Moment Matching

Bayesian Models for Ecologists

Tom Hobbs

June 4, 2024

DEB 2042028

Work flow: probability distributions

- General properties and definitions
  - discrete random variables
  - continuous random variables
- Specific distributions (cheat sheet and Probability Lab 2)
- Marginal distributions (Probability Lab 3)
- Moment matching (Probability Lab 4)

Why moment matching?

Remember yesterday's Probability lab 2 where you were told to use a normal distribution to model the data and then asked why that was a bad choice? A better choice would be a distribution with non-negative support. How do you model the data using these distributions? Same problem for modeling proportions (zero *to* one support).

Motivation: flexibility in analysis

Deterministic models	Types of data
general linear	real numbers
nonlinear	non-negative real numbers
differential equations	counts
difference equations	0 to 1
auto-regressive	0 or 1
occupancy	counts in categories
state-transition	proportions in categories
integral-projection	ordinal categories

univariate and
multivariate

Motivation: flexibility in analysis

Probability model	Support for random variable
normal	all real numbers
lognormal	non-negative real numbers
gamma	non-negative real numbers
beta	0 to 1 real numbers
Bernoulli	0 or 1
binomial	counts in 2 categories
Poisson	counts
multinomial	counts in >2 categories
negative binomial	counts
Dirichlet	proportions in ≥ 2 categories
Cauchy	real numbers

$\boxed{\text{Board work on normal data model}}$

The problem

All distributions have parameters:

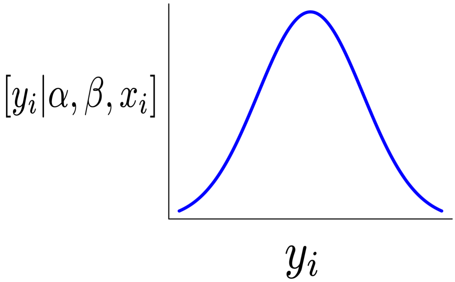

α and β are parameters of the distribution of the random variable y_i.

Types of parameters

Parameter name	Function
intensity, centrality, location	sets position on x axis
shape	controls dispersion and skew
scale, dispersion parameter	shrinks or expands width
rate	scale^{-1}

The problem

The normal and the Poisson are commonly used distributions for which the parameters of the distribution are the *same* as the moments. For all other distributions we will use, the parameters (e.g., α, β) are *functions* of the moments.

$$\begin{aligned} \alpha &= f_1(\mu, \sigma^2) \\ \beta &= f_2(\mu, \sigma^2) \end{aligned}$$

We can use these functions to "match"' the moments to the parameters,allowing use to "insert the mean" predicted by our model into the distribution.

Moment matching the gamma distribution

The gamma distribution: $[z|\alpha,\beta] = \frac{\beta^{\alpha} z^{\alpha-1} e^{-\beta z}}{\Gamma(\alpha)}$
The mean of the gamma distribution is

$$\mu = \frac{\alpha}{\beta}$$

and the variance is

$$\sigma^2 = \frac{\alpha}{\beta^2}.$$

Discover functions for α and β in terms of μ and σ^2.

Answer

$1)\mu = \frac{\alpha}{\beta}$
$2)\sigma^2 = \frac{\alpha}{\beta^2}$
Solve 1 for β, substitue for β in 2), solve for α :
$3)\ \alpha = \frac{\mu^2}{\sigma^2}$
Substitute rhs 3) for α in 2), solve for β :
$4)\ \beta = \frac{\mu}{\sigma^2}$

Example

Lets reconsider the mean above ground biomass problem from yesterday. How would model the data using a gamma distribution?

$$y_i \sim \text{gamma}\left(\frac{\mu^2}{\sigma^2}, \frac{\mu}{\sigma^2}\right)$$

Or, if you had a model that predicts the mean as a function of the covariate x_i and the parameters θ, $\mu_i = g(\theta, x_i)$

$$y_i \sim \text{gamma}\left(\frac{\mu_i^2}{\sigma^2}, \frac{\mu_i}{\sigma^2}\right)$$

Moment matching two parameters in the beta distribution

The beta distribution gives the probability density of random variables with support on $0, ..., 1$.

$$[z|\alpha,\beta] = \frac{z^{\alpha-1}(1-z)^{\beta-1}}{B(\alpha,\beta)}$$

$$B = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)}$$

$$\mu = \frac{\alpha}{\alpha+\beta}$$

$$\sigma^2 = \frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}$$

$$\alpha = \frac{\mu^2 - \mu^3 - \mu\sigma^2}{\sigma^2}$$

$$\beta = \frac{\mu - 2\mu^2 + \mu^3 - \sigma^2 + \mu\sigma^2}{\sigma^2}$$

You need some functions...

```
#BetaMomentMatch.R
# Function for parameters from moments
shape_from_stats <- function(mu, sigma){
 a <-(mu^2-mu^3-mu*sigma^2)/sigma^2
 b <- (mu-2*mu^2+mu^3-sigma^2+mu*sigma^2)/sigma^2
shape_ps <- c(a,b)
return(shape_ps)
}
# Functions for moments from parameters
beta.mean=function(a,b)a/(a+b)
beta.var = function(a,b)a*b/((a+b)^2*(a+b+1))
```

Moment matching for a single parameter illustrated with the beta distribution

We can solve for α in terms of μ and β,

$$\mu = \frac{\alpha}{\alpha+\beta} \tag{1}$$

$$\alpha = \frac{\mu\beta}{1-\mu}, \tag{2}$$

which allows us to use

$$\mu_i = g(\theta, x_i) \tag{3}$$

$$y_i \sim \text{beta}\left(\frac{\mu_i\beta}{1-\mu_i}, \beta\right) \tag{4}$$

to moment match the mean alone.

Moment matching for a single parameter illustrated with the lognormal distribution

The first parameter of the lognormal $= \alpha$, the mean of the random variable on the log scale. The second parameter $= \sigma^2_{\log}$, the variance of the random variable on the log scale
We often moment match the *median* the lognormal distribution:

$$\begin{aligned} \text{median} = \mu_i &= g(\theta, x_i) & (5) \\ \mu &= e^{\alpha} & (6) \\ \alpha &= \log(\mu_i) & (7) \\ y_i &\sim \text{lognormal}(\log(\mu_i), \sigma^2_{\text{log}}) & (8) \end{aligned}$$

In this case, σ^2 remains on log scale.

Data transformation: An alternative to moment matching

proportions (zero to one)

$$\mu_i = \frac{\exp(\beta_0 + \beta_1 x_i)}{1 + \exp(\beta_0 + \beta_1 x_i)}$$
$$\text{logit}(y_i) \sim \text{normal}(\text{logit}(\mu_i), \sigma^2)$$
$$\text{logit}(y_i) = \log\left(\frac{y_i}{1 - y_i}\right)$$

strictly non-negative (zero to infinity)

$$\mu_i = \exp(\beta_0 + \beta_1 x_i)$$
$$\log(y_i) \sim \text{normal}(\log(\mu_i), \sigma^2)$$

Transforming as an alternative

- Can speed MCMC convergence particularly if you use a glm module.
- May not fit as well as moment matched approach (or might fit better).
- Error prone interpretation.

Problems continued

Do Probability Lab 4