Chapter 2 Linear Transformations

2.1 Introduction

We were able to understand the limitations of exact inversion in the deblurring example because, in the Fourier representation, the forward problem could be written as a set of uncoupled scalar equations and therefore consisted of componentwise multiplication. Each Fourier component in the image is multiplied by the corresponding value of the Fourier transform of the point-spread function and then has noise added to give the data. The 'inverse' of the forward problem is then componentwise division of the Fourier transform of the data by the appropriate scalar from the Fourier transform of the point-spread function; problems arose when the divisor had small magnitude. We would like to have a similar description of other linear inverse problems so that we can similarly predict when problems will arise..

In general, a linear forward problem will not necessarily be shift-invariant and a Fourier representation of image and data space will not lead to the forward mapping being a set of uncoupled scalar equations. Usually the size of image and data spaces are different (an extreme case is continuous-discrete problems) in which case the matrix representing the forward mapping is not even square. Yet, remarkably, in all these cases there is an analogue of Fourier space – given by the *singular value decomposition* (SVD) – in which the forward operator reduces to componentwise scalar multiplication and the straightforward inverse is by scalar division.

The simplest case to understand is when both image and data space are finite-dimensional, i.e., for discrete-discrete problems. We will develop the SVD for that case first and then go on to the case where image or data space is continuous.

2.2 A Linear Algebra Primer

Linear algebra forms the mathematical basis for the vector and matrix analysis that we use to analyze linear inverse problems where both image and data space is discrete. Furthermore, when viewed from a certain perspective to be explained in Section 2.11.1, it has a natural generalization to functional analysis, which is the basic tool in analysis of continuous inverse problems.

M-dimensional (real) vector space is denoted $\mathbb{R}^M$ and consists of the set of all ordered M-tuples of real numbers, usually written as a column

$$\mathbf{u} = \begin{pmatrix} u_1 \\ u_2 \\ \vdots \\ u_M \end{pmatrix}$$

so that the subscript denotes a *row* index. Addition of vectors is defined by $(\mathbf{u}+\mathbf{v})_k = u_k + v_k$ and scalar multiplication by $(c\mathbf{u})_k = cu_k$, $c \in \mathbb{R}$. We will sometimes write the scalar to the right of the vector: $\mathbf{u}c = c\mathbf{u}$. This will facilitate our "star" notation for resolutions of unity. In fact, we can view the expression $\mathbf{u}c$ as the matrix product of the column vector $\mathbf{u}$ ($M \times 1$ matrix) and the scalar c (1×1 matrix). This merely amounts to a change of point of view: Rather than thinking of c as multiplying $\mathbf{u}$ to obtain $c\mathbf{u}$, we can think of $\mathbf{u}$ as multiplying c to get $\mathbf{u}c$. The result is, of course, the same.

Similarly, the set of all column vectors with M complex entries is denoted by $\mathbb{C}^M$. $\mathbb{C}^M$ is a complex vector space, with vector addition and scalar multiplication defined exactly as for $\mathbb{R}^M$, the only difference being that now the scalar c may be complex. A subspace of a vector space V is a subset $S \subset V$ that is "closed" under vector addition and scalar multiplication. That is, the sum of any two vectors in S also belongs S, as does any scalar multiple of any vector in S. For example, any plane through the origin is a subspace of $\mathbb{R}^3$, whereas the unit sphere is not. $\mathbb{R}^M$ is a subspace of $\mathbb{C}^M$, provided we use only real scalars in the

scalar multiplication. We therefore say that $\mathbb{R}^M$ is a real subspace of $\mathbb{C}^M$. It often turns out that even when analyzing real objects, complex methods are simpler than real methods. (For example, the complex exponential form of Fourier series is formally simpler than the real form using sines and cosines.)

A basis for $\mathbb{C}^M$ is a collection of vectors $\{\mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_N\}$ such that

1. *any* vector $\mathbf{u} \in \mathbb{C}^M$ can be written as a linear combination of the $\mathbf{b}_n$'s, i.e., $\mathbf{u} = \sum_n c_n \mathbf{b}_n$, and
2. there is just one set of coefficients $\{c_n\}$ for which this can be done.

The scalars c_n are called the *components* of $\mathbf{u}$ *with respect to the basis* $\{\mathbf{b}_n\}$, and they necessarily depend on the choice of basis. It can be shown that any basis for $\mathbb{C}^M$ has exactly M vectors, i.e., we must have $N = M$ above. If $N < M$, then some vectors in $\mathbb{C}^M$ cannot be expressed as linear combinations of the $\mathbf{b}_n$'s and we say that the collection $\{\mathbf{b}_n\}$ is *incomplete*. If $N > M$, then every vector can be expressed in an infinite number of different ways, hence the c_n's are not unique. We then say that the collection $\{\mathbf{b}_n\}$ is *overcomplete*. Condition (1) and (2), above, are equivalent to demanding that the M vectors in a basis are *complete*. Of course, not every collection of M vectors in $\mathbb{C}^M$ is a basis! In order to form a basis, a set of M vectors must be linearly independent, which means that no vector in it can be expressed as a linear combination of all the other vectors.

The standard basis $\{\mathbf{e}_1, \mathbf{e}_2, \ldots, \mathbf{e}_M\}$ in $\mathbb{C}^M$ (as well as in $\mathbb{R}^M$) is defined by

$$(\mathbf{e}_m)_k = \delta_{mk} \equiv \begin{cases} 1, & k = m \\ 0, & k \neq m \end{cases} .$$

in which δ_{mk} is the *Kronecker delta*.

2.2.1 Systems of linear equations

A function A from $\mathbb{C}^N$ to $\mathbb{C}^M$ is said to be linear if

$$A(c\mathbf{u} + \mathbf{v}) = cA(\mathbf{u}) + A(\mathbf{v}) \text{ for all } \mathbf{u}, \mathbf{v} \in \mathbb{C}^N \text{ and all } c \in \mathbb{C}.$$

Thus A preserves the vector space structure of $\mathbb{C}^N$. When A is linear we write $A\mathbf{u}$ instead of $A(\mathbf{u})$ because linear operators act in a way similar to multiplication.

Now let $\{\mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_N\}$ be a basis for $\mathbb{C}^N$ and $\{\mathbf{d}_1, \mathbf{d}_2, \ldots, \mathbf{d}_N\}$ be a basis for $\mathbb{C}^M$. Any vector $\mathbf{u} \in \mathbb{C}^N$ can be written as $\mathbf{u} = \sum_{n=0}^{N} u_n \mathbf{b}_n$ ($\{u_n\}$ are the components of $\mathbf{u}$ in this basis) and so

$$A\mathbf{u} = A \sum_{n=0}^{N} u_n \mathbf{b}_n = \sum_{n=0}^{N} u_n (A\mathbf{b}_n).$$

Since each $A\mathbf{b}_n \in \mathbb{C}^M$ it can be written as $A\mathbf{b}_n = \sum_{m=0}^{M} a_{mn}\mathbf{d}_m$, for some complex numbers a_{mn}, and therefore the components of $A\mathbf{u}$ with respect to the basis $\{\mathbf{d}_m\}$ are

$$(A\mathbf{u})_m = \sum_{n=0}^{N} a_{mn} u_n.$$

Hence the linear function A, given a particular choice of bases, can be represented by the $M \times N$ matrix $[a_{mn}]$. Conversely, any $M \times N$ matrix defines a linear operator from $\mathbb{C}^N$ to $\mathbb{C}^M$ (or possibly $\mathbb{R}^N$ to $\mathbb{R}^M$ if the matrix is real). Note that the operator is basis independent whereas matrices are basis dependent[1].

[1] A simple example of a basis-independent object is a vector which we think of as a geometric arrow. The components of the vector depend on the basis we use to represent it, but the geometric object exists independent of our choice of basis. Operations on the vector such as 'double its length' or 'rotate about some axis by 45 degrees' are linear; their martix representation depends on the basis we choose to describe the space.

The matrix is applied to a vector $\mathbf{f}$ in $\mathbb{R}^N$ by the process of matrix multiplication, e.g.,

$$\mathbf{d} = \mathbf{A}\mathbf{f}. \tag{2.1}$$

There are two common ways of interpreting this matrix product, each is useful in different circumstances. The one which we shall most commonly think about is where the matrix $\mathbf{A}$ is considered as a collection of n column vectors $\mathbf{c}_1, \mathbf{c}_2, ..., \mathbf{c}_n$ written next to each other as shown

$$\mathbf{A} = \begin{pmatrix} \vdots & \vdots & & \vdots \\ \mathbf{c}_1 & \mathbf{c}_2 & \cdots & \mathbf{c}_n \\ \vdots & \vdots & & \vdots \end{pmatrix}. \tag{2.2}$$

The product $\mathbf{Af}$ may then be seen as the calculation of a linear combination of the vectors $\mathbf{c}_1, \mathbf{c}_2, ..., \mathbf{c}_n$ with the components of $\mathbf{f}$ giving the coefficients.

$$\mathbf{Af} = \begin{pmatrix} \vdots & \vdots & & \vdots \\ \mathbf{c}_1 & \mathbf{c}_2 & \cdots & \mathbf{c}_n \\ \vdots & \vdots & & \vdots \end{pmatrix} \begin{pmatrix} f_1 \\ f_2 \\ \vdots \\ f_n \end{pmatrix} \tag{2.3}$$

$$= f_1\mathbf{c}_1 + f_2\mathbf{c}_2 + \cdots + f_n\mathbf{c}_n. \tag{2.4}$$

In this interpretation of the matrix vector product, we highlight the fact that the image of the linear transformation is spanned by the columns of the matrix $\mathbf{A}$. When we now consider the system of equations

$$\mathbf{Af} = \mathbf{d} \tag{2.5}$$

we are asking how we can synthesize the vector $\mathbf{d}$ by taking the appropriate linear combination (specified by the solution $\mathbf{f}$) of the columns of the matrix $\mathbf{A}$.

An alternative way of viewing the system of equations (2.5) is as a set of M simultaneous equations in N unknowns. From this point of view, we focus on the *rows* rather than the columns of $\mathbf{A}$ and write the matrix product as

$$\begin{pmatrix} \cdots & \mathbf{r}_1^{\mathrm{T}} & \cdots \\ \cdots & \mathbf{r}_2^{\mathrm{T}} & \cdots \\ & \vdots & \\ \cdots & \mathbf{r}_m^{\mathrm{T}} & \cdots \end{pmatrix} \mathbf{f} = \begin{pmatrix} d_1 \\ d_2 \\ \vdots \\ d_m \end{pmatrix} \tag{2.6}$$

where $\mathbf{r}_i^{\mathrm{T}}$ is the i^{th} row of the matrix $\mathbf{A}$. The i^{th} simultaneous equation that is to be satisfied is the inner product relationship $\mathbf{r}_i^{\mathrm{T}}\mathbf{x} = d_i$ which defines in hyperplane in $\mathbb{R}^n$. The solution process may be thought of as finding the intersection of all of these hyperplanes.

Traditionally, a *system* of equations $\mathbf{Af} = \mathbf{d}$ is said to be *overdetermined* if there is no solution and *underdetermined* if there are infinitely many solutions. In the theory of inverse problems, we are trying to reconstruct $\mathbf{f}$ from a measurement of $\mathbf{d}$. As we shall see below, however, we need to be concerned with more than just the system of equations for the *specific* measurement result $\mathbf{d}$. Due to measurement errors, the measured value of $\mathbf{d}$ is necessarily different from $\mathbf{Af}$, and it is necessary to obtain good reconstructions even in the presence of noise. These considerations make it necessary to study the properties of the *operator* $\mathbf{A}$ in more detail.

2.2.2 Inner Products

The standard inner product in $\mathbb{C}^N$ generalizes the dot product in $\mathbb{R}^N$ and is defined as

$$\langle \mathbf{u}, \mathbf{v} \rangle = \sum_{n=1}^{N} \bar{u}_n v_n, \qquad \mathbf{u}, \mathbf{v} \in \mathbb{C}^N$$

where u_n and v_n are the components of $\mathbf{u}$ and $\mathbf{v}$ with respect to the standard basis and $\bar{u}_n$ denotes the complex conjugate[2] of u_n. When $\mathbf{u}$ and $\mathbf{v}$ are real vectors, this reduces to the usual inner product in $\mathbb{R}^N$. The standard inner product in $\mathbb{C}^N$ between $\mathbf{u}$ and $\mathbf{v}$ can also be written in matrix form as $\mathbf{u}^{\mathrm{H}}\mathbf{v}$ where the superscript H denotes the Hermitian conjugate of the matrix. This reduces to $\mathbf{u}^{\mathrm{T}}\mathbf{v}$ for $\mathbb{R}^N$. Using this definition, the *norm* of $\mathbf{u}$ defined by

$$||\mathbf{u}|| = \sqrt{\langle \mathbf{u}, \mathbf{u} \rangle} = \left(\sum_{n=1}^{N} |u_n|^2 \right)^{\frac{1}{2}} \tag{2.7}$$

is real and nonnegative, which is sensible. Thus the inner-product and the norm it induces satisfy the following important properties:

(P) *Positivity:* $||\mathbf{u}|| > 0$ for all $\mathbf{u} \neq \mathbf{0}$ in $\mathbb{C}^N$, and $||\mathbf{0}|| = 0$.

(H) *Hermiticity:* $\overline{\langle \mathbf{u}, \mathbf{v} \rangle} = \langle \mathbf{v}, \mathbf{u} \rangle$ for all $\mathbf{u}, \mathbf{v}$ in $\mathbb{C}^N$.

(L) *Linearity:* $\langle \mathbf{u}, c\mathbf{v} + \mathbf{w} \rangle = c \langle \mathbf{u}, \mathbf{v} \rangle + \langle \mathbf{u}, \mathbf{w} \rangle$ for all $\mathbf{u}, \mathbf{v}, \mathbf{w}$ in $\mathbb{C}^N$ and $c \in \mathbb{C}$.

Note that linearity only holds for the second argument. From (L) and (H) it follows that the inner product is *anti-linear* in the first argument, i.e.

$$\langle c\mathbf{u} + \mathbf{w}, \mathbf{v} \rangle = \bar{c} \langle \mathbf{u}, \mathbf{v} \rangle + \langle \mathbf{w}, \mathbf{v} \rangle .$$

Inner-products other than the standard one are often used, particularly when a non-standard inner-product is more natural. Such a circumstance occurs when the different coordinates correspond to quantities with different units or scales and we end up with the inner-product

$$\langle \mathbf{u}, \mathbf{v} \rangle = \sum_{n=1}^{N} \mu_n \bar{u}_n v_n,$$

where the μ_n are positive weights. This inner-product, like all others, also satisfies properties (P), (H), and (L).

While the norm is defined in terms of the inner-product, it is possible to recover the inner-product from the norm using the *polarization identity*:

$$\langle \mathbf{u}, \mathbf{v} \rangle = \frac{1}{4} \left(||\mathbf{u} + \mathbf{v}||^2 - ||\mathbf{u} - \mathbf{v}||^2 \right) + \frac{1}{4\mathrm{i}} \left(||\mathbf{u} + \mathrm{i}\mathbf{v}||^2 - ||\mathbf{u} - \mathrm{i}\mathbf{v}||^2 \right). \tag{2.8}$$

For real vectors, the second term is zero and the identity reduces to $\langle \mathbf{u}, \mathbf{v} \rangle = \frac{1}{4} \left(||\mathbf{u} + \mathbf{v}||^2 - ||\mathbf{u} - \mathbf{v}||^2 \right)$.

Vectors $\mathbf{u}, \mathbf{v}$ in $\mathbb{C}^N$ are said to be *orthogonal* (w.r.t. the inner-product) if $\langle \mathbf{u}, \mathbf{v} \rangle = 0$. A basis $\{\mathbf{b}_n\}$ is *orthonormal* (w.r.t. the inner-product) if $\langle \mathbf{b}_k, \mathbf{b}_l \rangle = \delta_{kl}$. Hence the concept of orthogonality is relative to the choice of inner-product, as is the concept of length. The standard inner-product has the useful property that the standard basis is orthonormal with respect to that inner-product.

2.2.3 Resolution of the Identity

Given an **orthonormal** basis $\{\mathbf{b}_k\}_{k=1}^N$ with respect to some inner product, the **projection** of a vector $\mathbf{v}$ along the basis vector $\mathbf{b}_k$ is given by:

$$\mathbf{b}_k \langle \mathbf{b}_k, \mathbf{v} \rangle \tag{2.9}$$

This is a vector of length $\langle \mathbf{b}_k, \mathbf{v} \rangle$ in the direction of the unit vector $\mathbf{b}_k$.

[2] We use the convention of conjugating the first variable which is common in physics. In mathematics the convention is to conjugate the second variable. Our choice has the benefit of making the "star" operator linear.

If the projections of $\mathbf{v}$ along all of the basis vectors are added together, the result must be equal to $\mathbf{v}$, since the basis vectors span the entire space, i.e.,

$$\mathbf{v} = \sum_{k=1}^{N} \mathbf{b}_k \, \langle \mathbf{b}_k, \mathbf{v} \rangle \tag{2.10}$$

For a basis $\{\mathbf{b}_k\}_{k=1}^N$ which is not necessarily orthonormal, it is still possible to write any vector $\mathbf{v}$ as a linear combination of the basis vectors $\mathbf{b}_k$, but the coefficients are not simply given by inner products of the form $\langle \mathbf{b}_k, \mathbf{v} \rangle$.

If we now consider the standard inner product on $\mathbb{C}^N$, the projection of $\mathbf{v}$ on $\mathbf{b}_k$ may be written as $\mathbf{b}_k \left(\mathbf{b}_k^{\mathrm{H}} \mathbf{v} \right)$, and

$$\mathbf{v} = \sum_{k=1}^{N} \mathbf{b}_k \, \left(\mathbf{b}_k^{\mathrm{H}} \mathbf{v} \right) = \left(\sum_{k=1}^{N} \mathbf{b}_k \, \mathbf{b}_k^{\mathrm{H}} \right) \mathbf{v}. \tag{2.11}$$

Since this is true for all vectors $\mathbf{v}$,

$$\mathbf{I} = \sum_{k=1}^{N} \mathbf{b}_k \, \mathbf{b}_k^{\mathrm{H}} \tag{2.12}$$

where $\mathbf{I}$ is the $N \times N$ identity matrix. Note that since $\mathbf{b}_k$ is a column N-vector, $\mathbf{b}_k \, \mathbf{b}_k^{\mathrm{H}}$ is an $N \times N$ matrix for each k, called the outer product of $\mathbf{b}_k$ with itself. The representation of $\mathbf{I}$ in terms of the outer products of vectors in an orthonormal basis is called a *resolution of the identity*. The matrix $\mathbf{b}_k \, \mathbf{b}_k^{\mathrm{H}}$ is called a *projection operator* along $\mathbf{b}_k$.

2.3 The Linear Inverse Problem

We consider several specific examples which may be regarded as prototypical of the linear inverse problems which are encountered in practice. These should be kept in mind as we discuss linear transformations more abstractly in the subsequent sections. Schematically, all linear inverse problems may be represented by the block diagram shown. The input (or "image") $\mathbf{f}$ is the collection of quantities we wish to reconstruct and the output $\mathbf{d}$ are the data we measure. In a linear inverse problem, the relationship between $\mathbf{f}$ and $\mathbf{d}$ is

$$\mathbf{d} = \mathbf{A}\mathbf{f} + \mathbf{n} \tag{2.13}$$

where $\mathbf{A}$ is a *linear* transformation, and $\mathbf{n}$ represents an *additive* noise process which prevents us from knowing the noise-free data $\mathbf{y} = \mathbf{A}\mathbf{f}$ precisely. Different applications give rise to image and data spaces of different sizes, but we shall wish to present a unified treatment of all of these.

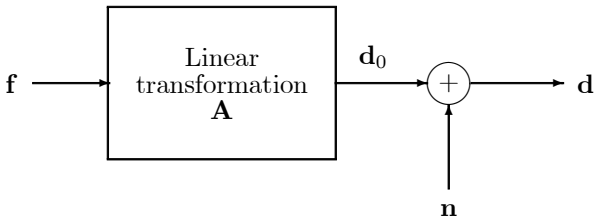

2.3.1 Model Fitting to Data

In the simplest case, we measure a collection of data $\{(x_i, y_i)\}_{i=1}^m$ and attempt to fit a straight line $y = f_0 + f_1 x$ to the data, where f_1 is the gradient and f_0 is the intercept. The dimension of the "image" space is only

two, but the data space is m dimensional. In terms of the general picture, we have

$$\underbrace{\begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_m \end{pmatrix}}_{\mathbf{d}} = \underbrace{\begin{pmatrix} 1 & x_1 \\ 1 & x_2 \\ \vdots & \vdots \\ 1 & x_m \end{pmatrix}}_{\mathbf{A}} \underbrace{\begin{pmatrix} f_0 \\ f_1 \end{pmatrix}}_{\mathbf{f}} + \underbrace{\begin{pmatrix} n_1 \\ n_2 \\ \vdots \\ n_m \end{pmatrix}}_{\mathbf{n}} \tag{2.14}$$

The noise $\mathbf{n}$ here represents the uncertainty in our measurements of y_i. We normally would make $m \gg 2$ measurements, and find (unless we are very lucky) that the points do not lie exactly on any straight line. Nevertheless, we usually "solve" for f_0 and f_1 by finding the line which is best in the least-squares sense. i.e., we find $\hat{\mathbf{f}}$ so as to minimize the value of $\left\|\mathbf{d} - \mathbf{A}\hat{\mathbf{f}}\right\|^2$. So long that $m \geq 2$, and that $x_1 \neq x_2$ there is a unique solution for $\hat{\mathbf{f}}$.

Instead of a straight line, we may fit a polynomial of the form $y = f_0 + f_1 x + \cdots + f_n x^n$ to the data. Even more generally, we need not use powers of x, but may be given a collection of n functions $g_1(x)$, $g_2(x), \ldots, g_n(x)$ and be required to fit $y = f_1 g_1(x) + f_2 g_2(x) + \cdots + f_n g_n(x)$. In this case, we write

$$\underbrace{\begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_m \end{pmatrix}}_{\mathbf{d}} = \underbrace{\begin{pmatrix} g_1(x_1) & g_2(x_1) & \ldots & g_n(x_1) \\ g_1(x_2) & g_2(x_2) & \ldots & g_n(x_2) \\ \vdots & \vdots & \ddots & \vdots \\ g_1(x_m) & g_2(x_m) & \ldots & g_n(x_m) \end{pmatrix}}_{\mathbf{A}} \underbrace{\begin{pmatrix} f_1 \\ f_1 \\ \vdots \\ f_n \end{pmatrix}}_{\mathbf{f}} + \underbrace{\begin{pmatrix} n_1 \\ n_2 \\ \vdots \\ n_m \end{pmatrix}}_{\mathbf{n}} \tag{2.15}$$

Once again, we would normally make $m \gg n$ measurements in order to find the values of $f_1, \ldots, f_n$. The situation in which there are much more data than "image" parameters to fit is generally called "model-fitting". It is usually the case in model fitting that the model is not a perfect fit to the data for any choice of the parameters, and we have to use some criterion such as least squares to obtain a "best" solution. It is also usually the case that there is a unique best solution in the least squares sense. One can see how things can go wrong, if for example the functions $g_1(x), \ldots, g_n(x)$ are poorly chosen. For example, if the functions are not linearly independent, there may be an infinite number of choices of $\mathbf{f}$, all of which give the same value of $\mathbf{Af}$. In such a situation, solutions will not be unique. In the following we shall consider more precisely the conditions under which solutions exist and to what extent they are unique.

2.3.2 Indirect Imaging

In model fitting, the size of the "image" is taken to be much smaller than the size of the data. If this is not the case, the parameters of the model are not well-determined by the data. For example, if we are fitting a parabola to data which consist only of two points, there are an infinite number of parabolas which pass exactly through the data. There are however problems in which the size of the image space is comparable to or larger than the size of data space. This situation often occurs in practice because images are often functions of continuous variables such as time or (physical) space, so that the dimensionality of image space is infinite. On the other hand, only a finite number of data can be measured, so that data space is necessarily finite dimensional. An example of a problem of this type is that of image deconvolution discussed in the previous chapter. The true image is actually a function defined on a plane and hence lies in an infinite dimensional space, but we made the space finite dimensional by discretizing the image into pixels. As discussed previously, the data consist of linear combinations of the values in the image. In the example we chose to discretize the blurred picture (the data) so that it has the same number of points as the image. This makes it convenient from the point of view of using discrete Fourier transforms, but there is no fundamental reason for doing this. As a second example, consider again the problem of X-ray tomography, in which line integrals through the patient are recorded at various angles in order to reconstruct the distribution of absorption within the body. The dimension of the image space is again infinite, but can be made finite by suitable discretizing the cross sectional area of the patient. The dimension of the data space depends on the number of angles used, and the number of points at which the X-ray intensity is measured.

The fineness of the discretization of the image is to some extent arbitrary, so long as it is fine enough that details can be represented adequately. As we refine the discretization, the number of image points increases, and at some stage, there will be an infinite number of images which will all map to exactly the same data under the linear transformation. (This will certainly happen once the number of image points exceeds the number of data points, but it can also happen well before that). Choosing a good reconstruction from among the infinity of possibilities then becomes an issue, since we would not like the appearance of the solution to change markedly as the discretization is refined beyond a certain point.

We shall refer to problems in which the number of image points is comparable to or greater than the number of data points as "indirect imaging." The distinction between model fitting and indirect imaging is qualitative rather than quantitative and we shall thus study them as special cases of the linear inverse problem.

2.4 Anatomy of a linear transformation

One way to treat the system of equations

$$\mathbf{d} = \mathbf{A}\mathbf{f} \tag{2.16}$$

where $\mathbf{A}$ is rectangular (and hence certainly not invertible) is to consider the operator $\mathbf{A}^{\mathrm{T}}\mathbf{A}$ (or $\mathbf{A}\mathbf{A}^{\mathrm{T}}$) which is square and potentially invertible. This idea of analyzing the properties of $\mathbf{A}^{\mathrm{T}}\mathbf{A}$ (or $\mathbf{A}\mathbf{A}^{\mathrm{T}}$) in order to give information about $\mathbf{A}$ leads to a very important way of characterizing the behaviour of any finite dimensional linear transformation called the *singular value decomposition.* We shall investigate this in greater detail after a preliminary review of the eigenspace properties of real symmetric matrices. The advantage of first considering symmetric (and thus square) matrices is that the linear transformation defined by such a matrix maps a space into itself, whereas the domain and range spaces of the transformation defined by a rectangular matrix are different.

2.4.1 Eigenvalues and eigenvectors of real symmetric matrices

In elementary courses on linear algebra, it is shown that a real, symmetric (and hence square) $m \times m$ matrix $\mathbf{M}$ always has *real* eigenvalues and the eigenvectors of such a matrix may always be chosen to form an *orthonormal basis* of $\mathbb{R}^m$. If the eigenvalues are denoted μ_i and the corresponding eigenvectors are denoted $\mathbf{u}_i$, the statement that $\mathbf{u}_i$ is an eigenvector associated with eigenvalue μ_i may be written

$$\mathbf{M}\mathbf{u}_i = \mathbf{u}_i\mu_i. \tag{2.17}$$

If we now write the column vectors $\mathbf{u}_1, ..., \mathbf{u}_m$ next to each other to form the square matrix

$$\mathbf{U} = \begin{pmatrix} \vdots & \vdots & & \vdots \\ \mathbf{u}_1 & \mathbf{u}_2 & \cdots & \mathbf{u}_m \\ \vdots & \vdots & & \vdots \end{pmatrix}, \tag{2.18}$$

the relations (2.17) for $i = 1, ..., m$ may be written as

$$\begin{aligned} \mathbf{MU} &= \begin{pmatrix} \vdots & \vdots & & \vdots \\ \mathbf{M}\mathbf{u}_1 & \mathbf{M}\mathbf{u}_2 & \cdots & \mathbf{M}\mathbf{u}_m \\ \vdots & \vdots & & \vdots \end{pmatrix} = \begin{pmatrix} \vdots & \vdots & & \vdots \\ \mu_1\mathbf{u}_1 & \mu_2\mathbf{u}_2 & \cdots & \mu_m\mathbf{u}_m \\ \vdots & \vdots & & \vdots \end{pmatrix} \\ &= \begin{pmatrix} \vdots & \vdots & & \vdots \\ \mathbf{u}_1 & \mathbf{u}_2 & \cdots & \mathbf{u}_m \\ \vdots & \vdots & & \vdots \end{pmatrix} \begin{pmatrix} \mu_1 & & & \\ & \mu_2 & & \\ & & \ddots & \\ & & & \mu_m \end{pmatrix} = \mathbf{UD}, \end{aligned} \tag{2.19}$$

where $\mathbf{D}$ is the diagonal matrix with the eigenvalues on the diagonal. Since $\mathbf{U}$ is an orthogonal matrix, it is invertible and so

$$\mathbf{M} = \mathbf{U}\mathbf{D}\mathbf{U}^{-1} = \mathbf{U}\mathbf{D}\mathbf{U}^{\mathrm{T}}. \tag{2.20}$$

Note that the identity $\mathbf{U}^{-1} = \mathbf{U}^{\mathrm{T}}$ follows from the columns of $\mathbf{U}$ being orthonormal; such a matrix is called *unitary*. The decomposition, in equation 2.20, of the original matrix $\mathbf{M}$ in terms of its eigenvectors may also be written as

$$\mathbf{M} = \sum_{k=1}^{m} \mu_k \mathbf{u}_k \mathbf{u}_k^{\mathrm{T}} \tag{2.21}$$

in which we are highlighting the representation of $\mathbf{M}$ as the weighted sum of outer-products of its eigenvectors. We can verify that the right hand side is indeed equal to $\mathbf{M}$ by considering its action on the eigenvectors $\mathbf{u}_i$. Since these eigenvectors form a basis, a verification which shows that they are transformed correctly ensures (by linearity) that all vectors are also transformed correctly.

The decomposition (2.21) means that the action of the real symmetric matrix $\mathbf{M}$ on an input vector $\mathbf{x} \in \mathbb{R}^m$ may be understood in terms of three steps:

1. It resolves the input vector along each of the eigenvectors $\mathbf{u}_k$, the component of the input vector along the i^{th} eigenvector being given by $\mathbf{u}_k^{\mathrm{T}}\mathbf{x}$,
2. The amount along the k^{th} eigenvector is multiplied by the eigenvalue μ_k,
3. The product tells us how much of the k^{th} eigenvector $\mathbf{u}_k$ is present in the product $\mathbf{M}\mathbf{x}$.

Thus, the eigenvectors of $\mathbf{M}$ define a "privileged" basis in which the action of $\mathbf{M}$ is particularly simple. Each of the components of $\mathbf{x}$ along the m eigenvectors is stretched *independently* by an amount given by the eigenvalue. Figure 2.1 is a schematic representation of the process. Only two of the m orthogonal eigenvectors are shown.

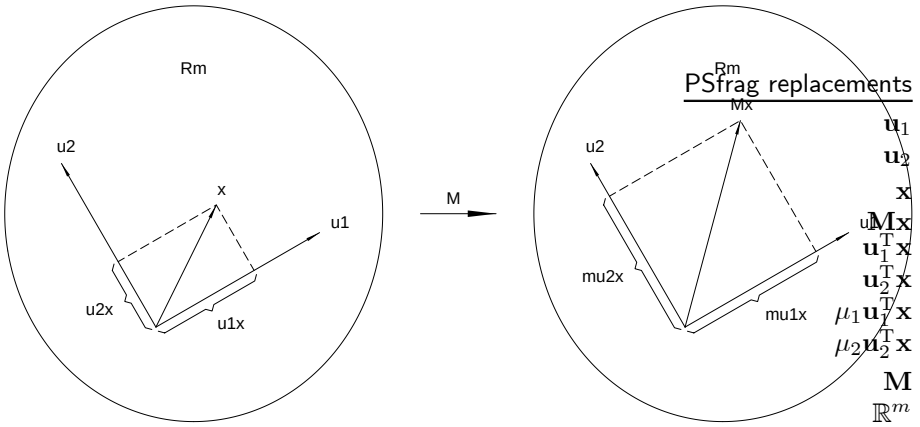

Figure 2.1 Effect of a real symmetric matrix $\mathbf{M}$ on a vector $\mathbf{x}$.

2.4.2 Functions of a real symmetric matrix

Since a real symmetric matrix $\mathbf{M}$ may be written as in (2.20), it is easy to compute any power of $\mathbf{M}$

$$\mathbf{M}^n = \left(\mathbf{U}\mathbf{D}\mathbf{U}^{\mathrm{T}}\right)^n = \mathbf{U}\mathbf{D}^n\mathbf{U}^{\mathrm{T}}, \tag{2.22}$$

since $\mathbf{U}^{\mathrm{T}}\mathbf{U} = \mathbf{U}\mathbf{U}^{\mathrm{T}} = \mathbf{I}$ as $\mathbf{U}$ is unitary. Since $\mathbf{D}$ is diagonal, raising $\mathbf{D}$ to the n^{th} power simply raises each of its (diagonal) elements to the n^{th} power. If we define arbitrary functions of a matrix in terms of the

associated power series, we see that

$$f(\mathbf{M}) = \mathbf{U} f(\mathbf{D}) \mathbf{U}^{\mathrm{T}} \tag{2.23}$$

$$= \sum_{k=1}^{m} f(\mu_k) \mathbf{u}_k \mathbf{u}_k^{\mathrm{T}}. \tag{2.24}$$

In particular, the inverse of the matrix $\mathbf{M}$ is

$$\mathbf{M}^{-1} = \sum_{k=1}^{m} \frac{1}{\mu_k} \mathbf{u}_k \mathbf{u}_k^{\mathrm{T}}. \tag{2.25}$$

The matrix is invertible provided that no eigenvalue is equal to zero. The eigenvectors of $\mathbf{M}^{-1}$ are the same as those of $\mathbf{M}$, only the eigenvalues are reciprocated. Each direction which is stretched when $\mathbf{M}$ is applied contracted by $\mathbf{M}^{-1}$, and vice versa.

2.4.3 Singular value decomposition of a real rectangular matrix

Let us suppose that $\mathbf{A} \in \mathbb{R}^{m \times n}$ is a real rectangular matrix which maps vectors in $\mathbb{R}^n$ to vectors in $\mathbb{R}^m$. We may consider the two square symmetric matrices $\mathbf{A}^{\mathrm{T}}\mathbf{A}$ and $\mathbf{A}\mathbf{A}^{\mathrm{T}}$ which may be formed from $\mathbf{A}$ and which are of size $n \times n$ and $m \times m$ respectively.

Since each of these matrices are square and symmetric, we may obtain the eigenvectors and eigenvalues of each. The eigenvectors may be chosen to form orthonormal bases of the respective spaces. We note that each of the matrices is *positive semidefinite*, which means that all their eigenvalues are non-negative. This is easy to see for if $\mathbf{v}$ is an eigenvector of $\mathbf{A}^{\mathrm{T}}\mathbf{A}$, belonging to eigenvalue λ, then

$$\mathbf{A}^{\mathrm{T}}\mathbf{A}\mathbf{v} = \lambda \mathbf{v}. \tag{2.26}$$

Multiplying on the left by $\mathbf{v}^{\mathrm{T}}$ and grouping the terms,

$$\left(\mathbf{v}^{\mathrm{T}}\mathbf{A}^{\mathrm{T}}\right)(\mathbf{A}\mathbf{v}) = \lambda \left(\mathbf{v}^{\mathrm{T}}\mathbf{v}\right). \tag{2.27}$$

On the left hand side we have a non-negative quantity, the square of the norm of $\mathbf{A}\mathbf{v}$. On the right, $\mathbf{v}^{\mathrm{T}}\mathbf{v}$ is positive and so λ must be non-negative.

Label the n orthonormal eigenvectors of $\mathbf{A}^{\mathrm{T}}\mathbf{A}$ as $\mathbf{v}_i$ with associated eigenvalues λ_i and assume that we have sorted them so that

$$\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n \geq 0. \tag{2.28}$$

Similarly, label the m orthonormal eigenvectors of $\mathbf{A}\mathbf{A}^{\mathrm{T}}$ as $\mathbf{u}_i$ with associated eigenvalues μ_i and sort them so that

$$\mu_1 \geq \mu_2 \geq \cdots \geq \mu_m \geq 0. \tag{2.29}$$

Consider the first eigenvector $\mathbf{v}_1$ of $\mathbf{A}^{\mathrm{T}}\mathbf{A}$ and suppose that λ_1 is not equal to zero. The vector $\mathbf{A}\mathbf{v}_1$ is then non-zero. We wish to show that $\mathbf{A}\mathbf{v}_1$ is in fact an eigenvector of $\mathbf{A}\mathbf{A}^{\mathrm{T}}$. To check this, we notice that

$$\left(\mathbf{A}\mathbf{A}^{\mathrm{T}}\right)(\mathbf{A}\mathbf{v}_1) = \mathbf{A}\left(\mathbf{A}^{\mathrm{T}}\mathbf{A}\right)\mathbf{v}_1 = \lambda_1 (\mathbf{A}\mathbf{v}_1). \tag{2.30}$$

This shows that $\mathbf{A}\mathbf{v}_1$ is indeed an eigenvector of $\mathbf{A}\mathbf{A}^{\mathrm{T}}$ which belongs to the eigenvalue λ_1. If we normalize $\mathbf{A}\mathbf{v}_1$ to have unit length by forming

$$\frac{\mathbf{A}\mathbf{v}_1}{\|\mathbf{A}\mathbf{v}_1\|}, \tag{2.31}$$

this is a normalized eigenvector of $\mathbf{A}\mathbf{A}^{\mathrm{T}}$ and so must be one of the $\mathbf{u}_i$ mentioned above provided that the eigenvalues of $\mathbf{A}\mathbf{A}^{\mathrm{T}}$ are not degenerate.

Continuing the above argument, we see that each non-zero eigenvalue λ_i of $\mathbf{A}^{\mathrm{T}}\mathbf{A}$ must also be an eigenvalue of $\mathbf{A}\mathbf{A}^{\mathrm{T}}$. A similar argument starting with an eigenvector $\mathbf{u}_i$ of $\mathbf{A}\mathbf{A}^{\mathrm{T}}$ belonging to a non-zero eigenvalue μ_i shows that the vector $\mathbf{A}^{\mathrm{T}}\mathbf{u}_1/\left\|\mathbf{A}^{\mathrm{T}}\mathbf{u}_1\right\|$ is a normalized eigenvector of $\mathbf{A}^{\mathrm{T}}\mathbf{A}$ with the same eigenvalue.

The conclusion to the above is that the non-zero eigenvalues of $\mathbf{A}^{\mathrm{T}}\mathbf{A}$ are the same as the non-zero eigenvalues of $\mathbf{A}\mathbf{A}^{\mathrm{T}}$ and vice versa. If there are r non-zero eigenvalues, this means that $\lambda_1 = \mu_1, ..., \lambda_r = \mu_r$ and that all subsequent eigenvalues must be zero, i.e., $\lambda_{r+1} = ... = \lambda_n = 0$ and that $\mu_{r+1} = ... = \mu_m = 0$.

It turns out to be possible to arrange that for all $k = 1, ..., r$,

$$\mathbf{u}_k = \frac{\mathbf{A}\mathbf{v}_k}{\|\mathbf{A}\mathbf{v}_k\|} \text{ and } \mathbf{v}_k = \frac{\mathbf{A}^{\mathrm{T}}\mathbf{u}_k}{\|\mathbf{A}^{\mathrm{T}}\mathbf{u}_k\|}. \tag{2.32}$$

This happens automatically if the non-zero eigenvalues of $\mathbf{A}^{\mathrm{T}}\mathbf{A}$ and $\mathbf{A}\mathbf{A}^{\mathrm{T}}$ are non-degenerate. Even if there are degeneracies, it is possible to choose the appropriate linear combinations in the degenerate eigenspaces so that these are true. The value of r is known as the *rank* of the matrix $\mathbf{A}$ (or of the matrix $\mathbf{A}^{\mathrm{T}}$). It should be clear that $r \leq m$ and $r \leq n$.

The norms in (2.32) may be evaluated,

$$\|\mathbf{A}\mathbf{v}_k\|^2 = (\mathbf{A}\mathbf{v}_k)^{\mathrm{T}}(\mathbf{A}\mathbf{v}_k) = \mathbf{v}_k^{\mathrm{T}}\left(\mathbf{A}^{\mathrm{T}}\mathbf{A}\right)\mathbf{v}_k = \lambda_k, \tag{2.33}$$

where the last equality holds because $\mathbf{v}_k$ is an eigenvalue of $\mathbf{A}^{\mathrm{T}}\mathbf{A}$ belonging to λ_k, and because $\mathbf{v}_k$ is normalized so that $\mathbf{v}_k^{\mathrm{T}}\mathbf{v}_k = 1$. Similarly, $\left\|\mathbf{A}^{\mathrm{T}}\mathbf{u}_k\right\|^2 = \mu_k$. Since $\lambda_k = \mu_k > 0$, we may define σ_k to be the square root of the eigenvalue and write

$$\|\mathbf{A}\mathbf{v}_k\| = \left\|\mathbf{A}^{\mathrm{T}}\mathbf{u}_k\right\| = \sigma_k = \sqrt{\lambda_k} = \sqrt{\mu_k}, \tag{2.34}$$

for $k = 1, 2, ..., r$. Equation (2.32) then takes the simple form

$$\mathbf{A}\mathbf{v}_k = \sigma_k\mathbf{u}_k \tag{2.35}$$

$$\mathbf{A}^{\mathrm{T}}\mathbf{u}_k = \sigma_k\mathbf{v}_k. \tag{2.36}$$

The effect of the linear transformation $\mathbf{A}$ on the unit vector $\mathbf{v}_k \in \mathbb{R}^n$ is to take it to the vector $\sigma_k\mathbf{u}_k \in \mathbb{R}^m$ of length σ_k in the direction of the unit vector $\mathbf{u}_k \in \mathbb{R}^m$. The effect of the linear transformation $\mathbf{A}^{\mathrm{T}}$ on the unit vector $\mathbf{u}_k \in \mathbb{R}^m$ is to take it to the vector $\sigma_k\mathbf{v}_k \in \mathbb{R}^n$ of length σ_k in the direction of the unit vector $\mathbf{v}_k \in \mathbb{R}^n$.

On the other hand for $k > r$, the eigenvalue of $\mathbf{A}^{\mathrm{T}}\mathbf{A}$ associated with $\mathbf{v}_k$ is zero and so $\mathbf{A}^{\mathrm{T}}\mathbf{A}\mathbf{v}_k = 0$. Premultiplying this by $\mathbf{v}_k^{\mathrm{T}}$ shows that $\|\mathbf{A}\mathbf{v}_k\| = 0$ and hence that $\mathbf{A}\mathbf{v}_k = 0$. We thus have that

$$\mathbf{A}\mathbf{v}_k = 0 \text{ for } k = r+1, ..., n \tag{2.37}$$

and similarly,

$$\mathbf{A}^{\mathrm{T}}\mathbf{u}_k = 0 \text{ for } k = r+1, ..., m. \tag{2.38}$$

Equations (2.35) and (2.37) together describe how $\mathbf{A}$ acts on the vectors in the basis $\{\mathbf{v}_k\}$ for $k = 1, ..., n$. By linearity, any operator which has the same action as $\mathbf{A}$ on each of the vectors of the basis must be the same as $\mathbf{A}$. Thus we may write

$$\mathbf{A} = \sum_{k=1}^{r} \sigma_k\mathbf{u}_k\mathbf{v}_k^{\mathrm{T}}, \tag{2.39}$$

and it is easy to check (using the orthonormality of the basis $\{\mathbf{v}_k\}$) that the right hand side does have the same action as $\mathbf{A}$ on the basis.

Taking the transpose of (2.39) gives

$$\mathbf{A}^{\mathrm{T}} = \sum_{k=1}^{r} \sigma_k\mathbf{v}_k\mathbf{u}_k^{\mathrm{T}} \tag{2.40}$$

and it is again easy to check that this is consistent with (2.36) and (2.38).

The orthonormal vectors $\{\mathbf{v}_k\}$ are known as the *right singular vectors*, the vectors $\{\mathbf{u}_k\}$ are known as the *left singular vectors*, and the scalars $\{\sigma_k\}$ are called the *singular values* of the matrix $\mathbf{A}$.

We may write the column vectors $\mathbf{u}_k$ next to each other to form an orthogonal $m \times m$ matrix $\mathbf{U}$ and stack the row vectors $\mathbf{v}_k^{\mathrm{T}}$ on top of each other to form the orthogonal $n \times n$ matrix $\mathbf{V}^{\mathrm{T}}$. The equation (2.39) may then be written in matrix form as

$$\mathbf{A} = \mathbf{U}\mathbf{S}\mathbf{V}^{\mathrm{T}}$$

where $\mathbf{S}$ is an $m \times n$ matrix whose only non-zero elements are the first r entries on the diagonal, i.e., $s_{kk} = \sigma_k$.

2.5 Interpretation of the singular value decomposition of a matrix

In this section, we discuss the interpretation of the singular value decomposition (2.39). When the matrix $\mathbf{A}$ acts on a vector $\mathbf{f}$, we may write the product as

$$\mathbf{A}\mathbf{f} = \sum_{k=1}^{r} \mathbf{u}_k \sigma_k \left(\mathbf{v}_k^{\mathrm{T}} \mathbf{f}\right), \tag{2.41}$$

this may again be understood as the sequence:

1. It resolves the input vector along each of the right singular vectors $\mathbf{v}_k$, the component of the input vector along the k^{th} singular vector being given by $\mathbf{v}_k^{\mathrm{T}}\mathbf{f}$,
2. The amount along the k^{th} direction is multiplied by the singular value σ_k,
3. The product tells us how much of the k^{th} left singular vector $\mathbf{u}_k$ is present in the product $\mathbf{A}\mathbf{f}$.

In effect the decomposition shows how a complicated operation such as matrix multiplication can be split into r *independent* multiplications, each of which takes a component along a vector in $\mathbb{R}^n$ and converts it into a component along a vector in $\mathbb{R}^m$. This result is all the more remarkable since $\{\mathbf{v}_k\}_{k=1}^r$ can be extended to an orthonormal basis $\{\mathbf{v}_k\}_{k=1}^n$ for $\mathbb{R}^n$ and $\{\mathbf{u}_k\}_{k=1}^r$ can be extended to an orthonormal basis $\{\mathbf{u}_k\}_{k=1}^m$ for $\mathbb{R}^m$. The action of $\mathbf{A}$ on a vector $\mathbf{f}$ is shown schematically in Figure 2.2. For convenience, we only show two of the n dimensions in the domain and two of the m dimensions in the range. This figure is rather similar to Figure 2.1, but notice that the bases in the two spaces are now *different*, even though they can both be chosen to be orthonormal.

The action of the transpose of $\mathbf{A}$ can also be worked out using the singular value decomposition. If $\mathbf{y} \in \mathbb{R}^m$, we see that,

$$\mathbf{A}^{\mathrm{T}}\mathbf{y} = \sum_{k=1}^{r} \mathbf{v}_k \sigma_k \left(\mathbf{u}_k^{\mathrm{T}} \mathbf{y}\right), \tag{2.42}$$

which may be understood as the sequence:

1. Resolve the vector $\mathbf{y}$ along each of the left singular vectors $\mathbf{u}_k$, the component of the input vector along the k^{th} singular vector being given by $\mathbf{u}_k^{\mathrm{T}}\mathbf{y}$,
2. The amount along the k^{th} direction is multiplied by the singular value σ_k,
3. The product tells us how much of the k^{th} right singular vector $\mathbf{v}_k$ is present in the product $\mathbf{A}^{\mathrm{T}}\mathbf{y}$.

These steps are shown in Figure 2.3. Note that the singular values for $\mathbf{A}^{\mathrm{T}}$ are the same as those for $\mathbf{A}$.

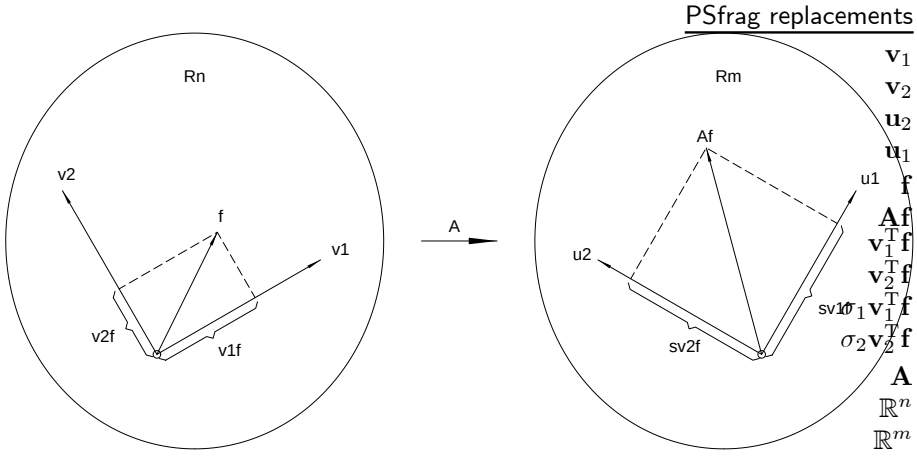

Figure 2.2 Effect of a rectangular matrix $\mathbf{A} \in \mathbb{R}^{m\times n}$ on a vector $\mathbf{f} \in \mathbb{R}^n$.

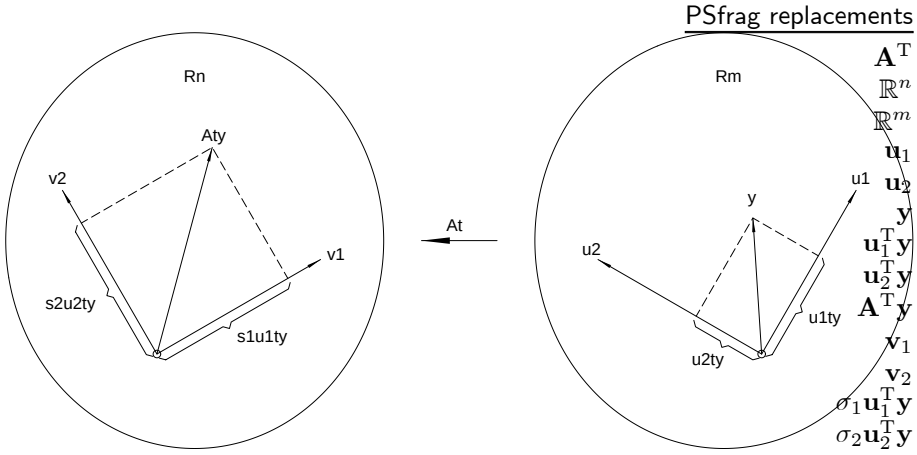

Figure 2.3 Effect of a rectangular matrix $\mathbf{A}^{\mathrm{T}} \in \mathbb{R}^{n\times m}$ on a vector $\mathbf{y} \in \mathbb{R}^m$.

2.6 Geometry of a linear transformation

Equations (2.35) and (2.38) tell us that the image of $\mathbf{A}$ is spanned by $\mathbf{u}_1, ..., \mathbf{u}_r$ and the null space of $\mathbf{A}^{\mathrm{T}}$ (i.e., the space which is mapped to zero under the action of $\mathbf{A}^{\mathrm{T}}$) is spanned by $\mathbf{u}_{r+1}, ..., \mathbf{u}_m$. Together they span all of $\mathbb{R}^m$, and we write $\mathbb{R}^m = \text{image}\,(\mathbf{A}) \oplus \text{null}\left(\mathbf{A}^{\mathrm{T}}\right)$. Similarly, from equations (2.37) and (2.36), the image of $\mathbf{A}^{\mathrm{T}}$ is spanned by $\mathbf{v}_1, ..., \mathbf{v}_r$ and the null space of $\mathbf{A}$ is spanned by $\mathbf{v}_{r+1}, ..., \mathbf{v}_n$. Together they span all of $\mathbb{R}^n$, and we write $\mathbb{R}^n = \text{image}\left(\mathbf{A}^{\mathrm{T}}\right) \oplus \text{null}\,(\mathbf{A})$. The symbol $\oplus$ denotes a *direct sum* of the spaces. If we write $C = A \oplus B$ where A, B and C are vector spaces, this means that any vector $\mathbf{c} \in C$ can be written in a unique way as the sum of a vector $\mathbf{a} \in A$ and a vector $\mathbf{b} \in B$. The above direct sums are also *orthogonal*, so that the vectors $\mathbf{a}$ and $\mathbf{b}$ are at right angles to each other. The dimensions satisfy $\dim C = \dim A + \dim B$.

It is convenient to visualize these relationships using the diagram of Figure 2.4. In the image space $\mathbb{R}^n$, there are r dimensions associated with the image of $\mathbf{A}^{\mathrm{T}}$ and $n - r$ dimensions associated with the null space of $\mathbf{A}$. Since we cannot draw more than two dimensions on a page, we represent each of these spaces by a single axis. Since the spaces are orthogonal, the axes are drawn at right angles to each other. Similarly, in data space $\mathbb{R}^m$, there are r dimensions associated with the image of $\mathbf{A}$ and $m - r$ dimensions associated with the null space of $\mathbf{A}^{\mathrm{T}}$. In the figure these are also represented schematically as two orthogonal axes.

The action of the linear transformation $\mathbf{A}$ is to map non-zero vectors in the space image $\left(\mathbf{A}^{\mathrm{T}}\right)$ into non-zero vectors in the space image $(\mathbf{A})$. All vectors which are orthogonal to image $\left(\mathbf{A}^{\mathrm{T}}\right)$ (i.e., those orthogonal to every *row* of $\mathbf{A}$) are necessarily in null $(\mathbf{A})$ and are mapped to zero under the action of $\mathbf{A}$.

Similarly, the action of the linear transformation $\mathbf{A}^{\mathrm{T}}$ is to map non-zero vectors in the space image $(\mathbf{A})$

into non-zero vectors in the space image $\left(\mathbf{A}^{\mathrm{T}}\right)$. All vectors which are orthogonal to image $(\mathbf{A})$ (i.e., those orthogonal to every *column* of $\mathbf{A}$) are necessarily in null $\left(\mathbf{A}^{\mathrm{T}}\right)$ and are mapped to zero under the action of $\mathbf{A}^{\mathrm{T}}$.

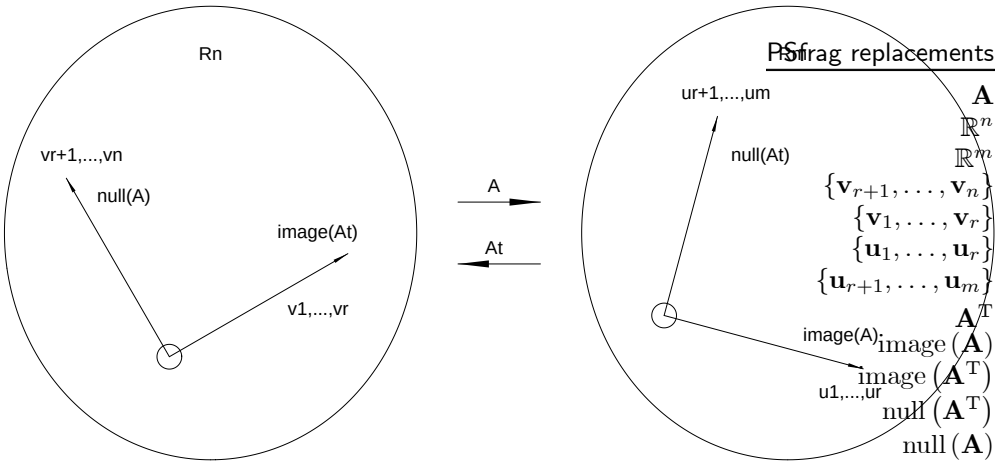

Figure 2.4 The image and null spaces of $\mathbf{A}$ and $\mathbf{A}^{\mathrm{T}}$ for a linear transformation of rank r.

Let us now see how the singular value decomposition is useful in revealing how a linear transformation converts an "image" vector $\mathbf{f}$ into a "data" vector $\mathbf{Af}$. Any image vector $\mathbf{f}$ may be written as the sum of right singular vectors $\mathbf{v}_k$, the length of the component being given by the projection along the singular vector,

$$\mathbf{f} = \sum_{k=1}^{n} f_k \mathbf{v}_k \text{ where } f_k = \mathbf{v}_k^{\mathrm{T}} \mathbf{f}. \tag{2.43}$$

After this is converted into "data " by multiplication by $\mathbf{A}$, the result is

$$\mathbf{Af} = \sum_{k=1}^{r} (\sigma_k f_k)\, \mathbf{u}_k. \tag{2.44}$$

Information about the projection of $\mathbf{f}$ along $\mathbf{v}_k$ is "encoded" in the data as the component along the direction $\mathbf{u}_k$. The size of the projection is multiplied by σ_k to give the coefficient of $\mathbf{u}_k$, namely $\sigma_k f_k$. The singular value is the *factor* which tells us by how much each component defining the image is amplified or attenuated when it is converted into the data.

Notice that only the projections of the image $\mathbf{f}$ along the first r right singular vectors $\mathbf{v}_k$ play a role in determining the data $\mathbf{Af}$. If $r < n$, this means that the data are "blind" to certain aspects of the image: the data do not allow us to distinguish between images which have the same projections along the first r right singular vectors. Such images will look different, since they may differ in their projections along the remaining singular vectors. Equivalently, we may add to an image *any* element of the null space of $\mathbf{A}$ and the data will not be changed in any way. If we now think of solving the inverse problem of reconstructing $\mathbf{f}$ from a measurement of $\mathbf{d}$, it is clear that at best, those components of $\mathbf{f}$ along the first r right singular vectors are determined by the data. However the data tell us nothing at all about the components of $\mathbf{f}$ along the remaining $n-r$ right singular vectors, and we have to use some other means for determining these components.

The above indicates that by calculating and plotting the right singular vectors, we can get an idea of what types of structure in the image will be visible in the data and also the types of structure which are invisible in the data.

2.7 The Singular Value Decomposition in Model Fitting Problems

In model-fitting problems, such as fitting of a straight line $y = f_0 + f_1 x$ to a collection of m data points $\{(x_k, y_k)\}_{k=1}^m$, the dimensionality of the image space is very low. The forward problem is

$$\underbrace{\begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_m \end{pmatrix}}_{\mathbf{d}} = \underbrace{\begin{pmatrix} 1 & x_1 \\ 1 & x_2 \\ \vdots & \vdots \\ 1 & x_m \end{pmatrix}}_{\mathbf{A}} \underbrace{\begin{pmatrix} f_0 \\ f_1 \end{pmatrix}}_{\mathbf{f}}, \tag{2.45}$$

so image space is two-dimensional ($n = 2$), while data space is m dimensional. For model-fitting to give well-defined answers, the rank r of the matrix $\mathbf{A}$ is equal to n, so that the image of $\mathbf{A}$ is a two-dimensional subspace of $\mathbb{R}^m$. The data sets which lie in the image of $\mathbf{A}$ are those for which all the points lie exactly on some straight line. For most data sets, the points will not all lie on a straight line, and in these cases $\mathbf{d} \notin \text{image}(\mathbf{A})$. In the least-squares approach, the model parameters $\hat{\mathbf{f}}$ are chosen so that the $\mathbf{A}\hat{\mathbf{f}}$ is as close as possible to $\mathbf{d}$, i.e.,

$$\hat{\mathbf{f}} = \arg\min \|\mathbf{d} - \mathbf{A}\mathbf{f}\|^2 \tag{2.46}$$

Suppose we compute the singular value decomposition of $\mathbf{A}$, i.e., we find the left singular vectors $\{\mathbf{u}_k\}_{k=1}^m$, the right singular vectors $\{\mathbf{v}_k\}_{k=1}^n$ and the singular values σ_k such that:

$$\mathbf{A} = \sum_{k=1}^{r} \sigma_k \mathbf{u}_k \mathbf{v}_k^{\mathrm{T}} \tag{2.47}$$

Since $\{\mathbf{u}_k\}_{k=1}^m$ form a basis of data space, we may write the data $\mathbf{d}$ as a the linear combination:

$$\mathbf{d} = \sum_{k=1}^{m} \mathbf{u}_k \left(\mathbf{u}_k^{\mathrm{T}} \mathbf{d}\right). \tag{2.48}$$

Then, given any $\mathbf{f}$, we see that

$$\|\mathbf{d} - \mathbf{A}\mathbf{f}\|^2 = \left\| \sum_{k=1}^{m} \mathbf{u}_k \left(\mathbf{u}_k^{\mathrm{T}} \mathbf{d}\right) - \sum_{k=1}^{r} \sigma_k \mathbf{u}_k \left(\mathbf{v}_k^{\mathrm{T}} \mathbf{f}\right) \right\|^2 \tag{2.49}$$

$$= \left\| \sum_{k=1}^{r} \mathbf{u}_k \left\{\mathbf{u}_k^{\mathrm{T}} \mathbf{d} - \sigma_k \left(\mathbf{v}_k^{\mathrm{T}} \mathbf{f}\right)\right\} + \sum_{k=r+1}^{m} \mathbf{u}_k \left(\mathbf{u}_k^{\mathrm{T}} \mathbf{d}\right) \right\|^2. \tag{2.50}$$

Using the theorem of Pythagorus (since the vectors $\{\mathbf{u}_k\}$ are orthogonal),

$$\|\mathbf{d} - \mathbf{A}\mathbf{f}\|^2 = \sum_{k=1}^{r} \left|\mathbf{u}_k^{\mathrm{T}} \mathbf{d} - \sigma_k \left(\mathbf{v}_k^{\mathrm{T}} \mathbf{f}\right)\right|^2 + \sum_{k=r+1}^{m} \left|\mathbf{u}_k^{\mathrm{T}} \mathbf{d}\right|^2. \tag{2.51}$$

Choosing $\hat{\mathbf{f}}$ so as to minimize $\|\mathbf{d} - \mathbf{A}\mathbf{f}\|^2$ is now straightforward. The second term on the right-hand side is the square of the perpendicular distance from $\mathbf{d}$ to the image of $\mathbf{A}$, and is completely unaffected by the choice of $\mathbf{f}$. The first term on the right hand side can be reduced to zero (its minimum possible value) by choosing $\hat{\mathbf{f}}$ such that

$$\mathbf{v}_k^{\mathrm{T}} \hat{\mathbf{f}} = \frac{\mathbf{u}_k^{\mathrm{T}} \mathbf{d}}{\sigma_k} \text{ for } k = 1, 2, \ldots, r. \tag{2.52}$$

Whether or not this completely determines $\hat{\mathbf{f}}$ depends on whether $r = n$ or $r < n$. For model fitting, $r = n$, and so the unique solution to the model fitting problem is:

$$\hat{\mathbf{f}} = \sum_{k=1}^{n} \mathbf{v}_k \left(\mathbf{v}_k^{\mathrm{T}} \hat{\mathbf{f}}\right) = \sum_{k=1}^{n} \mathbf{v}_k \left(\frac{\mathbf{u}_k^{\mathrm{T}} \mathbf{d}}{\sigma_k}\right) = \left(\sum_{k=1}^{n} \frac{1}{\sigma_k} \mathbf{v}_k \mathbf{u}_k^{\mathrm{T}}\right) \mathbf{d}. \tag{2.53}$$

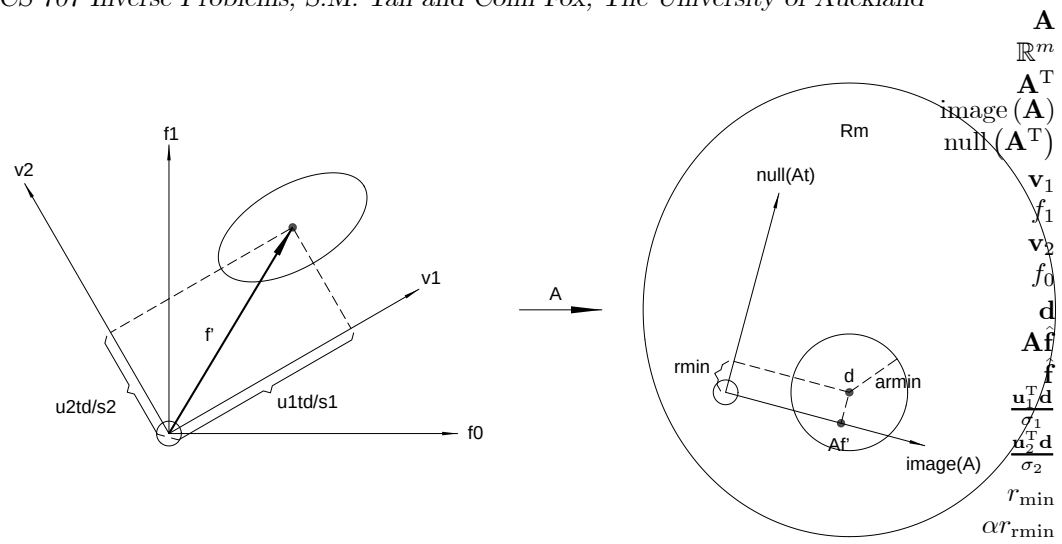

Figure 2.5 Geometry of a model-fitting problem

This process is illustrated for the problem of fitting a straight line in Figure 2.5. Note that the two-dimensional image space is depicted in its entirety, but that image $(\mathbf{A})$, which is a two dimensional subspace in data space, is only depicted schematically by a line. The other $m-2$ dimensions in data space are also depicted as a line perpendicular to image $(\mathbf{A})$. The data $\mathbf{d}$ are shown as being slightly off image $(\mathbf{A})$ due to noise. The reconstucted model parameters $\hat{\mathbf{f}}$ are chosen so that $\mathbf{A}\hat{\mathbf{f}}$ is as close as possible to $\mathbf{d}$ in data space. The components of $\hat{\mathbf{f}}$ along the right singular vectors $\mathbf{v}_k$ are given by $\left(\mathbf{u}_k^{\mathrm{T}}\mathbf{d}\right)/\sigma_k$ as indicated by (2.53). The singular vectors $\mathbf{v}_1$ and $\mathbf{v}_2$ are at right angles to each other, but may make an angle to the f_0 and f_1 axes which represent the intercept and gradient of the fitted straight line, respectively. The ellipse in image space depicts the uncertainty in the fitted parameter values, and will be discussed in a subsequent section.

2.7.1 Relationship to the Moore-Penrose inverse

The Moore-Penrose inverse was introduced in the stage 2 experiment on data fitting using Matlab. In order to minimize the misfit $C = \|\mathbf{d}-\mathbf{A}\mathbf{f}\|^2$, we may write

$$C = \|\mathbf{d}-\mathbf{A}\mathbf{f}\|^2 = \sum_{k=1}^{m}\left(d_k - \sum_{l=1}^{n} a_{kl}f_l\right)^2 \tag{2.54}$$

Then

$$\frac{\partial C}{\partial f_i} = \sum_{k=1}^{m} 2\left(d_k - \sum_{l=1}^{n} a_{kl}f_l\right)(-a_{ki}) = 0 \text{ for } i = 1,\ldots,n \tag{2.55}$$

for an extremum. This may be written as

$$\sum_{l=1}^{n}\left(\sum_{k=1}^{m} a_{ki}a_{kl}\right) f_l = \sum_{k=1}^{m} a_{ki}d_k, \tag{2.56}$$

which in matrix form is

$$\left(\mathbf{A}^{\mathrm{T}}\mathbf{A}\right)\mathbf{f} = \mathbf{A}^{\mathrm{T}}\mathbf{d}. \tag{2.57}$$

These are known as the *normal equations* of the least-squares problem. We obtain a unique solution provided that $\mathbf{A}^{\mathrm{T}}\mathbf{A}$ is invertible, and find the best fit parameters $\hat{\mathbf{f}}$ using

$$\hat{\mathbf{f}} = \left(\mathbf{A}^{\mathrm{T}}\mathbf{A}\right)^{-1}\mathbf{A}^{\mathrm{T}}\mathbf{d} \tag{2.58}$$

The matrix $\left(\mathbf{A}^{\mathrm{T}}\mathbf{A}\right)^{-1}\mathbf{A}^{\mathrm{T}}$ is known as the Moore-Penrose inverse of $\mathbf{A}$.

It is easy to relate this to the singular-value decomposition. Using (2.39) and (2.40),

$$\mathbf{A}^{\mathrm{T}}\mathbf{A} = \left(\sum_{k=1}^{r}\sigma_k\mathbf{v}_k\mathbf{u}_k^{\mathrm{T}}\right)\left(\sum_{l=1}^{r}\sigma_l\mathbf{u}_l\mathbf{v}_l^{\mathrm{T}}\right) = \sum_{k=1}^{r}\sum_{l=1}^{r}\sigma_k\sigma_l\mathbf{v}_k\left(\mathbf{u}_k^{\mathrm{T}}\mathbf{u}_l\right)\mathbf{v}_l^{\mathrm{T}} = \sum_{k=1}^{r}\sigma_k^2\mathbf{v}_k\mathbf{v}_k^{\mathrm{T}} \tag{2.59}$$

since $\mathbf{u}_k^{\mathrm{T}}\mathbf{u}_l = \delta_{kl}$. So $\left\{\sigma_k^2\right\}_{k=1}^{r}$ are the nonzero eigenvalues of $\mathbf{A}^{\mathrm{T}}\mathbf{A}$. Since $\mathbf{A}^{\mathrm{T}}\mathbf{A}$ is an $n\times n$ matrix, it is invertible iff $r=n$. If the matrix is invertible, then

$$\left(\mathbf{A}^{\mathrm{T}}\mathbf{A}\right)^{-1} = \sum_{k=1}^{r}\frac{1}{\sigma_k^2}\mathbf{v}_k\mathbf{v}_k^{\mathrm{T}}, \tag{2.60}$$

and

$$\left(\mathbf{A}^{\mathrm{T}}\mathbf{A}\right)^{-1}\mathbf{A}^{\mathrm{T}} = \left(\sum_{k=1}^{r}\frac{1}{\sigma_k^2}\mathbf{v}_k\mathbf{v}_k^{\mathrm{T}}\right)\left(\sum_{l=1}^{r}\sigma_l\mathbf{v}_l\mathbf{u}_l^{\mathrm{T}}\right) = \sum_{k=1}^{r}\frac{1}{\sigma_k}\mathbf{v}_k\mathbf{u}_k^{\mathrm{T}} \tag{2.61}$$

Comparing this with (2.53) shows that the least squares solution as calculated using the singular value decomposition is identical to that using the Moore Penrose inverse.

2.7.2 Effects of Noise on Model Parameter Estimates

It is instructive to consider how well the model parameters $\hat{\mathbf{f}}$ are determined in a least squares fitting procedure. The method is based on the idea that the value of $\|\mathbf{d}-\mathbf{A}\mathbf{f}\|^2$ is a measure of how unlikely $\mathbf{f}$ is when the measured data are $\mathbf{d}$. The "best" parameter estimate $\hat{\mathbf{f}}$ is thus the one which miimizes $\|\mathbf{d}-\mathbf{A}\mathbf{f}\|^2$, i.e.,

$$\left\|\mathbf{d}-\mathbf{A}\hat{\mathbf{f}}\right\|^2 = \min_{\mathbf{f}\in\mathbb{R}^n}\left\|\mathbf{d}-\mathbf{A}\hat{\mathbf{f}}\right\|^2 = r_{\min}^2 \tag{2.62}$$

where $r_{\min}$ is the distance between the data $\mathbf{d}$ and the image of $\mathbf{A}$, (in which the data should lie, if noise were absent). The value of $r_{\min}$ allows us to estimate the amount of noise likely to be present. In order to get some idea of how confident we are about $\hat{\mathbf{f}}$, we consider the set of probable $\mathbf{f}$ values for the given data and noise level. This is known as the "feasible set",

$$F = \left\{\mathbf{f} : \|\mathbf{d}-\mathbf{A}\mathbf{f}\|^2 \le \alpha r_{\min}^2\right\}, \tag{2.63}$$

where α is chosen according to the confidence level required. We shall show that this set is an ellipse in image space centred about $\hat{\mathbf{f}}$ as shown in Figure.2.5.

Using the singular value decomposition of $\mathbf{A}$, we find from (2.51) that

$$\|\mathbf{d}-\mathbf{A}\mathbf{f}\|^2 = \sum_{k=1}^{r}\left|\mathbf{u}_k^{\mathrm{T}}\mathbf{d}-\sigma_k\left(\mathbf{v}_k^{\mathrm{T}}\mathbf{f}\right)\right|^2 + r_{\min}^2 \tag{2.64}$$

Substituting the values of $\mathbf{v}_k^{\mathrm{T}}\hat{\mathbf{f}}$ from (2.52),

$$\begin{aligned}\|\mathbf{d}-\mathbf{A}\mathbf{f}\|^2 &= r_{\min}^2 + \sum_{k=1}^{r}\left|\sigma_k\left(\mathbf{v}_k^{\mathrm{T}}\hat{\mathbf{f}}\right)-\sigma_k\left(\mathbf{v}_k^{\mathrm{T}}\mathbf{f}\right)\right|^2 \\ &= r_{\min}^2 + \sum_{k=1}^{r}\sigma_k^2\left|\mathbf{v}_k^{\mathrm{T}}\left(\mathbf{f}-\hat{\mathbf{f}}\right)\right|^2\end{aligned} \tag{2.65}$$

From the definition of the set F, we want those $\mathbf{f}$ which satisfy

$$\sum_{k=1}^{r}\sigma_k^2\left|\mathbf{v}_k^{\mathrm{T}}\left(\mathbf{f}-\hat{\mathbf{f}}\right)\right|^2 \le (\alpha-1)\,r_{\min}^2 \tag{2.66}$$

The boundary of this set is an ellipse centred at $\hat{\mathbf{f}}$ with principal axes aligned with the singular vectors $\mathbf{v}_k$. The length of the k'th principal semi-axis is

$$\frac{r_{\min}\sqrt{\alpha-1}}{\sigma_k}, \tag{2.67}$$

from which it is apparent that we are less confident about the parameters along the directions corresponding to small singular values.

As is apparent from Figure 2.5, the *independent* uncertainties along the singular vector directions $\mathbf{v}_k$ translate into *correlated* uncertaities along the axes f_0 and f_1 of the estimated parameters. We shall later consider in more detail how to quantify such correlated uncertainties which arise from parameter fitting.

2.8 The Singular Value Decomposition in General

As the number n of parameters in the model becomes large, it becomes difficult to ensure that they are all independent. Once there are dependent parameters, it becomes possible to achieve the same data vector from more than one set of parameters. Formally, the forward map is not one-one and $\exists \mathbf{f}_1, \mathbf{f}_2$ such that $\mathbf{A}\mathbf{f}_1 = \mathbf{A}\mathbf{f}_2$, but $\mathbf{f}_1 \neq \mathbf{f}_2$. As n becomes large, the problem of model fitting merges into that of indirect imaging. In the regime of indirect imaging, the number of points in image space is chosen to be so large that it can adequately represent the object of interest. Since there are now many images which map to the same data, it becomes necessary to choose from among them using an additional criterion of optimality.

In terms of the singular value decomposition, the rank r of the forward map $\mathbf{A}$ is less than n when the parameters of the model are not independent. This means that the null space of $\mathbf{A}$ contains some non-zero vectors. In fact, all vectors which are linear combinations of $\mathbf{v}_{r+1}, \ldots, \mathbf{v}_n$ are in the null space of $\mathbf{A}$, so that

$$\mathbf{A}\left(c_{r+1}\mathbf{v}_{r+1} + \cdots + c_n\mathbf{v}_n\right) = 0 \tag{2.68}$$

for every choice of coefficients $c_{r+1}, \ldots, c_n$.

Let us now consider what happens if we try to use the principle of least squares to reconstruct the image, when $r < n$. For measured data $\mathbf{d}$, and a trial image $\mathbf{f}$, (2.51) states that

$$\|\mathbf{d} - \mathbf{A}\mathbf{f}\|^2 = \sum_{k=1}^{r} \left|\mathbf{u}_k^{\mathrm{T}}\mathbf{d} - \sigma_k\left(\mathbf{v}_k^{\mathrm{T}}\mathbf{f}\right)\right|^2 + \sum_{k=r+1}^{m} \left|\mathbf{u}_k^{\mathrm{T}}\mathbf{d}\right|^2. \tag{2.69}$$

In order to minimize $\|\mathbf{d} - \mathbf{A}\mathbf{f}\|^2$ over all possible $\mathbf{f}$, the best that we can do is to ensure that the first term on the right hand side is equal to zero, i.e., that

$$\mathbf{v}_k^{\mathrm{T}}\mathbf{f} = \frac{\mathbf{u}_k^{\mathrm{T}}\mathbf{d}}{\sigma_k} \text{ for } k = 1, 2, \ldots, r. \tag{2.70}$$

If $r = n$, this completely defines $\mathbf{f}$, but if $r < n$, we see that only the projections of $\mathbf{f}$ along the first r right singular vectors $\mathbf{v}_1, \ldots, \mathbf{v}_r$ are determined. The projections on the remaining $n - r$ are completely arbitrary. Thus, instead of a single "best" solution $\hat{\mathbf{f}}$, all images of the form:

$$\sum_{k=1}^{r} \mathbf{v}_k \left(\frac{\mathbf{u}_k^{\mathrm{T}}\mathbf{d}}{\sigma_k}\right) + c_{r+1}\mathbf{v}_{r+1} + \cdots + c_n\mathbf{v}_n \tag{2.71}$$

for every choice of coefficients $c_{r+1}, \ldots, c_n$ will give the same minimum value for $\|\mathbf{d} - \mathbf{A}\mathbf{f}\|^2$. Needless to say, if some of the arbitrary coefficients are large, the reconstructed image will look terrible. As mentioned above, we need to select the arbitrary coefficients according to some other critereon of optimality, because the data do *not* determine them at all.

The situation is depicted schematically in Figure 2.6. If $r < m$, it is very likely that noise will cause the data $\mathbf{d}$ to lie outside the image of $\mathbf{A}$. If $\mathbf{y}$ is the point which is closest to $\mathbf{d}$ in image$(\mathbf{A})$, we find that an

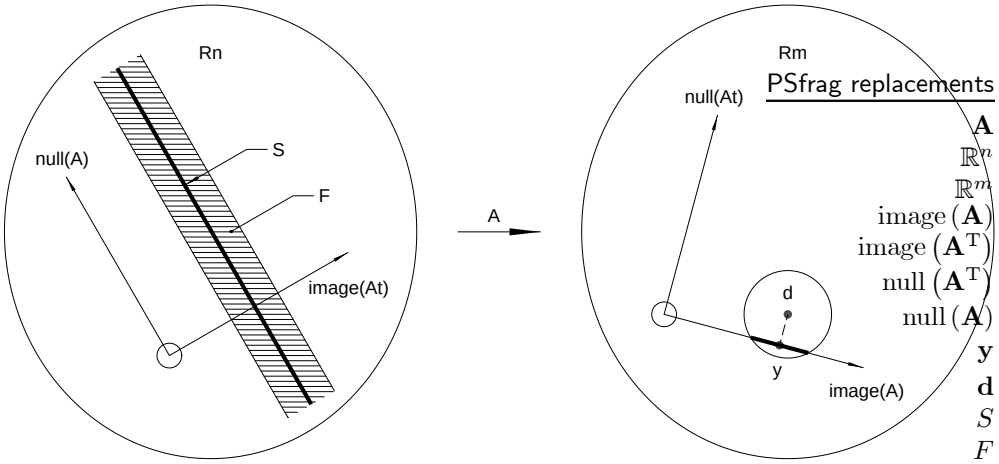

Figure 2.6 Geometry of a linear inverse problem.

infinite number of images $\mathbf{f}$ map to the point $\mathbf{d}$. All such images differ by some vector in null $(\mathbf{A})$, and so the set S in image space which maps to $\mathbf{y}$ is represented schematically by a line parallel to the subspace null $(\mathbf{A})$. When we include the uncertainty in $\mathbf{d}$, as represented by the m-dimensional sphere about $\mathbf{d}$, we are led to consider the feasible set F (as defined by (2.63)) which maps to the intersection of the sphere and image $(\mathbf{A})$. In the image space, the feasible set is represented by a strip as shown. Although feasible in the sense of fitting the data to some specified precision, most of the images in F have wildly oscillatory and non-physical properties (e.g., negative intensities, etc.) The essential goal of regularization methods to be discussed in the next chapter is to choose a "reasonable" solution within the feasible set.

2.9 Classifying linear operators

In the theory of inverse problems it is useful to classify operators $\mathbf{A}$ according to whether the system of equations $\mathbf{Af} = \mathbf{d}$ has solutions for various values of $\mathbf{d}$.

1. If the image of the operator $\mathbf{A}$ has smaller dimension than the data space, i.e., if $r < m$, there are vectors $\mathbf{d} \in \mathbb{R}^m$ which lie outside the range of $\mathbf{A}$. This means that for some possible $\mathbf{d}$, the equations $\mathbf{Af} = \mathbf{d}$ have no solution and it is necessary to use some condition such as least-squares to select a vector $\hat{\mathbf{f}}$ which minimizes the discrepancy between $\mathbf{Af}$ and $\mathbf{d}$. Whether or not this least-squares solution is unique or not depends on the next condition.

2. If the image of the operator $\mathbf{A}$ has smaller dimension than the image space, i.e., if $r < n$, then there are an infinite number of vectors $\mathbf{x}$ all of which map under $\mathbf{A}$ to the same vector $\mathbf{d}$. This means that for some choices of $\mathbf{d}$, the equations $\mathbf{Af} = \mathbf{d}$ have infinitely many solutions. Whether or not solutions exist for all values of $\mathbf{d}$ depends on the previous condition.

3. The only situation in which $\mathbf{Af} = \mathbf{d}$ has a unique solution for every $\mathbf{d}$ is if $r = m = n$ so that the matrix $\mathbf{A}$ is square and invertible. This situation almost *never* holds in practice and it is almost always a bad idea to try to force an inverse problem into a form with a square invertible matrix with a view of solving the problem by a solution of these equations. This reasons for this should become clearer in the following discussions.

Whenever the second condition holds, it is necessary to use additional information over and above the data collected in order to select a good reconstruction from among the possible reconstructions. One way of doing this is by using the process of *regularization* which we shall examine in more detail later. In practice, it is often the case that both the first two conditions hold and r is strictly less than both m and n. When this

is the case, there is no solution to $\mathbf{Af} = \mathbf{d}$ for some $\mathbf{d}$ while there are infinitely many solutions for other possible $\mathbf{d}$. If there is no solution for some $\mathbf{d}$, it makes sense to find $\mathbf{f}$ to minimize $\|\mathbf{d} - \mathbf{Af}\|^2$. It turns out that this minimization gives a unique $\mathbf{f}$ if $r = n$ but there are an infinity of vectors all of which give exactly the same minimum norm if $r < n$.

2.10 The effects of noise and small singular values

In the theory discussed so far, we have drawn a sharp distinction between the eigenvalues of $\mathbf{A}^{\mathrm{T}}\mathbf{A}$ which are non-zero and those which are zero. Indeed the rank r of $\mathbf{A}$ may be defined as the number of non-zero eigenvalues of $\mathbf{A}^{\mathrm{T}}\mathbf{A}$. In practice of course, when the eigenvalues of $\mathbf{A}^{\mathrm{T}}\mathbf{A}$ are sorted in decreasing order, there is a smooth transition from the large eigenvalues through the small eigenvalues to the tiny eigenvalues and the actual rank is always equal to the smaller of m or n. A more useful concept is the effective rank, which depends on the threshold below which we consider the eigenvalue (or the corresponding singular value) to be negligible.

For typical measurement processes, large or slowly-varying portions in the image are well-represented in the data while structures on fine scales or with high frequency components tend to be poorly represented. This is because measurements can usually only be made with a certain spatial or temporal resolution. For example, if we are taking a photograph of an object with visible light, structures on the object with a scale smaller than the wavelength are invisible. The singular vectors in image space associated with the large singular values for such an operator will tend to be smooth, while those associated with the small singular values will tend to be highly irregular or oscillatory.

In order to understand how the small singular values affect the reconstruction process, we consider a simple model for the measurement uncertainty or noise that is always present in data. Let us suppose that the measured data $\mathbf{d}$ may be written as the sum of the transformed image $\mathbf{y} = \mathbf{Af}$ and a noise vector $\mathbf{n}$ so that

$$\mathbf{d} = \mathbf{Af} + \mathbf{n}. \tag{2.72}$$

The vector $\mathbf{f}$ represents the true underlying image and $\mathbf{y}$ is the data that would have been obtained in the absence of noise. Neither of these quantities is known in practice, but the aim of reconstruction is to find a vector approximating to $\mathbf{f}$. Substituting the singular-value decomposition of $\mathbf{A}$ into this yields

$$\mathbf{d} = \left(\sum_{k=1}^{n} \sigma_k \mathbf{u}_k \mathbf{v}_k^{\mathrm{T}} \right) \mathbf{f} + \mathbf{n} \tag{2.73}$$

where the rank has been taken to be n, the size of the image space (assumed to be smaller than m). The forward mapping is strictly 1-1, and so there is a unique least-squares solution which we have seen is given by:

$$\hat{\mathbf{f}} = \sum_{k=1}^{n} \left(\frac{\mathbf{u}_k^{\mathrm{T}} \mathbf{d}}{\sigma_k} \right) \mathbf{v}_k. \tag{2.74}$$

Substituting the expansion (2.73) gives

$$\hat{\mathbf{f}} = \sum_{k=1}^{n} \left(\mathbf{v}_k^{\mathrm{T}} \mathbf{f} + \frac{\mathbf{u}_k^{\mathrm{T}} \mathbf{n}}{\sigma_k} \right) \mathbf{v}_k \tag{2.75}$$

$$= \mathbf{f} + \sum_{k=1}^{n} \frac{\left(\mathbf{u}_k^{\mathrm{T}} \mathbf{n} \right)}{\sigma_k} \mathbf{v}_k \tag{2.76}$$

where we have made use of the fact that the $\{\mathbf{v}_k\}$ form a basis for $\mathbb{R}^n$. We see that the reconstruction is the sum of the true image and terms due to the noise. The error term along the direction of $\mathbf{v}_k$ in image space arises from the component of the noise in the direction of $\mathbf{u}_k$ in data space, *divided by* the singular value σ_k.

If we now suppose that some of the singular values σ_k are small, this division will give a very large random component, often completely swamping the component of $\mathbf{f}$ in that direction. Another way of thinking about this is to see that the small singular values correspond to directions in image space for which the data contain very little information about the image. In attempting to reconstruct those aspects of $\mathbf{f}$ which lie along these directions, we have to amplify the small signal buried in the data by dividing by the small singular value. Such a scheme is risky because the noise which is inevitably present in the data is also going to be amplified by a large factor, corrupting the reconstruction.

Thus when there are small singular values, the least squares method can give bad reconstructions. It is better to consider small singular values as being effectively zero, and to regard the components along such directions as being free parameters which are not determined by the data.

When the singular values of the measurement operator $\mathbf{A}$ are ranked in non-increasing order, the rate at which they decrease with index gives valuable information about how much we can hope to reconstruct from data taken using that measurement process for a given amount of noise in the data. The more rapid is the decrease, the less we can reconstruct reliably for a given noise level. Equivalently, in order to get good reconstructions when the singular values decrease rapidly, an extremely high signal-to-noise ratio in the data is required.

2.11 Continuous transformations

Our analysis of the forward map A via the SVD required that both image space and data space to be finite dimensional so the A is represented by a $m \times n$ matrix. When one or both spaces are continuous we can still perform a singular-value decomposition with the aid of a little more mathematical machinery, as presented in the next section. We will only deal fully with the case where just one of the spaces is continuous and investigate some of the issues for the continuous-continuous problem.in an assignment.

We deal with continuous spaces by first considering how the ideas of matrices may be extended to infinite dimensional spaces and then show how to construct the matrix (conjugate) transpose so that we are able to form the analogs of $\mathbf{A}^{\mathrm{T}}\mathbf{A}$ and $\mathbf{A}\mathbf{A}^{\mathrm{T}}$. Both these constructions are achieved by using the "star" operator defined in the next section.

2.11.1 Adjoint operators and the star notation

2.11.1.1 Adjoints of Operators

One of the most powerful and pervasive concepts in linear algebra is the concept of adjoints. Its generalization to function spaces will allow the use of matrix notation for continuous-discrete and continuous-continuous forward operators.

For fixed inner-products in $\mathbb{C}^N$ and $\mathbb{C}^M$, the adjoint of a linear operator $F : \mathbb{C}^N \to \mathbb{C}^M$ is a another linear operator

$$F^* : \mathbb{C}^M \to \mathbb{C}^N$$

satisfying the relation

$$\langle \mathbf{u}, F^*\mathbf{v} \rangle = \langle F\mathbf{u}, \mathbf{v} \rangle , \quad \text{for all } \mathbf{u} \in \mathbb{C}^N, \mathbf{v} \in \mathbb{C}^M. \tag{2.77}$$

Note that the first inner-product is in $\mathbb{C}^N$ while the second is in $\mathbb{C}^M$. Strictly, equation 2.77 is not a definition of F^* as it is actually a (non-trivial) consequence of the properties of inner-products that F^* must exist and be unique.

A concrete picture of adjoints can be gained by choosing bases $\{\mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_N\}$ of $\mathbb{C}^N$ and $\{\mathbf{d}_1, \mathbf{d}_2, \ldots, \mathbf{d}_M\}$ of $\mathbb{C}^M$ which are orthonormal with respect to the respective inner-products. Recall that the operator F is

represented by the $M \times N$ matrix $[F_{mn}]$ and F^* is represented by the $N \times M$ matrix $[(F^*)_{nm}]$. The relationship between these two matrices can be found by considering the defining expressions

$$F\mathbf{b}_n = \sum_{m=0}^{M} \mathbf{d}_m F_{mn}, \quad F^*\mathbf{d}_m = \sum_{n=0}^{N} \mathbf{b}_n \left(F^*\right)_{nm}.$$

Since both bases are orthonormal, the inner-product of both sides of the first equation with $\mathbf{d}_k$, and the inner-product of both sides of the second equation with $\mathbf{b}_l$ gives

$$\langle \mathbf{d}_k, F\mathbf{b}_n \rangle = F_{kn}, \quad \langle \mathbf{b}_l, F^*\mathbf{d}_m \rangle = (F^*)_{lm}.$$

Setting $k = m$ and $l = n$ and using the Hermiticity property of the inner-product,

$$(F^*)_{nm} = \langle F\mathbf{b}_n, \mathbf{d}_m \rangle = \overline{\langle \mathbf{d}_m, F\mathbf{b}_n \rangle} = \overline{F_{mn}}.$$

Hence the matrix $[(F^*)_{nm}]$ representing F^* is the complex conjugate of the transpose[3] of the matrix $[F_{mn}]$ representing F. So a simple way of thinking of the adjoint operator is to think of the conjugate transpose of the representing matrix. In the general case of non-orthonormal bases, the relationship between matrices is more complicated. However we will typically only be considering orthonormal bases and so the simple picture is sufficient.

2.11.1.2 Adjoints of Vectors

The notion of the adjoint also applies to vectors when we associate the vector with its natural operation as a matrix. Think of the vector $\mathbf{u} \in \mathbb{C}^N$ as a simple operator that takes a complex number and produces a vector, i.e.,

$$\mathbf{u} : \mathbb{C} \to \mathbb{C}^N,$$

defined in the obvious way: for $c \in \mathbb{C}$,

$$\mathbf{u}c = c\mathbf{u}.$$

The left-hand side describes $\mathbf{u}$ acting on c while the right-hand side defines the result and is just the scalar multiplication of $\mathbf{u}$ by c.

Then the adjoint of $\mathbf{u}$ is

$$\mathbf{u}^* : \mathbb{C}^N \to \mathbb{C}$$

defined by equation 2.77, which implies[4] that for $\mathbf{v} \in \mathbb{C}^N$

$$\mathbf{u}^*\mathbf{v} = \langle \mathbf{u}, \mathbf{v} \rangle$$

So $\mathbf{u}^*$ is the operator that takes the inner product with $\mathbf{u}$, or equivalently, is the process of multiplying by the conjugate transpose of the representation of the vector. So, again, the adjoint may be thought of as the conjugate transpose. Since $\mathbf{u}$ is represented by a column vector, $\mathbf{u}^*$ is a row vector.

In the Physics literature another notation is often used in which one writes

$$\begin{array}{ccll} \mathbf{u}^* & \text{as} & \langle \mathbf{u}| & \text{(called a bra)} \\ \mathbf{v} & \text{as} & |\mathbf{v}\rangle & \text{(called a ket)} \end{array}$$

Finally, note that every linear operator that maps $\mathbb{C}^N \to \mathbb{C}$ can be written as $\mathbf{u}^*$ for some $\mathbf{u} \in \mathbb{C}^N$.

[3] The conjugate transpose is also called the Hermitian conjugate.

[4] Consider $\langle \mathbf{u}c, \mathbf{v} \rangle$ for any $c \in \mathbb{C}$ and $\mathbf{v} \in \mathbb{C}^N$. As the inner product is anti-linear in the first argument we have $\langle \mathbf{u}c, \mathbf{v} \rangle = \bar{c} \langle \mathbf{u}, \mathbf{v} \rangle$ while the definition of the $\mathbf{u}^*$ is $\langle \mathbf{u}c, \mathbf{v} \rangle = \langle c, \mathbf{u}^*\mathbf{v} \rangle = \bar{c}\mathbf{u}^*\mathbf{v}$. The implication follows.

2.11.1.3 Adjoints of Functions

So far we appear to have done little more than observe that the inner product may be written as $\langle \mathbf{u}, \mathbf{v} \rangle = \mathbf{u}^* \mathbf{v}$ where $\mathbf{u}^*$ is usually the conjugate transpose of $\mathbf{u}$ – though our notation does work for representations in non-orthonornmal bases as well. But the real power of the notation is that it works for infinite-dimensional vectors as well – such as those represented by functions of a continuous variable.

For example, a vector space of functions that we often use is the space of all square integrable functions defined on the real line. That is, the set of functions

$$L^2(\mathbb{R}) = \left\{ f : \mathbb{R} \to \mathbb{C} \text{ and } \|f\| \equiv \int_{-\infty}^{\infty} |f(x)|^2 \, \mathrm{d}x < \infty \right\}.$$

For $f, g \in L^2(\mathbb{R})$, the usual inner product is given by

$$\langle f, g \rangle = \int_{-\infty}^{\infty} \overline{f(x)} g(x) \, \mathrm{d}x.$$

In the same way as for finite dimensional vectors, we can think of $f(x)$ as a simple operator

$$f{:}\mathbb{C} \to L^2(\mathbb{R})$$

defined by

$$fc = cf(x).$$

The adjoint of this operator is

$$f^*{:}L^2(\mathbb{R}) \to \mathbb{C}$$

which is given by, for any $g \in L^2(\mathbb{R})$,

$$f^* g = \langle f, g \rangle.$$

Again we note that any bounded linear functional[5] that maps $L^2(\mathbb{R})$ to $\mathbb{C}$ can be written as f^* for some function f.

Example – the Fourier Transform

The star notation provides a handy shorthand notation for the Fourier transform. For any frequency let

$$e_\nu(t) = \exp(2\pi \mathrm{i} \nu t)$$

then the Fourier transform of a function $f(t) \in L^2(\mathbb{R})$ is

$$F(\nu) = e_\nu^* f.$$

Example – first N moments

Consider reconstructing a function $f(x)$ defined on the interval $[0, 1]$ from measurements of its first N moments

$$d_l = \int_0^1 x^{l-1} f(x) \, \mathrm{d}x \qquad l = 1, 2, \ldots, N.$$

By defining the functions

$$a_l(x) = x^{l-1} \quad \text{for } x \in [0, 1], \; l = 1, 2, \ldots, N$$

we can write the forward map as

$$d_l = a_l^* f, \qquad l = 1, 2, \ldots, N$$

[5] A functional is just a linear operator that takes a function as its argument.

or in the "matrix-vector" form –

$$\mathbf{d} = Af$$

where $\mathbf{d}$ is the N-dimensional vector of data. Here A is the discrete×continuous "matrix"

$$A = \left(a_l^*(x)\right) = \begin{pmatrix} \cdots & 1 & \cdots \\ \cdots & x & \cdots \\ \cdots & x^2 & \cdots \\ & \vdots & \\ \cdots & x^{N-1} & \cdots \end{pmatrix}$$

in which the rows are the indicated functions defined on $[0,1]$. The generalization of the matrices $\mathbf{A}^{\mathrm{T}}\mathbf{A}$ and $\mathbf{A}\mathbf{A}^{\mathrm{T}}$ are the representations of the operators $\mathbf{A}^*\mathbf{A}$ and $\mathbf{A}\mathbf{A}^*$ – the first of which is continuous×continuous while the later is discrete×discrete. So naturally we choose the smaller matrix which is the $N \times N$ matrix which has kl component

$$(AA^*)_{kl} = a_k^* a_l = \int_0^1 x^{l-1} x^{k-1}\,\mathrm{d}x = \frac{1}{k+l-1},$$

where we have used the important relation that $\left(a_l^*\right)^* = a_l$.

First the simplest case, $N = 1$. Now $AA^* = (1)$ which has one eigenvalue of 1 with eigenvector (1). Consequently, A has one non-zero singular value, equalling 1, and the corresponding singular vector in image space is

$$A^*(1) = 1(x),$$

i.e., the constant function with value 1. The left-hand side is the matrix A^* acting on the data vector (1) ,and since $A = a_1^*(x) = 1^*$, in this case, then $A^* = a_1(x) = 1(x)$ and the function $1(x)$ acting on the number 1 simply results in the function $1(x)$ which is the constant function with value 1.

A bit more substantial is the case $N = 3$: Now

$$AA^* = \begin{pmatrix} 1 & \frac{1}{2} & \frac{1}{3} \\ \frac{1}{2} & \frac{1}{3} & \frac{1}{4} \\ \frac{1}{3} & \frac{1}{4} & \frac{1}{5} \end{pmatrix}$$

$$\begin{pmatrix} 1 & \frac{1}{2} & \frac{1}{3} \\ \frac{1}{2} & \frac{1}{3} & \frac{1}{4} \\ \frac{1}{3} & \frac{1}{4} & \frac{1}{5} \end{pmatrix}$$

which has the normalized eigenvectors: $\mathbf{u}_1 = \begin{pmatrix} 0.82706 \\ 0.46039 \\ 0.32248 \end{pmatrix}$, $\mathbf{u}_2 = \begin{pmatrix} 0.54744 \\ -0.52822 \\ -0.64909 \end{pmatrix}$, and $\mathbf{u}_3 = \begin{pmatrix} 0.12766 \\ -0.71375 \\ 0.68868 \end{pmatrix}$, with eigenvalues 1.4083, 0.12233, and 2.6873×10^{-3}, respectively. The vectors are also the left singular vectors of A while the singular values are the square-roots of the eigenvalues and are: $\sigma_1 = 1.1867$, $\sigma_2 = 0.34976$, and $\sigma_3 = 5.1839 \times 10^{-2}$. The right singular vectors may be found, as in equation 2.32 or 2.36, by operating on each of the left singular vectors by A^* and dividing by the corresponding singular value. The resulting functions are

$$v_1(x) = \frac{A^* \begin{pmatrix} 0.82706 \\ 0.46039 \\ 0.32248 \end{pmatrix}}{1.1867} = 0.69694 + 0.38796x + 0.27175x^2$$

$$v_2(x) = \frac{A^* \begin{pmatrix} 0.54744 \\ -0.52822 \\ -0.64909 \end{pmatrix}}{0.34976} = 1.5652 - 1.5102x - 1.8558x^2$$

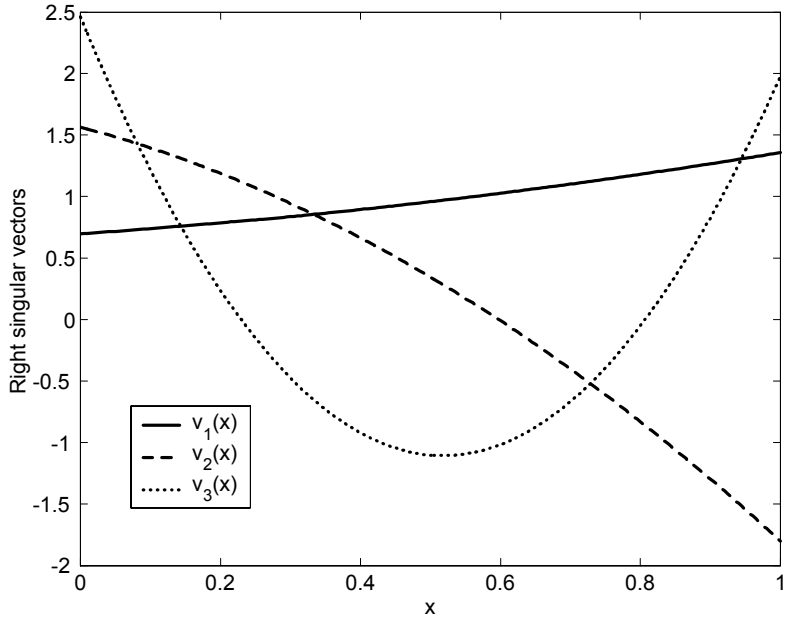

Figure 2.7 The three right singular vectors

$$v_3(x) = \frac{A^* \begin{pmatrix} 0.12766 \\ -0.71375 \\ 0.68868 \end{pmatrix}}{5.1839 \times 10^{-2}} = 2.4626 - 13.769x + 13.285x^2$$

With a little work you should be able to show that these functions are indeed orthonormal over the interval $[0,1]$. Just for the record, here is a graph of the three right singular vectors

These are the three structures in the image which contribute to the data. Note that the measurement process is about 20 times less sensitive to $v_3(x)$ than to $v_1(x)$.

In the case where 10 moments are measured we have $N = 10$ and

$$AA^* = \begin{pmatrix} 1 & \frac{1}{2} & \frac{1}{3} & \frac{1}{4} & \frac{1}{5} & \frac{1}{6} & \frac{1}{7} & \frac{1}{8} & \frac{1}{9} & \frac{1}{10} \\ \frac{1}{2} & \frac{1}{3} & \frac{1}{4} & \frac{1}{5} & \frac{1}{6} & \frac{1}{7} & \frac{1}{8} & \frac{1}{9} & \frac{1}{10} & \frac{1}{11} \\ \frac{1}{3} & \frac{1}{4} & \frac{1}{5} & \frac{1}{6} & \frac{1}{7} & \frac{1}{8} & \frac{1}{9} & \frac{1}{10} & \frac{1}{11} & \frac{1}{12} \\ \frac{1}{4} & \frac{1}{5} & \frac{1}{6} & \frac{1}{7} & \frac{1}{8} & \frac{1}{9} & \frac{1}{10} & \frac{1}{11} & \frac{1}{12} & \frac{1}{13} \\ \frac{1}{5} & \frac{1}{6} & \frac{1}{7} & \frac{1}{8} & \frac{1}{9} & \frac{1}{10} & \frac{1}{11} & \frac{1}{12} & \frac{1}{13} & \frac{1}{14} \\ \frac{1}{6} & \frac{1}{7} & \frac{1}{8} & \frac{1}{9} & \frac{1}{10} & \frac{1}{11} & \frac{1}{12} & \frac{1}{13} & \frac{1}{14} & \frac{1}{15} \\ \frac{1}{7} & \frac{1}{8} & \frac{1}{9} & \frac{1}{10} & \frac{1}{11} & \frac{1}{12} & \frac{1}{13} & \frac{1}{14} & \frac{1}{15} & \frac{1}{16} \\ \frac{1}{8} & \frac{1}{9} & \frac{1}{10} & \frac{1}{11} & \frac{1}{12} & \frac{1}{13} & \frac{1}{14} & \frac{1}{15} & \frac{1}{16} & \frac{1}{17} \\ \frac{1}{9} & \frac{1}{10} & \frac{1}{11} & \frac{1}{12} & \frac{1}{13} & \frac{1}{14} & \frac{1}{15} & \frac{1}{16} & \frac{1}{17} & \frac{1}{18} \\ \frac{1}{10} & \frac{1}{11} & \frac{1}{12} & \frac{1}{13} & \frac{1}{14} & \frac{1}{15} & \frac{1}{16} & \frac{1}{17} & \frac{1}{18} & \frac{1}{19} \end{pmatrix}.$$

The square root of the eigenvalues give the singular values

index	**singular value**
1	1.3236
2	0.5856
3	0.18906
4	5.0308×10^{-2}
5	1.1347×10^{-2}
6	2.1748×10^{-3}
7	3.5057×10^{-4}
8	4.634×10^{-5}
9	4.761×10^{-6}
10	3.3064×10^{-7}

Notice how the singular values fall off rapidly. With a measurement accuracy of 1 part in 10^3, for example, we would expect to measure the first 6 components (to varying degrees) while the remaining components of the data will typically be smaller than the noise.