

Variance, covariance and correlation

Ten problems on variance, covariance and correlation: the shortcut formulas, the variance of a sum, bilinearity and the quadratic form $a^T \Sigma a$, the bounds on a correlation, a dependent pair with zero covariance, the covariance matrix as a positive semidefinite matrix, the covariance of a linear map and whitening, the variance of an average of correlated variables, the sample covariance as $X^T H X$, and the correlation matrix with the bound on a common correlation, with worked solutions and the mistakes that add variances of correlated terms or read zero covariance as independence.

BEFORE YOU START

Variance and covariance are the second moments that nearly every calculation in machine learning runs on: the variance of a weighted sum is what initialisation schemes control, the covariance matrix is what PCA diagonalises and what the multivariate Gaussian is built from, and the variance of an average of correlated predictions is why ensembles stop improving. All of it follows from one algebraic fact, that covariance is bilinear, and from one inequality, that a variance is never negative. These ten problems derive the shortcut formulas, the variance of a sum and of an arbitrary linear combination, the bounds on a correlation, a pair of variables that are dependent but uncorrelated, the properties of a covariance matrix, how it transforms under a linear map, the variance of a mean of correlated terms, the sample covariance in matrix form and the correlation matrix, ending with the smallest correlation that d variables can all share. The five mistakes at the end are the ones that produce a plausible number: variances added across correlated terms, a linear map's covariance with its transposes swapped, zero covariance read as independence, a variance scaled by a factor instead of its square, and a sample covariance built from uncentred data.

- $\mathbb{E}$ is expectation, which is linear: $\mathbb{E}[au + bv] = a \mathbb{E}[u] + b \mathbb{E}[v]$ for constants a, b . A constant c has $\mathbb{E}[c] = c$.
- The variance of a random variable u with mean $m_u = \mathbb{E}[u]$ is $\text{Var}(u) = \mathbb{E}[(u - m_u)^2]$, and its standard deviation is $\sigma_u = \sqrt{\text{Var}(u)}$. The covariance of u and v is $\text{Cov}(u, v) = \mathbb{E}[(u - m_u)(v - m_v)]$, so $\text{Var}(u) = \text{Cov}(u, u)$. When $\sigma_u, \sigma_v > 0$ the correlation is $\rho = \text{Cov}(u, v) / (\sigma_u \sigma_v)$.
- u and v are uncorrelated if $\text{Cov}(u, v) = 0$. They are independent if every event about u is independent of every event about v ; independence gives $\mathbb{E}[g(u)h(v)] = \mathbb{E}[g(u)] \mathbb{E}[h(v)]$ for any functions g, h .
- A random vector $x \in \mathbb{R}^d$ has mean $\mu = \mathbb{E}[x]$ (entry by entry) and covariance matrix $\Sigma = \mathbb{E}[(x - \mu)(x - \mu)^T]$, with $\Sigma_{ij} = \text{Cov}(x_i, x_j)$ and $\Sigma_{ii} = \text{Var}(x_i)$.

- Data are $X \in \mathbb{R}^{N \times d}$ with rows $x_n^\top$ (examples), sample mean $\bar{x} = \frac{1}{N} \sum_n x_n = \frac{1}{N} X^\top \mathbf{1}$ with $\mathbf{1}$ the all-ones vector, and sample covariance $S = \frac{1}{N-1} \sum_n (x_n - \bar{x})(x_n - \bar{x})^\top$, which is what `np.cov` computes; the maximum-likelihood page's estimate divides by N instead.
- A symmetric matrix M is positive semidefinite, written $M \succeq 0$, if $a^\top M a \geq 0$ for every a , equivalently if its eigenvalues are all ≥ 0 (the eigenvalues page).

Problems

1. Show that $\text{Var}(u) = \mathbb{E}[u^2] - (\mathbb{E}[u])^2$, that $\text{Cov}(u, v) = \mathbb{E}[uv] - \mathbb{E}[u] \mathbb{E}[v]$, difficulty $\cdot$ and that $\text{Var}(au + b) = a^2 \text{Var}(u)$ for constants a, b .

2. Show that $\text{Var}(u + v) = \text{Var}(u) + \text{Var}(v) + 2 \text{Cov}(u, v)$ and $\text{Var}(u - v) = \text{Var}(u) + \text{Var}(v) - 2 \text{Cov}(u, v)$. When do the variances simply add? difficulty $\cdot$

3. Show that $\text{Cov}(\sum_i a_i u_i, \sum_j b_j v_j) = \sum_{i,j} a_i b_j \text{Cov}(u_i, v_j)$. Deduce that difficulty $\cdot \cdot$ for a random vector x with covariance Σ and constant vectors a, b , $\text{Cov}(a^\top x, b^\top x) = a^\top \Sigma b$ and $\text{Var}(a^\top x) = a^\top \Sigma a$.

4. Show that $-1 \leq \rho \leq 1$, that $\rho = \pm 1$ exactly when $v = au + b$ for difficulty $\cdot \cdot$ constants with $a \neq 0$ having the sign of ρ , and that $\text{corr}(au + b, cv + d) = \text{sgn}(ac) \rho$ for $ac \neq 0$.

5. Let u be uniform on $\{-1, 0, 1\}$ and $v = u^2$. Compute $\text{Cov}(u, v)$, show difficulty $\cdot \cdot$ that u and v are not independent, and compute $\text{Var}(u + v)$ directly and by Problem 2.

6. For a random vector $x \in \mathbb{R}^d$ with mean μ and covariance Σ , show that difficulty $\cdot \cdot$ $\Sigma = \mathbb{E}[xx^\top] - \mu\mu^\top$, that Σ is symmetric and positive semidefinite, and that Σ is singular exactly when some $a \neq 0$ makes $a^\top x$ almost surely constant.

7. Let $y = Ax + b$ with $A \in \mathbb{R}^{m \times d}$ and $b \in \mathbb{R}^m$ constant. Show that $\mathbb{E}[y] =$ difficulty $\cdot \cdot$ $A\mu + b$ and $\text{Cov}(y) = A\Sigma A^\top$. Then, with $\Sigma = Q\Lambda Q^\top$ and every eigenvalue positive, show that $W = \Lambda^{-1/2}Q^\top$ gives $\text{Cov}(Wx) = I$.

8. Let $x_1, \dots, x_N$ each have variance σ^2 , and $\bar{x} = \frac{1}{N} \sum_n x_n$. Show that difficulty ..
 $\text{Var}(\bar{x}) = \frac{1}{N^2} \mathbf{1}^\top \Sigma \mathbf{1}$ in general, that it is σ^2/N when the x_n are uncorrelated, and that it is $\sigma^2(\rho + \frac{1-\rho}{N})$ when every pair has correlation ρ . What happens as $N \rightarrow \infty$ in the last case?

9. Let $H = I - \frac{1}{N} \mathbf{1}\mathbf{1}^\top$ ($N \times N$). Show that $HX = X - \mathbf{1}\bar{x}^\top$, the centred difficulty ..
data, that H is a symmetric projection ($H^\top = H$ and $H^2 = H$), and that $S = \frac{1}{N-1} X^\top H X$.

10. The correlation matrix is $R = D^{-1/2} \Sigma D^{-1/2}$ with $D =$ difficulty ...
 $\text{diag}(\Sigma_{11}, \dots, \Sigma_{dd})$. Show that $R_{ij} = \text{corr}(x_i, x_j)$, that R has unit diagonal and $R \succeq 0$. Then,
for $R = (1-r)I + r\mathbf{1}\mathbf{1}^\top$, where every pair has the same correlation r , find the eigenvalues
of R and deduce that it is a valid correlation matrix exactly when $-\frac{1}{d-1} \leq r \leq 1$.

Answers

1. $\text{Var}(u) = \mathbb{E}[u^2] - (\mathbb{E}[u])^2$; $\text{Cov}(u, v) = \mathbb{E}[uv] - \mathbb{E}[u] \mathbb{E}[v]$; $\text{Var}(au + b) = a^2 \text{Var}(u)$
2. $\text{Var}(u \pm v) = \text{Var}(u) + \text{Var}(v) \pm 2 \text{Cov}(u, v)$; the variances add exactly when u and v are uncorrelated, which independent variables are
3. $\text{Cov}(\sum_i a_i u_i, \sum_j b_j v_j) = \sum_{i,j} a_i b_j \text{Cov}(u_i, v_j)$; $\text{Cov}(a^\top x, b^\top x) = a^\top \Sigma b$; $\text{Var}(a^\top x) = a^\top \Sigma a$
4. $-1 \leq \rho \leq 1$; $\rho = \pm 1$ exactly when $v = au + b$ with $\text{sgn}(a) = \rho$; $\text{corr}(au + b, cv + d) = \text{sgn}(ac) \rho$
5. $\text{Cov}(u, u^2) = 0$ although $v = u^2$ is a function of u ; $\text{Var}(u) = \frac{2}{3}$, $\text{Var}(v) = \frac{2}{9}$ and $\text{Var}(u + v) = \frac{8}{9} = \text{Var}(u) + \text{Var}(v)$, the cross term being 0
6. $\Sigma = \mathbb{E}[xx^\top] - \mu\mu^\top$; $\Sigma = \Sigma^\top$; $a^\top \Sigma a = \text{Var}(a^\top x) \geq 0$ for all a ; Σ is singular exactly when some $a \neq 0$ has $a^\top x$ almost surely constant
7. $\mathbb{E}[Ax + b] = A\mu + b$; $\text{Cov}(Ax + b) = A\Sigma A^\top$; $W = \Lambda^{-1/2}Q^\top$ whitens: $\text{Cov}(Wx) = I$
8. $\text{Var}(\bar{x}) = \frac{1}{N^2} \mathbf{1}^\top \Sigma \mathbf{1}$; σ^2/N when uncorrelated; $\sigma^2(\rho + \frac{1-\rho}{N})$ under a common correlation ρ , which tends to $\rho\sigma^2$ rather than 0 as $N \rightarrow \infty$
9. $HX = X - \mathbf{1}\bar{x}^\top$; $H^\top = H = H^2$; $S = \frac{1}{N-1} X^\top HX$, and the maximum-likelihood page's $\frac{1}{N}$ version is $\frac{1}{N} X^\top HX$
10. $R_{ij} = \text{corr}(x_i, x_j)$, $R_{ii} = 1$, $R \succeq 0$; $(1-r)I + r\mathbf{1}\mathbf{1}^\top$ has eigenvalues $1 + (d-1)r$ (once, eigenvector $\mathbf{1}$) and $1-r$ ($d-1$ times), so it is a correlation matrix exactly when $-\frac{1}{d-1} \leq r \leq 1$

Worked solutions

Problem 1

Show that $\text{Var}(u) = \mathbb{E}[u^2] - (\mathbb{E}[u])^2$, that $\text{Cov}(u, v) = \mathbb{E}[uv] - \mathbb{E}[u] \mathbb{E}[v]$, and that $\text{Var}(au + b) = a^2 \text{Var}(u)$ for constants a, b .

With $m = \mathbb{E}[u]$: $(u - m)^2 = u^2 - 2mu + m^2$.

Expand the square; m is a number.

$\text{Var}(u) = \mathbb{E}[u^2] - 2m \mathbb{E}[u] + m^2 = \mathbb{E}[u^2] - 2m^2 + m^2 = \mathbb{E}[u^2] - m^2$.

Linearity of $\mathbb{E}$, then $\mathbb{E}[u] = m$.

$(u - m_u)(v - m_v) = uv - m_u v - m_v u + m_u m_v$, with expectation $\mathbb{E}[uv] - m_u m_v - m_v m_u + m_u m_v = \mathbb{E}[uv] - m_u m_v$.

Expand, then linearity with $\mathbb{E}[v] = m_v$ and $\mathbb{E}[u] = m_u$.

$\mathbb{E}[au + b] = am + b$, so $au + b - (am + b) = a(u - m)$ and $\text{Var}(au + b) = \mathbb{E}[a^2(u - m)^2] = a^2 \text{Var}(u)$.

The shift b cancels in the deviation, and the constant a^2 comes out of the expectation.

$\text{Var}(u) = \mathbb{E}[u^2] - (\mathbb{E}[u])^2$; $\text{Cov}(u, v) = \mathbb{E}[uv] - \mathbb{E}[u] \mathbb{E}[v]$; $\text{Var}(au + b) = a^2 \text{Var}(u)$

The variance formula is the covariance formula with $v = u$. A shift leaves spread unchanged, and scaling by a scales the standard deviation by $|a|$, so the variance by a^2 (Mistake 4). Since a variance is an expectation of a square, $\mathbb{E}[u^2] \geq (\mathbb{E}[u])^2$ always, which is the initialisation page's starting point.

Problem 2

Show that $\text{Var}(u + v) = \text{Var}(u) + \text{Var}(v) + 2 \text{Cov}(u, v)$ and $\text{Var}(u - v) = \text{Var}(u) + \text{Var}(v) - 2 \text{Cov}(u, v)$. When do the variances simply add?

$$\mathbb{E}[u + v] = m_u + m_v, \text{ so } (u + v) - (m_u + m_v) = (u - m_u) + (v - m_v).$$

Linearity of $\mathbb{E}$.

$$\left((u - m_u) + (v - m_v)\right)^2 = (u - m_u)^2 + 2(u - m_u)(v - m_v) + (v - m_v)^2.$$

Expand the square.

$$\text{Var}(u + v) = \text{Var}(u) + 2 \text{Cov}(u, v) + \text{Var}(v).$$

Take expectations term by term; each is a definition.

$\text{Cov}(u, -v) = -\text{Cov}(u, v)$ and $\text{Var}(-v) = \text{Var}(v)$, so step 3 with $-v$ in place of v gives the difference.

$-v$ has mean $-m_v$ and deviation $-(v - m_v)$, which flips the sign of the cross term and leaves the square alone; Problem 1 with $a = -1$.

$\text{Var}(u \pm v) = \text{Var}(u) + \text{Var}(v) \pm 2 \text{Cov}(u, v)$; **the variances add exactly when u and v are uncorrelated, which independent variables are**

Independence gives $\mathbb{E}[uv] = \mathbb{E}[u] \mathbb{E}[v]$, so $\text{Cov} = 0$ by Problem 1; the converse fails (Problem 5). A sum of correlated variables can have variance anywhere between $(\sigma_u - \sigma_v)^2$ and $(\sigma_u + \sigma_v)^2$ (Problem 4), and adding variances regardless is Mistake 1. Problem 8 is this formula applied to N terms.

Problem 3

Show that $\text{Cov}(\sum_i a_i u_i, \sum_j b_j v_j) = \sum_{i,j} a_i b_j \text{Cov}(u_i, v_j)$. Deduce that for a random vector x with covariance Σ and constant vectors a, b , $\text{Cov}(a^\top x, b^\top x) = a^\top \Sigma b$ and $\text{Var}(a^\top x) = a^\top \Sigma a$.

Let $U = \sum_i a_i u_i$ and $V = \sum_j b_j v_j$. Then $U - \mathbb{E}[U] = \sum_i a_i (u_i - \mathbb{E}[u_i])$ and likewise $V - \mathbb{E}[V] = \sum_j b_j (v_j - \mathbb{E}[v_j])$.

Linearity of $\mathbb{E}$ gives $\mathbb{E}[U] = \sum_i a_i \mathbb{E}[u_i]$; subtract.

$(U - \mathbb{E}[U])(V - \mathbb{E}[V]) = \sum_i \sum_j a_i b_j (u_i - \mathbb{E}[u_i])(v_j - \mathbb{E}[v_j]).$	Multiply the two sums term by term.
$\text{Cov}(U, V) = \sum_{i,j} a_i b_j \mathbb{E}[(u_i - \mathbb{E}[u_i])(v_j - \mathbb{E}[v_j])] = \sum_{i,j} a_i b_j \text{Cov}(u_i, v_j).$	Linearity of $\mathbb{E}$ over the double sum.
With $u = v = x$: $\text{Cov}(a^\top x, b^\top x) = \sum_{i,j} a_i \Sigma_{ij} b_j = a^\top \Sigma b$, and $b = a$ gives $\text{Var}(a^\top x) = a^\top \Sigma a$.	$\Sigma_{ij} = \text{Cov}(x_i, x_j)$, and $a^\top \Sigma b = \sum_{i,j} a_i \Sigma_{ij} b_j$ is the definition of the matrix product.
$\text{Cov}(\sum_i a_i u_i, \sum_j b_j v_j) = \sum_{i,j} a_i b_j \text{Cov}(u_i, v_j); \text{Cov}(a^\top x, b^\top x) = a^\top \Sigma b; \text{Var}(a^\top x) = a^\top \Sigma a$	Covariance is bilinear, so it behaves like an inner product on random variables, and Σ is the matrix of that inner product on the coordinates of x. The quadratic form $a^\top \Sigma a$ is never negative because it is a variance, which is Problem 6, and it is the quantity PCA maximises over unit vectors a on the PCA page.

Problem 4

Show that $-1 \leq \rho \leq 1$, that $\rho = \pm 1$ exactly when $v = au + b$ for constants with $a \neq 0$ having the sign of ρ , and that $\text{corr}(au + b, cv + d) = \text{sgn}(ac) \rho$ for $ac \neq 0$.

Let $\tilde{u} = (u - m_u)/\sigma_u$ and $\tilde{v} = (v - m_v)/\sigma_v$. Then $\mathbb{E}[\tilde{u}] = \mathbb{E}[\tilde{v}] = 0$, $\text{Var}(\tilde{u}) = \text{Var}(\tilde{v}) = 1$ and $\text{Cov}(\tilde{u}, \tilde{v}) = \text{Cov}(u, v)/(\sigma_u \sigma_v) = \rho$.	Problem 1 with $a = 1/\sigma_u$ and $b = -m_u/\sigma_u$ for the variance; Problem 3 for the covariance.
$0 \leq \text{Var}(\tilde{u} - \tilde{v}) = 1 + 1 - 2\rho$, so $\rho \leq 1$.	A variance is the expectation of a square; Problem 2 for the difference.
$0 \leq \text{Var}(\tilde{u} + \tilde{v}) = 2 + 2\rho$, so $\rho \geq -1$.	Problem 2 for the sum.

If $\rho = 1$ then $\text{Var}(\tilde{u} - \tilde{v}) = 0$, so $\tilde{u} - \tilde{v}$ equals its mean 0 and $v = m_v + (\sigma_v/\sigma_u)(u - m_u)$, a line of positive slope; $\rho = -1$ gives $\tilde{v} = -\tilde{u}$, a line of negative slope. Conversely, if $v = au + b$ then $\text{Cov}(u, v) = a \text{Var}(u)$ and $\sigma_v = |a|\sigma_u$, so $\rho = a/|a| = \text{sgn}(a)$.

A random variable with zero variance is constant; Problems 3 and 1 for the converse.

$\text{Cov}(au + b, cv + d) = ac \text{Cov}(u, v)$ and $\sigma_{au+b} \sigma_{cv+d} = |a||c|\sigma_u \sigma_v$, so the correlation is $\frac{ac}{|ac|} \rho$.

Problem 3 for the covariance, with the shifts dropping out as in Problem 1; Problem 1 for the standard deviations.

$-1 \leq \rho \leq 1$; $\rho = \pm 1$ **exactly when** $v = au + b$ **with** $\text{sgn}(a) = \rho$; $\text{corr}(au + b, cv + d) = \text{sgn}(ac) \rho$

This is the Cauchy–Schwarz inequality for the inner product of Problem 3. Correlation measures only linear association and ignores units: metres to feet or Celsius to Fahrenheit leaves it unchanged, and negating one variable flips its sign. $\rho = 0$ says nothing about nonlinear dependence (Problem 5).

Problem 5

Let u be uniform on $\{-1, 0, 1\}$ and $v = u^2$. Compute $\text{Cov}(u, v)$, show that u and v are not independent, and compute $\text{Var}(u + v)$ directly and by Problem 2.

$\mathbb{E}[u] = 0$ and $\mathbb{E}[u^2] = \frac{1}{3}(1 + 0 + 1) = \frac{2}{3}$, so $\text{Var}(u) = \frac{2}{3}$.

Each value has probability $\frac{1}{3}$; Problem 1.

v is 1 with probability $\frac{2}{3}$ and 0 with probability $\frac{1}{3}$, so $\mathbb{E}[v] = \frac{2}{3}$, $\mathbb{E}[v^2] = \mathbb{E}[v] = \frac{2}{3}$ and $\text{Var}(v) = \frac{2}{3} - \frac{4}{9} = \frac{2}{9}$.

$v^2 = v$ for a 0–1 variable; Problem 1.

$\mathbb{E}[uv] = \mathbb{E}[u^3] = \frac{1}{3}(-1 + 0 + 1) = 0$, so $\text{Cov}(u, v) = 0 - 0 \cdot \frac{2}{3} = 0$.	Problem 1's shortcut.
$P(v = 1 \mid u = 0) = 0$ while $P(v = 1) = \frac{2}{3}$, so u and v are dependent.	Knowing $u = 0$ fixes $v = 0$; independence would need every conditional probability to equal the marginal (the Bayes page).
$u + v$ takes the values 0 (at $u = -1$), 0 (at $u = 0$) and 2 (at $u = 1$), so $\mathbb{E}[u + v] = \frac{2}{3}$, $\mathbb{E}[(u + v)^2] = \frac{4}{3}$ and $\text{Var}(u + v) = \frac{4}{3} -$ $\frac{4}{9} = \frac{8}{9}$.	Enumerate the three cases; Problem 1.
$\text{Cov}(u, u^2) = 0$ although $v = u^2$ is a function of u; $\text{Var}(u) = \frac{2}{3}$, $\text{Var}(v) = \frac{2}{9}$ and $\text{Var}(u + v) = \frac{8}{9} = \text{Var}(u) + \text{Var}(v)$, the cross term being 0	Covariance sees only linear association, and v is an even function of a symmetric u , so the positive and negative sides cancel. Zero correlation is enough for variances to add (Problem 2) but not for independence (Mistake 3): $\mathbb{E}[u^2v] =$ $\mathbb{E}[u^4] = \frac{2}{3}$ while $\mathbb{E}[u^2] \mathbb{E}[v] = \frac{4}{9}$. For jointly Gaussian variables the two notions coincide, which is the case people generalise from.

Problem 6

For a random vector $x \in \mathbb{R}^d$ with mean μ and covariance Σ , show that $\Sigma = \mathbb{E}[xx^\top] - \mu\mu^\top$, that Σ is symmetric and positive semidefinite, and that Σ is singular exactly when some $a \neq 0$ makes $a^\top x$ almost surely constant.

$(x - \mu)(x - \mu)^\top = xx^\top - x\mu^\top - \mu x^\top +$ $\mu\mu^\top$, with expectation $\mathbb{E}[xx^\top] - \mu\mu^\top -$ $\mu\mu^\top + \mu\mu^\top = \mathbb{E}[xx^\top] - \mu\mu^\top$.	Expand the outer product; $\mathbb{E}$ applies entry by entry and $\mathbb{E}[x] = \mu$.
$\Sigma^\top = \mathbb{E}[(x - \mu)(x - \mu)^\top]^\top = \Sigma$.	$(ab^\top)^\top = ba^\top$, and here $a = b$.
For every $a \in \mathbb{R}^d$, $a^\top \Sigma a = \text{Var}(a^\top x) \geq 0$.	Problem 3; a variance is nonnegative.

If $\Sigma a = 0$ for some $a \neq 0$ then $a^\top \Sigma a = 0$, so $\text{Var}(a^\top x) = 0$ and $a^\top x$ equals its mean almost surely. Conversely, if $\text{Var}(a^\top x) = 0$ then $a^\top \Sigma a = 0$; writing $\Sigma = Q\Lambda Q^\top$, that is $\sum_j \lambda_j (q_j^\top a)^2 = 0$ with every $\lambda_j \geq 0$, so $\lambda_j q_j^\top a = 0$ for each j and $\Sigma a = \sum_j \lambda_j q_j (q_j^\top a) = 0$.

A singular matrix has a nonzero null vector; the eigenvalues page's decomposition of a symmetric matrix, with step 3 giving $\lambda_j \geq 0$.

$\Sigma = \mathbb{E}[xx^\top] - \mu\mu^\top$; $\Sigma = \Sigma^\top$; $a^\top \Sigma a = \text{Var}(a^\top x) \geq 0$ **for all** a ; **Σ is singular exactly when some $a \neq 0$ has $a^\top x$ almost surely constant**

A singular covariance says the data lie on a hyperplane. One-hot encoding a category with all K columns does this, since the columns sum to 1 on every example, which is the dummy-variable trap of linear regression. The multivariate-Gaussian page needs Σ invertible to write its density, and the PCA page reads the eigenvalues of Σ as variances along the q_j .

Problem 7

Let $y = Ax + b$ with $A \in \mathbb{R}^{m \times d}$ and $b \in \mathbb{R}^m$ constant. Show that $\mathbb{E}[y] = A\mu + b$ and $\text{Cov}(y) = A\Sigma A^\top$. Then, with $\Sigma = Q\Lambda Q^\top$ and every eigenvalue positive, show that $W = \Lambda^{-1/2}Q^\top$ gives $\text{Cov}(Wx) = I$.

$$\mathbb{E}[y] = A \mathbb{E}[x] + b = A\mu + b.$$

Linearity, entry by entry: $y_i = \sum_j A_{ij}x_j + b_i$.

$$y - \mathbb{E}[y] = A(x - \mu).$$

Subtract step 1 from y ; b cancels.

$$\text{Cov}(y) = \mathbb{E}[A(x - \mu)(x - \mu)^\top A^\top] = A \mathbb{E}[(x - \mu)(x - \mu)^\top] A^\top = A\Sigma A^\top.$$

$(Av)(Av)^\top = Avv^\top A^\top$, and the constant matrices come out of the expectation.

$$\text{Cov}(Wx) = W\Sigma W^\top = \Lambda^{-1/2}Q^\top Q\Lambda Q^\top Q\Lambda^{-1/2} = \Lambda^{-1/2}\Lambda\Lambda^{-1/2} = I.$$

Step 3 with $A = W$; $Q^\top Q = I$ because Q is orthogonal, and diagonal matrices multiply entry by entry.

$$\begin{aligned}\mathbb{E}[Ax + b] &= A\mu + b; \text{Cov}(Ax + b) = \\ &A\Sigma A^\top; W = \Lambda^{-1/2}Q^\top \text{ whitens:} \\ \text{Cov}(Wx) &= I\end{aligned}$$

Shapes: $(m \times d)(d \times d)(d \times m) = m \times m$, symmetric and positive semidefinite as a covariance must be; the version with the transposes swapped is Mistake 2. The multivariate-Gaussian page proves the same for Gaussian x and adds that y is Gaussian; nothing here needed a distribution. Any W with $W\Sigma W^\top = I$ whitens, for instance L^{-1} from the Cholesky factor $\Sigma = LL^\top$; the PCA page's scores are $Q^\top(x - \mu)$, whitened by $\Lambda^{-1/2}$.

Problem 8

Let $x_1, \dots, x_N$ each have variance σ^2 , and $\bar{x} = \frac{1}{N} \sum_n x_n$. Show that $\text{Var}(\bar{x}) = \frac{1}{N^2} \mathbf{1}^\top \Sigma \mathbf{1}$ in general, that it is σ^2/N when the x_n are uncorrelated, and that it is $\sigma^2(\rho + \frac{1-\rho}{N})$ when every pair has correlation ρ . What happens as $N \rightarrow \infty$ in the last case?

$$\begin{aligned}\bar{x} &= a^\top x \text{ with } a = \frac{1}{N} \mathbf{1}, \text{ so } \text{Var}(\bar{x}) = \\ a^\top \Sigma a &= \frac{1}{N^2} \mathbf{1}^\top \Sigma \mathbf{1} = \frac{1}{N^2} \sum_{n,m} \Sigma_{nm}.\end{aligned}$$

Problem 3; $\mathbf{1}^\top \Sigma \mathbf{1}$ adds up every entry of Σ .

Uncorrelated: $\Sigma = \sigma^2 I$, so $\mathbf{1}^\top \Sigma \mathbf{1} = N\sigma^2$ and $\text{Var}(\bar{x}) = \sigma^2/N$.

Only the N diagonal entries are nonzero.

Common correlation: $\Sigma_{nn} = \sigma^2$ and $\Sigma_{nm} = \rho\sigma^2$ for $n \neq m$, so $\mathbf{1}^\top \Sigma \mathbf{1} = N\sigma^2 + N(N-1)\rho\sigma^2$.

N diagonal entries and $N(N-1)$ off-diagonal ones.

$$\begin{aligned}\text{Var}(\bar{x}) &= \sigma^2 \frac{N + N(N-1)\rho}{N^2} = \\ \sigma^2 \frac{1 + (N-1)\rho}{N} &= \sigma^2 \left(\rho + \frac{1-\rho}{N} \right).\end{aligned}$$

Divide by N^2 , then write $1 + (N-1)\rho = N\rho + (1-\rho)$.

$\text{Var}(\bar{x}) = \frac{1}{N^2} \mathbf{1}^\top \Sigma \mathbf{1}; \sigma^2/N$ **when uncorrelated;** $\sigma^2(\rho + \frac{1-\rho}{N})$ **under a common correlation ρ , which tends to $\rho\sigma^2$ rather than 0 as $N \rightarrow \infty$**

Averaging removes only the uncorrelated part of the variance. This is the ensemble calculation: averaging N models whose errors have correlation ρ can reduce the variance by at most a factor ρ however many models are used, which is why random forests decorrelate their trees by subsampling features, and why Mistake 1 at scale, σ^2/N for correlated models, overstates the gain. Problem 10 says which values of ρ are possible.

Problem 9

Let $H = I - \frac{1}{N} \mathbf{1} \mathbf{1}^\top$ ($N \times N$). Show that $HX = X - \mathbf{1} \bar{x}^\top$, the centred data, that H is a symmetric projection ($H^\top = H$ and $H^2 = H$), and that $S = \frac{1}{N-1} X^\top HX$.

$\mathbf{1}^\top X = N \bar{x}^\top$, so $\frac{1}{N} \mathbf{1} \mathbf{1}^\top X = \mathbf{1} \bar{x}^\top$ and $HX = X - \mathbf{1} \bar{x}^\top$.

$\mathbf{1}^\top X$ is the row of column sums, which is N times the row of column means; $\mathbf{1} \bar{x}^\top$ repeats $\bar{x}^\top$ in every row, so row n of HX is $(x_n - \bar{x})^\top$.

$H^\top = I - \frac{1}{N} (\mathbf{1} \mathbf{1}^\top)^\top = H$.

$(\mathbf{1} \mathbf{1}^\top)^\top = \mathbf{1} \mathbf{1}^\top$.

$H^2 = I - \frac{2}{N} \mathbf{1} \mathbf{1}^\top + \frac{1}{N^2} \mathbf{1} (\mathbf{1}^\top \mathbf{1}) \mathbf{1}^\top = I - \frac{2}{N} \mathbf{1} \mathbf{1}^\top + \frac{1}{N} \mathbf{1} \mathbf{1}^\top = H$.

Expand; $\mathbf{1}^\top \mathbf{1} = N$.

$\sum_n (x_n - \bar{x})(x_n - \bar{x})^\top = (HX)^\top (HX) = X^\top H^\top HX = X^\top HX$.

$M^\top M$ is the sum of the outer products of the rows of M , and row n of HX is $(x_n - \bar{x})^\top$; then steps 2 and 3.

$HX = X - \mathbf{1} \bar{x}^\top$; $H^\top = H = H^2$; $S = \frac{1}{N-1} X^\top HX$, **and the maximum-likelihood page's $\frac{1}{N}$ version is $\frac{1}{N} X^\top HX$**

H projects onto the vectors whose entries sum to zero (the least-squares page's projections), and centering is that projection applied to each column. The two normalisations differ by a scalar, so they share eigenvectors; `np.cov(X.T)` divides by $N - 1$, `bias=True` by N . Skipping the H is Mistake 5.

Problem 10

The correlation matrix is $R = D^{-1/2}\Sigma D^{-1/2}$ with $D = \text{diag}(\Sigma_{11}, \dots, \Sigma_{dd})$. Show that $R_{ij} = \text{corr}(x_i, x_j)$, that R has unit diagonal and $R \succeq 0$. Then, for $R = (1 - r)I + r\mathbf{1}\mathbf{1}^\top$, where every pair has the same correlation r , find the eigenvalues of R and deduce that it is a valid correlation matrix exactly when $-\frac{1}{d-1} \leq r \leq 1$.

$$(D^{-1/2}\Sigma D^{-1/2})_{ij} = \frac{\Sigma_{ij}}{\sqrt{\Sigma_{ii}}\sqrt{\Sigma_{jj}}} = \frac{\text{Cov}(x_i, x_j)}{\sigma_i\sigma_j}, \text{ which is 1 when } i = j.$$

A diagonal matrix on the left scales row i by its entry and on the right scales column j .

$$R = \text{Cov}(D^{-1/2}x) \succeq 0.$$

Problem 7 with $A = D^{-1/2}$, which is symmetric: R is the covariance of the standardised variables x_i/σ_i , and every covariance is positive semidefinite by Problem 6.

$$R\mathbf{1} = (1 - r)\mathbf{1} + r\mathbf{1}(\mathbf{1}^\top\mathbf{1}) = (1 + (d - 1)r)\mathbf{1}.$$

$\mathbf{1}^\top\mathbf{1} = d$; so $\mathbf{1}$ is an eigenvector with eigenvalue $1 + (d - 1)r$.

$$\text{For any } a \text{ with } \mathbf{1}^\top a = 0: Ra = (1 - r)a + r\mathbf{1}(\mathbf{1}^\top a) = (1 - r)a.$$

The zero-sum vectors form a $(d - 1)$ -dimensional space, so $1 - r$ is an eigenvalue with multiplicity $d - 1$, and with step 3 all d eigenvalues are found.

$$R \succeq 0 \text{ exactly when } 1 - r \geq 0 \text{ and } 1 + (d - 1)r \geq 0, \text{ that is when } -\frac{1}{d-1} \leq r \leq 1.$$

A symmetric matrix is positive semidefinite exactly when its eigenvalues are nonnegative (the eigenvalues page).

$$R_{ij} = \text{corr}(x_i, x_j), R_{ii} = 1, R \succeq 0; (1 - r)I + r\mathbf{1}\mathbf{1}^\top \text{ has eigenvalues } 1 + (d - 1)r \text{ (once, eigenvector } \mathbf{1}) \text{ and } 1 - r \text{ (} d - 1 \text{ times), so it is a correlation matrix exactly when } -\frac{1}{d-1} \leq r \leq 1$$

Three variables cannot all be pairwise correlated at -0.9 : the lower bound is $-\frac{1}{2}$ for $d = 3$ and tends to 0 as d grows, so a large set of variables can be only slightly negatively correlated on average. At $r = -\frac{1}{d-1}$ the sum $\mathbf{1}^\top x$ has variance $\mathbf{1}^\top R\mathbf{1} = d(1 + (d - 1)r) = 0$, Problem 6's singular case. Problem 8's variance of the mean is nonnegative on exactly this range.

Where this goes wrong

1. Adding the variances of correlated terms

Variances add for independent variables, which is the case every first example uses.

$$\text{Var}(u + v) = \text{Var}(u) + \text{Var}(v) + 2 \text{Cov}(u, v)$$

Right so far: Problem 2.

“Variances add.”

The habit that causes the mistake: additivity learned on independent variables and carried over to a pair that is correlated.

$$\text{Var}(u + v) = \text{Var}(u) + \text{Var}(v) \text{ ✗}$$

With $\rho = \frac{1}{2}$ and equal variances σ^2 the true variance is $3\sigma^2$, not $2\sigma^2$, and with $v = c - u$ the sum is constant while the formula says $2\sigma^2$. At scale it is Problem 8 with the correlations dropped: an average of N correlated models is credited with variance σ^2/N , which it cannot reach, because $\rho\sigma^2$ remains however large N is.

2. Covariance of a linear map written with the transposes swapped

Quadratic forms are written $a^\top \Sigma a$, and the matrix version looks like the same thing with A in place of a .

$$y = Ax + b \text{ with } \text{Cov}(y) = \mathbb{E}[(y - \mathbb{E}[y])(y - \mathbb{E}[y])^\top]$$

Right so far: the definition, as in Problem 7.

$$\text{“Var}(a^\top x) = a^\top \Sigma a, \text{ so Cov}(Ax) = A^\top \Sigma A. \text{”}$$

The analogy that causes the mistake: the vector formula has its transpose on the left because the map is $a^\top x$, a row times x ; for Ax the map's transpose lands on the right.

$$\text{Cov}(Ax + b) = A^\top \Sigma A \quad \text{✗}$$

For $A \in \mathbb{R}^{m \times d}$ with $m \neq d$ the product $A^\top \Sigma A$ is not even defined, since $A^\top$ is $d \times m$ and Σ is $d \times d$. When $m = d$ it is a different matrix: for $A = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$ and $\Sigma = I$, $A \Sigma A^\top = \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}$ while $A^\top \Sigma A = \begin{pmatrix} 1 & 1 \\ 1 & 2 \end{pmatrix}$. The vector case is the $m = 1$ instance of Problem 7 with $A = a^\top$, and $(a^\top) \Sigma (a^\top)^\top = a^\top \Sigma a$ agrees.

3. Zero covariance read as independence

Independence implies zero covariance, and the implication is usually learned in that direction only.

$$\text{Cov}(u, v) = 0$$

Right so far: a computed fact, as for u and u^2 in Problem 5.

“Zero covariance means no relation, so u and v are independent.”

The converse that causes the mistake: true for jointly Gaussian variables, and generalised from there.

$$P(v = 1 \mid u = 0) = P(v = 1) \text{ for } v = u^2 \quad \text{✗}$$

Problem 5: the left side is 0 and the right is $\frac{2}{3}$. Covariance is one number measuring linear association, and a symmetric nonlinear dependence leaves it at 0. The consequences are not only probabilistic: $\mathbb{E}[uv]$ factorises but $\mathbb{E}[u^2v] = \frac{2}{3} \neq \mathbb{E}[u^2] \mathbb{E}[v] = \frac{4}{9}$, so any calculation that uses independence beyond the second moment is wrong.

4. Scaling a variance by the factor instead of its square

Expectations scale linearly, and a variance is an expectation.

$$\mathbb{E}[au] = a \mathbb{E}[u]$$

Right so far: linearity.

“Variance is an expectation, so it scales the same way.”

The analogy that causes the mistake: variance is the expectation of a square, and the factor sits inside the square.

$$\text{Var}(au) = a \text{Var}(u) \quad \text{✗}$$

Problem 1: $\text{Var}(au) = a^2 \text{Var}(u)$.
 Converting metres to centimetres multiplies a variance by 10^4 , not 100, and halving a layer's outputs quarters their variance, which is the square in the initialisation page's $\text{Var}(wx) = \sigma_w^2 \mathbb{E}[x^2]$.
 For $a < 0$ the wrong formula makes a variance negative. It is correct for the standard deviation, which scales by $|a|$, and that is where it comes from.

5. Sample covariance from uncentred data

$X^\top X$ is the Gram matrix of least squares and of every normal equation, and the covariance looks like it with a $\frac{1}{N-1}$ in front.

$$S = \frac{1}{N-1} \sum_n (x_n - \bar{x})(x_n - \bar{x})^\top$$

Right so far: the definition, as in Problem 9.

“The sum of the outer products of the rows is $X^\top X$.”

The shortcut that causes the mistake: that is right for the rows of X , but the rows summed here are those of the centred HX .

$$S = \frac{1}{N-1} X^\top X \quad \text{✗}$$

Expanding $x_n = (x_n - \bar{x}) + \bar{x}$ gives $\sum_n x_n x_n^\top = \sum_n (x_n - \bar{x})(x_n - \bar{x})^\top + N\bar{x}\bar{x}^\top$, so the line computes $S + \frac{N}{N-1}\bar{x}\bar{x}^\top$: the covariance plus a rank-one term that dominates whenever the means are large next to the spread. Its leading eigenvector then points towards $\bar{x}$ rather than along the direction of greatest spread, which is the PCA page's first mistake. The fix is $X^\top HX$, or subtracting the column means before forming the product.