Profiling and marginalization

A NUISANCE PARAMETER

$L_p(s) = \max_{b \geq 0} L(s, b)$

COMPLETE EXAMPLE

One signal and one control count

```
%%%
Two treatments of an unknown background yield.
%%%
flatppl_compat = "0.1"
counts = load_module(
    "https://raw.githubusercontent.com/
    flatppl-community/flatppl-stats-counts/
    4f0a6a8eedf0595d1d779fe6ba6a5579fa96eaf4/
    main.flatppl"
)
n = 7
m = 12
ratio = 5.0
data = record(on = n, off = m)
law(s, b) = joint(on = Poisson(s + b),
    off = Poisson(ratio * b))
joint_at(s, b) = logdensityof(law(s, b), data)
b_hat(s) = counts.onoff_background_mle(
    n, m, s, ratio)
profile(s) = joint_at(s, b_hat(s))
shape = 2.0
rate = 1.0
marginal(s) = counts.onoff_gamma_logdensity(
    n, m, s, ratio, shape, rate)
s = 3.0
bg_at_s = b_hat(s)
log_profile = profile(s)
log_marginal = marginal(s)
```

Evaluate a signal scan

```
scan_p(t) = counts.onoff_profile_loglikelihood(
    n, m, t, ratio)
grid = linspace(0.0, 12.0, 61)
lp = scan_p.(grid)
lm = marginal.(grid)
```

Evaluate the module's scalar profile formula and the Gamma marginal on the same grid. The dots broadcast both calls over the signal values.

The integral over background

At fixed signal the marginal likelihood is $\int_0^\infty L(s, b)\pi(b)\,\mathrm{d}b$. The Gamma density weights each background interval in the integral.

ANNOTATED LINES

A Signal and control counts 13–15

```
law(s, b) = joint(on = Poisson(s + b),
    off = Poisson(ratio * b))
joint_at(s, b) = logdensityof(law(s, b), data)
```

The signal region has Poisson mean `s + b`. The control exposure is five times larger, with mean `5*b`. The observed counts are 7 and 12, respectively.

B Profile at fixed signal 17–19

```
b_hat(s) = counts.onoff_background_mle(
    n, m, s, ratio)
profile(s) = joint_at(s, b_hat(s))
```

At a fixed `s`, maximize the joint likelihood over nonnegative background. Inserting the maximizer `b_hat(s)` gives the profile likelihood as a function of `s`.

C Integrate a background prior 21–24

```
shape = 2.0
rate = 1.0
marginal(s) = counts.onoff_gamma_logdensity(
    n, m, s, ratio, shape, rate)
```

For integer counts the module integrates exactly against a Gamma prior for `b`. Shape 2 and rate 1 give a prior mean of 2.

D Evaluation at a fixed signal 26–29

```
s = 3.0
bg_at_s = b_hat(s)
log_profile = profile(s)
log_marginal = marginal(s)
```

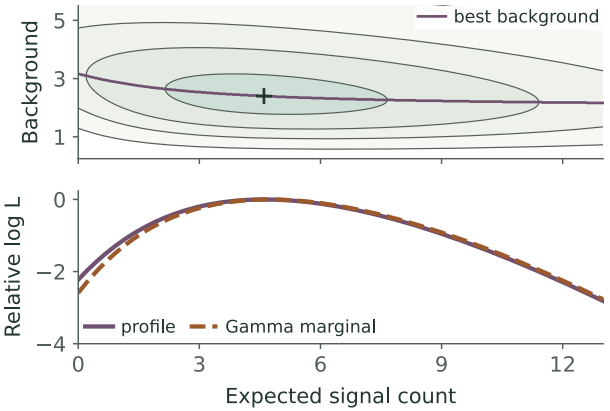
At `s = 3` the fitted background is 2.53424. Before shifting the curves to their peaks, the log likelihoods are −4.26916 for profiling and −5.37638 for marginalization.

Check the joint maximum

The control count gives $\hat{b} = \frac{12}{5} = 2.4$. The signal region then gives $\hat{s} = 7 - 2.4 = 4.6$. Both means reproduce their observed counts.

COMPARING SIGNAL LIKELIHOODS

Profile and marginal signal likelihoods



The ridge shows the fitted background at each signal. Each log-likelihood curve in the lower panel has its own maximum at zero. The dashed curve integrates over the Gamma prior.

Locations of the maxima

The profile peaks at 4.60 and the marginal near 4.73 for this prior. Measure changes in log likelihood relative to each curve's own maximum.

Background fit at zero signal

At `s = 0` the two regions have means b and $5b$. The total of 19 counts across six exposure units gives $\hat{b}(0) = \frac{19}{6}$. This fixes the first point on the profile ridge.

Interpret the exposure ratio

`ratio` is the control exposure divided by the signal-region exposure. Since `b` denotes the expected signal-region background count, multiplying it by `ratio` gives the expected control count.

From a marginal likelihood to a posterior

The background prior affects only the marginal curve. Multiply that curve by a signal prior and normalize to obtain a signal posterior. The counts module supplies both scan functions.

Jessica Cox · bcox@mpp.mpg.de