2.2 Matrix and Vector Algebra

This subject is very large and well-developed and it is not my intention to repeat material better found elsewhere[9]. Only a brief survey of essential results is provided.

Continued on next page...

A matrix is an $M \times N$ array of elements of the form

$$\mathbf{A} = \{A_{ij}\}, \quad 1 \le i \le M,\ 1 \le j \le N\,.$$

Normally a matrix is denoted by a bold-faced capital letter. A vector is a special case of an $M \times 1$ matrix, written as a bold-face lower case letter, for example, $\mathbf{q}$. Corresponding capital or lower case letters for Greek symbols are also indicated in bold-face. Unless otherwise stipulated, vectors are understood to be columnar. The transpose of a matrix interchanges its rows and columns. Transposition applied to vectors is sometimes used to save space in printing, for example, $\mathbf{q} = [q_1, q_2, ..., q_N]^T$ is the same as

$$\mathbf{q} = \begin{bmatrix} q_1 \\ q_2 \\ \vdots \\ q_N \end{bmatrix}\,.$$

Matrices and Vectors

A conventional measure of length of a vector is $\sqrt{\mathbf{a}^T\mathbf{a}} = \sqrt{\sum_i^N a_i^2} = \|\mathbf{a}\|$. The inner, or dot, product between two $L \times 1$ vectors $\mathbf{a}$, $\mathbf{b}$ is written $\mathbf{a}^T\mathbf{b} \equiv \mathbf{a} \cdot \mathbf{b} = \sum_{i=1}^L a_i b_i$ and is a scalar. Such an inner product is the "projection" of $\mathbf{a}$ onto $\mathbf{b}$ (or vice-versa). It is readily shown that $|\mathbf{a}^T\mathbf{b}| \le \|\mathbf{a}\|\,\|\mathbf{b}\|$; the magnitude of this projection can be measured as,

$$|\mathbf{a}^T\mathbf{b}| = \|\mathbf{a}\|\ \|\mathbf{b}\|\,|\cos\phi|\,,$$

where the magnitude of $\cos\phi$ ranges between zero, when the vectors are orthogonal, and one, when they are parallel.

Suppose we have a collection of N vectors, $\mathbf{e}_i$, each of dimension N. If it is possible to represent perfectly an arbitrary N–dimensional vector $\mathbf{f}$ as the linear sum,

{31001}
$$\mathbf{f} = \sum_{i=1}^{N} \alpha_i \mathbf{e}_i\,, \tag{2.1}$$

then $\mathbf{e}_i$ are said to be a "basis." A necessary and sufficient condition for them to have that property is that they should be "independent," that is, no one of them should be perfectly representable by the others:

{31002}
$$\mathbf{e}_j - \sum_{i=1,\, i\neq j}^{N} \beta_i \mathbf{e}_i \neq 0, \quad 1 \le j \le N\,. \tag{2.2}$$

A subset of the $\mathbf{e}_j$ are said to span a subspace (all vectors perfectly representable by the subset). For example, $[1,-1,0]^T, [1,1,0]^T$ span the subspace of all vectors $[v_1, v_2, 0]^T$. A "spanning set" completely describes the subspace too, but might have additional, redundant vectors. Thus the vectors $[1,-1,0]^T, [1,1,0]^T, [1,1/2,0]$ span the subspace but are not a basis for it.

The expansion coefficients α_i in (2.1) are obtained by taking the dot product of (2.1) with each of the vectors in turn:

$$\sum_{i=1}^{N} \alpha_i \mathbf{e}_k^T \mathbf{e}_i = \mathbf{e}_k^T \mathbf{f}, \quad 1 \le k \le N\,, \tag{2.3}$$ {31003}

a system of N equations in N unknowns. The α_i are most readily found if the $\mathbf{e}_i$ are a mutually orthonormal set, that is, if

$$\mathbf{e}_i^T \mathbf{e}_j = \delta_{ij}\,,$$

but this requirement is not a necessary one. With a basis, the information contained in the set of projections, $\mathbf{e}_i^T \mathbf{f} = \mathbf{f}^T \mathbf{e}_i$, is adequate then to determine the α_i and thus all the information required to reconstruct $\mathbf{f}$ is contained in the dot products..

The concept of "nearly-dependent" vectors is helpful and can be understood heuristically. Consider figure 3.1, in which the space is two-dimensional. Then the two vectors $\mathbf{e}_1$, $\mathbf{e}_2$, as depicted there, are independent and can be used to expand an arbitrary two-dimensional vector $\mathbf{f}$ in the plane. The simultaneous equations become,

$$\begin{aligned} \alpha_1 \mathbf{e}_1^T \mathbf{e}_1 + \alpha_2 \mathbf{e}_1^T \mathbf{e}_2 &= \mathbf{e}_1^T \mathbf{f} \\ \alpha_1 \mathbf{e}_2^T \mathbf{e}_1 + \alpha_2 \mathbf{e}_2^T \mathbf{e}_2 &= \mathbf{e}_2^T \mathbf{f}. \end{aligned} \tag{2.4}$$

The vectors become nearly parallel as the angle ϕ in Fig. 3.1 goes to zero; as long as they are not identically parallel, they can still be used mathematically to represent $\mathbf{f}$ perfectly. An important feature is that even if the lengths of $\mathbf{e}_1$,$\mathbf{e}_2, \mathbf{f}$ are all order-one, the expansion coefficients $a_{1,2}$ can have extremely large magnitudes when the angle ϕ becomes small and $\mathbf{f}$ is nearly orthogonal to both (measured by angle η).

That is to say, we find readily from (2.4),

$$\alpha_1 = \frac{\left(\mathbf{e}_1^T \mathbf{f}\right)\left(\mathbf{e}_2^T \mathbf{e}_2\right) - \left(\mathbf{e}_2^T \mathbf{f}\right)\left(\mathbf{e}_1^T \mathbf{e}_2\right)}{\left(\mathbf{e}_1^T \mathbf{e}_1\right)\left(\mathbf{e}_2^T \mathbf{e}_2\right) - \left(\mathbf{e}_1^T \mathbf{e}_2\right)^2}, \tag{2.5}$$

$$\alpha_2 = \frac{\left(\mathbf{e}_2^T \mathbf{f}\right)\left(\mathbf{e}_1^T \mathbf{e}_1\right) - \left(\mathbf{e}_1^T \mathbf{f}\right)\left(\mathbf{e}_2^T \mathbf{e}_1\right)}{\left(\mathbf{e}_1^T \mathbf{e}_1\right)\left(\mathbf{e}_2^T \mathbf{e}_2\right) - \left(\mathbf{e}_1^T \mathbf{e}_2\right)^2}. \tag{2.6}$$

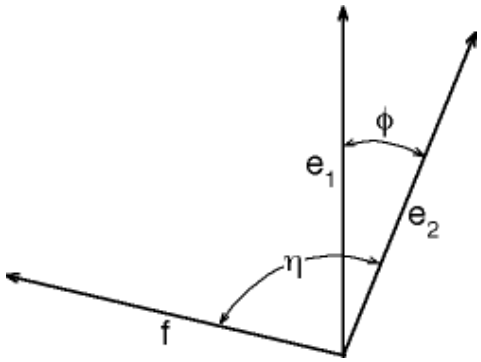

Figure 2.1: Schematic of expansion of an arbitrary vector $\mathbf{f}$ in two vectors $\mathbf{e}_1, \mathbf{e}_2$ which may nearly coincide in direction. {ocip3.1new.ep

Suppose for simplicity, that $\mathbf{f}$ has unit length, and that the $\mathbf{e}_i$ have also been normalized to unit length as shown in Figure 3.1. We have then,

$$\alpha_1 = \frac{\cos(\eta - \phi) - \cos\phi\cos\eta}{1 - \cos^2\phi} = \frac{\sin\eta}{\sin\phi} \tag{2.7}$$

$$\alpha_2 = \cos\eta - \sin\eta\cot\phi \tag{2.8}$$

and whose magnitudes can become arbitrarily large as $\phi \to 0$. One can imagine a situation in which $\alpha_1\mathbf{e}_1$ and $\alpha_2\mathbf{e}_2$ were separately measured and found to be very large. One could then erroneously infer that the sum vector, $\mathbf{f}$, was equally large. This property of the expansion in non-orthogonal vectors potentially producing large coefficients becomes important later (Chapter 4) as a way of gaining insight into the behavior of so-called non-normal operators. The generalization to higher dimensions is left to the reader's intuition. One anticipates that as ϕ becomes very small, numerical problems can arise in using these "almost parallel" vectors.

gramschmidt}

Gram-Schmidt Process

One often has a set of p-independent, but non-orthonormal vectors, $\mathbf{h}_i$, and it is convenient to find a new set $\mathbf{g}_i$, which are orthonormal. The "Gram-Schmidt process" operates by induction. Suppose the first k of the $\mathbf{h}_i$ have been orthonormalized to a new set, $\mathbf{g}_i$. To generate vector $k+1$, let

{31030}

$$\mathbf{g}_{k+1} = \mathbf{h}_{k+1} - \sum_j^k \gamma_j \mathbf{g}_j \,. \tag{2.9}$$

Because $\mathbf{g}_{k+1}$ must be orthogonal to the preceding $\mathbf{g}_i$, $i = 1, ..., k$, take the dot products of (2.9) with each of these vectors, producing a set of simultaneous equations for determining the unknown γ_j. The resulting $\mathbf{g}_{k+1}$ is easily given unit norm by division by its length.

Given the first k of N necessary vectors, an additional $N-k$ independent vectors, $\mathbf{h}_i$ are needed. There are several possibilities. The necessary extra vectors might be generated by filling their elements with random numbers. Or a very simple trial set like $\mathbf{h}_{k+1} = [1,0,0,...,0]^T$, $\mathbf{h}_{k+2} = [0,1,0,...0], \ldots$ might be adequate. If one is unlucky, the set chosen might prove not to be independent of the existing $\mathbf{g}_i$. But a simple numerical perturbation usually suffices to render them so. In practice, the algorithm is changed to what is usually called the "modified Gram-Schmidt process" for purposes of numerical stability.[10]

2.2.1 Matrix Multiplication and Identities

It has been found convenient and fruitful to usually define multiplication of two matrices $\mathbf{A}, \mathbf{B}$, written as $\mathbf{C} = \mathbf{AB}$, such that

$$C_{ij} = \sum_{p=1}^{P} A_{ip} B_{pj} \,. \tag{2.10}$$ {31004}

For the definition (2.10) to make sense, $\mathbf{A}$ must be a $M \times P$ matrix and $\mathbf{B}$ must be $P \times N$ (including the special case of $P \times 1$, a column vector). That is, the two matrices must be "conformable." If two matrices are multiplied, or a matrix and a vector are multiplied, conformability is implied—otherwise one can be assured that an error has been made. Note that $\mathbf{AB} \neq \mathbf{BA}$ even where both products exist, except under special circumstance. Define $\mathbf{A}^2 = \mathbf{AA}$, etc. Other definitions of matrix multiplication exist, and we will later define the Hadamard product when it is needed..

The mathematical operation in (2.10) may appear arbitrary, but a physical interpretation is available: Matrix multiplication is the dot product of all of the rows of $\mathbf{A}$ with all of the columns of $\mathbf{B}$. Thus multiplication of a vector by a matrix represents the projections of the rows of the matrix onto the vector.

Define a matrix, $\mathbf{E}$, each of whose columns is the corresponding vector $\mathbf{e}_i$, and a vector, $\boldsymbol{\alpha} = \{\alpha_i\}$, in the same order. Then the expansion (2.1) can be written compactly as,

$$\mathbf{f} = \mathbf{E}\boldsymbol{\alpha} \,. \tag{2.11}$$ {31005}

The transpose of a matrix $\mathbf{A}$ is written $\mathbf{A}^T$ and is defined as $\{A^T\}_{ij} = A_{ji}$, an interchange of the rows and columns of $\mathbf{A}$. Thus $\left(\mathbf{A}^T\right)^T = \mathbf{A}$. A "symmetric matrix" is one for which $\mathbf{A}^T = \mathbf{A}$. The product $\mathbf{A}^T\mathbf{A}$ represents the array of all the dot products of the columns of $\mathbf{A}$ with themselves, and similarly, $\mathbf{AA}^T$ represents the set of all dot products of all the rows of $\mathbf{A}$ with themselves. It follows that $(\mathbf{AB})^T = \mathbf{B}^T\mathbf{A}^T$. Because we have $(\mathbf{AA}^T)^T = \mathbf{AA}^T$, $(\mathbf{A}^T\mathbf{A})^T = \mathbf{A}^T\mathbf{A}$, both of these matrices are symmetric.

The "trace" of a square $M \times M$ matrix $\mathbf{A}$ is defined as $\text{trace}(\mathbf{A}) = \sum_i^M A_{ii}$. A "diagonal matrix" is square and zero except for the terms along the main diagonal, although we will later generalize this definition. The operator $\text{diag}(\mathbf{q})$ forms a square diagonal matrix with $\mathbf{q}$ along the main diagonal.

The special $L \times L$ diagonal matrix $\mathbf{I}_L$, with $I_{ii} = 1$, is the "identity." Usually, when the dimension of $\mathbf{I}_L$ is clear from the context, the subscript is omitted. $\mathbf{IA} = \mathbf{A}$, $\mathbf{AI} = \mathbf{I}$, for any $\mathbf{A}$ for which the products make sense. If there is a matrix $\mathbf{B}$, such that $\mathbf{BE} = \mathbf{I}$, then $\mathbf{B}$ is the "left-inverse" of $\mathbf{E}$. If $\mathbf{B}$ is the left inverse of $\mathbf{E}$ and $\mathbf{E}$ is square, a standard result is that it must also be a right inverse: $\mathbf{EB} = \mathbf{I}$, $\mathbf{B}$ is then called "*the* inverse of $\mathbf{E}$" and is usually written $\mathbf{E}^{-1}$. Square matrices with inverses are "non-singular." Analytical expressions exist for some inverses, and numerical linear algebra books explain how to find them, when they exist. If $\mathbf{E}$ is not square, one must distinguish left and right inverses, sometimes written $\mathbf{E}^+$ and referred to as "generalized inverses." Some of them will be encountered later. A useful result is that $(\mathbf{AB})^{-1} = \mathbf{B}^{-1}\mathbf{A}^{-1}$, if the inverses exist. A useful notational shorthand is $\left(\mathbf{A}^{-1}\right)^T = \left(\mathbf{A}^T\right)^{-1} \equiv \mathbf{A}^{-T}$.

A definition of the "length," or norm of a vector has already been introduced. But several choices are possible; for present purposes, the conventional l_2 norm already defined,

{31006}
$$\|\mathbf{f}\|_2 \equiv (\mathbf{f}^T\mathbf{f})^{1/2} = \left(\sum_{i=1}^{N} f_i^2\right)^{1/2}, \tag{2.12}$$

is most useful; often the subscript will be omitted. Eq. (2.12) leads in turn to the measure of distance between two vectors, $\mathbf{a}$, $\mathbf{b}$ as,

{31007}
$$\|\mathbf{a}-\mathbf{b}\|_2 = \sqrt{(\mathbf{a}-\mathbf{b})^T(\mathbf{a}-\mathbf{b})}, \tag{2.13}$$

the familiar Cartesian distance. Distances can also be measured in such a way that deviations of certain elements of $\mathbf{c} = \mathbf{a} - \mathbf{b}$ count for more than others—that is, a metric, or set of weights can be introduced with a definition,

{31008}
$$\|\mathbf{c}\|_W = \sqrt{\sum\nolimits_i c_i W_{ii} c_i}, \tag{2.14}$$

depending upon the importance to be attached to magnitudes of different elements, stretching and shrinking various coordinates. Finally, in the most general form, distance can be measured in a coordinate system both stretched and rotated relative to the original one

{31009}
$$\|\mathbf{c}\|_W = \sqrt{\mathbf{c}^T\mathbf{W}\mathbf{c}} \tag{2.15}$$

where $\mathbf{W}$ is an arbitrary matrix (but usually, for physical reasons, symmetric and positive definite[11], implying that $\mathbf{c}^T\mathbf{W}\mathbf{c} \geq 0$).

2.2.2 Linear Simultaneous Equations

Consider a set of M-linear equations in N-unknowns,

{31010}
$$\mathbf{E}\mathbf{x} = \mathbf{y}. \tag{2.16}$$

Because of the appearance of simultaneous equations in situations in which the y_i are observed, and where $\mathbf{x}$ are parameters we wish to determine, it is often convenient to refer to (2.16) as a set of measurements of $\mathbf{x}$ which produced the observations or data, $\mathbf{y}$. If $M > N$, the system is said to be "formally overdetermined." If $M < N$, it is "underdetermined," and if $M = N$, it is "formally just-determined." The use of the word "formally" has a purpose we will come to. Knowledge of the matrix inverse to $\mathbf{E}$ would make it easy to solve a set of L equations in L unknowns, by left-multiplying (2.16) by $\mathbf{E}^{-1}$:

$$\mathbf{E}^{-1}\mathbf{E}\mathbf{x} = \mathbf{I}\mathbf{x} = \mathbf{x} = \mathbf{E}^{-1}\mathbf{y}$$

The reader is cautioned that although matrix inverses are a very powerful theoretical tool, one is usually ill-advised to solve large sets of simultaneous equations by employing $\mathbf{E}^{-1}$; better numerical methods are available for the purpose[12].

There are several ways to view the meaning of any set of linear simultaneous equations. If the columns of $\mathbf{E}$ continue to be denoted $\mathbf{e}_i$, then (2.16) is,

$$x_1\mathbf{e}_1 + x_2\mathbf{e}_2 + \cdots + x_N\mathbf{e}_N = \mathbf{y}\,. \tag{2.17}$$ {31011}

The ability to so describe an arbitrary $\mathbf{y}$, or to solve the equations, would thus depend upon whether the $M \times 1$, vector $\mathbf{y}$ can be specified by a sum of N-column vectors, $\mathbf{e}_i$. That is, it would depend upon their being a spanning set. In this view, the elements of $\mathbf{x}$ are simply the corresponding expansion coefficients. Depending upon the ratio of M to N, that is, the number of equations compared to the number of unknown elements, one faces the possibility that there are fewer expansion vectors $\mathbf{e}_i$ than elements of $\mathbf{y}$ ($M > N$), or that there are more expansion vectors available than elements of $\mathbf{y}$ ($M < N$). Thus the overdetermined case corresponds to having *fewer* expansion vectors, and the underdetermined case corresponds to having *more* expansion vectors, than the dimension of $\mathbf{y}$. It is possible that in the overdetermined case, the too-few expansion vectors are not actually independent, so that there are even fewer vectors available than is first apparent. Similarly, in the underdetermined case, there is the possibility that although it appears we have more expansion vectors than required, fewer may be independent than the number of elements of $\mathbf{y}$, and the consequences of that case need to be understood as well.

An alternative interpretation of simultaneous linear equations denotes the rows of $\mathbf{E}$ as $\mathbf{r}_i^T$, $1 \le i \le M$. Then Eq.(2.16) is a set of M-inner products,

$$\mathbf{r}_i^T \mathbf{x} = y_i, \quad 1 \le i \le M\,. \tag{2.18}$$ {31012}

That is, the set of simultaneous equations is also equivalent to being provided with the value of M–dot products of the N–dimensional unknown vector, $\mathbf{x}$, with M known vectors, $\mathbf{r}_i$. Whether that is sufficient information to determine $\mathbf{x}$ depends upon whether the $\mathbf{r}_i$ are a spanning set. In this view, in the overdetermined case, one has *more* dot products available than unknown elements x_i, and in the underdetermined case, there are *fewer* such values than unknowns.

A special set of simultaneous equations for square matrices, $\mathbf{A}$, is labelled the "eigenvalue/eigenvector problem,"

{eigen1}
$$\mathbf{A}\mathbf{e} = \lambda \mathbf{e}. \tag{2.19}$$

In this set of linear simultaneous equations one seeks a special vector, $\mathbf{e}$, such that for some as yet unknown scalar eigenvalue, λ, there is a solution. An $N \times N$ matrix will have up to N solutions $(\lambda_i, \mathbf{e}_i)$, but the nature of these elements and their relations require considerable effort to deduce. We will look at this problem more later; for the moment, it again suffices to say that numerical methods for solving Eq. (2.19) are well-known. {pageeigen}

2.2.3 Matrix Norms

A number of useful definitions of a matrix size, or norm, exist. For present purposes the so-called "spectral norm" or "2–norm" defined as

$$\|\mathbf{A}\|_2 = \sqrt{\text{maximum eigenvalue of } (\mathbf{A}^T\mathbf{A})} \tag{2.20}$$

is usually adequate. Without difficulty, it may be seen that this definition is equivalent to

$$\|\mathbf{A}\|_2 = \max \frac{\mathbf{x}^T\mathbf{A}^T\mathbf{A}\mathbf{x}}{\mathbf{x}^T\mathbf{x}} = \max \frac{\|\mathbf{A}\mathbf{x}\|_2}{\|\mathbf{x}\|_2} \tag{2.21}$$

where the maximum is defined over all vectors $\mathbf{x}$.[13] Another useful measure is the "Frobenius norm,"

$$\|\mathbf{A}\|_F = \sqrt{\sum_{i=1}^{M}\sum_{j=1}^{N} A_{ij}^2} = \sqrt{\text{trace}(\mathbf{A}^T\mathbf{A})}\,. \tag{2.22}$$

Neither norm requires $\mathbf{A}$ to be square. These norms permit one to derive various useful results. Consider one illustration. Suppose $\mathbf{Q}$ is square, and $\|\mathbf{Q}\| < 1$, then

{31029}
$$(\mathbf{I} + \mathbf{Q})^{-1} = \mathbf{I} - \mathbf{Q} + \mathbf{Q}^2 - \cdots\,, \tag{2.23}$$

which may be verified by multiplying both sides by $\mathbf{I} + \mathbf{Q}$, doing term-by-term multiplication and measuring the remainders with either norm.

Nothing has been said about actually finding the numerical values of either the matrix inverse or the eigenvectors and eigenvalues. Computational algorithms for obtaining them have been developed by experts, and are discussed in many good textbooks.[14] Software systems like MATLAB, Maple, IDL and Mathematica implement them in easy-to-use form. For purposes of this book, we assume the reader has at least a rudimentary knowledge of these techniques and access to a good software implementation.

2.2.4 Identities. Differentiation.

There are some identities and matrix/vector definitions which prove useful.

A square "positive definite" matrix $\mathbf{A}$, is one for which the scalar "quadratic form,"

$$J = \mathbf{x}^T \mathbf{A} \mathbf{x} \tag{2.24}$$ {31013}

is positive for all possible vectors $\mathbf{x}$. (It suffices to consider only symmetric $\mathbf{A}$ because for a general matrix, $\mathbf{x}^T\mathbf{A}\mathbf{x} = \mathbf{x}^T[(\mathbf{A} + \mathbf{A}^T)/2]\mathbf{x}$, which follows from the scalar property of the quadratic form.) If $J \geq 0$ for all $\mathbf{x}$, $\mathbf{A}$ is "positive semi-definite," or "non-negative definite." Linear algebra books show that a necessary and sufficient requirement for positive definiteness is that $\mathbf{A}$ have only positive eigenvalues (Eq. 2.19) and a semi-definite one must have all non-negative eigenvalues.

We end up doing a certain amount of differentiation and other operations with respect to matrices and vectors. A number of formulas are very helpful, and save a lot of writing. They are all demonstrated by doing the derivatives term-by-term. Let $\mathbf{q}$, $\mathbf{r}$ be $N \times 1$ column vectors, and $\mathbf{A}$, $\mathbf{B}$, $\mathbf{C}$ be matrices. The derivative of a matrix by a scalar is just the matrix of element by element deriviatives. Alternatively, if s is any scalar, its derivative by a vector,

$$\frac{\partial s}{\partial \mathbf{q}} = \left[\frac{\partial s}{\partial q_1} \cdots \frac{\partial s}{\partial q_N} \right]^T \tag{2.25}$$ {31014a}

is a column vector (the gradient; some authors define it to be a row vector). The derivative of

one vector by another is defined as a matrix:

$$\frac{\partial \mathbf{r}}{\partial \mathbf{q}} = \left\{ \frac{\partial r_i}{\partial q_j} \right\} = \begin{Bmatrix} \frac{\partial r_1}{\partial q_1} & \frac{\partial r_2}{\partial q_1} & . & \frac{\partial r_M}{\partial q_1} \\ \frac{\partial r_1}{\partial q_2} & . & . & \frac{\partial r_M}{\partial q_2} \\ . & . & . & . \\ \frac{\partial r_1}{\partial q_N} & . & . & \frac{\partial r_M}{\partial q_N} \end{Bmatrix} \equiv \mathbf{B}. \tag{2.26}$$ {31015}

If $\mathbf{r}$, $\mathbf{q}$ are of the same dimension, the determinant of $\mathbf{B} = \det(\mathbf{B})$ is the "Jacobian" of $\mathbf{r}$.[15]

The second derivative of a scalar,

$$\frac{\partial^2 s}{\partial \mathbf{q}^2} = \left\{ \frac{\partial}{\partial \mathbf{q}_i} \frac{\partial s}{\partial \mathbf{q}_j} \right\} = \begin{Bmatrix} \frac{\partial^2 s}{\partial q_1^2} & \frac{\partial^2 s}{\partial q_1 q_2} & . & . & \frac{\partial^2 s}{\partial q_1 q_N} \\ . & . & . & . & . \\ \frac{\partial^2 s}{\partial q_N \partial q_1} & . & . & . & \frac{\partial^2 s}{\partial q_N^2} \end{Bmatrix} \tag{2.27}$$ {31014b}

is the "Hessian" of s and is the derivative of the gradient of s.

Assuming conformability, the inner product, $J = \mathbf{r}^T\mathbf{q} = \mathbf{q}^T\mathbf{r}$, is a scalar. The differential of J is,

$$dJ = d\mathbf{r}^T\mathbf{q} + \mathbf{r}^T d\mathbf{q} = d\mathbf{q}^T\mathbf{r} + \mathbf{q}^T d\mathbf{r}, \tag{2.28}$$

and hence the partial derivatives are,

$$\frac{\partial(\mathbf{q}^T\mathbf{r})}{\partial \mathbf{q}} = \frac{\partial\left(\mathbf{r}^T\mathbf{q}\right)}{\partial \mathbf{q}} = \mathbf{r}\,, \tag{2.29}$$ {31016}

$$\frac{\partial\left(\mathbf{q}^T\mathbf{q}\right)}{\partial \mathbf{q}} = 2\mathbf{q}\,. \tag{2.30}$$ {31017}

It follows immediately that for matrix/vector products,

$$\frac{\partial}{\partial \mathbf{q}}(\mathbf{B}\mathbf{q}) = \mathbf{B}^T,\ \frac{\partial}{\partial \mathbf{q}}\left(\mathbf{q}^T\mathbf{B}\right) = \mathbf{B}. \tag{2.31}$$ {31017a}

The first of these is used repeatedly, and attention is called to the apparently trivial fact that differentiation of $\mathbf{B}\mathbf{q}$ with respect to $\mathbf{q}$ produces the transpose of $\mathbf{B}$—the origin, as seen later, of so-called adjoint models. For a quadratic form,

$$\begin{aligned} J &= \mathbf{q}^T\mathbf{A}\mathbf{q} \\ \frac{\partial J}{\partial \mathbf{q}} &= \left(\mathbf{A} + \mathbf{A}^T\right)\mathbf{q}\,, \end{aligned} \tag{2.32}$$ {31018}

and its Hessian is $2\mathbf{A}$ if $\mathbf{A} = \mathbf{A}^T$.

Differentiation of a scalar function (e.g., J in Eq. 2.32) or a vector by a matrix, $\mathbf{A}$, is readily defined.[16] Differentiation of a matrix by another matrix function by another matrix results in a third, very large, matrix. One special case of the *differential* of a matrix function proves useful later on. It can be shown[17] that

$$d\mathbf{A}^n = (d\mathbf{A})\,\mathbf{A}^{n-1} + \mathbf{A}\,(d\mathbf{A})\,\mathbf{A}^{n-2} + ... + \mathbf{A}^{n-1}\,(d\mathbf{A})\,, \tag{2.33}$$ {deltaA1}

where $\mathbf{A}$ is square. Thus the derivative with respect to some scalar, k, is

$$\frac{d\mathbf{A}^n}{dk} = \frac{(d\mathbf{A})}{dk}\mathbf{A}^{n-1} + \mathbf{A}^{n-2}\frac{(d\mathbf{A})}{dk}\mathbf{A} + ... + \mathbf{A}^{n-1}\left(\frac{d\mathbf{A}}{dk}\right). \tag{2.34}$$ {deltaA2}

There are a few, unfortunately unintuitive, matrix inversion identities which are essential later. They are derived by considering the square, partitioned matrix,

$$\begin{Bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{B}^T & \mathbf{C} \end{Bmatrix} \tag{2.35}$$ {31022}

where $\mathbf{A}^T = \mathbf{A}$, $\mathbf{C}^T = \mathbf{C}$, but $\mathbf{B}$ can be rectangular of conformable dimensions in (2.35).[18] The most important of the identities, sometimes called the "matrix inversion lemma" is, in one form,

$$\begin{aligned} \{\mathbf{C} - \mathbf{B}^T\mathbf{A}^{-1}\mathbf{B}\}^{-1} &= \{\mathbf{I} - \mathbf{C}^{-1}\mathbf{B}^T\mathbf{A}^{-1}\mathbf{B}\}^{-1}\mathbf{C}^{-1} \\ &= \mathbf{C}^{-1} - \mathbf{C}^{-1}\mathbf{B}^T(\mathbf{B}\mathbf{C}^{-1}\mathbf{B}^T - \mathbf{A})^{-1}\mathbf{B}\mathbf{C}^{-1} \end{aligned} \tag{2.36}$$ {31023}

where it is assumed that the inverses exist.[19] A variant is,

$$\mathbf{A}\mathbf{B}^T(\mathbf{C} + \mathbf{B}\mathbf{A}\mathbf{B}^T)^{-1} = (\mathbf{A}^{-1} + \mathbf{B}^T\mathbf{C}^{-1}\mathbf{B})^{-1}\mathbf{B}^T\mathbf{C}^{-1}. \tag{2.37}$$ {31024}

Eq. (2.37) is readily confirmed by left-multiplying both sides by $(\mathbf{A}^{-1} + \mathbf{B}^T\mathbf{C}^{-1}\mathbf{B})$, and right-multiplying by $(\mathbf{C} + \mathbf{B}\mathbf{A}\mathbf{B}^T)$ and showing that the two sides of the resulting equation are equal. Another identity, found by "completing the square," is demonstrated by directly multiplying it out, and requires $\mathbf{C} = \mathbf{C}^T$ ($\mathbf{A}$ is unrestricted, but the matrices must be conformable as shown):

$$\mathbf{A}\mathbf{C}\mathbf{A}^T - \mathbf{B}\mathbf{A}^T - \mathbf{A}\mathbf{B}^T = (\mathbf{A} - \mathbf{B}\mathbf{C}^{-1})\mathbf{C}(\mathbf{A} - \mathbf{B}\mathbf{C}^{-1})^T - \mathbf{B}\mathbf{C}^{-1}\mathbf{B}^T\,. \tag{2.38}$$ {31025}