Exercise: Laplace approximation to $p(\mu, \log\sigma|\mathcal{D})$ for a univariate Gaussian.

In this exercise, we will compute a Laplace approximation of $p(\mu, \log\sigma|\mathcal{D})$ for a Gaussian, using an uninformative prior $p(\mu, \log\sigma) \propto 1$.

First, let $\ell = \log\sigma$ and $s^2 = \frac{1}{n}\sum_{i=1}^{n}(x_i - \overline{x})^2$. The log posterior is

$$\log p(\mu, \ell|\mathcal{D}) = -n\log\sigma - \frac{1}{2\sigma^2}[ns^2 + n(\overline{x} - \mu)^2] + \text{const} \tag{1}$$

where we ignore the dependence of $\log p(\mathcal{D})$ on the parameters.

1. Show that the first derivatives are

$$\frac{\partial}{\partial\mu}\log p(\mu, \ell|\mathcal{D}) = \frac{n(\overline{x} - \mu)}{\sigma^2} \tag{2}$$

$$\frac{\partial}{\partial\ell}\log p(\mu, \ell|\mathcal{D}) = -n + \frac{ns^2 + n(\overline{x} - \mu)^2}{\sigma^2} \tag{3}$$

2. Show that the Hessian matrix is given by

$$\mathbf{H} = \begin{pmatrix} \frac{\partial^2}{\partial\mu^2}\log p(\mu, \ell|\mathcal{D}) & \frac{\partial^2}{\partial\mu\partial\ell}\log p(\mu, \ell|\mathcal{D}) \\ \frac{\partial^2}{\partial\ell^2}\log p(\mu, \ell|\mathcal{D}) & \frac{\partial^2}{\partial\ell^2}\log p(\mu, \ell|\mathcal{D}) \end{pmatrix} \tag{4}$$

$$= \begin{pmatrix} -\frac{n}{\sigma^2} & -2n\frac{\overline{x}-\mu}{\sigma^2} \\ -2n\frac{\overline{x}-\mu}{\sigma^2} & -\frac{2}{\sigma^2}(ns^2 + n(\overline{x} - \mu)^2) \end{pmatrix} \tag{5}$$

3. Use this to derive a Laplace approximation to the posterior $p(\mu, \ell|\mathcal{D})$.