STA 6351, Report.1.11

Carson Slater *Baylor University*

1.11 Coupled Poisson–normal model—computational example

- Suppose we have $n = 8$ observations:

i	1	2	3	4	5	6	7	8
x_i	0.0	3.9	1.5	0.0	4.2	2.1	0.7	0.0

  Note the structural zeros at $i \in \{1, 4, 8\}$. Positive counts were at $i \in \{2, 3, 5, 6, 7\}$.

- Use the starting values:

$$\lambda^{(0)} = 1.2000, \quad \mu^{(0)} = 1.8000, \quad \sigma^{2(0)} = 0.9000.$$

  Truncate the series at $n_{\max} = \lfloor \lambda^{(0)} + 8\sqrt{\lambda^{(0)}} \rfloor = 10$. The Poisson tail $\Pr(N > 10 \mid \lambda^{(0)}) \approx 6.22 \times 10^{-8}$ is negligible.

- For each nonzero x_i, define

$$w_{in} = \Pr(N_i = n \mid X_i = x_i; \lambda^{(0)}, \mu^{(0)}, \sigma^{2(0)}), \quad n = 1, \dots, 10.$$

  The largest weights (top 4) and the remaining mass are:

i	Top responsibilities $(n : w_{in})$	remainder
2	$(2 : 0.6899, 3 : 0.1523, 1 : 0.1439, 4 : 0.0132)$;	rest ≈ 0.0007
3	$(1 : 0.8771, 2 : 0.1149, 3 : 0.0076, 4 : 0.0004)$;	rest ≈ 0.0000
5	$(2 : 0.7872, 3 : 0.1910, 4 : 0.0157, 1 : 0.0052)$;	rest ≈ 0.0012
6	$(1 : 0.7942, 2 : 0.1894, 3 : 0.0154, 4 : 0.0009)$;	rest ≈ 0.0000
7	$(1 : 0.9216, 2 : 0.0741, 3 : 0.0042, 4 : 0.0002)$;	rest ≈ 0.0000

  Conditional means of N_i for the nonzeros (useful checks):

$$\Big[\mathbb{E}[N_i \mid x_i]\Big] = \begin{cases} 2.0371 & (i = 2), \\ 1.1313 & (i = 3), \\ 2.1622 & (i = 5), \\ 1.2230 & (i = 6), \\ 1.0830 & (i = 7). \end{cases}$$

  For zeros, $\mathbb{E}[N_i \mid x_i = 0] = 0$ exactly.

- The starting scores are as follows. Let $\varphi \equiv \sigma^2$.

$$U_\lambda = \sum_{i=1}^{n} \left(-1 + \frac{\mathbb{E}[N_i \mid x_i]}{\lambda} \right) = -1.6361,$$

$$U_\mu = \sum_{i:x_i \neq 0} \sum_{n=1}^{10} w_{in} \frac{x_i - n\mu}{\varphi} = -1.4956, \quad U_\varphi = \sum_{i:x_i \neq 0} \sum_{n=1}^{10} w_{in} \left(-\frac{1}{2\varphi} + \frac{(x_i - n\mu)^2}{2n\varphi^2} \right) = -0.4921.$$

- Fisher (complete-data expected) information at the start:

$$\mathcal{I}_{\lambda\lambda} = \sum_{i=1}^{n} \frac{\mathbb{E}[N_i \mid x_i]}{\lambda^2} = 5.3033, \quad \mathcal{I}_{\mu\mu} = \sum_{i:x_i\neq 0} \sum_{n} w_{in} \frac{n}{\varphi} = 8.4852,$$

$$\mathcal{I}_{\varphi\varphi} = \sum_{i:x_i\neq 0} \sum_{n} w_{in} \left(\frac{1}{2\varphi^2} + \frac{(x_i - n\mu)^2}{n\varphi^3} \right) = 8.1656, \quad \mathcal{I}_{\mu\varphi} = \sum_{i:x_i\neq 0} \sum_{n} w_{in} \frac{x_i - n\mu}{\varphi^2} = -1.6618.$$

  Cross-terms involving λ are 0 in this complete-data information.

- To compute one Fisher–scoring update,

$$\boldsymbol{\theta}^{(1)} = \boldsymbol{\theta}^{(0)} + \mathcal{I}(\boldsymbol{\theta}^{(0)})^{-1} U(\boldsymbol{\theta}^{(0)}),$$

  we solve the 3×3 linear system (block-diagonal in λ vs (μ, φ)) yielding the increment

$$\Delta\lambda = -0.3085, \quad \Delta\mu = -0.1959, \quad \Delta\varphi = -0.1001.$$

  Thus,

$$\lambda^{(1)} = 0.8915, \quad \mu^{(1)} = 1.6041, \quad \sigma^{2(1)} = 0.7999.$$

- Next, recompute w_{in} at $\boldsymbol{\theta}^{(1)}$, update the scores and information, and iterate. A simple line search (or step halving) safeguards monotone increase of ℓ while maintaining $\sigma^{2(t)} > 0$.

- **Checks & comments.**
  - The weights align with intuition: large x_i (e.g., 4.2) place most mass on $n = 2, 3$; small x_i (e.g., 0.7) favor $n = 1$.
  - Because the normal part depends on n, μ, and σ^2 couple (nonzero $\mathcal{I}_{\mu\varphi}$), closed-form MLEs no longer exist—iterations are essential.
  - The Poisson truncation at $n_{\max} = 10$ is numerically harmless here; students should verify the tail bound for their own starts.

Solution

After updating the parameter vector $\boldsymbol{\theta}^{(1)} = (\lambda^{(1)}, \mu^{(1)}, \sigma^{2(1)})$, we recompute the latent responsibilities

$$w_{in}^{(1)} = \frac{\Pr(N_i = n \mid \lambda^{(1)})\, f(x_i \mid N_i = n; \mu^{(1)}, \sigma^{2(1)})}{\sum_{m=1}^{n_{\max}} \Pr(N_i = m \mid \lambda^{(1)})\, f(x_i \mid N_i = m; \mu^{(1)}, \sigma^{2(1)})}.$$

Here, $\Pr(N_i = n \mid \lambda^{(1)}) = e^{-\lambda^{(1)}} (\lambda^{(1)})^n / n!$ and $f(x_i \mid N_i = n; \mu^{(1)}, \sigma^{2(1)}) = \mathcal{N}(x_i \mid n\mu^{(1)}, n\sigma^{2(1)})$ for $x_i > 0$, while zeros correspond to structural mass at δ_0.

Once $w_{in}^{(1)}$ is obtained, we update

$$\mathbb{E}[N_i \mid x_i] = \sum_{n=1}^{n_{\max}} n\, w_{in}^{(1)},$$

and recompute the Fisher information and score functions for the next iteration. A simple line search (or step halving) ensures a monotone increase of the observed-data log-likelihood ℓ while maintaining $\sigma^{2(t)} > 0$.

Checks & comments.

- The recomputed responsibilities w_{in} at $\boldsymbol{\theta}^{(1)} = (\lambda^{(1)}, \mu^{(1)}, \sigma^{2(1)}) = (0.8915, 1.6041, 0.7999)$ confirm the expected structure: large observations (e.g., $x_5 = 4.2$) place most of their probability mass on $n = 2, 3$, while smaller ones (e.g., $x_7 = 0.7$) favor $n = 1$.
- The updated conditional expectations of the latent counts are

$$\mathbb{E}[N_i \mid x_i] = (0.000,\ 2.103,\ 1.127,\ 0.000,\ 2.218,\ 1.230,\ 1.076,\ 0.000),$$

  showing that the fitted Poisson–normal model continues to capture increasing latent means with increasing signal intensity x_i.
- Because the normal component depends jointly on n, μ, and σ^2, the parameters remain coupled through the nonzero cross–information term $\mathcal{I}_{\mu\varphi}$; hence closed–form MLEs do not exist, and Fisher–scoring iterations are required.
- The truncation at $n_{\max} = 10$ remains numerically safe: the largest omitted Poisson tail probability is below 10^{-10} under $\lambda^{(1)} = 0.8915$.

Appendix

```r
knitr::opts_chunk$set(
  dev = "cairo_pdf",
  fig.width = 5,
  fig.height = 5,
  fig.align = 'center',
  echo = FALSE,
  message = FALSE,
  warning = FALSE,
  error = FALSE,
  results = 'markup'
)

# Load required libraries
library("tidyverse")
library("patchwork")
library("glue")
library("scales", warn.conflicts = FALSE)
library("extrafont")
library("tinytex")
library("knitr")
library("tidyr")
library("latex2exp")
library("MASS")
library("kableExtra")

theme_set(theme_minimal(base_family = "Roboto Condensed"))

conflicted::conflicts_prefer(
  readr::col_factor(),
  purrr::discard(),
  dplyr::lag(),
  readr::parse_date(),
  kableExtra::group_rows(),
  dplyr::select
)
# Given data
x <- c(0.0, 3.9, 1.5, 0.0, 4.2, 2.1, 0.7, 0.0)
n <- length(x)
n_max <- 10

# Updated parameters after one Fisher-scoring step
lambda <- 0.8915
mu <- 1.6041
sigma2 <- 0.7999
```

```r
# Compute weights w_in = P(N_i = n | X_i = x_i; theta)
w <- matrix(0, nrow = n, ncol = n_max)

for (i in 1:n) {
  for (k in 1:n_max) {
    # Poisson prior on N_i
    pois_part <- dpois(k, lambda)

    # Conditional normal likelihood
    if (x[i] == 0) {
      norm_part <- ifelse(k == 0, 1, 0) # structural zero
    } else {
      norm_part <- dnorm(x[i], mean = k * mu, sd = sqrt(k * sigma2))
    }
    w[i, k] <- pois_part * norm_part
  }
  # Normalize
  w[i, ] <- w[i, ] / sum(w[i, ])
}

# Expected counts for nonzero observations
E_Ni <- rowSums(w * matrix(rep(1:n_max, each = n), nrow = n))
E_Ni[x == 0] <- 0

# Display updated responsibilities and conditional expectations
w
E_Ni
```