points satisfying this equation is unchanged if (t, ω) is replaced by $(-t, -\omega)$ it follows that, as sets, $l_{t,\omega} = l_{-t,-\omega}$.

Very often it is convenient to parametrize the points on the unit circle by a real number, to that end we set

$$\omega(\theta) = (\cos(\theta), \sin(\theta)). \tag{1.7}$$

Since cos and sin are 2π-periodic it clear that $\omega(\theta)$ and $\omega(\theta + 2\pi)$ are the same point on the unit circle. Using this notation the line $l_{t,\theta} = l_{t,\omega(\theta)}$ is the set of solutions to the equation

$$\cos(\theta)x + \sin(\theta)y = t.$$

Both notations are used in the sequel. The vector

$$\hat{\omega}(\theta) = (-\sin(\theta), \cos(\theta)),$$

is perpendicular to $\omega(\theta)$. For any real number s,

$$\omega \cdot (t\omega + s\hat{\omega}) = t$$

and therefore we can describe $l_{t,\omega}$ parametrically as the set of points

$$l_{t,\omega} = \{t\omega + s\hat{\omega} \mid s \in (-\infty, \infty)\}.$$

Both $\hat{\omega}$ and $-\hat{\omega}$ are unit vectors which are perpendicular to ω; $\hat{\omega}$ is singled out by the condition that the 2×2 matrix $(\omega\hat{\omega})$ has determinant $+1$. This shows that the pair (t, ω) determines an *oriented line* The vector $\hat{\omega}$ is the "positive" direction along the line $l_{t,\omega}$.

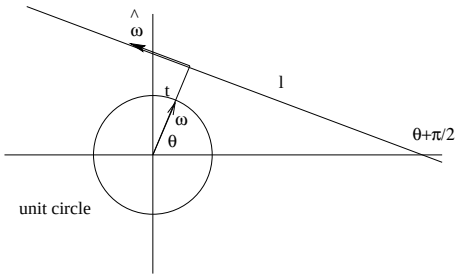

Figure 1.10: Parameterization of oriented lines in the plane.

The vector ω is the direction perpendicular to the line and the number t is called the *affine parameter* of the line, $|t|$ is the distance from the line to the origin of the coordinate system. The pair (t, ω) defines two half planes

$$H^+_{t,\omega} = \{(x, y) \mid (x, y) \cdot \omega > t\} \text{ and } H^-_{t,\omega} = \{(x, y) \mid (x, y) \cdot \omega < t\},$$

the line $l_{t,\omega}$ is the common boundary of these half planes. Facing along the line $l_{t,\omega}$ in the direction specified by $\hat{\omega}$, the half plane $H^-_{t,\omega}$ lies to the left. To summarize, the pairs $(t, \omega) \in \mathbb{R}^1 \times \mathbb{S}^1$ parametrize the *oriented* lines in the plane, which we sometimes call the *space of oriented lines.*

Exercise 1.2.1. Show that

$$|t| = \min\{\sqrt{x^2 + y^2} \ : \ (x, y) \in l_{t,\omega}\}.$$

Exercise 1.2.2. Show that if ω is fixed then the family of lines $\{l_{t,\omega} \ : \ t \in \mathbb{R}\}$ are parallel.

Exercise 1.2.3. Show that every line in the family $\{l_{t,\hat{\omega}} \ : \ t \in \mathbb{R}\}$ is orthogonal to every line in the family $\{l_{t,\omega} \ : \ t \in \mathbb{R}\}$.

Exercise 1.2.4. Each choice of direction ω defines a coordinate system on $\mathbb{R}^2$,

$$(x, y) = t\omega + s\hat{\omega}.$$

Find the inverse, expressing (t, s) as functions of (x, y). Show that the the area element in the plane satisfies

$$dxdy = dtds.$$

1.2.2 Reconstructing an object from its shadows

Now we can quantitatively describe the shadow. Because there are two lines in each family of parallel lines which are tangent to the boundary of D we need a way to select one of them. To do this we choose an orientation for the boundary of D; this operation is familiar from Green's theorem in the plane. The positive direction on the boundary is selected so that, when facing in that direction the region lies to the left; the counterclockwise direction is, by convention the positive direction, see figure 1.11.

Fix a direction $\omega = (\cos(\theta), \sin(\theta))$. In the family of parallel lines $l_{t,\omega}$ there are two values of t, $t_0 < t_1$ such that the lines $l_{t_0,\omega}$ and $l_{t_1,\omega}$ are tangent to the boundary of D,

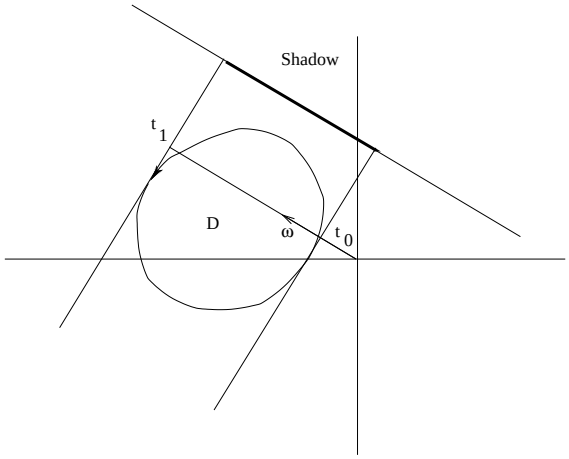

Figure 1.11: The measurement of the shadow

see figure 1.11. Examining the diagram it is clear that the orientation of the boundary at the point of tangency and that of the line agree, for t_1 and are opposite for t_0. For $\omega(\theta) = (\cos(\theta), \sin(\theta))$ define $h_D(\theta) = t_1$. We call $h_D(\theta)$ the *shadow function* for D. The mathematical formulation of reconstruction problem is: Can the boundary of the region D be determined from its shadow function?

The line $l_{h_D(\theta),\omega(\theta)}$ is given parametrically by

$$\{h_D(\theta)(\cos(\theta), \sin(\theta)) + s(-\sin(\theta), \cos(\theta)) \mid s \in (-\infty, \infty)\}.$$

To determine the boundary of D it would suffice to determine the point of tangency of $l_{h_D(\theta),\omega(\theta)}$ with the boundary of D, in other words we would like to find the function $s(\theta)$ so that for each θ,

$$(x(\theta), y(\theta)) = h_D(\theta)(\cos(\theta), \sin(\theta)) + s(\theta)(-\sin(\theta), \cos(\theta)) \tag{1.8}$$

is a point on the boundary of D.

The function $s(\theta)$ is found by recalling that, at the point of tangency, the direction of the tangent line to D is $\hat{\omega}(\theta)$. For a curve in the plane given parametrically by $(x(\theta), y(\theta))$ the direction of the tangent line at a point θ_0 is the same as that of the vector $(x'(\theta_0), y'(\theta_0))$. Differentiating the expression given in (1.8) and using the fact that $\partial_\theta \omega = \hat{\omega}$ we find that

$$(x'(\theta), y'(\theta)) = (h'_D(\theta) - s(\theta))\omega(\theta) + (h_D(\theta) + s'(\theta))\hat{\omega}(\theta). \tag{1.9}$$

Since the tangent line at $(x(\theta), y(\theta))$ is parallel to $\hat{\omega}(\theta)$ it follows from (1.9) that

$$h'_D(\theta) - s(\theta) = 0. \tag{1.10}$$

This gives a parametric representation for the boundary of a convex region in terms of its shadow function: If the shadow function is $h_D(\theta)$ then the boundary of D is given parametrically by

$$(x(\theta), y(\theta)) = h_D(\theta)\omega(\theta) + h'_D(\theta)\hat{\omega}(\theta). \tag{1.11}$$

Note that we have assumed that $h_D(\theta)$ is a differentiable function. This is not always true, for example if the region D is a polygon then the shadow function is not everywhere differentiable.

Let D denote a convex region and h_D its shadow function. We can think of $D \to h_D$ as a mapping from convex regions in the plane to 2π periodic functions. It is reasonable to enquire if every 2π periodic function is the shadow function of a convex region. The answer to this question is **no.** For strictly convex regions with smooth boundaries we are able to characterize the range of this mapping. If h is twice differentiable then the tangent vector to the curve defined by

$$(x(\theta), y(\theta)) = h(\theta)\omega(\theta) + h'(\theta)\hat{\omega}(\theta) \tag{1.12}$$

is given by

$$(x'(\theta), y'(\theta)) = (h''(\theta) + h(\theta))\hat{\omega}(\theta).$$

In our construction of the shadow function we observed that the tangent vector to the curve at $(x(\theta), y(\theta))$ and $\hat{\omega}(\theta)$ point in the same direction. From our formula for the tangent vector we see that this implies that

$$h''(\theta) + h(\theta) > 0 \text{ for all } \theta \in [0, 2\pi]. \tag{1.13}$$

This gives a necessary condition for a twice differentiable function h to be the shadow function for a strictly convex region with a smooth boundary. Mathematically we are determining the of the map that takes a convex body $D \subset \mathbb{R}^2$ to its shadow function h_D, under the assumption that h_D is twice differentiable. This is a convenient mathematical assumption, though in an applied context it is likely to be overly restrictive.

Exercise 1.2.5. Suppose that D_n is a regular n-gon. Determine the shadow function $h_{D_h}(\theta)$.

Exercise 1.2.6. Suppose that $h(\theta)$ is 2π-periodic, twice differentiable function which satisfies (1.13). Show that the curve given by (1.12) is the boundary of a strictly convex region.

Exercise 1.2.7. * Find a characterizations of those functions which are shadow functions of convex regions without assuming that they are twice differentiable or that the region is strictly convex.

Exercise 1.2.8. If $h(\theta)$ is any differentiable function then equation (1.12) defines a curve, by plotting examples, determine what happens if the condition (1.13) is not satisfied.

Exercise 1.2.9. Suppose that h_D is a function satisfying (1.13). Show that the area enclosed by Γ_h is given by the

$$\text{Area}(D_h) = \frac{1}{2}\int_0^{2\pi} [(h(\theta))^2 - (h'(\theta))^2] d\theta.$$

Explain why this implies that a function satisfying (1.13) also satisfies the estimate

$$\int_0^{2\pi} (h'(\theta))^2 d\theta < \int_0^{2\pi} (h(\theta))^2 d\theta.$$

Exercise 1.2.10. Let h be a smooth 2π-periodic function which satisfies (1.13). Prove that the curvature of the boundary of the region with this shadow function, at the point $h(\theta)\omega(\theta) + h'(\theta)\hat{\omega}(\theta)$ is given by

$$\kappa(\theta) = \frac{1}{h(\theta) + h''(\theta)}. \tag{1.14}$$

Exercise 1.2.11. Suppose that h is a function satisfying (1.13). Show that another parametric representation for boundary of the region with this shadow function is

$$\theta \mapsto \left(-\int_0^\theta (h(s) + h''(s)) \sin(s) ds, \int_0^\theta (h(s) + h''(s)) \cos(s) ds \right).$$

Exercise 1.2.12. Which positive functions $\kappa(\theta)$ defined on S^1 are the curvatures of closed convex curves? Prove the following result: A positive function $\kappa(\theta)$ on S^1 is the curvature of a closed, strictly convex curve (parametrized by its tangent direction) if and only if

$$\int_0^\infty \frac{\sin(s) ds}{\kappa(s)} = 0 = \int_0^\infty \frac{\cos(s) ds}{\kappa(s)}.$$

Exercise 1.2.13. Let D be a convex region with shadow function h_D. For a vector $v \in \mathbb{R}^2$ define the translated region

$$D^v = \{(x, y) + v \ : \ (x, y) \in D\}.$$

Find the relation between h_D and h_{D^v}. Explain why this answer is inevitable in light of the formula (1.14), for the curvature.

Exercise 1.2.14. Let D be a convex region with shadow function h_D. For a rotation $A \in SO(2)$ define the rotated region

$$D^A = \{A(x, y) \ : \ (x, y) \in D\}.$$

Find the relation between h_D and h_{D^A}.

Exercise 1.2.15. * If h_1 and h_2 are 2π-periodic functions satisfying (1.13) then they are the shadow functions of convex regions D_1 and D_2. The sum, $h_1 + h_2$ also satisfies (1.13) and so is the shadow function of a convex region, D_3. Describe, geometrically how D_3 is determined by D_1 and D_2.

1.2.3 Approximate reconstructions

See: A.7.2.

In a realistic situation we can only make finitely many measurements. The shadow function is measured at a finite set of angles $\{h_D(\theta_1), \ldots, h_D(\theta_m)\}$. How can this data

be used to construct an approximation to the region D, which cast these shadows? We consider two different strategies both of which rely on the special geometric properties of convex regions. Recall that a convex region always lies in one of the half planes determined by the support line at any point of its boundary. The half plane lying "below" the oriented line $l_{t,\omega}$ is the set defined by

$$H^-_{t,\omega} = \{(x,y) \mid (x,y)\cdot\omega < t\}.$$

Since the boundary of D and $l_{h(\theta),\omega(\theta)}$ have the same orientation at the point of contact, it follows that D lies in each of the half planes

$$H^-_{h(\theta_j),\omega(\theta_j)}, \quad j = 1,\ldots,m.$$

As D lies in each of these half planes it also lies in their intersection. This defines a polygon

$$P_m = \bigcap_{j=1}^{m} H_{h(\theta_j),\omega(\theta_j)}.$$

P_m is a convex polygon which contains D, this then provides one sort of approximation for D from the measurement of a finite set of shadows. This is a stable approximation to D as small changes in the measurements of either the angles θ_j or the corresponding affine parameters $h(\theta_j)$ lead to small changes in the approximating polygon.

The difficulty with using the exact reconstruction formula (1.11) is that h is only known at finitely many values, $\{\theta_j\}$. From this information it is not possible to exactly compute the derivatives, $h'(\theta_j)$. We could use a finite difference approximation for the derivative to determine a finite set of points which approximate points on the boundary of D :

$$(x(\theta_j), y(\theta_j)) = h(\theta_j)\omega(\theta_j) + \frac{h(\theta_j) - h(\theta_{j+1})}{\theta_j - \theta_{j+1}}\hat{\omega}(\theta_j).$$

If the measurements were perfect, the boundary of D smooth and the numbers $\{|\theta_j - \theta_{j+1}|\}$ small then the finite difference approximations to $h'(\theta_j)$ would be accurate and these points would lie close to points on the boundary of D. Joining these points, in the given order gives a polygon, P' which approximates D. If the points could be computed exactly then P' would be contained in D. With approximate values this cannot be asserted with certainty, though under the assumptions above, P' should be largely contained within D.

This gives a different way to reconstruct an approximation to D from a finite set of measurements. This method is not as robust as the first technique because it requires the measured data to be differentiated. In order for the finite difference $\frac{h(\theta_j)-h(\theta_{j+1})}{\theta_j-\theta_{j+1}}$ to be a good approximation to $h'(\theta_j)$ it is generally necessary for $|\theta_j - \theta_{j+1}|$ to be small. Moreover the errors in the measurements of $h(\theta_j)$ and $h(\theta_{j+1})$ must also be small compared to $|\theta_j - \theta_{j+1}|$. This difficulty arises in solution of the reconstruction problem in X-ray CT, the exact reconstruction formula calls for the measured data to be differentiated. In general, measured data is corrupted by random noise, and random noise is usually "non-differentiable."

This means that measurements of a function must be regularized before they can used approximate derivatives. One way to handle this is to improve the quality of individual

measurements. One assumes that the errors in individual measurements have "mean zero:" if the same measurement is repeated many times then the average of the individual measurements should approach the true value. This is the approach taken in magnetic resonance imaging. Another possibility is to make a large number of measurements at closely spaced angles $\{(h_j, j\Delta\theta) \ : \ j = 1, \ldots, N\}$ which are then "averaged" to give less noisy approximations on a coarser grid. There are many ways to do the averaging. One way is to find a differentiable function, H belonging to a family of functions of dimension $M < N$ and minimizes the *square error*

$$e(H) = \sum_{j=1}^{N} (h_j - H(j\Delta\theta))^2.$$

For example H could be taken to be a polynomial of degree $M - 1$, or a continuously differentiable, piecewise cubic function. Using values of H one can find an approximation to the boundary of D which is hopefully less corrupted by noise. Fine structure in the boundary is also blurred by such a procedure. This is closer to the approach used in X-ray CT.

Exercise 1.2.16. Suppose that the angles $\{\theta_j\}$ can be measured exactly but there is an uncertainty of size ϵ in the measurement of the affine parameters, $h(\theta_j)$. Find a polygon $P_{m,\epsilon}$ which gives the best possible approximation to D which certainly contains D.

Exercise 1.2.17. Suppose that we know that $|h''(\theta)| < M$, and the measurement errors are bounded by $\epsilon > 0$. For what angle spacing is the error in using a finite difference approximation for h' due to the uncertainty in the measurements equal to that caused by the non-linearity of h itself.

1.2.4 Can an object be reconstructed from its width?

To measure the location of the shadow requires an expensive detector which can accurately locate a transition from light to dark. It would be much cheaper to build a device, similar to the exposure meter in a camera, to measure the length of the shadow region without determining its precise location. It is therefore an interesting question whether or not the boundary of a region can be reconstructed from measurements of the *widths* of its shadows. Let $w_D(\theta)$ denote the width of the shadow in direction θ, a moments consideration shows that

$$w_D(\theta) = h_D(\theta) + h_D(\theta + \pi).$$

From this formula it follows that w_D does not determine D. Let $e(\theta) \not\equiv 0$ be a function that satisfies

$$e(\theta) + e(\theta + \pi) = 0. \tag{1.15}$$

From the discussion in section 1.2.2 we know that if h_D has two derivatives such that $h_D'' + h_D > 0$ then $h_D(\theta)$ is the shadow function of a strictly convex region. Let e be a smooth function satisfying (1.15) such that

$$h_D'' + h_D + e'' + e > 0$$

as well, then $h_D + e$ is also the shadow function for a different strictly convex region. Observe that $h_D(\theta) + e(\theta)$ is the shadow function for a different region, D' which has the same *width* of shadow for each direction as D. That is

$$w_D(\theta) = (h_D(\theta) + e(\theta)) + (h_D(\theta + \pi) + e(\theta + \pi)) = w_{D'}(\theta).$$

To complete this discussion note that any function with a Fourier representation of the form

$$e(\theta) = \sum_{j=0}^{\infty} [a_j \sin(2j+1)\theta + b_j \cos(2j+1)\theta]$$

satisfies (1.15). This is an infinite dimensional space of functions. This implies that if $w_D(\theta)$ is the "width of the shadow" function for a convex region D then there is an infinite dimensional set of regions with the same "width of the shadow" function. Consequently the simpler measurement is inadequate to reconstruct the boundary of a convex region. The figure below shows the unit disk and another region which has constant "shadow width" equal to 2.

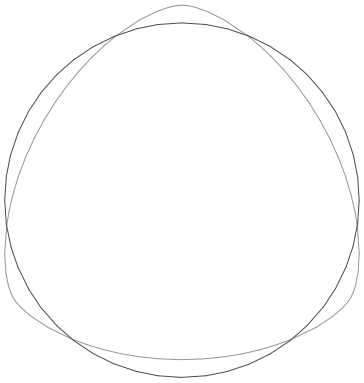

Figure 1.12: Two regions of constant width 2

Exercise 1.2.18. Show that the width function satisfies $w_D'' + w_D > 0$.

Exercise 1.2.19. Is it true that every twice differentiable, π-periodic function, w satisfying $w'' + w > 0$ is the width function of a convex domain?

Exercise 1.2.20. Our motivation for considering whether or not a convex body is determined by the width of its shadows was to replace our expensive detector, which can determine where a shadow begins and ends, with a less expensive detector. The cheaper detector can only measure the width of the covered region. Can you find a way to use a detector which only measures the length of an illuminated region to locate the edge of the shadow? Hint: Only cover half of the detector with photosensitive material.

1.3 Linearity

See: A.2.

As we have seen, using measurements to determine other quantities requires the solution of systems of equations. Sometimes we need to solve differential equations and sometimes algebraic equations. In almost all practical applications one is eventually reduced to solving systems of linear equations. This is true even for physical systems which are described by non-linear equations. Non-linear equations are usually solved iteratively where the iteration step involves the solution of linear equations. There are many reasons why linear equations and linear models are ubiquitous. From the pragmatic point of view, there is a "complete" mathematical theory for systems of linear equations. One has necessary and sufficient conditions for linear equations to have solutions, a description of the space of solutions when they exist and practical algorithms for finding them. This is not true even for the simplest families on non-linear equations. On a more conceptual level, for systems with some sort of intrinsic smoothness, a linear model often suffices to describe small deviations from an equilibrium state.

We give a quick overview of the theory of systems of linear, algebraic equations. This is not intended to serve as a text on linear algebra, merely a review of some important concepts. Detailed discussions of various aspect of this material can be found in [43] or [78].

Consider a system of m equations for n unknowns, $(x_1, \ldots, x_n)$:

$$\begin{array}{ccc} a_{11}x_1 + a_{12}x_2 + & \cdots + a_{1n}x_n & = y_1 \\ a_{21}x_1 + a_{22}x_2 + & \cdots + a_{2n}x_n & = y_2 \\ \vdots & \vdots & \vdots \\ a_{m1}x_1 + a_{m2}x_2 + & \cdots + a_{mn}x_n & = y_m. \end{array} \tag{1.16}$$

There are four questions which require answers:

Existence:
: For a given m-vector $(y_1, \ldots, y_m)$ does there exist an n-vector $(x_1, \ldots, x_n)$ which satisfies the equations in (1.16)?

Uniqueness:
: When a solution exists is it unique? More generally, describe the space of solutions.

Solve in practice:
: Give an algorithm to find approximations for the solutions of (1.16) and criteria to select a solution when there is more than one.

Stability:
: How sensitive is the solution to small variations in the coefficients (a_{ij}) or the right hand side (y_j)?

It is a somewhat unexpected, but very important fact that these issues are in practice, rather independent of one another.

Before proceeding with our analysis we first need to simplify the notation. It is very cumbersome to have to work with complicated expressions like (1.16), instead we use standard matrix and vector notation. Let $\mathbf{a}$ denote the $m \times n$ array of numbers $(a_{ij})_{i=1\ldots m, j=1\ldots n}$ and $\mathbf{x}$ denote the column n-vector. The result of multiplying $\mathbf{x}$ by $\mathbf{a}$ is a column m-vector whose i^{th} entry is

$$(\mathbf{ax})_i = a_{i1}x_1 + a_{i2}x_2 + \cdots + a_{in}x_n.$$

The system of equations (1.16) is concisely expressed as

$$\mathbf{a}\mathbf{x} = \mathbf{y},$$

Here y is a column m-vector with entries $(y_1, \dots, y_m)$. We briefly recall the properties of matrix multiplication, let $\mathbf{x}^1$ and $\mathbf{x}^2$ be n-vectors, then

$$\mathbf{a}(\mathbf{x}^1 + \mathbf{x}^2) = \mathbf{a}\mathbf{x}^1 + \mathbf{a}\mathbf{x}^2$$

and for any number c

$$\mathbf{a}(c\mathbf{x}^1) = c(\mathbf{a}\mathbf{x}^1).$$

These are the properties that characterize linearity.

1.3.1 Solving linear equations

Suppose that $\mathbf{x}^0$ is a solution of the equation $\mathbf{a}\mathbf{x} = 0$ and $\mathbf{x}^1$ is a solution of the equation $\mathbf{a}\mathbf{x}^1 = \mathbf{y}$ then the rules above show that for any number c we have

$$\mathbf{a}(c\mathbf{x}^0 + \mathbf{x}^1) = c\mathbf{a}\mathbf{x}^0 + \mathbf{a}\mathbf{x}^1 = \mathbf{a}\mathbf{x}^1 = \mathbf{y}.$$

If $\mathbf{y} = 0$ as well then we conclude that the set of solutions to the equation

$$\mathbf{a}\mathbf{x} = 0$$

is a linear space, that is if $\mathbf{x}^0$ and $\mathbf{x}^1$ solve this equation then so does $\mathbf{x}^0 + \mathbf{x}^1$ as well as $c\mathbf{x}^0$, for any number c. This space is called the *null space or kernel*of $\mathbf{a}$. It is denoted by $\ker(\mathbf{a})$ and always contains, at least the zero vector $\mathbf{0} = (0, \dots, 0)$. These observations answer the question above about uniqueness.

Theorem 1.3.1. *Let* $\mathbf{a}$ *be an* $m \times n$ *matrix. Given a vector* $\mathbf{y}$, *if* $\mathbf{x}^1$ *satisfies* $\mathbf{a}\mathbf{x}^1 = \mathbf{y}$ *then every other solution to this equation is of the form* $\mathbf{x}^1 + \mathbf{x}^0$ *where* $\mathbf{x}^0 \in \ker(\mathbf{a})$. *Moreover, every vector of this form solves the equation* $\mathbf{a}\mathbf{x} = \mathbf{y}$.

As a simple corollary it follows that the solution of the equation $\mathbf{a}\mathbf{x} = \mathbf{y}$ is unique only if the null space of $\mathbf{a}$ contains only the 0-vector.

In order to answer the question of existence it is convenient to introduce the notion of a "dot" or inner product. If $\mathbf{x}$ and $\mathbf{y}$ are two n-vectors then define

$$\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{j=1}^{n} x_j y_j = x \cdot y.$$

Suppose that $\mathbf{a}$ is an $m \times n$-matrix, $\mathbf{x}$ is an n-vector and $\mathbf{y}$ is an m-vector then $\mathbf{a}\mathbf{x}$ is an m-vector and

$$\langle \mathbf{a}\mathbf{x}, \mathbf{y} \rangle = \sum_{i=1}^{m} \sum_{j=1}^{n} y_i a_{ij} x_j.$$

The transpose of the matrix $\mathbf{a}$ is the $n \times m$ matrix $\mathbf{a}^t$ whose ij-entry is a_{ji}. From the previous formula it follows that that

$$\langle \mathbf{a}\mathbf{x}, \mathbf{y} \rangle = \langle \mathbf{x}, \mathbf{a}^t \mathbf{y} \rangle.$$

Suppose that $\mathbf{y}$ is a non-zero vector in the null space of $\mathbf{a}^t$ (note that here we are using the transpose!) and the equation $\mathbf{ax} = \mathbf{b}$ has a solution. Using the calculations above we see that

$$\langle \mathbf{b}, \mathbf{y} \rangle = \langle \mathbf{ax}, \mathbf{y} \rangle = \langle \mathbf{x}, \mathbf{a}^t\mathbf{y} \rangle = 0.$$

The last equality follows from the fact that $\mathbf{a}^t\mathbf{y} = 0$. This gives a necessary condition for existence of a solution to the equation $\mathbf{ax} = \mathbf{b}$, the vector $\mathbf{b}$ must satisfy $\langle \mathbf{b}, \mathbf{y} \rangle = 0$ for every solution of the homogeneous equation $\mathbf{a}^t\mathbf{y} = 0$. This also turns out to be sufficient.

Theorem 1.3.2. *Let* $\mathbf{a}$ *be an* $m \times n$*-matrix and* $\mathbf{b}$ *and* m*-vector. The equation* $\mathbf{ax} = \mathbf{b}$ *has a solution if and only if*

$$\langle \mathbf{b}, \mathbf{y} \rangle = 0$$

for every vector $\mathbf{y}$ *satisfying the homogeneous equation*

$$\mathbf{a}^t\mathbf{y} = 0.$$

Putting these two results together we obtain that

Corollary 1.3.1. *Let* $\mathbf{a}$ *be an* $m \times n$*-matrix the equation* $\mathbf{ax} = \mathbf{b}$ *has a* **unique** *solution for any vector* $\mathbf{b}$ *if and only if* $\ker(\mathbf{a}) = \{0\}$ *and* $\ker(\mathbf{a}^t) = \{0\}$.

In a physical situation the vector $\mathbf{x}$ describes the state of a system and the entries of the vector $\mathbf{b}$ are results of measurements made on the system while it is in this state. The matrix $\mathbf{a}$ is a model for the measurement process: it is the assertion that if the physical object is described by the parameters, $\mathbf{x}$ then the results of the experiments performed should be the vector $\mathbf{b} = \mathbf{ax}$. This is a *linear model* because the map from the state of the system to the measurements is a linear map. The problem of determining the state of the system from the measurements is precisely the problem of solving this system of linear equations.

Example 1.3.1. Suppose we have a collection of photons sources, labeled by $1 \leq i \leq n$ and an array of detectors, labeled by$1 \leq j \leq m$. The matrix $\mathbf{P}$ has entries $0 \leq p_{ij} \leq 1$. The ij-entry is the probability that a particle emitted from source i is detected by detector j. Since a given photon can be detected by at most one detector it follows that

$$\sum_{j=1}^{m} p_{ij} \leq 1 \text{ for } i = 1, \ldots n.$$

If $d_j, j = 1, \ldots, m$ is the number of photons detected at detector j and $s_i, i = 1, \ldots, n$ is the number of photons emitted by source i then our model predicts that

$$\mathbf{Ps} = \mathbf{d}.$$

If $m = n$ and $\mathbf{P}$ is an invertible matrix then we can use the measurements $\mathbf{d}$ to obtain a unique vector $\mathbf{s}$. Since the model is probabilistic this should be regarded as an expected value for the distribution of sources. If $m > n$ then we have more measurements than unknowns, so any measurement errors or flaws in the model could make it impossible to find a vector $\mathbf{s}$ so that $\mathbf{Ps} = \mathbf{d}$. This is a frequent situation in image reconstruction problems. One

chooses a way to measure the error, usually a function of the form $e(\mathbf{Ps} - \mathbf{d})$ and seeks a vector s which minimizes the error. Finally we may have more sources than detectors. The measurements are then inadequate, in principle to determine their distribution. This is also a common circumstance in image reconstruction problems and is resolved by making some a priori assumptions about the allowable distribution of sources to obtain a determined (or even overdetermined) problem.

As illustrated by this example and explained in the theorem there are essentially 3 types of linear models for systems with finitely many degrees of freedom.

Determined:

The simplest case arises when the number of *independent* measurements and parameters describing the state of the system are the same. This implies that $n = m$. In this case the measurements uniquely determine the state of the system. Mathematically we say that the matrix, $\mathbf{a}$ is invertible. In the situation that $n = m$ this is equivalent to the statement that the homogeneous equation, $\mathbf{ax} = 0$ has only the trivial solution, $\mathbf{x} = 0$. The inverse matrix is denoted by $\mathbf{a}^{-1}$, it is both a left and a right inverse to $\mathbf{a}$,

$$\mathbf{a}^{-1}\mathbf{a} = \mathrm{Id}_n = \mathbf{a}\mathbf{a}^{-1}.$$

Here Id_n denotes the $n \times n$ identity matrix, that in

$$(\mathrm{Id}_n)_{ij} = \begin{cases} 1 \text{ if } i = j, \\ 0 \text{ if } i \neq j. \end{cases}$$

From the mathematical point of view, the unique solution is obtained by setting

$$\mathbf{x} = \mathbf{a}^{-1}\mathbf{y}.$$

Except in special cases, the inverse matrix $\mathbf{a}^{-1}$ is not computed directly.

Overdetermined:

In this case we have more measurements than parameters, i.e. $m > n$. If the model and measurements are perfect then there should be a unique $\mathbf{x}$ with $\mathbf{ax} = \mathbf{y}$. In general, neither is true and there will not exist any $\mathbf{x}$ exactly satisfying this equation. Having more measurements than parameters can be used to advantage in several different ways. In example 1.3.2 we explain how to use the conditions for solvability given in Theorem 1.3.2 to determine physical parameters. Often times measurements are noisy. A model for the noise in the measurements can be used to select a criterion for a "best approximate solution." The error function is usually defined by picking a norm $\|\cdot\|$ on the space of measurements. We then try to find the vector $\mathbf{x}$ which minimizes the error, $e(\mathbf{x}) = \|\mathbf{ax} - \mathbf{y}\|$. The most commonly used error function is that defined by the *square norm*

$$\|\mathbf{y}\|_2^2 = \sum_{j=1}^{m} y_j^2.$$

There are two reasons why this measure of the error is often employed: 1. It is a natural choice if the noise is normally distributed, 2. The problem of minimizing $\|\mathbf{ax} - \mathbf{y}\|_2$ can be reduced to the problem of solving a system of linear equations.

Underdetermined:

Most of the problems in image reconstruction are underdetermined, that is we do not have enough data to uniquely determine a solution. In mathematical tomography a "perfect reconstruction" requires an infinite number of exact measurements. These are, of course never available. In a linear algebra problem, this is the case where $m < n$. When the measurements $\mathbf{y}$ do not uniquely determine the state $\mathbf{x}$, additional criteria are needed to determine which solution to actually use, for example one might use the solution to $\mathbf{ax} = \mathbf{y}$ which is of smallest norm. Another approach is to assume that $\mathbf{x}$ belongs to a subspace whose dimension is equal to the number of independent measurements. Both of these approaches are used in medical imaging.

Example 1.3.2. In the refraction problem considered in example 1.1.4 we remarked that the refractive index of the lower fluid n_2 could be determined by an additional measurement. Suppose that we shine a beam of light in at a different angle, so that the upper angle is ϕ_1 and the lower angle is ϕ_2. This light beam is displaced by l_2 as it passes through the fluid. We now have 3 equations for the two unknowns:

$$\begin{pmatrix} 1 & 1 \\ \tan(\theta_1) & \tan(\theta_2) \\ \tan(\phi_1) & \tan(\phi_2) \end{pmatrix} \begin{pmatrix} h_1 \\ h_2 \end{pmatrix} = \begin{pmatrix} h \\ l_1 \\ l_2 \end{pmatrix}. \tag{1.17}$$

In order for this equation to have a solution the measurements (h, l_1, l_2) must satisfy the condition

$$\begin{pmatrix} 1 \\ \tan(\theta_1) \\ \tan(\phi_1) \end{pmatrix} \times \begin{pmatrix} 1 \\ \tan(\theta_2) \\ \tan(\phi_2) \end{pmatrix} \cdot \begin{pmatrix} h \\ l_1 \\ l_2 \end{pmatrix} = 0.$$

Here $\times$ is the vector cross product. Since

$$\frac{\sin(\theta_1)}{\sin(\theta_2)} = \frac{\sin(\phi_1)}{\sin(\phi_2)} = \frac{n_2}{n_1}$$

and the angles θ_1 and ϕ_1 as well as (h, l_1, l_2) are assumed known, this solvability conditions gives a non-linear equation which allows the determination of $\frac{n_2}{n_1}$ from the measured data.

This brings us to the final questions of giving practical algorithms for finding the solutions of linear equations and their stability. We leave a detailed discussion of algorithms to later chapters for, in practice one needs to select a method that is well adapted to the equations under consideration. We conclude this section with a discussion of the problem of stability. In practical situations many things conspire to limit the accuracy achievable in using measurements to predict the state of a system. There are errors in the model itself as well as noise in the actual measurements. Once the model is made and the measurements are taken one needs to solve systems of equations. In finite time one can only work with numbers having finitely many decimal places, so rounding errors are also an unavoidable problem.

Given all these sources of error one would like to have a *stable* algorithm for solving the system of equations. Suppose that $\mathbf{a}$ is an $n \times n$ invertible matrix that models a

measurement process. If $\mathbf{x}_1$ and $\mathbf{x}_2$ are two states of our system then, because the model is linear the difference in the measurements can easily be computed

$$\mathbf{y}_1 - \mathbf{y}_2 = \mathbf{a}\mathbf{x}_1 - \mathbf{a}\mathbf{x}_2 = \mathbf{a}(\mathbf{x}_1 - \mathbf{x}_2).$$

From this formula we see that nearby states result in nearby measurements. However the reverse is often not true. There may exist states $\mathbf{x}_1$ and $\mathbf{x}_2$ which are not nearby, in the sense that $\|\mathbf{x}_1 - \mathbf{x}_2\|$ is large but $\|\mathbf{a}(\mathbf{x}_1 - \mathbf{x}_2)\|$ is small. Physically, the measurements performed are not sufficiently independent to distinguish certain pairs of states, which are not, in fact very close together. In numerical analysis this is known as an *ill-conditioned* equation. Briefly, a small error in the measurement process can be magnified by applying $\mathbf{a}^{-1}$ to the measurement vector. For an ill-conditioned problem even a good algorithm for solving linear equations can produce meaningless results.

Example 1.3.3. For example, consider the system with $m = n = 2$ and

$$\mathbf{a} = \begin{pmatrix} 1 & 0 \\ 1 & 10^{-5} \end{pmatrix}.$$

Then $\mathbf{x}$ is given by $\mathbf{a}^{-1}\mathbf{y}$ where

$$\mathbf{a} = \begin{pmatrix} 1 & 0 \\ -10^5 & 10^5 \end{pmatrix}.$$

If the actual data is $\mathbf{y} = (1, 1)$ but we make an error in measurement and measure, $\mathbf{y}_m = (1, 1 + \epsilon)$ then the relative error is

$$\frac{|\mathbf{y}_m - \mathbf{y}|}{|\mathbf{y}|} = \epsilon 10^5.$$

Even though the measurements uniquely determine the state of the system, a small error in measurement is vastly amplified.

In image reconstruction the practical problems of solving systems of linear equations are considerable. It is not uncommon to have $10,000$-equations in $10,000$-unknowns. These huge systems arise as finite dimensional approximations to linear equations for functions of continuous variables. We close this section with a short discussion of linear algebra in in infinite dimensional spaces. This is a theme which occupies a large part of this book.

Exercise 1.3.1. Let $\mathbf{a}$ be an $m \times n$ matrix. Show that if $\ker \mathbf{a} = \ker \mathbf{a}^t = \mathbf{0}$ then $n = m$. Is the converse true?

Exercise 1.3.2. Suppose that the state of a system is described by the vector $\mathbf{x}$. The measurements are modeled as inner products $\{\mathbf{a}_j \cdot \mathbf{x} : j = 1, \ldots, m\}$. However the measurements are noisy and each is repeated m_j times leading to measured values $\{y_j^1, \ldots, y_j^{m_j}\}$. Define an error function by

$$e(\mathbf{x}) = \sum_{j=1}^{m} \sum_{i=1}^{m_j} (\mathbf{a}_j \cdot \mathbf{x} - y_j^i)^2.$$

Show that $e(\mathbf{x})$ is minimized by the vector which satisfies the averaged equations

$$\mathbf{a}_j \cdot \mathbf{x} = \frac{1}{m_j} \sum_{i=1}^{m_j} y_j^i.$$

1.3.2 Infinite dimensional linear algebra

The state of a 'system' in medical imaging is described by a function of continuous variables. In this introductory section we consider real valued functions defined on the real line. Let $f(x)$ describe the state of the system. A linear measurement of the state is usually described as an integral

$$\mathcal{M}(f)(x) = \int_{-\infty}^{\infty} m(x,y)f(y)dy.$$

Here $m(x,y)$ is a function on $\mathbb{R} \times \mathbb{R}$ which provides a model for the measurement process. It can be thought of as an infinite 'matrix' with indices x and y. A linear transformation of an infinite dimensional space is called a *linear operator*. A linear transformation which can be expressed as an integral is called an *integral operator*.

Suppose that the function $g(x)$ is the output of the measurement process, to reconstruct f means solving the linear equation

$$\mathcal{M}f = g.$$

This is a concise way to write a system of infinitely many equations in infinitely many unknowns. Theorems 1.3.1 and 1.3.2 contain the complete theory for the existence and uniqueness of solutions to linear equations in finitely many variables. These theorems are entirely algebraic in character. No such theory exists for equations in infinitely many variables. It is usually a very complicated problem to describe both the domain and range of such a transformation. We close this section with a few illustrative examples.

Example 1.3.4. Perhaps the simplest linear operator is the indefinite integral

$$\mathcal{I}(f)(x) = \int_{0}^{x} f(y)dy.$$

If we use the continuous functions on $\mathbb{R}$ as the domain of $\mathcal{I}$ then every function in the range is continuously differentiable. Moreover the null-space of $\mathcal{I}$ is the zero function. Observe that the domain and range of $\mathcal{I}$ are fundamentally different spaces. Because $\mathcal{I}(f)(0) = 0$ not every continuously differentiable function is in the range of $\mathcal{I}$. The derivative is a left inverse to $\mathcal{I}$ as the Fundamental Theorem of Calculus states that if f is continuous then

$$\frac{d}{dx} \circ \mathcal{I}(f)(x) = f(x).$$

On the other hand it is not quite a right inverse because

$$\mathcal{I}(\frac{df}{dx})(x) = f(x) - f(0).$$

The domain of $\mathcal{I}$ can be enlarged to include all *locally integrable* functions. These are functions such that

$$\int_{0}^{x} |f(y)|dy < \infty$$

for every $x \in \mathbb{R}$. Enlarging the domain also enlarges the range. For example the function $|x|$ lies in the enlarged range of $\mathcal{I}$,

$$|x| = \int_0^x \text{sign}(y)dy,$$

where $\text{sign}(y) = 1$ if $y \geq 0$ and -1 if $y < 0$. Even though $|x|$ is not differentiable at $x = 0$ it is still the indefinite integral of a locally integrable function, however the formula

$$\frac{d|x|}{dx} = \text{sign}(x)$$

does not make sense at $x = 0$.

Example 1.3.5. Changing the lower limit of integration to $-\infty$ leads to a very different sort of linear transformation. Initially $\mathcal{I}_\infty$ is defined for continuous functions f, vanishing for sufficiently negative x by

$$\mathcal{I}_\infty(f)(x) = \int_{-\infty}^x f(y)dy.$$

Once again the null-space of $\mathcal{I}_\infty$ consists of the zero function alone. The domain can be enlarged to include locally integrable functions such that

$$\lim_{R\to\infty} \int_{-R}^0 |f(y)|dy < \infty. \tag{1.18}$$

If f is continuous then we can apply the F.T.C. to obtain

$$\frac{d}{dx} \circ \mathcal{I}(f) = f.$$

If a function g belongs to the range of $\mathcal{I}$ then

$$\lim_{x\to-\infty} g(x) = 0. \tag{1.19}$$

There are once differentiable functions satisfying this condition which do not belong to the range of $\mathcal{I}_\infty$. For example,

$$f(x) = \frac{x\cos x - \sin x}{x^2} = \frac{d}{dx}\frac{\sin x}{x}$$

satisfies (1.19) but $\frac{\cos x}{x}$ does not satisfy (1.18). With the domain defined by (1.18) the precise range of $\mathcal{I}_\infty$ is rather difficult to describe.

This example illustrates how a integral operator may have a simple definition on a certain domain, which by a limiting process can be extended to a larger domain. The domain of such an operator is often characterized by a size condition like (1.18).

Example 1.3.6. A real physical measurement is always some sort of an average. If the state of the system is described by a function f of a single variable x then the average of f over an interval of length 2δ is

$$\mathcal{M}_\delta(f)(x) = \frac{1}{2\delta}\int_{x-\delta}^{x+\delta} f(y)dy.$$

A natural domain for $\mathcal{M}_\delta$ is all locally integrable functions. To what extent is f determined by $\mathcal{M}_\delta(f)$? Suppose that f and g are two states, then, because the integral is linear

$$\mathcal{M}_\delta(f) - \mathcal{M}_\delta(g) = \mathcal{M}_\delta(f-g).$$

The extent to which $\mathcal{M}_\delta(f)$ determines f is characterized by the null-space of $\mathcal{M}_\delta$,

$$\mathcal{N}_\delta = \{f \,:\, \mathcal{M}_\delta(f) = 0\}.$$

Proceeding formally, we can differentiate $\mathcal{M}_\delta(f)$ to obtain

$$\frac{d\mathcal{M}_\delta(f)}{dx} = f(x+\delta) - f(x-\delta) \tag{1.20}$$

If $f \in \mathcal{N}_\delta$ then $\mathcal{M}_\delta(f)$ is surely constant and therefore

$$f \in \mathcal{N}_\delta \Rightarrow f(x+\delta) - f(x-\delta) = 0,$$

in other words f is periodic with periodic 2δ. A periodic function has an expansion in terms of sines and cosines, that is

$$f(x) = a_0 + \sum_{j=1}^{\infty}\left[a_j\cos(\frac{\pi jx}{\delta}) + b_j\sin(\frac{\pi jx}{\delta})\right].$$

If $a_0 = 0$ then $\mathcal{M}_\delta(f) = 0$. This shows that the null-space of $\mathcal{M}_\delta$ is infinite dimensional.

In applications one often has additional information about the state of the system, for example one might know that

$$\lim_{|x|\to\infty} f(x) = 0. \tag{1.21}$$

A periodic function that tends to zero at infinity must be identically zero, so among such functions the measurements $\mathcal{M}_\delta(f)$ would appear to determine f completely. To *prove* this statement we need to know somewhat more about f than (1.21). With a more quantitative condition like

$$\|f\|_p = \left[\int_{-\infty}^{\infty}|f(y)|^p dy\right]^{\frac{1}{p}} < \infty, \tag{1.22}$$

for a p between 1 and 2, it is possible to show that $\mathcal{M}_\delta(f) = 0$ implies that $f = 0$. For such functions the measurement $\mathcal{M}_\delta(f)$ uniquely determines f. However, f cannot be *stably* reconstructed from $\mathcal{M}_\delta(f)$. A small error in measurement can lead to a very large error in the reconstructed state.

The integral in (1.22) defines a measure for the size of f called the L^p-norm. It is a generalization of the notion of a norm on a finite dimensional vector space and satisfies the familiar conditions for a norm:

$$\|af\|_p = |a|\|f\|_p \text{ and } \|f+g\|_p \leq \|f\|_p + \|g\|_p.$$

The first step in analyzing linear transformations of infinite dimensional spaces is the introduction of norms on the domain and range. This was not necessary in finite dimensions but is absolutely essential in the infinite dimensional case. In medical image reconstruction there is a small list of linear transformations that are very important, the Fourier transform, Radon transform and Abel transform. A large part of this text is devoted to the analysis of these operators.

Exercise 1.3.3. Prove that the null-space of $\mathcal{I}$ acting on $\mathcal{C}^0(\mathbb{R})$ is the zero function.

1.4 Conclusion

By examining a large collection of examples we have seen how physical systems can be described using mathematical models. The models suggest measurements which one can make to determine the state of the system. It is important to keep in mind that mathematical models are just that, models, often toy models. A good model must satisfy two opposing requirements: the model should accurately depict the system under study while at the same time being simple enough to be usable. Which models are "simple enough to be useful" depends on what you know, one of our goals, in the succeeding chapters is to develop some sophisticated mathematical tools to work with models. The workhorse throughout this book and in most applications of mathematics to problems in measurement and signal processing is the Fourier transform.

The models used in medical imaging usually involve infinitely many degrees of freedom. The state of the system is described by a function of continuous variables. Ultimately of course only a finite number of measurements can be made and only a finite amount of time is available to process them. Our analysis of the reconstruction process in X-ray CT passes through several stages, beginning with a description of the complete, perfect data situation and concluding with an analysis of the effects of noise on the quality of an approximate image, reconstructed from finitely many measurements.

In mathematics, problems of determining the state of a physical system from feasible measurements are gathered under the rubric of *inverse problems*. The division of problems into inverse problems and *direct problems* is often a matter of history. Usually a physical theory which models how the state of the system determines feasible measurements preceded a description of the inverse process: how the state can be determined from measurements. Example 1.1.6 is typical though very simple example. Formula (1.22) describes the solution to the direct problem: the determination of the transit time from a knowledge of the sound speed *and* the depth. The inverse problem asks for a determination of the depth from a knowledge of the sound speed and the transit time. While many of the problems which arise in medical imaging are considered to be inverse problems, we do not give any systematic development of this subject. The curious reader is referred to the very nice article by Joe Keller which contains analyses of many classical inverse problems, see [41].