

Optimization Exercise

Download the R routines `lzip` and `gzip` (they are contained in the same file) and read them into R using the `source` command. Add `gzip` as a gradient attribute to `lzip` using the command `attr(lzip,"gradient")=gzip`. R routines that can take advantage of the analytical evaluation of the gradient will now be able to use this user-supplied gradient function. Generate some zero-inflated Poisson data using mixing proportion .5 and a Poisson mean of 1:

```
x=0+rbinom(100,1,.5)*rpois(100,lambda=1)
```

Fit the zero-inflated Poisson model by calling `nlm`:

```
nlm(lzip,c(.5,1),hessian=T,xsam=x)
```

The solution will be listed in the output as the attribute “estimate”; the estimated proportion of observations from the degenerate distribution (all mass placed at 0) will be listed first and the estimated Poisson mean will be listed next. Repeat this exercise with sample sizes of 20, 50 and 100 and a mean of 1, 2 and 3. Comment on your results.

We can also fit these models in SAS. Download the SAS file for fitting zero-inflated Poisson data on the website and run it for some of the same scenarios you used for the R code. The SAS code uses PROC NLMIXED to evaluate log likelihoods; it generates its own zero-inflated Poisson data, so results will not be strictly comparable unless you change the program to accept user-supplied data.