Invariances

A blog about research methods, statistics, infrastructure, cognition, memory, and perception

Sunday, January 24, 2016

Roll Your Own: How to Compute Bayes Factors For Your Priors

Sometimes people ask what prior they should use in computing Bayes factor. It's a great question, and trying to answer it leads to a deeper understanding of Bayesian methods. Here I provide some R code so that you can visualize and compute Bayes factors for (almost) any prior you wish. At the end of this blog I provide all the R code in one block. You might even want to go down there and cut-and-paste it into your R editor now.

The Bayes factor compares two models. Let's take one of them to be a point null. The other, the alternative, is up to you.

Models

Let's first start with a data model. Data are independent normal draws with mean μ and variance σ^2 . It is more convenient re-express the normal as a function of effect size, δ and σ^2 where $\delta = \mu/\sigma$. Here is the formal specification:

$$Y_i \mid \delta, \sigma^2 \stackrel{iid}{\sim} \text{Normal}(\sigma\delta, \sigma^2).$$

A null model is implemented by setting $\delta = 0$. The alternative model on effect size is up to you. Here is how you do it. Below is the specification for the R function altDens, the density of the alternative. I have chosen a normal with mean of .5 and a standard deviation of .3. I have also truncated this distribution at zero, and negative values are not allowed. You can see that in the specification of the function altDens. Set lo and hi, the support of the alternative. Here I set them to 0 and infinity, respectively.

```
#Specify Alternative (up to constant of proportionality)
lo=0 #lower bound of support
hi=Inf #upper bound of support
altDens=function(delta)
dnorm(delta,.5,.3)*as.integer(delta>lo)*as.integer(delta<hi)
```

You may notice that this alternative is not quite a proper density because it does not integrate to 1 That's ok, the following function scales the alternative so the density does indeed integrate to 1.0.

```
#Normalize alternative density in case user does not,
K=1/integrate(altDens,lower=lo,upper=hi)$value
f=function(delta) K*altDens(delta)
```

Visualizing the Models

Let's now take a look at the competing models on δ , the effect size. Here is the code and the graphs:

```
delta=seq(-2,3,.01)
maxAlt=max(f(delta))
plot(delta,f(delta),typ='n',xlab="Effect Size Parameter Delta",
ylab="Density",ylim=c(0,1.4*maxAlt),main="Models")
```

About Me

Jeff Rouder

[View my complete profile](#)

Blog Archive

► [2017](#) (2)

▼ [2016](#) (11)

► [October](#) (1)

► [April](#) (2)

► [March](#) (4)

► [February](#) (2)

▼ [January](#) (2)

Roll Your Own: How to Compute Bayes Factors For Yo...

[Bayesian Modeling to Improve The Spatial and Tempo...](#)

► [2015](#) (12)

```
arrows(0,0,0,1.3*maxAlt,col='darkblue',lwd=2)
lines(delta,f(delta),col='darkgreen',lwd=2)
legend("topright",legend=c("Null","Alternative"),col=c('darkblue','darkgreen'),lwd=2)
```

Bayes Factor Computation

Bayes factor is based on prediction. It is the ratio of the predicted density of the data under one model relative to that under another. The predicted density of data for the null is easy to compute and is related to the central t distribution. Here is the code.

```
nullPredF=function(obs,N) dt(sqrt(N)*obs,N-1)
```

We can compute this predicted density for any observed effect size or for all of them. The following code does it for a reasonable range of effect sizes for a sample size of 30. You can change N, the sample size, as needed.

```
obs=seq(-2,3,.01)
N=30
nullPred=nullPredF(obs,N)
```

Getting the predictive density for the alternative is a bit harder. For each nonzero effect size parameter, the distribution of the observed effect follows a noncentral t distribution. Hence, to obtain predictions across all nonzero effect size parameters, we need to integrate the alternative model against the noncentral t distribution. Here is the code with a simple loop:

```
altPredIntegrand=function(delta,obs,N)
dt(sqrt(N)*obs,N-1,sqrt(N)*delta)*f(delta)

altPredF=function(obs,N)
integrate(altPredIntegrand,lower=lo,upper=hi,obs=obs,N=N)$value

I=length(obs)
altPred=1:I
for (i in 1:I) altPred[i]=altPredF(obs[i],N)
```

Now we can plot the predictions for all observed effect sizes:

```
top=max(altPred,nullPred)
plot(type='l',obs,nullPred,ylim=c(0,top),xlab="Observed Effect
Size",ylab="Density",main="Predictions",col='darkblue',lwd=2)
```

```
lines(obs,altPred,col='darkgreen',lwd=2)
legend("topright",legend=c("Null","Alternative"),col=c('darkblue',
'e','darkgreen'),lwd=2)
```

Let's Run an Experiment

Suppose we just ran an experiment and observed a sample effect size of .4. Let's look at the predictions for this value.

```
plot(type='l',obs,nullPred,ylim=c(0,top),xlab="Observed Effect
Size",ylab="Density",main="Predictions",col='darkblue',lwd=2)
lines(obs,altPred,col='darkgreen',lwd=2)
legend("topright",legend=c("Null","Alternative"),col=c('darkblue',
'e','darkgreen'),lwd=2)
```

```
my.es=.4
abline(v=my.es,lty=2,lwd=2,col='red')
valNull=nullPredF(my.es,N)
valAlt=altPredF(my.es,N)
points(pch=19,c(my.es,my.es),c(valNull,valAlt))
```

```
#cat("Predictive Density under the null is ",valNull)
#cat("Predictive Density under the specified alternative is ",v
alAlt)
#cat("Bayes factor (alt/null) is ",valAlt/valNull)
```

How well did the models do? Well, this value was clearly better predicted under the specified alternative than under the null. The density under the null is .04; the density under the alternative .205, and the ratio between them is 5.16-to-1. This ratio is the Bayes factor. In this case the Bayes factor indicate support of slightly more than 5-to-1 for the specified alternative.

Other Alternatives

The alternative is specified with just a bit of R code that can be changed. Suppose, for example, one wished a flat or uniform alternative from -.2 to 1.2. Here is the bit of code

```
#Specify Alternative (up to constant of proportionality)
lo=-.2 #lower bound of support
hi=1.2 #upper bound of support
altDens=function(delta)
dunif(delta,lo,hi)
```

Run this first, and then run the rest to see what happens.
Here is another specification; this alternative is a gamma

```
#Specify Alternative (up to constant of proportionality)
lo=0 #lower bound of support
hi=Inf #upper bound of support
altDens=function(delta)
dgamma(delta,shape=2,scale=.5)
```

The R Code in One Block

```
N=30
my.es=.4

#####
#Specify Alternative (up to constant of proportionality)
#Change this section to change alternative
lo=0 #lower bound of support
```

```
hi=Inf #upper bound of support
altDens=function(delta)
dnorm(delta,.5,.3)*as.integer(delta>lo)*as.integer(delta<hi)
#####

#Normalize alternative density in case user does not,
K=1/integrate(altDens,lower=lo,upper=hi)$value
f=function(delta) K*altDens(delta)

delta=seq(-2,3,.01)

#Plot Alternative as a density and Null as a point arrow
maxAlt=max(f(delta))
plot(delta,f(delta),typ='n',xlab="Effect Size Parameter Delta",
ylab="Density",ylim=c(0,1.4*maxAlt),main="Models")
arrows(0,0,0,1.3*maxAlt,col='darkblue',lwd=2)
lines(delta,f(delta),col='darkgreen',lwd=2)
legend("topright",legend=c("Null","Alternative"),col=c('darkblue',
'darkgreen'),lwd=2)

#Prediction Function Under Null
nullPredF=function(obs,N) dt(sqrt(N)*obs,N-1)

#Prediction Function Under the Alternative
altPredIntegrand=function(delta,obs,N)
dt(sqrt(N)*obs,N-1,sqrt(N)*delta)*f(delta)
altPredF=function(obs,N)
integrate(altPredIntegrand,lower=lo,upper=hi,obs=obs,N=N)$value

obs=delta
I=length(obs)
nullPred=nullPredF(obs,N)
altPred=1:I
for (i in 1:I) altPred[i]=altPredF(obs[i],N)

#Plot The Predictions
top=max(altPred,nullPred)
plot(type='l',obs,nullPred,ylim=c(0,top),xlab="Observed Effect
Size",ylab="Density",main="Predictions",col='darkblue',lwd=2)
lines(obs,altPred,col='darkgreen',lwd=2)
legend("topright",legend=c("Null","Alternative"),col=c('darkblue',
'darkgreen'),lwd=2)

#Evaluate Predicted Density at Observed Value my.es
abline(v=my.es,lty=2,lwd=2,col='red')
valNull=nullPredF(my.es,N)
valAlt=altPredF(my.es,N)
points(pch=19,c(my.es,my.es),c(valNull,valAlt))
cat("Bayes factor (alt/null) is ",valAlt/valNull)
```

Posted by [Jeff Rouder](#) at 12:26 AM

Recommend this on Google

1 comment:

Jessy Ryder said...

Most employers require a masters or a management degree for a top job placement, as a result the need for an MBA is essential in landing a lucrative job. [science homework help online](#)

November 17, 2016 at 8:25 AM

[Post a Comment](#)

Links to this post

[Create a Link](#)

[Newer Post](#)

[Home](#)

[Older Post](#)

Subscribe to: [Post Comments \(Atom\)](#)

Invariances

A blog about research methods, statistics, infrastructure, cognition, memory, and perception

Monday, March 21, 2016

Roll Your Own II: Bayes Factors With Null Intervals

The Bayes factors we develop compare the null model to an alternative model. This null model is almost always a single point---the true effect is identically zero. People sometimes confuse our advocacy for Bayes factor with that for point-null-hypothesis testing. They even critique Bayes factor with the Cohenesque claim that the point null is never true.

Bayes factor is a general way of measuring the strength of evidence from data for competing models. It is not tied to the point null.

We develop for the point null because we think it is a useful, plausible, theoretically meaningful model. Others might disagree, and these disagreements are welcome as part of the exchange of viewpoints in science.

In the blog post **Roll Your Own: How to Compute Bayes Factors for Your Priors**, I provided **R** code to compute a Bayes factor between a point-null and a user-specified alternative for a simple setup motivated by the one-sample t-test. I was heartened by the reception and I hope a few of you are using the code (or the comparable code provided by Richard Morey). There have been some requests to generalize the code for non-point nulls. Here, let's explore the Bayes factor for any two models in a simple setup. As it turns out, the generalization is instructive and computationally *trivial*. We have all we need from the previous posts.

Using Interval Nulls: An Example

Consider the following two possibilities:

I. Perhaps you feel the point null is too constrained and would rather adopt a null model with mass on a small region around zero rather than at the point. John Kruschke calls these regions ROPEs (regions of posterior equivalence).

II. Perhaps you are more interested in the direction in an effect rather than whether it is zero or not. In this case, you might consider testing two one-sided models against each other.

For this blog, I am going to retain four different priors. Let's start with a data model. Data are independent normal draws with mean μ and variance σ^2 . It is more convenient re-express the normal as a function of effect size, δ and σ^2 where $\delta = \mu/\sigma$). Here is the formal specification:

$$Y_i \mid \delta, \sigma^2 \stackrel{iid}{\sim} \text{Normal}(\sigma\delta, \sigma^2).$$

Now, the substantive positions as prior models on effect size:

1. M_0 , A Point Null Model: $\delta = 0$
2. M_1 , A ROPE Model: $\delta \sim \text{Unif}(-.25, .25)$
3. M_2 , A Positive Model: $\delta \sim \text{Gamma}(3, 2.5)$

About Me

Jeff Rouder

[View my complete profile](#)

Blog Archive

► [2017](#) (2)

▼ [2016](#) (11)

► [October](#) (1)

► [April](#) (2)

▼ [March](#) (4)

[The Effect-Size Puzzler, The Answer](#)

[The Effect-Size Puzzler](#)

[Roll Your Own II: Bayes Factors With Null Interva...](#)

[Statistical Difficulties from the Outer Limits](#)

► [February](#) (2)

► [January](#) (2)

► [2015](#) (12)

4. M_3 , A Negative Model: $-\delta \sim \text{Gamma}(3,2.5)$

Here are these four models expressed graphically as distributions:

I picked these four models, but you can pick as many ones as you wish. For example, you can include a normal if you wish.

Oh, let's look at some data. Suppose the observed effect size is .35 for an N of 60.

Going Transitive

Bayes factors are the comparison between two models. Hence we would like to compute the Bayes factors between any of these models. Let B_{ij} be the comparison between the i th and j th model. We want a Table like this:

B_{00}	B_{01}	B_{02}	B_{03}
B_{10}	B_{11}	B_{12}	B_{13}
B_{20}	B_{21}	B_{22}	B_{23}
B_{30}	B_{31}	B_{32}	B_{33}

Off the bat, we know the Bayes factor between a model and itself is 1 and that $B_{ij} = 1/B_{ji}$. So we only need to worry about the lower corner.

1			
B_{10}	1		
B_{20}	B_{21}	1	
B_{30}	B_{31}	B_{32}	1

We can use the code below, from the previous post to figure out the null vs. all the other models.

$$B_{10} = 4.9, \quad B_{20} = 4.2, \quad B_{30} = .0009$$

Here we see that the point null is not as attractive or the ROPE null or the positive model. It is more attractive, however, than the negative model.

Suppose, however, that you are most interested in the ROPE null and its comparison to the positive and negative model. The missing Bayes factors are B_{12} , B_{13} , and B_{23} .

The key application of transitivity is as follows:

$$B_{ij} = B_{ik} \times B_{kj}.$$

So, we can compute B_{12} as follows:

$$B_{12} = B_{10} \times B_{02} = B_{10}/B_{20} = 4.9/4.2 = 1.2.$$

The other two Bayes factors are computed likewise: $B_{13} = 5444$ and $B_{23} = 4667$

So what have we learned. Clearly, if you were pressed to choose a direction, it is in the positive direction. That said, the evidence for a positive effect is slight when compared to a ROPE null.

Snippets of R Code

```
#First, Define Your Models as a List
```

```
#lo, lower bound of support
```

```
#hi, upper bound of support
```

```
#fun, density function
```

```
#here are Models M1, M2, M3
```

```
#add or change here for your models
```

```
mod1=list(lo=-.25,hi=.25,fun=function(x,lo,hi) dunif(x,lo,hi))
```

```
mod2=list(lo=0,hi=Inf,fun=function(x,lo,hi) dgamma(x,shape=3,rate=2.5))
```

```
mod3=list(lo=-Inf,hi=0,fun=function(x,lo,hi) dgamma(-x,shape=3,rate=2.5))
```

```
#note, we dont need to specify the point null, it is built into the code
```

```
#Lets make sure the densities are proper, here is a function to do so:
```

```
normalize=function(mod)
```

```
return(c(mod,K=1/integrate(mod$fun,lower=mod$lo,upper=mod$hi,lo=mod$lo,hi=mod$hi)$value))
```

```
#and now we normalize the three models
```

```
mod1=normalize(mod1)
```

```
mod2=normalize(mod2)
```

```
mod3=normalize(mod3)
```

```
#Observed Data
es=.35
N=60

#Here is the key function that computes the Bayes factor between a model and the
point null
BF.mod.0=function(mod,es,N)
{
  f= function(delta) mod$fun(delta,mod$lo,mod$hi)*mod$K
  pred.null=dt(sqrt(N)*es,N-1)
  altPredIntegrand=function(delta,es,N)
    dt(sqrt(N)*es,N-1,sqrt(N)*delta)*f(delta)
  pred.alt=
    integrate(altPredIntegrand,lower=mod$lo,upper=mod$hi,es=es,N=N)$value
  return(pred.alt/pred.null)
}

B10=BF.mod.0(mod1,es,N)
B20=BF.mod.0(mod2,es,N)
B30=BF.mod.0(mod3,es,N)

print(paste("B10=",B10," B20=",B20," B30=",B30))

B12=B10/B20
B13=B10/B30
B23=B20/B30

print(paste("B12=",B12," B13=",B13," B23=",B23))
```

Posted by [Jeff Rouder](#) at 12:02 PM

Recommend this on Google

2 comments:

Matt Williams said...

Hi Jeff,

I think this is a neat post: Computing Bayes Factors for comparing more plausible models sounds like a better idea than focusing on a point null. And it's great that you've shown how to do this.

But to expand on a question I asked you on Twitter:

The BFs alone don't directly lead to conclusions about which *model* is actually most probable. For that we need posterior probabilities.

But finding the posterior probabilities require us to define the prior probability that each model is correct. We presumably wouldn't want to just assume that each model is equally plausible and define prior probability = 0.25 for each of the four models: That would lead to a rather strange and very informative prior probability distribution (sort of a hill plus tower plus radio antenna shape in the case of these four models!)

We could make a more considered decision about how much prior probability to place on each of the four models based on existing knowledge. And use that to convert the Bayes Factors to a posterior probability distribution that tells us the posterior probability that each model is correct.

But if we're going to go to the trouble of specifying priors, estimating models for the data, and calculating a posterior probability distribution... Why not just skip straight to Bayesian estimation, and simply look at the proportion of the posterior draws that fall in each interval? (Or that fall exactly on an effect size of zero, in the case of the point null?)

I.e., What does the Bayes Factor approach add that you can't achieve via Bayesian estimation?

One advantage I see is that it may make it a bit easier to deal with priors that include a point mass on zero, but I'm wondering if there are other advantages that lead you to be more enthusiastic about this approach? :)

March 29, 2016 at 9:02 PM

Jeff Rouder said...

Hi Matt, Thanks for writing and kind words.

I must admit, I am partial to the point null. I like it because it has solid theoretical meaning---it denotes an invariance. I think certain invariances, even if they do not hold exactly, are exceedingly useful. For example, the Earth doesn't follow an ellipse around the sun exactly as its orbit is perturbed ever so slightly by the moon. Yet, the Kepler Laws are a monumental achievement.

Also, I can appreciate an emphasis on posterior odds. I prefer to report BFs and let people multiply them by whatever prior odds they wish. The BFs are the evidence from the data, or how beliefs should be updated in light of data. If I wish to provide additional guidance, I will suggest prior odds, but by reporting BFs I leave the reader free to interpret the evidence as they wish. I am not wed to making decisions for myself or others; doing so to me is akin to asking how bananas is a lot of bananas. Why make a decision when we can just report the number of bananas at hand.

There is no real difference between BF and estimation; I view them as united. Consider a spike-and-slab model on effect size with mass on zero and density over a continuum. There are three prior settings: the prior mass on zero (the spike), and the prior center and dispersion of the slab. The spike-and-slab form is conjugate, so the posterior is also a spike-and-slab with updated values of the mass of the point, and the center and dispersion of the slab. So this posterior is estimate. The ratio of posterior spike mass to prior spike mass is also the Bayes factor.

I'll write a post about it one day.

April 5, 2016 at 10:35 PM

[Post a Comment](#)

Links to this post

[Create a Link](#)

[Newer Post](#)

[Home](#)

[Older Post](#)

Subscribe to: [Post Comments \(Atom\)](#)