UJEP

Department of Mathematics

Mathematics for Data Engineers

KMA/MDI

Part I

Linear algebra:
models, geometry, and modes

Lectures 1–4

Yaroslav Bazaikin

17 October 2026

Preface

This is Part I of *Mathematics for Data Engineers*, KMA/MDI, developed at the Department of Mathematics for students of the Department of Informatics, UJEP. Lectures 1–4 follow the revised syllabus of September 2026. They develop the linear algebra used later in optimization, Fourier series, and integral transforms.

Three questions for a linear model

Which observations can the model explain? Can its parameters be recovered uniquely? If the observations do not fit exactly, which approximation minimizes a specified error?

Each lecture ends with six review questions and eight exercises. The last two exercises, marked with an asterisk, extend the core material. Sections labelled *Further perspective* are supplementary.

The course in five parts

The lecture allocation below follows the revised syllabus. This volume contains Part I; the later parts use its language and methods.

Part	Lectures	Subject and purpose
I	1–4	**Linear algebra.** Vector spaces and coordinates; linear maps, kernel, image, and rank; inner products and orthonormal bases; Gram–Schmidt; orthogonal matrices; eigenvalues and eigenvectors.
II	5–7	**Differential calculus and optimization.** Partial and directional derivatives, gradient, Hessian, Taylor models, free extrema, constrained extrema, and Lagrange multipliers.
III	8–9	**Integration and function geometry.** Riemann integration, integration in several dimensions, substitution, function norms, and L^2 approximation.
IV	10	**Function series.** Orthonormal function systems, Fourier series, and Haar wavelet series as signal approximations.
V	11–13	**Integral transforms.** Fourier, Laplace, Hilbert, and wavelet transforms; frequency, dynamics, phase, and localization.

Notation

Vectors are columns. The transpose is denoted by $^{\mathsf{T}}$; I is an identity matrix of the required size. Scalars are real unless complex numbers are explicitly introduced. The Euclidean inner product is $\langle \boldsymbol{x}, \boldsymbol{y} \rangle = \boldsymbol{x}^{\mathsf{T}}\boldsymbol{y}$, and $\|\boldsymbol{x}\|^2 = \boldsymbol{x}^{\mathsf{T}}\boldsymbol{x}$. The symbol span denotes all linear combinations of the listed objects. A matrix whose columns form a basis converts coordinates in that basis to the original coordinates.

Contents

Preface **i**

The course in five parts **ii**

1 Vector spaces and coordinates **1**
1.1 Vectors . . . 1
1.2 Subspaces and linear models . . . 1
1.3 Span, redundancy, and independence . . . 2
1.4 Bases, dimension, and unique coordinates . . . 3
1.5 Changing the basis . . . 4
1.6 Further perspective: choosing function coordinates . . . 5
1.7 Questions and exercises . . . 6

2 Linear maps and information **7**
2.1 Linearity and the matrix of a map . . . 7
2.2 Kernel, image, and rank . . . 7
2.3 Existence and uniqueness of a reconstruction . . . 8
2.4 Why rank plus nullity equals input dimension . . . 9
2.5 Composition and coordinates of an operator . . . 10
2.6 Further perspective: sampling can hide a function . . . 10
2.7 Questions and exercises . . . 12

3 Inner products and approximation **13**
3.1 Measuring length and error . . . 13
3.2 Cauchy–Schwarz and orthogonality . . . 13
3.3 Orthonormal bases and Gram–Schmidt . . . 14
3.4 Projection is the best approximation . . . 15
3.5 Least squares and the calibration model . . . 16
3.6 Further perspective: unequal measurement tolerances . . . 18
3.7 Questions and exercises . . . 19

4 Orthogonal transforms and eigenmodes **20**
4.1 Orthogonal matrices preserve geometry . . . 20
4.2 Four samples: averages and local details . . . 21
4.3 Eigenvectors and the characteristic equation . . . 22
4.4 Symmetric matrices and quadratic geometry . . . 23
4.5 Further perspective: repeated filtering and complex modes . . . 25
4.6 Questions and exercises . . . 26

Lecture 1

Vector spaces, models, and coordinates

A linear model is a span of prescribed objects. A basis identifies the parameters needed to represent each member uniquely.

1.1 Vectors

A *real vector space* is a set V with addition $V \times V \to V$ and scalar multiplication $\mathbb{R} \times V \to V$. Addition is associative and commutative, has a zero vector $\mathbf{0}$, and gives each vector an additive inverse. For all $\boldsymbol{u}, \boldsymbol{v} \in V$ and $\alpha, \beta \in \mathbb{R}$,

$$\alpha(\boldsymbol{u} + \boldsymbol{v}) = \alpha\boldsymbol{u} + \alpha\boldsymbol{v}, \qquad (\alpha + \beta)\boldsymbol{u} = \alpha\boldsymbol{u} + \beta\boldsymbol{u},$$

$$\alpha(\beta\boldsymbol{u}) = (\alpha\beta)\boldsymbol{u}, \qquad 1\boldsymbol{u} = \boldsymbol{u}.$$

The operations take values in V. The axioms imply $0\boldsymbol{u} = \mathbf{0}$: expand $(0 + 0)\boldsymbol{u}$ and subtract $0\boldsymbol{u}$. Similarly, $(-1)\boldsymbol{u} = -\boldsymbol{u}$.

Examples are $\mathbb{R}^n$ with componentwise operations, fixed-size real matrices, and $\mathcal{P}_m$, the polynomials of degree at most m (including zero). The space $C[-1, 1]$ consists of continuous real functions with pointwise operations. A vector is an object in such a space; its coordinates depend on a chosen basis. Our basis and dimension arguments below concern finite-dimensional spaces.

Example 1.1. Calibration: offset, slope, and patterns

A temperature channel is tested at reference temperatures $T = 20, 25, 30$ (in degrees Celsius). Set $t = (T - 25)/5$, giving input levels $-1, 0, 1$, and model its output in volts by $a + bt$. The unknown *offset* a is the output at 25°C; the *slope* b is its change per unit of t, here a 5°C increase.

$$\boldsymbol{p} = (a - b, a, a + b)^{\mathsf{T}} = a\boldsymbol{u} + b\boldsymbol{v}, \qquad \boldsymbol{u} = (1, 1, 1)^{\mathsf{T}}, \quad \boldsymbol{v} = (-1, 0, 1)^{\mathsf{T}}.$$

A *pattern* is a prescribed response vector: increasing a by one adds $\boldsymbol{u}$, while increasing b by one adds $\boldsymbol{v}$. The patterns are known; their coefficients are to be calibrated. For $a = 3$, $b = 1$, the predicted readings are $(2, 3, 4)^{\mathsf{T}}$ volts, ordered by input level.

The model permits real a, b and forms a vector space. The response is affine in t and linear in the parameter pair (a, b).

1.2 Subspaces and linear models

A subset $W \subseteq V$ is a *subspace* if it is a vector space with the operations inherited from V. Lines and planes through the origin are examples of subspaces of Euclidean space.

A usable subspace test

A subset $W \subseteq V$ is a subspace precisely when $\mathbf{0} \in W$ and

$$\alpha\boldsymbol{u} + \beta\boldsymbol{v} \in W \quad \text{for every } \boldsymbol{u}, \boldsymbol{v} \in W \text{ and every } \alpha, \beta \in \mathbb{R}.$$

Reasoning: why the test suffices. The test gives closure under addition ($\alpha = \beta = 1$), scalar multiplication ($\beta = 0$), and additive inverses ($\alpha = -1$). The other axioms are inherited from V. Conversely, a vector space with the inherited operations contains zero and is closed under linear combinations. □

Example 1.2. An exact consistency check

An affine calibration response must have equal increments between equally spaced inputs. For a three-entry record this condition is $x_2 - x_1 = x_3 - x_2$, that is

$$W = \{\boldsymbol{x} \in \mathbb{R}^3 : x_1 - 2x_2 + x_3 = 0\}.$$

Zero satisfies the condition, and for $\boldsymbol{x}, \boldsymbol{y} \in W$,

$$(\alpha x_1 + \beta y_1) - 2(\alpha x_2 + \beta y_2) + (\alpha x_3 + \beta y_3) = 0.$$

Thus W is a subspace. Conversely, set $a = x_2$ and $b = x_3 - x_2$. The condition gives $x_1 = a - b$, hence

$$\boldsymbol{x} = (a - b, a, a + b)^{\mathsf{T}} = a\boldsymbol{u} + b\boldsymbol{v}.$$

So the check characterizes precisely the records our calibration model can explain. Prescribing instead $x_1 - 2x_2 + x_3 = 1$ gives a translated plane, an *affine set*: zero no longer satisfies the constraint. Such nonzero constraints will occur in optimization.

The same test applies to functions. Polynomials in $\mathcal{P}_2$ satisfying $p(0) = 0$ form a subspace: they are $bt + ct^2$. The condition $p(0) = 1$ defines an affine set. Polynomials of *exactly* degree two do not form a vector space: $t^2 + (-t^2) = 0$ leaves the set. The phrase "degree at most" therefore matters.

1.3 Span, redundancy, and independence

A *linear combination* of $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k$ is any sum $c_1\boldsymbol{v}_1 + \cdots + c_k\boldsymbol{v}_k$. Their *span* is the set of all these sums:

$$\operatorname{span}\{\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k\} = \left\{ \sum_{j=1}^{k} c_j\boldsymbol{v}_j : c_j \in \mathbb{R} \right\}.$$

It is a subspace: combining two such sums gives another such sum. It is the smallest subspace containing all the listed vectors, since any subspace containing them must contain their linear combinations.

The vectors are *linearly independent* if

$$c_1\boldsymbol{v}_1 + \cdots + c_k\boldsymbol{v}_k = \mathbf{0} \quad \Longrightarrow \quad c_1 = \cdots = c_k = 0.$$

Otherwise they are dependent. A nonzero relation means that at least one of the vectors can be expressed using the others. Its coefficient can be isolated and divided out. That vector contributes no new direction to the span.

Example 1.3. A redundant calibration feature

Suppose a calibration model already uses $\boldsymbol{u}, \boldsymbol{v}$ and a new feature $\boldsymbol{w} = (0, 1, 2)^{\mathsf{T}}$ is proposed. The original patterns are independent: in $c\boldsymbol{u} + d\boldsymbol{v} = \boldsymbol{0}$, the middle entry gives $c = 0$, then the last gives $d = 0$. But $\boldsymbol{w} = \boldsymbol{u} + \boldsymbol{v}$, so

$$a\boldsymbol{u} + b\boldsymbol{v} + c\boldsymbol{w} = (a + c)\boldsymbol{u} + (b + c)\boldsymbol{v}.$$

The extra coefficient adds no possible prediction. Replacing (a, b, c) by $(a - \delta, b - \delta, c + \delta)$ leaves the predicted record unchanged for every $\delta \in \mathbb{R}$. The data can identify only $a + c$ and $b + c$. More precise readings at these same inputs cannot identify all three coefficients; the model must remove the redundancy or impose additional information.

1.4 Bases, dimension, and unique coordinates

A space is *finite-dimensional* if it has a finite spanning list. A *basis* is a linearly independent spanning list. Thus a basis gives a representation for every vector, with unique coefficients. For a basis $B = (\boldsymbol{b}_1, \dots, \boldsymbol{b}_n)$, write

$$\boldsymbol{x} = c_1\boldsymbol{b}_1 + \cdots + c_n\boldsymbol{b}_n, \qquad [\boldsymbol{x}]_B = (c_1, \dots, c_n)^{\mathsf{T}}.$$

The coordinate column $[\boldsymbol{x}]_B$ depends on the ordered basis. The vector $\boldsymbol{x}$ does not change when we describe it differently.

Reasoning: uniqueness by subtraction. Suppose $\sum_j c_j\boldsymbol{b}_j = \sum_j d_j\boldsymbol{b}_j$. Subtract the two expressions to get $\sum_j (c_j - d_j)\boldsymbol{b}_j = \boldsymbol{0}$. Independence gives $c_j = d_j$ for every j. Thus independence gives unique coordinates. □

The *dimension* of a finite-dimensional space is the number of vectors in any basis. Why does the number not depend on our choice? If n vectors span the space, any list of more than n vectors is dependent: express them in the spanning list and solve the resulting homogeneous system of n equations in more than n unknowns. Elimination leaves a free variable and hence a nonzero relation. Comparing two bases in both directions shows that their lengths agree.

Deleting dependent vectors from a finite spanning list produces a basis. Conversely, an independent list can be extended to a basis: add a vector outside its span until it spans the space. The length bound just proved ensures termination. We will use this extension in the rank–nullity proof.

The standard columns $\boldsymbol{e}_1, \dots, \boldsymbol{e}_n$ form a basis of $\mathbb{R}^n$. The polynomials $1, t, t^2$ form a basis of $\mathcal{P}_2$, so $\dim \mathcal{P}_2 = 3$. The model space $W = \operatorname{span}\{\boldsymbol{u}, \boldsymbol{v}\}$ has dimension two even though its members have three entries. The ambient measurement dimension and the model dimension are different numbers.

Example 1.4. Does the calibration model explain the observations?

At $t = -1, 0, 1$, suppose the channel gives readings $\boldsymbol{y} = (1, 2, 4)^{\mathsf{T}}$. Recovering an exact offset and slope would require

$$a - b = 1, \qquad a = 2, \qquad a + b = 4.$$

The first two equations give $b = 1$; the last two give $b = 2$. Equivalently, $y_1 - 2y_2 + y_3 = 1 \neq 0$, so $\boldsymbol{y} \notin W$.

No parameter choice explains all three readings exactly. This calculation alone does not distinguish measurement error from inadequacy of the affine model. With a specified error measure, Lecture 3 will find the closest predicted record in W.

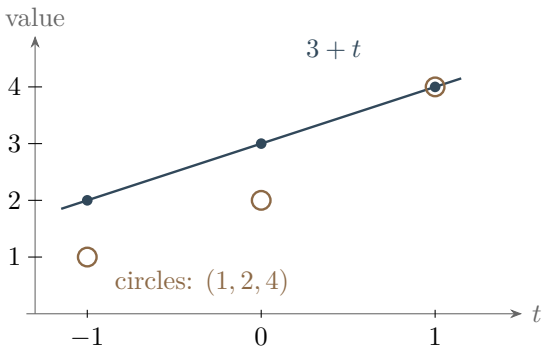


FIGURE 1.1. The filled points belong to one affine-response model. The open circles do not lie on a single straight line. The horizontal coordinate is the input level; it is not one of the three components of the record.

1.5 Changing the basis

Put the basis vectors of $\mathbb{R}^n$ into the columns of P:

$$P = \begin{pmatrix} | & & | \\ \boldsymbol{b}_1 & \cdots & \boldsymbol{b}_n \\ | & & | \end{pmatrix}.$$

Matrix multiplication is a compact notation for a linear combination of columns. Thus

$$\boldsymbol{x} = P[\boldsymbol{x}]_B, \qquad [\boldsymbol{x}]_B = P^{-1}\boldsymbol{x}. \tag{1.1}$$

The inverse exists because a basis gives exactly one coordinate column for every vector. For another basis C with basis matrix S,

$$[\boldsymbol{x}]_C = S^{-1}P[\boldsymbol{x}]_B.$$

The rightmost operation reconstructs the vector from its B-coordinates; the left operation expresses that vector in C.

Example 1.5. Separating a common level from disagreement

Two channels measure the same scalar quantity in the same units and report 3 and 1. Describe their record using $\boldsymbol{b}_1 = (1,1)^{\mathsf{T}}$ and $\boldsymbol{b}_2 = (1,-1)^{\mathsf{T}}$. The basis matrix and its inverse are

$$P = \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}, \qquad P^{-1} = \frac{1}{2}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}.$$

For any record, the coordinates are $(x_1 + x_2)/2$ and $(x_1 - x_2)/2$. In particular,

$$[(3,1)^{\mathsf{T}}]_B = (2,1)^{\mathsf{T}}, \qquad (3,1)^{\mathsf{T}} = 2\boldsymbol{b}_1 + \boldsymbol{b}_2.$$

The first coordinate is the common level; the second is half the disagreement. They summarize different questions without losing either reading. These basis vectors have length $\sqrt{2}$, so $3^2+1^2 = 10$ differs from $2^2+1^2 = 5$. Normalization in Lecture 4 will make squared-coordinate sums preserve length.

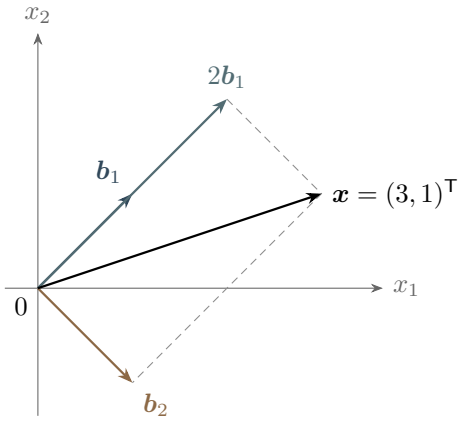


FIGURE 1.2. Two descriptions of the same vector. The dashed parallelogram shows $\boldsymbol{x} = 2\boldsymbol{b}_1 + \boldsymbol{b}_2$. Both coordinate axes use exactly the same scale.

1.6 Further perspective: choosing function coordinates

For the bases $B = (1, t)$ and $C = (1, t-1)$ of $\mathcal{P}_1$, the identity $t - 1 = -1 \cdot 1 + 1 \cdot t$ gives

$$[p]_B = \begin{pmatrix} 1 & -1 \\ 0 & 1 \end{pmatrix} [p]_C.$$

Thus $p(t) = 2 + 3t = 5 + 3(t-1)$ has C-coordinates $(5, 3)^\mathsf{T}$. The slope is unchanged; the constant coordinate now records the value at $t = 1$. Changing the reference input changes coordinates without changing the function.

In $\operatorname{span}\{\cos t, \sin t\}$, independence follows from $a\cos t + b\sin t = 0$ for all t: substituting $t = 0, \pi/2$ gives $a = b = 0$. Every function in this space therefore has unique coordinates in this basis.

For a sinusoidal signal, the addition formula gives

$$A\cos(t - \phi) = (A\cos\phi)\cos t + (A\sin\phi)\sin t.$$

For example, $2\cos(t - \pi/3) = \cos t + \sqrt{3}\sin t$ has coordinates $(1, \sqrt{3})^\mathsf{T}$. Amplitude and phase describe the same signal, but phase is not a linear coordinate. Fourier series will use a cosine–sine pair for each frequency.

1.7 Questions and exercises

Questions for review

T1.1. What operations and algebraic properties define a real vector space? Give one example whose vectors are functions.

T1.2. State the subspace test and explain why a nonzero right-hand side in a linear constraint usually produces an affine set.

T1.3. Explain span and linear independence in terms of a model's expressive power and redundant parameters.

T1.4. Why does a basis give both existence and uniqueness of coordinates? Reproduce the subtraction argument.

T1.5. Why do all bases of a finite-dimensional space have the same length? Distinguish model dimension from measurement dimension.

T1.6. In $\boldsymbol{x} = P[\boldsymbol{x}]_B$, what are the columns of P? Derive the conversion from basis B to basis C.

Exercises

1.1. For each subset of $\mathbb{R}^3$, decide whether it is a subspace. Justify your answer by a general argument or a counterexample: (a) $x_1 + x_2 + x_3 = 0$; (b) $x_1 + x_2 + x_3 = 1$; (c) $x_1, x_2, x_3 \geq 0$.

1.2. With $\boldsymbol{u} = (1, 1, 1)^\mathsf{T}$ and $\boldsymbol{v} = (-1, 0, 1)^\mathsf{T}$, find the coordinates of $(1, 3, 5)^\mathsf{T}$ in $(\boldsymbol{u}, \boldsymbol{v})$. Decide whether $(1, 3, 6)^\mathsf{T}$ belongs to their span. Interpret the result as an exact calibration test at inputs $-1, 0, 1$.

1.3. Determine whether $(1, 0, 1)^\mathsf{T}$, $(0, 1, 1)^\mathsf{T}$, and $(1, 1, 2)^\mathsf{T}$ are independent. Find a basis and the dimension of their span. Give an equation describing this span.

1.4. Find a basis and the dimension of $\{p \in \mathcal{P}_2 : p(1) = 0\}$. Express $p(t) = t^2 - 1$ in your basis.

1.5. For $B = ((1, 1)^\mathsf{T}, (1, -1)^\mathsf{T})$, compute $[(5, -1)^\mathsf{T}]_B$. Reconstruct the original vector and explain the meaning of both coordinates.

1.6. Let $C = ((1, 0)^\mathsf{T}, (1, 1)^\mathsf{T})$. Find the matrix converting B-coordinates from Exercise 1.5 to C-coordinates. Apply it to $(2, 3)^\mathsf{T}$ and check by reconstruction in both bases.

1.7. (*) Prove that the intersection of two subspaces is a subspace. Give two subspaces whose union is not a subspace. Explain why the union differs from their combined span.

1.8. (*) Use the bases $(1, t, t^2)$ and $(1, t-1, (t-1)^2)$ to obtain a coordinate-conversion matrix. Convert $p(t) = 1 + 2t + t^2$ and interpret its new constant coefficient.

Lecture 2

Linear maps and recoverable information

For a linear map, the kernel describes invisible input changes and the image describes attainable outputs. Together they determine existence and uniqueness of reconstruction.

2.1 Linearity and the matrix of a map

A map $T : V \to W$ is *linear* if

$$T(\alpha \boldsymbol{x} + \beta \boldsymbol{y}) = \alpha T\boldsymbol{x} + \beta T\boldsymbol{y}$$

for every pair of vectors and every pair of scalars. This is the superposition rule. In particular, $T\mathbf{0} = \mathbf{0}$. Remark, that the zero-input test is necessary but not sufficient. The scalar map $x \mapsto x^2$ sends zero to zero, yet $(1+1)^2 \neq 1^2 + 1^2$.

Reasoning: why columns determine a linear map. In $\mathbb{R}^n$, write $\boldsymbol{x} = \sum_j x_j \boldsymbol{e}_j$. Linearity gives

$$T\boldsymbol{x} = \sum_{j=1}^{n} x_j T\boldsymbol{e}_j = A\boldsymbol{x}, \qquad A = \begin{pmatrix} | & & | \\ T\boldsymbol{e}_1 & \cdots & T\boldsymbol{e}_n \\ | & & | \end{pmatrix}.$$

Thus the jth column is the output produced by the jth unit input. If the output is in $\mathbb{R}^m$, the matrix has m rows and n columns. □

Example 2.1. The calibration design matrix

Return to the model $a + bt$ at the three fixed input levels. Its parameter-to-prediction map is

$$\begin{pmatrix} a \\ b \end{pmatrix} \longmapsto A \begin{pmatrix} a \\ b \end{pmatrix}, \qquad A = \begin{pmatrix} 1 & -1 \\ 1 & 0 \\ 1 & 1 \end{pmatrix}.$$

This *design matrix* is known once the input levels and model features $1, t$ are chosen. Its columns are the responses to one unit of offset and one unit of slope. For $(a, b) = (3, 1)$, matrix multiplication produces $(2, 3, 4)^{\mathsf{T}}$.

Differentiation is another linear map: $(\alpha p + \beta q)' = \alpha p' + \beta q'$. For example, on $\mathcal{P}_2$, differentiation has columns equal to the coordinates of $0, 1, 2t$ in the chosen output basis. We can therefore use the same ideas for functions and for arrays of data.

2.2 Kernel, image, and rank

The *kernel* and *image* of T are

$$\ker T = \{\boldsymbol{x} : T\boldsymbol{x} = \mathbf{0}\}, \qquad \operatorname{im} T = \{T\boldsymbol{x} : \boldsymbol{x} \in V\}.$$

Here $\ker T \subseteq V$ and $\operatorname{im} T \subseteq W$. Both are subspaces: linear combinations of zero-output inputs still have zero output, and

$$\alpha T\boldsymbol{x} + \beta T\boldsymbol{y} = T(\alpha\boldsymbol{x} + \beta\boldsymbol{y}).$$

The *rank* is $\dim(\operatorname{im} T)$. For a matrix, its image is the span of its columns, so the rank counts independent output directions. The *nullity* is $\dim(\ker T)$; it counts independent input directions that are invisible to the output.

Example 2.2. Storing increments loses the initial level

A data pipeline keeps successive changes of a three-sample record instead of its values. Its encoder is

$$D = \begin{pmatrix} -1 & 1 & 0 \\ 0 & -1 & 1 \end{pmatrix}, \qquad D\boldsymbol{x} = (x_2 - x_1, x_3 - x_2)^{\mathsf{T}}.$$

Which records can be recovered from these two stored numbers? Solving $D\boldsymbol{x} = \boldsymbol{0}$ gives $x_1 = x_2 = x_3$, hence

$$\ker D = \operatorname{span}\{(1,1,1)^{\mathsf{T}}\}, \qquad \dim\ker D = 1.$$

Every increment pair $(d_1, d_2)^{\mathsf{T}}$ is attainable: use $(0, d_1, d_1 + d_2)^{\mathsf{T}}$. Therefore $\operatorname{im} D = \mathbb{R}^2$ and $\operatorname{rank} D = 2$.

The record $(1,2,4)^{\mathsf{T}}$ produces $(1,2)^{\mathsf{T}}$. Every record with these increments is $(c, c+1, c+3)^{\mathsf{T}}$, $c \in \mathbb{R}$. Storing the initial value $x_1 = 1$ alongside the increments selects the original record. Without that anchor, the absolute level cannot be reconstructed.

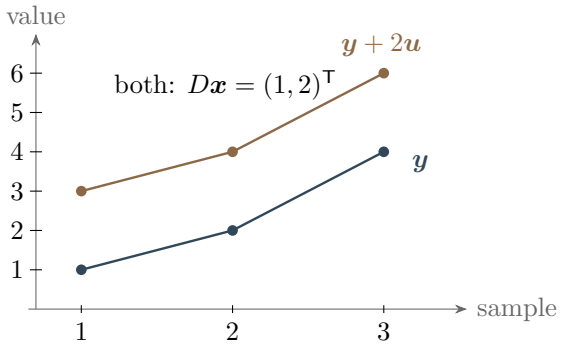


FIGURE 2.1. An entire family of vertically shifted records has the same first differences. The joining segments only help compare increments; no interpolation model is assumed by the difference operator.

A map is *injective* if equal outputs imply equal inputs, and *surjective* if every member of its stated codomain is an output. For a linear map,

$$T \text{ injective} \iff \ker T = \{\boldsymbol{0}\}, \qquad T \text{ surjective onto } W \iff \operatorname{im} T = W.$$

The first equivalence follows from $T\boldsymbol{x} = T\boldsymbol{y}$ if and only if $\boldsymbol{x} - \boldsymbol{y} \in \ker T$. The design map $A : \mathbb{R}^2 \to \mathbb{R}^3$ is injective because its columns are independent, but not surjective because its image is only a plane.

2.3 Existence and uniqueness of a reconstruction

Proposition. Let $A \in \mathbb{R}^{m\times n}$ and $\boldsymbol{b} \in \mathbb{R}^m$. The equation $A\boldsymbol{x} = \boldsymbol{b}$ is solvable if and only if $\boldsymbol{b} \in \operatorname{im} A$. If $\boldsymbol{x}_0$ is one solution, the full solution set is

$$\boldsymbol{x} = \boldsymbol{x}_0 + \boldsymbol{z}, \qquad \boldsymbol{z} \in \ker A. \tag{2.1}$$

Indeed, adding a kernel vector leaves the output unchanged. Conversely, subtracting any two solutions gives $A(\boldsymbol{x} - \boldsymbol{x}_0) = \mathbf{0}$. This is another use of the subtraction argument from Lecture 1.

Three separate questions

1. *Does a solution exist?* Check whether $\boldsymbol{b} \in \operatorname{im} A$.
2. *If it exists, is it unique?* Check whether $\ker A = \{\mathbf{0}\}$.
3. *What freedom remains?* Add every vector in $\ker A$ to one solution.

For an $m \times n$ matrix, zero kernel is equivalent to rank n, and image $\mathbb{R}^m$ is equivalent to rank m.

For a square matrix these two rank requirements coincide. An $n \times n$ matrix is invertible precisely when its rank is n; then $A^{-1}\boldsymbol{b}$ is the unique solution for every right-hand side. For a rectangular matrix one should use the statements about kernel and image directly, rather than speak loosely of "the inverse".

Example 2.3. Elimination tests exact calibration

For the readings $\boldsymbol{y} = (1, 2, 4)^{\mathsf{T}}$, exact calibration means solving $A(a, b)^{\mathsf{T}} = \boldsymbol{y}$. Elimination gives

$$\left(\begin{array}{cc|c} 1 & -1 & 1 \\ 1 & 0 & 2 \\ 1 & 1 & 4 \end{array}\right) \longrightarrow \left(\begin{array}{cc|c} 1 & -1 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{array}\right).$$

Subtract the first row from the other two, then subtract twice the new second row from the third. The last row reads $0 = 1$: the data are outside the image. The two pivot columns still show that any attainable record would determine the parameters uniquely.

For general readings $\boldsymbol{y}$, the final row is $0 = y_1 - 2y_2 + y_3$. If this condition holds, $a = y_2$, $b = y_2 - y_1$. Elimination thus gives both a consistency test and a recovery formula, with a precise condition for using it.

Elementary row operations preserve a system's solution set because each is reversible. After elimination, pivot columns count the rank, and nonpivot variables in a homogeneous system can be chosen freely. To obtain a basis of the *original column space*, take the original columns at the pivot positions; the row operations generally change the actual column vectors.

2.4 Why rank plus nullity equals input dimension

Rank–nullity theorem. For a linear map $T : V \to W$ with V finite-dimensional,

$$\dim V = \dim \ker T + \dim \operatorname{im} T. \tag{2.2}$$

Reasoning: extend a basis of what is lost. Choose a basis $\boldsymbol{k}_1, \dots, \boldsymbol{k}_r$ of $\ker T$. Extend it to a basis of V by adding $\boldsymbol{v}_1, \dots, \boldsymbol{v}_s$. We claim that $T\boldsymbol{v}_1, \dots, T\boldsymbol{v}_s$ form a basis of the image.

They span the image because every input is a combination of the $\boldsymbol{k}_i$ and $\boldsymbol{v}_j$, and all the $T\boldsymbol{k}_i$ vanish. To prove independence, suppose $\sum_j c_j T\boldsymbol{v}_j = \mathbf{0}$. Then $\sum_j c_j \boldsymbol{v}_j$ belongs to the kernel, so it is a combination of the $\boldsymbol{k}_i$. Moving that combination to the left gives a relation among the chosen basis vectors of V. All coefficients must be zero, including every c_j.

Thus $\dim \operatorname{im} T = s$ and $\dim V = r + s$, proving the formula. □

For differencing $D : \mathbb{R}^3 \to \mathbb{R}^2$, the formula is $3 = 1+2$; for the calibration map $A : \mathbb{R}^2 \to \mathbb{R}^3$, it is $2 = 0 + 2$. The extended basis separates lost directions from directions with independent images.

Thus $m < n$ prevents injectivity of a linear map $\mathbb{R}^n \to \mathbb{R}^m$. A known initial value or another constraint can reduce the admissible inputs, but supplies additional information.

2.5 Composition and coordinates of an operator

If T is represented by A and a second map S by B, then their composition is represented by BA:

$$\boldsymbol{x} \longmapsto A\boldsymbol{x} \longmapsto B(A\boldsymbol{x}).$$

The rightmost matrix acts first. For our two recurring maps,

$$DA = \begin{pmatrix} 0 & 1 \\ 0 & 1 \end{pmatrix}.$$

Taking differences of the predictions $a + bt$ eliminates the offset a and produces the slope b twice. The two output numbers are equal, so their count does not imply two independent pieces of information.

Now let P be the input basis matrix and S the output basis matrix. If $\boldsymbol{y} = A\boldsymbol{x}$, then

$$[\boldsymbol{y}]_{\text{out}} = S^{-1}AP[\boldsymbol{x}]_{\text{in}}.$$

For an operator on one space, using the same new basis at input and output gives $P^{-1}AP$. This is the same physical or mathematical operation written in different coordinates.

Example 2.4. What a smoothing rule preserves

To reduce disagreement between two adjacent readings of one variable, replace each by 3/4 of itself plus 1/4 of its neighbour. This is the linear rule

$$M = \begin{pmatrix} 3/4 & 1/4 \\ 1/4 & 3/4 \end{pmatrix}.$$

In the mean–contrast basis $P = \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$,

$$P^{-1}MP = \begin{pmatrix} 1 & 0 \\ 0 & 1/2 \end{pmatrix}.$$

Thus $(3, 1)^{\mathsf{T}}$, with coordinates $(2, 1)^{\mathsf{T}}$, becomes a record with coordinates $(2, 1/2)^{\mathsf{T}}$: the output is $(5/2, 3/2)^{\mathsf{T}}$. The mean is preserved and the contrast is halved.

The rule suppresses genuine variation as well as noise: regarding disagreement as unwanted is a modelling decision. The map is invertible, but reversing it doubles any error in the contrast coordinate. Exact recoverability therefore does not imply stable recovery from perturbed data.

2.6 Further perspective: sampling can hide a function

Sampling at fixed points is a linear map. On $\mathcal{P}_2$, define

$$E(p) = (p(-1), p(0), p(1))^{\mathsf{T}}.$$

For $p(t) = a + bt + ct^2$, its matrix in the basis $(1, t, t^2)$ is

$$E = \begin{pmatrix} 1 & -1 & 1 \\ 1 & 0 & 0 \\ 1 & 1 & 1 \end{pmatrix}.$$

If all three sampled values are zero, then $a = 0$, $-b+c = 0$, and $b+c = 0$, hence $a = b = c = 0$. Its kernel is zero; a quadratic polynomial is uniquely determined by these three samples.

Example 2.5. Interpolation depends on the permitted function class

A pipeline must interpolate the calibration readings between measured input levels. If its model is $\mathcal{P}_2$, solving $E(p) = (1, 2, 4)^{\mathsf{T}}$ yields the unique polynomial

$$p(t) = 2 + \frac{3}{2}t + \frac{1}{2}t^2.$$

If cubic polynomials are also permitted, every

$$p_\alpha(t) = p(t) + \alpha(t^3 - t), \qquad \alpha \in \mathbb{R},$$

has the same measurements. The added term vanishes at $-1, 0, 1$. On $\mathcal{P}_3$, sampling has rank three and kernel $\operatorname{span}\{t^3 - t\}$: its image contains that of $\mathcal{P}_2$, and rank–nullity gives nullity one.

The quadratic is a uniquely recoverable model only within the chosen class. Matching three samples does not establish the physical response between them. A fourth distinct input, or a justified restriction of the model, supplies information that the original samples lack.

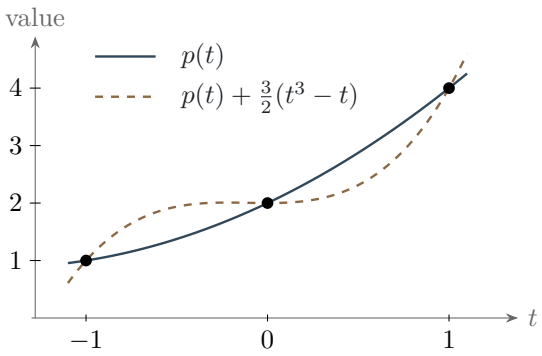


FIGURE 2.2. Two explicitly defined polynomials with identical samples. The dashed curve is allowed in $\mathcal{P}_3$ but not in $\mathcal{P}_2$. The curves are plotted from their formulas, not drawn as qualitative sketches.

2.7 Questions and exercises

Questions for review

T2.1. State linearity and derive the rule that the columns of a matrix are the images of unit input vectors. Why is $T\mathbf{0} = \mathbf{0}$ not a sufficient test?

T2.2. Define kernel, image, rank, and nullity. Explain their meanings for a measurement operator.

T2.3. Prove that all solutions of $A\boldsymbol{x} = \boldsymbol{b}$, when one exists, form $\boldsymbol{x}_0 + \ker A$.

T2.4. State and explain rank–nullity. Reconstruct the proof using a basis of the kernel extended to a basis of the domain.

T2.5. Distinguish injectivity, surjectivity, and invertibility. Which rank conditions apply to an $m \times n$ matrix?

T2.6. Derive $S^{-1}AP$ for the matrix of a map in new input and output bases. What does a diagonal representation reveal about the smoothing example?

Exercises

2.1. Which maps are linear? Give a proof or counterexample: (a) $(x_1, x_2) \mapsto (x_1 + x_2, x_1 - x_2)$; (b) $(x_1, x_2) \mapsto (x_1 + 1, x_2)$; (c) $(x_1, x_2) \mapsto x_1 x_2$.

2.2. For D in Example 2.2, find all records producing $(2, -1)^{\mathsf{T}}$. Select the unique record whose entries sum to zero. Identify the extra information that removes the ambiguity.

2.3. For A in Example 2.1, characterize its image and kernel. Determine whether $(2, 5, 8)^{\mathsf{T}}$ and $(2, 5, 9)^{\mathsf{T}}$ have preimages; find every preimage that exists.

2.4. Find the rank and a basis of the kernel of $B = \begin{pmatrix} 1 & 2 & 1 \\ 0 & 1 & 1 \end{pmatrix}$. Solve $B\boldsymbol{x} = (3, 1)^{\mathsf{T}}$ and verify the rank–nullity identity.

2.5. Compute DA and interpret its kernel as a set of parameter pairs. Explain why A is injective but DA is not.

2.6. Compute $P^{-1}MP$ for the matrices in Example 2.4. Use these coordinates to recover the input that produced $(4, 2)^{\mathsf{T}}$. Which type of measurement error is amplified by recovery?

2.7. (*) Write the matrix of differentiation $T : \mathcal{P}_2 \to \mathcal{P}_1$ in the bases $(1, t, t^2)$ and $(1, t)$. Find its image and kernel, and describe all polynomials with derivative $1 + 2t$.

2.8. (*) Find every polynomial in $\mathcal{P}_3$ with values $0, 1, 0$ at $-1, 0, 1$. Explain why a fourth measurement can remove the ambiguity if it is taken at $t = 2$, but not if it repeats $t = 0$.

Lecture 3

Inner products and best approximation

An inner product specifies how to measure error. Orthogonality then characterizes the unique best approximation within a finite-dimensional subspace.

3.1 Measuring length and error

An *inner product* on a real vector space V is a map $V \times V \to \mathbb{R}$ that is symmetric, linear in its first argument, and positive definite:

$$\langle \boldsymbol{x}, \boldsymbol{y} \rangle = \langle \boldsymbol{y}, \boldsymbol{x} \rangle , \qquad \langle \alpha \boldsymbol{x} + \beta \boldsymbol{z}, \boldsymbol{y} \rangle = \alpha \langle \boldsymbol{x}, \boldsymbol{y} \rangle + \beta \langle \boldsymbol{z}, \boldsymbol{y} \rangle ,$$

$$\langle \boldsymbol{x}, \boldsymbol{x} \rangle > 0 \quad (\boldsymbol{x} \neq \boldsymbol{0}).$$

Symmetry also gives linearity in the second argument. The induced length is $\|\boldsymbol{x}\| = \sqrt{\langle \boldsymbol{x}, \boldsymbol{x} \rangle}$. In $\mathbb{R}^n$ the *Euclidean* inner product is $\langle \boldsymbol{x}, \boldsymbol{y} \rangle = \boldsymbol{x}^{\mathsf{T}} \boldsymbol{y} = \sum_j x_j y_j$.

A *norm* is a function $N : V \to [0, \infty)$ with $N(\boldsymbol{x}) = 0$ exactly when $\boldsymbol{x} = \boldsymbol{0}$, $N(\alpha \boldsymbol{x}) = |\alpha| N(\boldsymbol{x})$, and $N(\boldsymbol{x} + \boldsymbol{y}) \leq N(\boldsymbol{x}) + N(\boldsymbol{y})$. The induced length above is a norm, as proved below. Other norms on $\mathbb{R}^n$ include

$$\|\boldsymbol{x}\|_1 = \sum_j |x_j|, \qquad \|\boldsymbol{x}\|_\infty = \max_j |x_j|.$$

Distance is the norm of a difference. For coordinate vectors, an unsubscripted norm is Euclidean unless a different inner product is specified.

Example 3.1. Choosing an error summary for a data record

Three readings of one variable are compared with reference values in a validation dataset. Suppose measured minus reference gives $\boldsymbol{e} = (0.1, -0.2, 0.2)^{\mathsf{T}}$, in a common unit. Then

$$\|\boldsymbol{e}\|_1 = 0.5, \qquad \|\boldsymbol{e}\| = 0.3, \qquad \|\boldsymbol{e}\|_\infty = 0.2, \qquad \mathrm{RMS}(\boldsymbol{e}) = \frac{0.3}{\sqrt{3}}.$$

The 1-norm reports total absolute discrepancy; the maximum norm answers whether every reading meets a componentwise tolerance. Euclidean length combines squared discrepancies; RMS (Root Mean Square) divides that length by $\sqrt{n}$ to report a typical scale across n entries.

For a tolerance of 0.15 per reading, this record fails because $0.2 > 0.15$. For a criterion based on squared error, its score is 0.09. The criterion must be specified before comparing models. If different channels use different units, divide by stated reference scales before combining their errors.

3.2 Cauchy–Schwarz and orthogonality

The inner product is controlled by the lengths:

$$| \langle \boldsymbol{x}, \boldsymbol{y} \rangle | \leq \|\boldsymbol{x}\| \, \|\boldsymbol{y}\| . \tag{3.1}$$

This is the Cauchy–Schwarz inequality. It justifies the angle formula

$$\cos\theta = \frac{\langle \boldsymbol{x}, \boldsymbol{y} \rangle}{\|\boldsymbol{x}\|\,\|\boldsymbol{y}\|}$$

for two nonzero vectors, since the right-hand side lies in $[-1, 1]$.

Reasoning: a nonnegative squared residual. If $\boldsymbol{y} = \boldsymbol{0}$, the inequality is immediate. Otherwise choose $\alpha = \langle \boldsymbol{x}, \boldsymbol{y} \rangle / \|\boldsymbol{y}\|^2$. Expanding a square gives

$$0 \le \|\boldsymbol{x} - \alpha \boldsymbol{y}\|^2 = \|\boldsymbol{x}\|^2 - \frac{\langle \boldsymbol{x}, \boldsymbol{y} \rangle^2}{\|\boldsymbol{y}\|^2}.$$

Rearrange and take square roots. For $\boldsymbol{y} \neq \boldsymbol{0}$, equality holds exactly when $\boldsymbol{x}$ is a multiple of $\boldsymbol{y}$. Including the zero case, equality is equivalent to linear dependence. □

The triangle inequality follows from the same ingredients:

$$\|\boldsymbol{x} + \boldsymbol{y}\|^2 = \|\boldsymbol{x}\|^2 + 2\,\langle \boldsymbol{x}, \boldsymbol{y} \rangle + \|\boldsymbol{y}\|^2 \le (\|\boldsymbol{x}\| + \|\boldsymbol{y}\|)^2.$$

Thus an inner product really does induce a norm.

Two vectors are *orthogonal* when their inner product is zero. In that case the cross term disappears and we obtain Pythagoras:

$$\|\boldsymbol{x} + \boldsymbol{y}\|^2 = \|\boldsymbol{x}\|^2 + \|\boldsymbol{y}\|^2 . \tag{3.2}$$

This identity holds in every real inner product space.

One application of Cauchy–Schwarz is error propagation. A linear summary $s = \boldsymbol{w}^{\mathsf{T}} \boldsymbol{y}$ changes by $\boldsymbol{w}^{\mathsf{T}} \boldsymbol{e}$ when the data change by $\boldsymbol{e}$. Therefore

$$|\Delta s| \le \|\boldsymbol{w}\|\,\|\boldsymbol{e}\| .$$

For the average of three readings, $\boldsymbol{w} = (1, 1, 1)^{\mathsf{T}}/3$ and $\|\boldsymbol{w}\| = 1/\sqrt{3}$. An error bound $\|\boldsymbol{e}\| \le \varepsilon$ gives $|\Delta s| \le \varepsilon/\sqrt{3}$.

3.3 Orthonormal bases and Gram–Schmidt

A list $\boldsymbol{q}_1, \dots, \boldsymbol{q}_k$ is *orthonormal* if

$$\left\langle \boldsymbol{q}_i, \boldsymbol{q}_j \right\rangle = \begin{cases} 1, & i = j, \\ 0, & i \neq j. \end{cases}$$

Such a list is independent. Indeed, taking the inner product of $\sum_j c_j \boldsymbol{q}_j = \boldsymbol{0}$ with $\boldsymbol{q}_i$ isolates $c_i = 0$. More generally, if $\boldsymbol{x}$ belongs to its span, its coordinates are immediately available:

$$\boldsymbol{x} = \sum_{j=1}^{k} \left\langle \boldsymbol{x}, \boldsymbol{q}_j \right\rangle \boldsymbol{q}_j . \tag{3.3}$$

Each coefficient is therefore a single inner product.

Starting with independent vectors $\boldsymbol{a}_1, \dots, \boldsymbol{a}_k$, the Gram–Schmidt procedure constructs an orthonormal basis of the same span. Set

$$\boldsymbol{r}_1 = \boldsymbol{a}_1, \qquad \boldsymbol{q}_1 = \frac{\boldsymbol{r}_1}{\|\boldsymbol{r}_1\|},$$

and successively use

$$\boldsymbol{r}_j = \boldsymbol{a}_j - \sum_{i=1}^{j-1} \langle \boldsymbol{a}_j, \boldsymbol{q}_i \rangle\, \boldsymbol{q}_i, \qquad \boldsymbol{q}_j = \frac{\boldsymbol{r}_j}{\|\boldsymbol{r}_j\|}. \tag{3.4}$$

Subtract the parts already explained, then normalize the new part.

Reasoning: why the procedure preserves information. Taking the inner product of $\boldsymbol{r}_j$ with any earlier $\boldsymbol{q}_i$ gives zero: the corresponding term in the sum cancels it. Also, $\boldsymbol{r}_j$ differs from $\boldsymbol{a}_j$ by a vector in the earlier span. Hence adding $\boldsymbol{r}_j$ adds exactly the same new direction as adding $\boldsymbol{a}_j$. Independence ensures $\boldsymbol{r}_j \neq \boldsymbol{0}$, so normalization is allowed. For a dependent input list, a zero residual signals redundancy and that vector is omitted. □

Example 3.2. Orthogonalizing calibration features

The affine model is also spanned by $1, t+1$. Add $(t+1)^2$ to permit curvature. At $t = -1, 0, 1$ these features give

$$\boldsymbol{a}_1 = (1,1,1)^{\mathsf{T}}, \quad \boldsymbol{a}_2 = (0,1,2)^{\mathsf{T}}, \quad \boldsymbol{a}_3 = (0,1,4)^{\mathsf{T}}.$$

They form a basis of sampled quadratic responses. Gram–Schmidt gives $\boldsymbol{q}_1 = (1,1,1)^{\mathsf{T}}/\sqrt{3}$ and

$$\boldsymbol{r}_2 = \boldsymbol{a}_2 - (1,1,1)^{\mathsf{T}} = (-1,0,1)^{\mathsf{T}}, \qquad \boldsymbol{q}_2 = \frac{(-1,0,1)^{\mathsf{T}}}{\sqrt{2}}.$$

Removing the constant and linear parts of the third feature gives

$$\boldsymbol{r}_3 = (0,1,4)^{\mathsf{T}} - \tfrac{5}{3}(1,1,1)^{\mathsf{T}} - 2(-1,0,1)^{\mathsf{T}} = \tfrac{1}{3}(1,-2,1)^{\mathsf{T}}.$$

Since $\|\boldsymbol{r}_3\| = \sqrt{6}/3$, we obtain $\boldsymbol{q}_3 = (1,-2,1)^{\mathsf{T}}/\sqrt{6}$. The new coordinates separate common level, linear response, and the second difference $y_1 - 2y_2 + y_3$. The first two describe the original model; the third measures the part it cannot explain. This is a change of features, not a change of the sampled quadratic model.

Nearly dependent input vectors give small residuals, so floating-point orthogonalization needs care. The exact calculation identifies which independent features the algorithm should retain.

3.4 Projection is the best approximation

For a subspace W of a real inner product space V, its *orthogonal complement* is

$$W^{\perp} = \{\boldsymbol{r} \in V : \langle \boldsymbol{r}, \boldsymbol{w} \rangle = 0 \text{ for every } \boldsymbol{w} \in W\}.$$

It is a subspace by linearity. If W is finite-dimensional, every $\boldsymbol{y} \in V$ has a unique decomposition $\boldsymbol{y} = \boldsymbol{p} + \boldsymbol{r}$ with $\boldsymbol{p} \in W$, $\boldsymbol{r} \in W^{\perp}$. Indeed, choose an orthonormal basis $\boldsymbol{q}_1, \ldots, \boldsymbol{q}_k$ of W and set

$$\boldsymbol{p} = \sum_{j=1}^{k} \left\langle \boldsymbol{y}, \boldsymbol{q}_j \right\rangle \boldsymbol{q}_j, \qquad \boldsymbol{r} = \boldsymbol{y} - \boldsymbol{p}. \tag{3.5}$$

Taking inner products with each $\boldsymbol{q}_j$ shows $\boldsymbol{r} \in W^{\perp}$. Two decompositions would differ by a vector in both W and $W^{\perp}$, whose squared norm must be zero. Thus the decomposition is unique, independently of the chosen basis. The vector $\boldsymbol{p}$ is the *orthogonal projection* of $\boldsymbol{y}$ onto W.

The best-approximation identity

For every candidate $\boldsymbol{z} \in W$,

$$\|\boldsymbol{y} - \boldsymbol{z}\|^2 = \|\boldsymbol{y} - \boldsymbol{p}\|^2 + \|\boldsymbol{p} - \boldsymbol{z}\|^2 . \tag{3.6}$$

Therefore $\boldsymbol{p}$ is the unique closest vector in W. The conclusion concerns the chosen inner product and its associated distance.

Reasoning: compare all candidates at once. Write $\boldsymbol{y} - \boldsymbol{z} = (\boldsymbol{y} - \boldsymbol{p}) + (\boldsymbol{p} - \boldsymbol{z})$. The first term is orthogonal to W, while the second belongs to W. Pythagoras gives the identity. The second squared norm vanishes only for $\boldsymbol{z} = \boldsymbol{p}$, proving the unique global minimum. □

Example 3.3. Combining two measurements of a common value

Two equally weighted channels report 3 and 1 for the same quantity. Seek a common estimate c minimizing $(3-c)^2 + (1-c)^2$. Candidate predictions form $W = \operatorname{span}\{(1,1)^\mathsf{T}\}$. Its unit direction is $\boldsymbol{q} = (1,1)^\mathsf{T}/\sqrt{2}$, so

$$\boldsymbol{p} = \left\langle (3,1)^\mathsf{T}, \boldsymbol{q} \right\rangle \boldsymbol{q} = (2,2)^\mathsf{T}, \qquad \boldsymbol{r} = (1,-1)^\mathsf{T}.$$

The projection identity becomes

$$(3-c)^2 + (1-c)^2 = 2 + 2(c-2)^2.$$

Thus $c = 2$ is the unique best common estimate, and the unavoidable squared discrepancy is 2. The residual describes disagreement that a common-value model cannot represent. Equal weighting gives the arithmetic mean; changing the stated error scales will change the estimate.

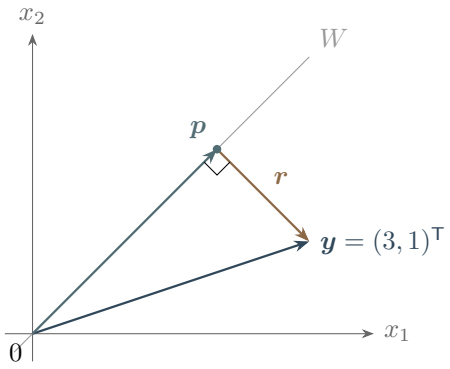


FIGURE 3.1. The residual from the projection is perpendicular to the model line. The three displayed points are exactly 0, $(2,2)$, and $(3,1)$, with equal coordinate scales.

For later constrained optimization, a useful consequence in $\mathbb{R}^n$ is

$$(\ker \boldsymbol{a}^\mathsf{T})^\perp = \operatorname{span}\{\boldsymbol{a}\} \qquad (\boldsymbol{a} \neq \boldsymbol{0}).$$

Indeed, $\ker \boldsymbol{a}^\mathsf{T}$ consists of vectors perpendicular to $\boldsymbol{a}$. Decompose any vector as $\alpha\boldsymbol{a} + \boldsymbol{v}$ with $\boldsymbol{v} \perp \boldsymbol{a}$. If it is also perpendicular to this kernel, its inner product with $\boldsymbol{v}$ is $\|\boldsymbol{v}\|^2 = 0$, so $\boldsymbol{v} = \boldsymbol{0}$. Thus a vector perpendicular to every direction of the hyperplane must be parallel to its normal.

3.5 Least squares and the calibration model

Let $A \in \mathbb{R}^{m\times n}$ and $\boldsymbol{y} \in \mathbb{R}^m$. The least-squares problem is

$$\min_{\boldsymbol{c}\in\mathbb{R}^n} J(\boldsymbol{c}), \qquad J(\boldsymbol{c}) = \|\boldsymbol{y} - A\boldsymbol{c}\|^2 .$$

The best prediction is the projection onto $\operatorname{im} A$. A vector $\boldsymbol{p} = A\boldsymbol{c}$ is that projection if and only if its residual $\boldsymbol{r} = \boldsymbol{y} - A\boldsymbol{c}$ is orthogonal to every column of A. Equivalently, the *normal equations* hold:

$$A^\mathsf{T}(\boldsymbol{y} - A\boldsymbol{c}) = \boldsymbol{0}, \qquad G\boldsymbol{c} = A^\mathsf{T}\boldsymbol{y}, \quad G = A^\mathsf{T}A. \tag{3.7}$$

Here G is the *Gram matrix*: G_{ij} is the inner product of model columns i, j. It is symmetric, and

$$\boldsymbol{h}^{\mathsf{T}} G \boldsymbol{h} = \|A\boldsymbol{h}\|^2 \geq 0, \qquad \ker G = \ker A. \tag{3.8}$$

For the kernel identity, $A\boldsymbol{h} = \mathbf{0}$ implies $G\boldsymbol{h} = \mathbf{0}$; conversely, $G\boldsymbol{h} = \mathbf{0}$ makes the squared norm zero. Thus G is always positive semidefinite and is positive definite precisely when the columns of A are independent; these terms are classified in Lecture 4.

The normal equations always have a solution because the projection exists. The fitted vector is unique. Its parameters are unique if and only if $\ker A = \{\mathbf{0}\}$; then G is invertible and $\boldsymbol{c} = G^{-1}A^{\mathsf{T}}\boldsymbol{y}$. Otherwise all minimizers are $\boldsymbol{c}_* + \ker A$.

Example 3.4. Calibrating by least squares

For the responses $(1, 2, 4)^{\mathsf{T}}$ at $t = -1, 0, 1$, choose the offset and slope that minimize the sum of squared response errors. Equal weights mean that the same response discrepancy has the same cost at every input. For the design matrix in Lecture 2,

$$A^{\mathsf{T}}A = \begin{pmatrix} 3 & 0 \\ 0 & 2 \end{pmatrix}, \qquad A^{\mathsf{T}}\boldsymbol{y} = \begin{pmatrix} 7 \\ 3 \end{pmatrix}.$$

The normal equations give $a = 7/3$, $b = 3/2$, hence

$$\boldsymbol{p} = \begin{pmatrix} 5/6 \\ 7/3 \\ 23/6 \end{pmatrix}, \qquad \boldsymbol{r} = \boldsymbol{y} - \boldsymbol{p} = \frac{1}{6}\begin{pmatrix} 1 \\ -2 \\ 1 \end{pmatrix}.$$

The calibrated response at the reference input is $7/3$; one unit of input increases the predicted response by $3/2$. The residual is observed minus predicted: the model underpredicts the outer readings and overpredicts the middle one.

Both $\boldsymbol{u}^{\mathsf{T}}\boldsymbol{r}$ and $\boldsymbol{v}^{\mathsf{T}}\boldsymbol{r}$ are zero, verifying the normal equations. The minimum loss is $J_* = 1/6$, and RMS is $\sqrt{J_*/3} = 1/(3\sqrt{2})$. More generally,

$$J(7/3 + \delta a,\, 3/2 + \delta b) = \frac{1}{6} + 3(\delta a)^2 + 2(\delta b)^2.$$

This proves uniqueness and quantifies the cost of parameter errors. The loss J is the sum of squares, with no factor $1/2$; this convention is retained in Part II.

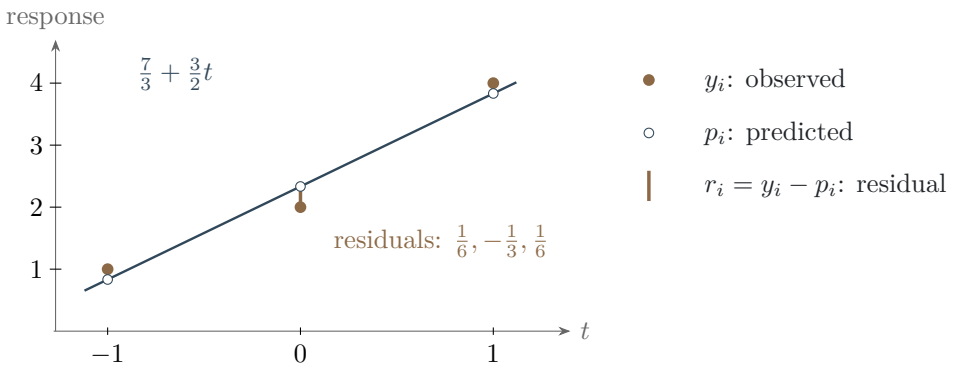


FIGURE 3.2. The least-squares line and its vertical data discrepancies. Orthogonality of the residual is a statement in the three-dimensional space of records, not a claim that the vertical segments are perpendicular to the line in this input–response plot.

The same calculation in the orthonormal basis from Example 3.2 gives

$$c_1 = \frac{7}{\sqrt{3}}, \qquad c_2 = \frac{3}{\sqrt{2}}, \qquad c_3 = \frac{1}{\sqrt{6}}.$$

The fit retains $c_1\boldsymbol{q}_1 + c_2\boldsymbol{q}_2$ and discards $c_3\boldsymbol{q}_3$. Pythagoras reads

$$\|\boldsymbol{y}\|^2 = 21 = \frac{49}{3} + \frac{9}{2} + \frac{1}{6}.$$

The discarded curvature coefficient accounts for the entire fitting error. The same identity will give errors of Fourier and wavelet truncation.

3.6 Further perspective: unequal measurement tolerances

If different measurements have different error scales, use a weighted inner product

$$\langle \boldsymbol{x}, \boldsymbol{y} \rangle_W = \boldsymbol{x}^{\mathsf{T}} W \boldsymbol{y}, \qquad W = \operatorname{diag}(w_1, \dots, w_n), \quad w_j > 0.$$

Strictly positive weights ensure positive definiteness. With $w_j = 1/s_j^2$, this is ordinary Euclidean geometry after dividing component j by its specified scale s_j. Projection and least squares still work, but orthogonality now refers to this weighted inner product. The normal equations become

$$A^{\mathsf{T}} W A \boldsymbol{c} = A^{\mathsf{T}} W \boldsymbol{y}.$$

Example 3.5. Combining readings with unequal tolerances

In a separate experiment, three channels measure one common quantity and report $1, 2, 4$. Their specified error scales are $1, 1, 1/2$ in the same units. Standardizing each residual gives

$$J(c) = (1-c)^2 + (2-c)^2 + 4(4-c)^2.$$

Weighted orthogonality to $(1,1,1)^{\mathsf{T}}$ requires

$$(1-c) + (2-c) + 4(4-c) = 0, \qquad c = \frac{19}{6}.$$

The third reading counts four times as much because its stated error scale is half as large. The estimate moves from the equal-weight mean $7/3$ towards 4.

3.7 Questions and exercises

Questions for review

T3.1. State the real inner product properties and the norm they induce. Distinguish Euclidean length, RMS, and maximum component error.

T3.2. Prove Cauchy–Schwarz by subtracting a suitable multiple of a vector. What does its equality case mean?

T3.3. Explain why an orthonormal list is independent and why its coordinates can be found one at a time by inner products.

T3.4. State Gram–Schmidt and explain both orthogonality and preservation of span. What does a zero residual mean?

T3.5. Define $W^{\perp}$. Prove the orthogonal decomposition and best-approximation identity for finite-dimensional W. Why is the projection unique?

T3.6. Derive the normal equations and $\ker(A^{\mathsf{T}}A) = \ker A$. When are the prediction and its parameters unique?

Exercises

3.1. For $\boldsymbol{e} = (2, -1, 2)^{\mathsf{T}}$, calculate the Euclidean norm, the 1-norm, the maximum norm, and RMS. If every component is divided by a common tolerance $s > 0$, how does each quantity change?

3.2. Use Cauchy–Schwarz to bound $|e_1 + 2e_2|$ when $\|\boldsymbol{e}\| \le 0.1$ in $\mathbb{R}^2$. Give an error vector attaining the positive upper bound.

3.3. Apply Gram–Schmidt in the stated order to $(1, 1, 0)^{\mathsf{T}}$ and $(1, 0, 1)^{\mathsf{T}}$. Check the two resulting unit lengths and their inner product.

3.4. Project $(2, 0, 1)^{\mathsf{T}}$ onto $W = \operatorname{span}\{(1, 1, 1)^{\mathsf{T}}\}$. Describe $W^{\perp}$ by an equation and verify that the residual lies in it. Find the squared error.

3.5. Fit $a + bt$ to values $(2, 2, 5)^{\mathsf{T}}$ at $t = -1, 0, 1$. Find a, b, all fitted values, the residual, and its squared norm. Verify both normal equations directly.

3.6. For $\boldsymbol{y} = (2, 2, 5)^{\mathsf{T}}$, calculate its coefficients in the three-vector orthonormal basis of Example 3.2. Use the discarded coefficient to reproduce the error in Exercise 3.5.

3.7. (*) Derive the weighted best constant for arbitrary data y_j and strictly positive weights w_j. Apply the formula to $(1, 2, 4)^{\mathsf{T}}$ with weights $(2, 1, 1)$.

3.8. (*) Let Q have orthonormal columns and set $P = QQ^{\mathsf{T}}$. Prove $P^{\mathsf{T}} = P$ and $P^2 = P$. Identify its image and kernel, and explain why applying the projection twice gives no further change.

Lecture 4

Orthogonal transforms and eigenmodes

Orthonormal coordinates preserve squared length; eigenvector coordinates separate an operator into scalar actions. For symmetric matrices, one basis does both.

4.1 Orthogonal matrices preserve geometry

A real square matrix Q is *orthogonal* when

$$Q^{\mathsf{T}}Q = I.$$

The (i, j) entry of $Q^{\mathsf{T}}Q$ is the inner product of columns i and j. Thus an orthogonal matrix is exactly a matrix whose columns form an orthonormal basis of $\mathbb{R}^n$.

Reasoning: the transpose is the inverse. If $Q\boldsymbol{x} = \mathbf{0}$, multiplying by Q^{T} gives $\boldsymbol{x} = \mathbf{0}$. The square matrix Q is therefore invertible, and $Q^{\mathsf{T}}Q = I$ implies $Q^{-1} = Q^{\mathsf{T}}$. Consequently $QQ^{\mathsf{T}} = I$ as well. Its rows, as well as its columns, are orthonormal. □

For any vectors $\boldsymbol{x}, \boldsymbol{y}$,

$$\langle Q\boldsymbol{x}, Q\boldsymbol{y}\rangle = \boldsymbol{x}^{\mathsf{T}}Q^{\mathsf{T}}Q\boldsymbol{y} = \langle \boldsymbol{x}, \boldsymbol{y}\rangle .$$

Hence orthogonal matrices preserve lengths, distances, and angles. For example, $\operatorname{diag}(1, -1)$ reflects the plane across its first coordinate axis. Taking determinants of $Q^{\mathsf{T}}Q = I$ gives $(\det Q)^2 = 1$: the determinant is 1 or -1.

Analysis, reconstruction, and energy

If the columns $\boldsymbol{q}_j$ of Q form an orthonormal basis, then

$$\boldsymbol{c} = Q^{\mathsf{T}}\boldsymbol{x}, \qquad \boldsymbol{x} = Q\boldsymbol{c}, \qquad \|\boldsymbol{x}\|^2 = \sum_j c_j^2.$$

For a subset S of the coefficients, set $\boldsymbol{p} = \sum_{j\in S} c_j \boldsymbol{q}_j$. Then

$$\|\boldsymbol{x} - \boldsymbol{p}\|^2 = \sum_{j\notin S} c_j^2.$$

The final identity gives an exact error budget for discarded coordinates.

For a tall $n \times k$ matrix with orthonormal columns, $Q^{\mathsf{T}}Q = I_k$ and $\|Q\boldsymbol{c}\| = \|\boldsymbol{c}\|$. If $k < n$, analysis followed by reconstruction gives $QQ^{\mathsf{T}}\boldsymbol{x}$, the projection onto $\operatorname{im} Q$. It equals $\boldsymbol{x}$ exactly when $\boldsymbol{x} \in \operatorname{im} Q$; $QQ^{\mathsf{T}} \neq I_n$.

Example 4.1. A normalized representation of two readings

For the two readings $(3,1)^{\mathsf{T}}$, use the unit mean and contrast directions

$$Q = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}, \qquad \boldsymbol{c} = Q^{\mathsf{T}}\boldsymbol{x} = (2\sqrt{2}, \sqrt{2})^{\mathsf{T}}.$$

The coefficients retain all information and satisfy $\|\boldsymbol{x}\|^2 = 10 = 8 + 2$. Their values are $\sqrt{2}$ times the mean and half-difference, because the patterns now have unit length.

If a compressed representation keeps only the common level, reconstruction gives $(2,2)^{\mathsf{T}}$. Its squared error is the discarded contrast energy $(\sqrt{2})^2 = 2$. Normalization makes this loss directly readable from the coefficient, without reconstructing and subtracting the records.

4.2 Four samples: averages and local details

In $\mathbb{R}^4$, consider the four unit patterns

$$\boldsymbol{h}_0 = \tfrac{1}{2}(1,1,1,1)^{\mathsf{T}}, \qquad \boldsymbol{h}_1 = \tfrac{1}{2}(1,1,-1,-1)^{\mathsf{T}},$$
$$\boldsymbol{h}_2 = \tfrac{1}{\sqrt{2}}(1,-1,0,0)^{\mathsf{T}}, \qquad \boldsymbol{h}_3 = \tfrac{1}{\sqrt{2}}(0,0,1,-1)^{\mathsf{T}}.$$

Their inner products vanish pairwise: in each overlap the positive and negative contributions cancel, and the last two patterns have disjoint nonzero entries. Their squared lengths are one. They therefore form an orthonormal basis.

These are finite Haar patterns. Their coefficients measure overall level, the difference between half-record averages, and two local differences. They separate variation at two scales.

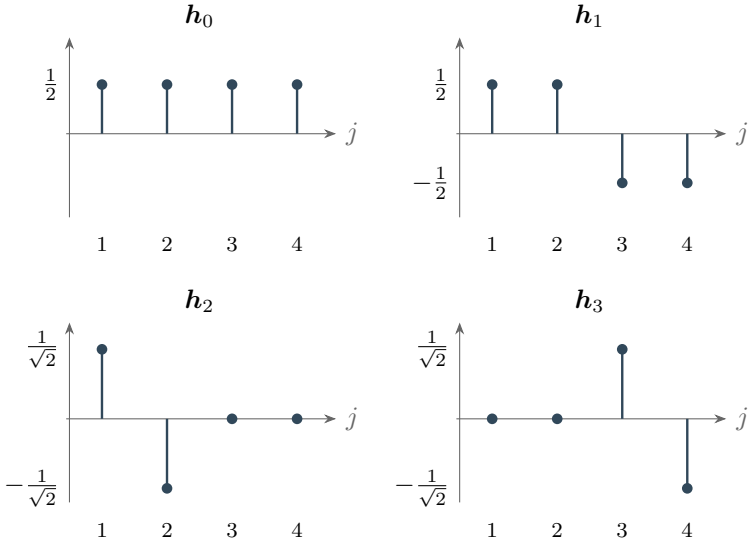


FIGURE 4.1. Four discrete orthonormal patterns. All panels have identical scales. The vertical stems show vector components; the normalization is a sum of four squares, not a continuous integral.

Example 4.2. Compressing a record while preserving half-record averages

A four-sample record is $\boldsymbol{x} = (1,3,2,6)^{\mathsf{T}}$. Store two coefficients so that reconstruction is constant on the first and the last adjacent pairs and preserves each pair's average. The required subspace is $\operatorname{span}\{\boldsymbol{h}_0, \boldsymbol{h}_1\}$. The full Haar transform (that is coordinates in Haar basis) gives

$$c_0 = 6, \quad c_1 = -2, \quad c_2 = -\sqrt{2}, \quad c_3 = -2\sqrt{2}.$$

For example, $c_1 = (1+3-2-6)/2 = -2$ and $c_3 = (2-6)/\sqrt{2} = -2\sqrt{2}$. Exact reconstruction and energy are

$$\boldsymbol{x} = 6\boldsymbol{h}_0 - 2\boldsymbol{h}_1 - \sqrt{2}\boldsymbol{h}_2 - 2\sqrt{2}\boldsymbol{h}_3, \qquad \|\boldsymbol{x}\|^2 = 50 = 36 + 4 + 2 + 8.$$

Keeping c_0, c_1 gives the projection

$$\boldsymbol{p} = 6\boldsymbol{h}_0 - 2\boldsymbol{h}_1 = (2, 2, 4, 4)^{\mathsf{T}}.$$

The pair averages are 2 and 4, as required. The discarded local contrasts give squared error $c_2^2 + c_3^2 = 10$. This is the best approximation within the stated subspace; selecting any two coefficients is a different compression problem.

The retained energy is $40/50 = 80\%$, whereas relative RMS error is

$$\frac{\mathrm{RMS}(\boldsymbol{x} - \boldsymbol{p})}{\mathrm{RMS}(\boldsymbol{x})} = \sqrt{\frac{10}{50}} = \frac{1}{\sqrt{5}} \approx 0.447.$$

Losing 20% of squared norm therefore means about 44.7% relative RMS error. The full transform loses no information; truncation does.

Here *energy* means squared Euclidean norm. A physical time-integrated energy additionally requires the sampling interval and physical scaling. Parts III–IV replace sums by integrals and state the corresponding normalization explicitly.

4.3 Eigenvectors and the characteristic equation

Let A be an $n \times n$ matrix. A nonzero vector $\boldsymbol{v}$ is an *eigenvector* with *eigenvalue* λ if

$$A\boldsymbol{v} = \lambda\boldsymbol{v}.$$

The operation scales the vector, reversing its direction if $\lambda < 0$. Zero is excluded because it satisfies the equation for every λ. Nonzero multiples of an eigenvector describe the same direction.

For a fixed eigenvalue, its eigenvectors together with zero form the *eigenspace* $\ker(A - \lambda I)$. A nonzero solution of $(A - \lambda I)\boldsymbol{v} = \boldsymbol{0}$ exists exactly when that square matrix is singular, giving the *characteristic equation*

$$\det(A - \lambda I) = 0. \tag{4.1}$$

For a 2×2 matrix, $\det \begin{pmatrix} a & b \\ c & d \end{pmatrix} = ad - bc$, so

$$\det(A - \lambda I) = \lambda^2 - (a + d)\lambda + (ad - bc).$$

First find the roots; then solve a homogeneous system for each root. A repeated root still requires finding the dimension of its eigenvector space.

Example 4.3. Finding the smoothing modes systematically

The smoothing rule from Lecture 2 has matrix $M = \begin{pmatrix} 3/4 & 1/4 \\ 1/4 & 3/4 \end{pmatrix}$. To identify directions whose shape is preserved, compute

$$\det(M - \lambda I) = (3/4 - \lambda)^2 - 1/16 = (\lambda - 1)(\lambda - 1/2).$$

For $\lambda = 1$, the eigenvector equation gives $v_1 = v_2$; for $\lambda = 1/2$, it gives $v_1 = -v_2$. Unit choices are

$$\boldsymbol{q}_+ = \frac{(1,1)^\mathsf{T}}{\sqrt{2}}, \qquad \boldsymbol{q}_- = \frac{(1,-1)^\mathsf{T}}{\sqrt{2}}.$$

Thus the mean and contrast patterns are *modes*: one application preserves their shapes and multiplies their coefficients by 1 and 1/2. These gains quantify exactly what the filter retains and attenuates. The characteristic equation recovers the useful basis without guessing it.

A real $n \times n$ matrix is *diagonalizable over* $\mathbb{R}$ if and only if it has a basis of real eigenvectors. With those vectors as the columns of P,

$$AP = P\Lambda, \qquad P^{-1}AP = \Lambda, \qquad \Lambda = \operatorname{diag}(\lambda_1, \ldots, \lambda_n).$$

Conversely, if P is invertible and $P^{-1}AP = \Lambda$ is diagonal, $AP = P\Lambda$ says that the columns of P form an eigenbasis. In these coordinates each component is scaled separately.

Not every real matrix has such a real basis. The shear $\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$ has only eigenvalue 1, but its eigenvectors satisfy $v_2 = 0$, giving only one independent direction. The quarter-turn $\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$ has characteristic equation $\lambda^2 + 1 = 0$ and no real eigenvectors. Orthogonality of a matrix does not guarantee real eigenvectors. Symmetry provides a different and particularly useful guarantee.

4.4 Symmetric matrices and quadratic geometry

A matrix is *symmetric* if $H^\mathsf{T} = H$. Its action interacts with the Euclidean inner product through

$$\langle H\boldsymbol{u}, \boldsymbol{v} \rangle = \langle \boldsymbol{u}, H\boldsymbol{v} \rangle .$$

This identity explains why its different eigenvalue directions are perpendicular.

Reasoning: distinct symmetric modes are orthogonal. Suppose $H\boldsymbol{u} = \lambda\boldsymbol{u}$ and $H\boldsymbol{v} = \mu\boldsymbol{v}$. Then

$$\lambda \langle \boldsymbol{u}, \boldsymbol{v} \rangle = \langle H\boldsymbol{u}, \boldsymbol{v} \rangle = \langle \boldsymbol{u}, H\boldsymbol{v} \rangle = \mu \langle \boldsymbol{u}, \boldsymbol{v} \rangle .$$

If $\lambda \neq \mu$, their inner product is zero. Symmetry is essential to the middle equality. □

Real spectral theorem

Every real symmetric matrix has an orthonormal basis of real eigenvectors. Equivalently,

$$H = Q\Lambda Q^\mathsf{T}, \qquad Q^\mathsf{T}Q = I, \qquad \Lambda = \operatorname{diag}(\lambda_1, \ldots, \lambda_n), \quad \lambda_j \in \mathbb{R}.$$

Repeated eigenvalues are allowed; an orthonormal basis can be chosen within each eigenspace. We don't give here existence proof.

For a real symmetric matrix H, the *quadratic form* $q(\boldsymbol{x}) = \boldsymbol{x}^\mathsf{T}H\boldsymbol{x}$ becomes, with $\boldsymbol{z} = Q^\mathsf{T}\boldsymbol{x}$,

$$q(\boldsymbol{x}) = \boldsymbol{z}^\mathsf{T}\Lambda\boldsymbol{z} = \sum_j \lambda_j z_j^2. \tag{4.2}$$

The table defines its sign classes and states the equivalent eigenvalue criteria. Each inequality for q must hold for every $\boldsymbol{x} \neq \mathbf{0}$.

Classification	Definition	Eigenvalue condition
Positive definite	$q(\boldsymbol{x}) > 0$	All $\lambda_j > 0$
Negative definite	$q(\boldsymbol{x}) < 0$	All $\lambda_j < 0$
Positive semidefinite	$q(\boldsymbol{x}) \geq 0$	All $\lambda_j \geq 0$
Negative semidefinite	$q(\boldsymbol{x}) \leq 0$	All $\lambda_j \leq 0$
Indefinite	Both signs occur	Some $\lambda_i > 0$, some $\lambda_j < 0$

Reasoning: signs in orthonormal coordinates. The sum in (4.2) has the indicated sign when its coefficients do; strict positivity follows because $\boldsymbol{x} \neq \boldsymbol{0}$ has a nonzero coordinate. Conversely, $q(\boldsymbol{q}_j) = \lambda_j$ on each unit eigenvector. This proves the definite and semidefinite criteria; the sum takes both signs exactly when its coefficients include both signs. □

Semidefinite includes definite as a special case; *singular semidefinite* means that a zero eigenvalue is present. A zero eigenvalue can also occur in an indefinite form. For every real symmetric H,

$$\lambda_{\min} \|\boldsymbol{x}\|^2 \leq q(\boldsymbol{x}) \leq \lambda_{\max} \|\boldsymbol{x}\|^2 . \tag{4.3}$$

Compare each λ_j in (4.2) with the extreme eigenvalues and use $\|\boldsymbol{z}\| = \|\boldsymbol{x}\|$. In the positive definite case the lower bound is strictly positive away from zero; this quantitative bound underlies the Hessian test in Part II.

Example 4.4. Reconciling inconsistent measurements

Two channels have unknown deviations x_1, x_2 from a reference level; either deviation may be negative. Separate readings give 1 and 0, while an aggregate channel reports $x_1 + x_2 = 0$. These observations are inconsistent. The individual error scales are 1; the aggregate scale is $\sqrt{2}$. Seek the best reconciliation by minimizing standardized squared discrepancies:

$$L(\boldsymbol{x}) = (x_1 - 1)^2 + x_2^2 + \tfrac{1}{2}(x_1 + x_2)^2 = 1 + \tfrac{1}{2}\boldsymbol{x}^{\mathsf{T}} H \boldsymbol{x} - \boldsymbol{b}^{\mathsf{T}} \boldsymbol{x},$$

$$H = \begin{pmatrix} 3 & 1 \\ 1 & 3 \end{pmatrix}, \qquad \boldsymbol{b} = (2, 0)^{\mathsf{T}}.$$

The mean and contrast eigenvectors have eigenvalues 4 and 2, so H is positive definite. Solve $H\boldsymbol{x}_* = \boldsymbol{b}$ to obtain $\boldsymbol{x}_* = (3/4, -1/4)^{\mathsf{T}}$. Symmetry gives

$$\begin{aligned} L(\boldsymbol{x}_* + \boldsymbol{h}) &= L(\boldsymbol{x}_*) + (H\boldsymbol{x}_* - \boldsymbol{b})^{\mathsf{T}} \boldsymbol{h} + \tfrac{1}{2}\boldsymbol{h}^{\mathsf{T}} H \boldsymbol{h} \\ &= \tfrac{1}{4} + \tfrac{1}{2}\boldsymbol{h}^{\mathsf{T}} H \boldsymbol{h}. \end{aligned}$$

The last term is positive for $\boldsymbol{h} \neq \boldsymbol{0}$, proving a unique global minimum. This is completion of the square in several variables.

The reconciled deviations are $0.75, -0.25$ and their sum is 0.5. The residuals (observed minus predicted) are $0.25, 0.25, -0.5$, with weighted squared sum $1/4$. They balance the stated tolerances. With $z_\pm = (h_1 \pm h_2)/\sqrt{2}$, the increase in loss is $2z_+^2 + z_-^2$: equal displacement along the mean mode costs twice as much as along the contrast mode.

Zero eigenvalues give flat directions of a quadratic form. For a twice continuously differentiable objective at a stationary point, a singular semidefinite Hessian leaves the extremum test inconclusive; higher-order terms matter. An indefinite Hessian still detects a saddle even if some eigenvalues are zero.

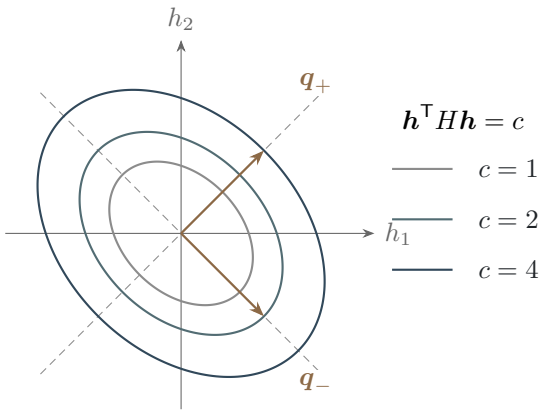


FIGURE 4.2. Contours in displacement coordinates $\boldsymbol{h} = \boldsymbol{x} - \boldsymbol{x}_*$. Their semiaxes are $\sqrt{c/4}$ and $\sqrt{c/2}$; the loss is $L = 1/4 + c/2$. Axis scales are equal.

4.5 Further perspective: repeated filtering and complex modes

Example 4.5. How repeated smoothing affects reconstruction

Starting from $\boldsymbol{x}^{(0)} = (3, 1)^\mathsf{T}$, after $k \geq 0$ smoothing steps the output is

$$\boldsymbol{x}^{(k)} = M^k \boldsymbol{x}^{(0)} = \begin{pmatrix} 2 + 2^{-k} \\ 2 - 2^{-k} \end{pmatrix}.$$

The distance to $(2, 2)^\mathsf{T}$ is $\sqrt{2}\, 2^{-k}$. Each finite power is invertible, but output contrast error ε is amplified to $2^k \varepsilon$ by its inverse. Additional smoothing increases reconstruction sensitivity to rounding and measurement errors.

The limiting projection onto constant records is not invertible: contrast is lost. It suppresses genuine variation as well as noise.

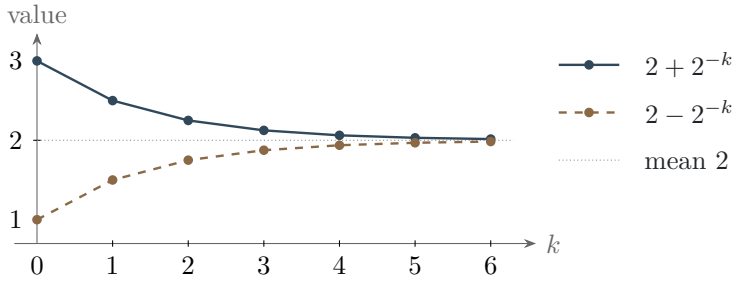


FIGURE 4.3. Exact discrete iterates of the smoothing rule. Lines join successive iterates only to guide the eye. The average stays at two while the contrast decays geometrically.

Over $\mathbb{C}$, the quarter-turn has eigenvalues $\pm i$. Euler's identity $e^{i\theta} = \cos\theta + i \sin\theta$ combines cosine and sine for Fourier analysis. The complex inner product convention will be

$$\langle \boldsymbol{z}, \boldsymbol{w} \rangle = \sum_j z_j \overline{w_j}, \qquad \|\boldsymbol{z}\|^2 = \sum_j |z_j|^2.$$

It is linear in the first argument and conjugate-linear in the second. An orthonormal basis matrix satisfies $Q^* Q = I$, with $Q^* = \overline{Q}^\mathsf{T}$; coefficients are $Q^* \boldsymbol{z}$, and discarded energy is a sum of $|c_j|^2$.

4.6 Questions and exercises

Questions for review

T4.1. Give equivalent descriptions of a real orthogonal matrix. Prove that its transpose is its inverse and that it preserves inner products.

T4.2. Derive the reconstruction and energy identities for orthonormal coordinates. How does the identity change when coefficients are discarded?

T4.3. Explain the four Haar patterns and calculate what each coefficient measures. Distinguish an energy fraction from a relative RMS error.

T4.4. Define eigenvalue, eigenvector, and eigenspace. Explain the characteristic equation and why finding its roots alone is not enough for diagonalization.

T4.5. State the real spectral theorem and prove orthogonality of eigenvectors belonging to distinct eigenvalues of a symmetric matrix.

T4.6. Define the five sign classes of symmetric forms and prove their eigenvalue criteria. How does completion of the square establish a unique minimum?

Exercises

4.1. Check that $Q = \frac{1}{5}\begin{pmatrix} 3 & -4 \\ 4 & 3 \end{pmatrix}$ is orthogonal. Transform $(1,2)^{\mathsf{T}}$ by Q^{T}, reconstruct it using Q, and verify preservation of the squared norm.

4.2. Find the four Haar coefficients of $(2,0,4,2)^{\mathsf{T}}$. Reconstruct the record. Retain only $\boldsymbol{h}_0$ and $\boldsymbol{h}_1$, and calculate both the approximation and its squared error.

4.3. Find the eigenvalues and unit eigenvectors of $\begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}$. Write an orthogonal diagonalization and check it by multiplication.

4.4. For $\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$, compute the characteristic polynomial and the eigenspace. Explain why the repeated eigenvalue does not provide a basis of eigenvectors.

4.5. Let $H = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}$. Express its form in eigenvector coordinates and find the semiaxes of $\boldsymbol{x}^{\mathsf{T}} H \boldsymbol{x} = 3$. Use eigenvalues to classify $H - \alpha I$ for $\alpha = 0, 1, 2, 3, 4$.

4.6. For the same H, minimize $F(\boldsymbol{x}) = \frac{1}{2}\boldsymbol{x}^{\mathsf{T}} H \boldsymbol{x} - (3,0)\boldsymbol{x}$ by solving $H\boldsymbol{x}_* = (3,0)^{\mathsf{T}}$ and completing the square. Give the minimum value and justify uniqueness.

4.7. (*) For the record in Example 4.2, retain the mean coefficient and only one detail coefficient. Which detail gives the smallest squared error? Compare the result with retaining only $\boldsymbol{h}_0$ and $\boldsymbol{h}_1$. Explain the general coefficient-selection principle for a fixed number of orthonormal terms.

4.8. (*) Starting from $(4,0)^{\mathsf{T}}$, apply the smoothing matrix M repeatedly. Find the smallest integer k for which the distance to $(2,2)^{\mathsf{T}}$ is below 0.05. Quantify the amplification of contrast errors when those k steps are inverted.

A compact set of connections

Linear algebra in this part	Later use
Basis and coordinate change	A representation adapted to a signal or a local model
Kernel and rank	Identifiability of linear models and local derivative models
Orthogonal complements	Tangent directions and normals in constrained optimization
Projection and Gram matrices	Least squares, its Hessian, and function approximation
Orthonormal energy identity	Error of truncating a Fourier or wavelet expansion
Spectral theorem and sign criteria	Hessian tests and principal curvature directions
Independent scalar modes	Gain, attenuation, and repeated action of a filter