

Computing discriminability and bias with the R software

Christophe Pallier*

June 26, 2002

Abstract

Psychologists often model detection or discrimination decisions as being influenced by two components: perceptual discriminability and response bias. In the signal detection framework, these are estimated by d' and β . Other, non-parametric parameters, A' and B''_D , have also been proposed. This note provides the code to compute those parameters with the statistical software R (<http://www.r-project.org>).

Introduction

Consider the following three situations :

- In a signal detection test, the subject is presented with a series of trials in which he must decide if a target signal (e.g. a faint tone in a noisy background) is present or absent.
- In a memory recognition test, the subject is first trained to memorize some stimuli. Then, in a subsequent test phase, he is presented with a series of stimuli and asked to decide if they are 'old' or 'new' items.
- In a discrimination task, the subject is presented with pairs of stimuli, and for each pair, asked to tell whether the stimuli are the same or different.

In all three cases, the trials belong to two categories (Y=target present; N=target absent). So do the responses of the subjects (Y=yes; N=no). Therefore, there are four possible combinations of trial type and response (see Table 1): A 'Hit' occurs when the subject correctly detects a signal (or

correctly recognize an old item, or correctly discriminate different stimuli...). A 'Miss' occurs when the subject fails to detect the presence of the target signal. A 'False alarm' (abbreviated as 'FA') occurs when the subject reports the target while it is not present. Finally, a 'Correct rejection' ('CR') occurs when the subject correctly reports no target.

Trial	Response	
	Y	N
Y	Hits	Misses
N	FA	CR

Table 1: Table of scores

The data of a single participant can be summarized in a table such as Table 1, which reports the counts of the different types of trial-response combinations. The four figures in the summary table are not independent:

$$\begin{cases} \text{Hits} + \text{Misses} = \text{total number of 'Y' trials} \\ \text{FA} + \text{CR} = \text{total number of 'N' trials} \end{cases}$$

Therefore, the data are often summarized by the two independent numbers:

$$\begin{cases} \text{Hit rate} = \text{Hits}/(\text{Hits}+\text{Misses}) \\ \text{FA rate} = \text{FA}/(\text{FA}+\text{CR}) \end{cases}$$

The percentage of correct responses is:

$$\% \text{correct} = (\text{Hits} + \text{CR})/(\text{Hits}+\text{Misses}+\text{FA}+\text{CR})$$

If there are as many trials of types 'Y' and 'N', the percentage of correct responses is simply the mean of the Hit and FA rates. The percentage of correct responses is a function of discriminability and bias. To introduce these notions, it is useful to plot the subject's performance on a graphics with

*www.pallier.org




Figure 1: The (fa, hit) unit square (left) and the percentage correct responses as a function of (fa, hit) (right)

FA and Hit rates as coordinates, as shown on Figure 1.

Psychologists model binary answers of human participants as being influenced by two distinct factors: (a) a perceptual discrimination component and (b) a response bias. For example, two subjects with similar perceptual discrimination capabilities can have different propensions to answer ‘yes’ or ‘no’. Or, in one individual, the response bias can be manipulated by varying the costs/benefits associated to the responses.

Consider the following three examples of extreme behaviors:

1. A subject (a) has Hit and FA rates at 100% and 0% respectively; he is considered to discriminate perfectly and have no response bias.
2. A subject (b) systematically answer ‘Yes’: his Hit and Fa rates are both at 100%. One would say that he has a strong response bias, and no evident discrimination capacity.
3. A subject (c) has Hit and FA rates both at 50%; he will be considered to have no response bias, and no discriminability either (his score falls on the “chance diagonal”).

To disentangle the two components (bias and discriminability) in a given subject, several measures have been proposed, that improve on the simple percentage of correct responses. The most used are d' and β and the non-parametric A' and B''_D [1].

d' and β originate from the signal detection theory framework. d' reflects the distance between the

two distributions: signal, and signal+noise and corresponds to the Z value of the hit-rate minus that of the false-alarm-rate. Though Z values can have any real value, normally distributed ones are between -2 and 2 about 95% of the time, so differences of twice that would be rare. The value for β is the ratio of the normal density functions at the criterion of the Z values used in the computation of d' . This reflects an observer’s bias to say ‘yes’ or ‘no’ with the unbiased observer having a value around 1.0. As the bias to say ‘yes’ increases, resulting in a higher hit-rate and false-alarm-rate, beta approaches 0.0. As the bias to say ‘no’ increases, resulting in a lower hit-rate and false-alarm-rate, beta increases over 1.0 on an open-ended scale [2].

A' and B''_D are non-parametric estimates of discriminability and bias. Given a point in the (fa, hit) unit square, A' is defined as the sum of the area of B and half the areas of A1 and A2 (cf. Fig. 4). The formula for computing A' and B''_D are:

If $hit > fa$,

$$A' = 1/2 + \frac{(hit - fa) * (1 + hit - fa)}{4 * hit * (1 - fa)}$$

If $fa > hit$,

$$A' = 1/2 - \frac{(fa - hit) * (1 + fa - hit)}{4 * fa * (1 - hit)}$$

and,

$$B''_D = \frac{(1 - hit) * (1 - fa) - hit * fa}{(1 - hit) * (1 - fa) + hit * fa}$$

We will not detail the rationale behind the construction of these parameter (see [1]). Suffice to say

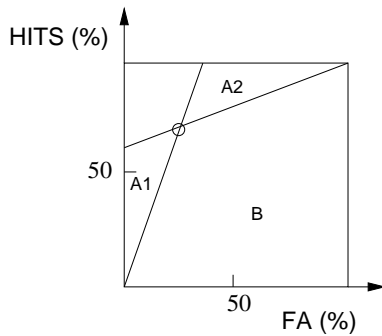


Figure 4: Areas involved in the computation of A'

that an A' near 1.0 indicates good discriminability, while a value near 0.5 means chance performance. A B''_D equal to 0.0 indicates no bias, positive numbers represent conservative bias (i.e. a tendency to answer 'no'), negative numbers represent liberal bias (i.e. a tendency to answer 'yes'). The maximum absolute value is 1.0.

The isolevel curves for the parameters that estimate discriminability are shown on figure 2. They can be compared with the right panel of Figure 1. The isolevel curves for the bias estimates d' and B''_D are shown on figure 3.

R functions

Here is the R code for four functions `dprime`, `beta`, `aprime` and `bppd` that compute the parameters as a function of Hit and FA rates:

```
dprime <- function(hit,fa) {
  qnorm(hit) - qnorm(fa)
}

beta <- function(hit,fa) {
  zhr <- qnorm(hit)
  zfar <- qnorm(fa)
  exp(-zhr*zhr/2+zfar*zfar/2)
}

aprime <-function(hit,fa) {
  a<-1/2+((hit-fa)*(1+hit-fa) /
    (4*hit*(1-fa)))
  b<-1/2-((fa-hit)*(1+fa-hit) /
    (4*fa*(1-hit)))
  a[fa>hit]<-b[fa>hit]
```

```
  a[fa==hit]<- .5
  a
}

bppd <-function(hit,fa) {
  ((1-hit)*(1-fa)-hit*fa) /
  ((1-hit)*(1-fa)+hit*fa)
}
```

Hit and FA rates can be passed as single numbers or as vectors. In the latter case, the function will return a vector of the same length. This is useful to get the parameters for a set of subjects. Note the trick in `aprime` to handle the two cases `fa>hit`, `hit>fa` without resorting to a loop on the vectors' elements.

Practical usage

We suppose that the code given above for the four functions is saved in a text file named, say, `discr.R`. To be able to access these functions, the first command to run is:

```
source('discr.R')
```

If you have already computed Hit and FA rates for each subject, you can enter them into two R vectors `hit` and `fa`, using the 'c' operator:

```
hit <- c(.7,.8,.6)
fa <- c(.2,.1,.3)
```

Then plot the performances and compute the parameters:

```
plot(fa, hit, xlim=c(0,1), ylim=c(0,1))
aprime(hit, fa)
bppd(hit, fa)
dprime(hit, fa)
beta(hit, fa)
```

If, rather, the rates are saved in a text file `data.txt` in a tabular format such as:

```
s1 0.7 0.2
s2 0.8 0.1
s3 0.6 0.3
```

You would then rather use:




Figure 2: Isolevel plots for the discriminability parameters

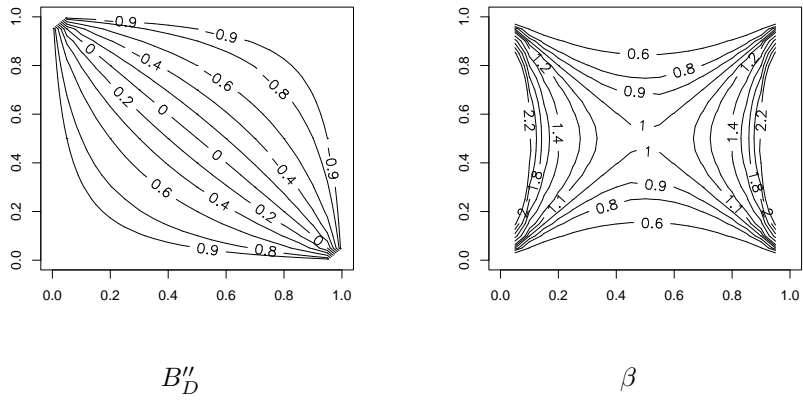


Figure 3: Isolevel plots for the bias parameters

```
a<-read.table('data.txt')
hits <- a$V2;
fa <- a$V3;
```

Finally, only raw data may be available, for example, as a table with one row per trial, with the three space-separated columns corresponding respectively to the subject code, the trial type, and the response of the subject:

```
s1 S Y
s1 S N
s1 N Y
s1 N N
s1 N Y
...
s2 S Y
s2 S Y
s2 N Y
s2 S N
...
```

In this case, to compute subjects' hit and fa rates:

```
a<-read.table('samp.dat',
              col.names=c('subj','target','resp'))
attach(a)
b<-table(target,resp,subj)
subj <- names(b[1,1,])
hit <- b['S','Y',]/(b['S','Y',]+b['S','N',])
fa <- b['N','Y',]/(b['N','Y',]+b['N','N',])
```

You can explore 'b' and check that $b['S','Y',]$ = Hits, $b['S','N',]$ = Misses, $b['N','Y',]$ = False alarms, and $b['N','N',]$ = Correct rejections. Whence, the formula for hit and fa.

Explorations

The isodiscriminability curves were obtained using the contour function:

```
a<-0:10/10
contour(z=outer(a,a,"aprime"))
contour(z=outer(a,a,"dprime"))
```

In figure 5, we looked at the behavior of percent correct, A' and d' when the fa rate was set to a fixed value of 5%, and the hit rate varied from 0% to 100%. The code was the following:

```
hit<-0:20/20
fa<-rep(.05,21)
pairs(cbind(hit,
            correct=correct(hit,fa),
            aprime=aprime(hit,fa),
            dprime=dprime(hit,fa)))
```

References

- [1] Wayne Donaldson (1992) Measuring Recognition Memory. *J. Exp. Psychol.: General*, 121, 3, 275–277
- [2] Gary Perlman. —STAT 5.4: data manipulation & analysis programs for unix and msdos (software package available on the web)

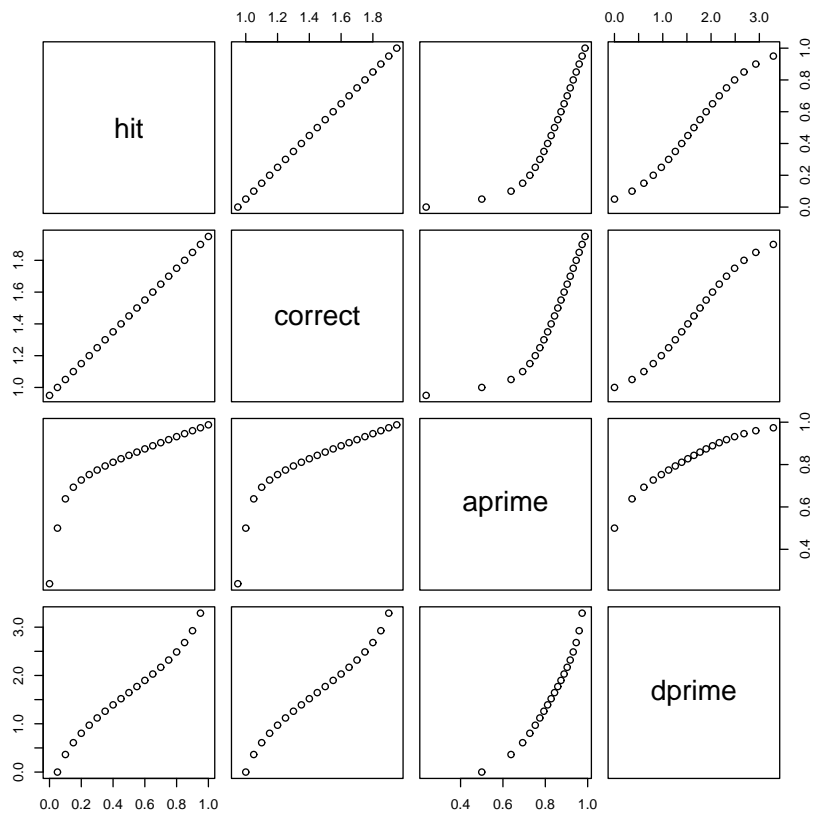


Figure 5: Relationship between percent correct, A' and d' for a constant FA rate=0.05