In-class exercise

Instructions

- **Don't look at the solution yet!** This is for your benefit.
- This exercise must be submitted within 48 hours of the lecture in which it was given.
- As long as you do the exercise on time, you get full credit—your performance does not matter.
- Without looking at the solution, take 5 minutes to try to solve the exercise.
- Pre-assessment: Write down how correct you think your answer is, from 0 to 100%.
- Post-assessment: Now, study the solution and give yourself a "grade" from 0 to 100%.
- Submit your work on the course website, including the pre- and post- assessments.

Exercise

Suppose $X_1, \ldots, X_n \overset{\text{iid}}{\sim} \mathcal{N}(\theta, \sigma^2)$, and θ is given a $\mathcal{N}(\mu_0, \sigma_0^2)$ prior.

(a) When $n = 1$, what is the marginal likelihood, $p(x_1)$?

(b) When $n > 1$, is it true that the marginal likelihood factors as $p(x_{1:n}) = p(x_1) \cdots p(x_n)$?

Solution

Part (a)

We could equivalently define the model in the following way: $\boldsymbol{\theta} \sim \mathcal{N}(\mu_0, \sigma_0^2)$ and $X_i = \boldsymbol{\theta} + Z_i$, where $Z_i \sim \mathcal{N}(0, \sigma^2)$ independently (that is, independently of each other and of $\boldsymbol{\theta}$) for $i = 1, \ldots, n$. Using the rule for linear combinations of independent normals, the marginal distribution of X_i is $\mathcal{N}(\mu_0, \sigma_0^2 + \sigma^2)$.

Part (b)

No. If it were true that $p(x_{1:n}) = p(x_1) \cdots p(x_n)$, this would imply that the X_i's are independent, but we know this is not the case, since, for instance, the posterior predictive $p(x_2|x_1)$ depends on x_1. (Note: The X_i's are *conditionally* independent given θ, but here we are talking about their marginal distribution.)