1. Let $X \sim N_n(\mu, \Sigma)$, and let A be an arbitrary $m \times n$ matrix. Prove directly from the definition that AX has an m-variate normal distribution $AX \sim N_m(A\mu, A\Sigma A^\top)$.

2. Let X and Y be $N(0,1)$ random variables such that $\text{cov}(X, Y) = 0$. Show by example that X and Y need not be independent. (Hint: Take $Y = XZ$ where X and Z are independent and the distribution of Z is chosen appropriately.)

3. Let $X \sim N_n(\mu, \sigma^2 I)$ where I is the $n \times n$ identity matrix, and let P be an $n \times n$ orthogonal projection matrix; i.e. $P^2 = P = P^\top$. Show that the random vectors PX and $(I - P)X$ are independent.

4. Let $X \sim N_n(\mu, \Sigma)$, and let X_1 denote the first n_1 components of X. Let μ_1 denote the first n_1 components of μ, and let Σ_{11} denote the upper left $n_1 \times n_1$ block of Σ. Show that $X_1 \sim N_{n_1}(\mu_1, \Sigma_{11})$.

5. Consider the simple linear regression model

$$Y_i = a + bx_i + \varepsilon_i, \quad i = 1, \ldots, n$$

where $\varepsilon_1, \ldots, \varepsilon_n \overset{\text{iid}}{\sim} N(0, \sigma^2)$ and $\sum_{i=1}^n x_i = 0$. Derive from first principles explicit expressions for the MLEs $\hat{a}$, $\hat{b}$ and $\hat{\sigma}^2$. Show that we can obtain the same expressions if we regard the simple linear regression model as a special case of the general linear model $Y = X\beta + \varepsilon$ and specialise the formulae $\hat{\beta} = (X^\top X)^{-1} X^\top Y$ and $\hat{\sigma}^2 = n^{-1} \|Y - X\hat{\beta}\|^2$.

6. The relationship between the range in metres, Y, of a howitzer with muzzle velocity v metres per second fired at angle of elevation α degrees is assumed to be

$$Y = \frac{v^2}{g} \sin(2\alpha) + \varepsilon,$$

where $g = 9.81$ and where $\varepsilon \sim N(0, \sigma^2)$. Estimate v from the following independent observations made on 9 shells.

α (deg)	5	10	15	20	25	30	35	40	45
$\sin 2\alpha$	0.1736	0.3420	0.5	0.6428	0.7660	0.8660	0.9397	0.9848	1
range (m)	4860	9580	14080	18100	21550	24350	26400	27700	28300

7. Consider the one-way analysis of variance (ANOVA) model

$$Y_{ij} = \mu_i + \varepsilon_{ij}, \quad i = 1, \ldots, I, j = 1, \ldots, n_i,$$

where $(\varepsilon_{ij}) \overset{\text{iid}}{\sim} N(0, \sigma^2)$. Derive from first principles explicit expressions for the MLEs $\hat{\mu}_1, \ldots, \hat{\mu}_I$ and $\hat{\sigma}^2$. Show that we can obtain the same expressions if we regard the ANOVA model as a special case of the general linear model $Y = X\beta + \varepsilon$ and specialise the formulae $\hat{\beta} = (X^\top X)^{-1} X^\top Y$ and $\hat{\sigma}^2 = n^{-1} \|Y - X\hat{\beta}\|^2$.

8. Consider the linear model $Y = X\beta + \varepsilon$, where $\mathbb{E}(\varepsilon) = 0$ and $\text{Cov}(\varepsilon) = \sigma^2 \Sigma$, for some unknown parameter $\sigma^2 > 0$ and known positive definite matrix Σ. Derive the form of the Generalised Least Squares estimator $\tilde{\beta}^{GLS}$, defined by

$$\tilde{\beta}^{GLS} = \text{argmin}_\beta (Y - X\beta)^\top \Sigma^{-1} (Y - X\beta).$$

State and prove a version of the Gauss–Markov theorem for $\tilde{\beta}^{GLS}$.

9. A medical scientist believes that a certain measurement of patient health (e.g. the blood cholesterol level) is influenced by whether a patient has a certain characteristic (e.g. green eyes). The scientist recruits n patients with various combinations of a list of p characteristics.

The data $(Y_i)_i$ is collected, where Y_i is the health measurement of the ith patient. The data is modelled as $Y = X\beta + \varepsilon$ where $\varepsilon \sim N_n(0, \sigma^2 I)$ for unknown parameters $\beta \in \mathbb{R}^p$ and σ^2, and where X is the $n \times p$ matrix such that $x_{ij} = 1$ if patient i has characteristic j and $x_{ij} = 0$ otherwise. Assume that the columns of X are linearly independent.

(a) For a certain characteristic j, interpret the hypotheses $H_0 : \beta_j = 0$ vs $H_1 : \beta_j \neq 0$. Find the likelihood ratio test of size α.

(b) The scientist repeats the test in part (a) for each characteristic $j = 1, \ldots, p$, and finds that there is one attribute j^* for which the null hypothesis is rejected. To what extent is there evidence that characteristic j^* affects the health measurement?

10. Consider the one-way ANOVA model of **7**. Letting $\bar{Y}_{i\cdot} = n_i^{-1} \sum_{j=1}^{n_i} Y_{ij}$ and $\bar{Y} = n^{-1} \sum_{i=1}^{I} \sum_{j=1}^{n_i} Y_{ij}$ with $n = n_1 + \ldots + n_I$, show from first principles that the size α likelihood ratio test of equality of means rejects H_0 if

$$\frac{\frac{1}{I-1} \sum_{i=1}^{I} n_i(\bar{Y}_{i\cdot} - \bar{Y})^2}{\frac{1}{n-I} \sum_{i=1}^{I} \sum_{j=1}^{n_i} (Y_{ij} - \bar{Y}_{i\cdot})^2} > k$$

where $F(k) = 1 - \alpha$ and F is the CDF of the $F_{I-1,n-I}$ distribution; i.e. if 'the ratio of the between groups sum of squares to the within groups sum of squares is large'.

11. In the standard linear model $Y = X\beta + \varepsilon$ with $\varepsilon \sim N_n(0, \sigma^2 I)$ and MLE $\hat{\beta}$, determine the distribution of the quadratic form $(\hat{\beta} - \beta)^\top X^\top X(\hat{\beta} - \beta)$. (Hint: Consider $(X^\top X)^{1/2}(\hat{\beta} - \beta)$, where the square root is a matrix square root.) Use the F-distribution to find a $(1 - \alpha)$-level confidence set for β, i.e. a random set $C(Y)$ such that $\mathbb{P}_{\beta,\sigma^2}(b \in C(Y)) = 1 - \alpha$ for all β, σ^2. What is the shape of this confidence set?

$^+$**12.** Download **R** from `http://cran.r-project.org/`. Use it to compute a 95% confidence set for the vector of mean chick weights for the different food supplements in the `chickwts` data set (one of the in-built data sets in **R**). (Hint: Type `?model.matrix` *to find out how to obtain the design matrix*.) Now use **R** to compute 95% confidence intervals for each of the individual mean chick weights. Which intervals exclude the estimate of the overall mean chick weight in the null model which assumes that the mean chick weight does not depend on the food supplement?