

Documentation for BAMROC Version 1.0

Thor Aspelund

The University of Iowa, Iowa City, Iowa 52246

March 27, 2001

SUMMARY. This is documentation to accompany the S-PLUS (2000) program BAMROC, Version 1.0. BAMROC fits, via maximum likelihood, Binormal Association Marginal (BAM) models for bivariate ordinal rating data, as described in Lang and Aspelund (2001).

KEY WORDS: binormal model; correlated ordinal rating data; loglinear model; receiver operating characteristic curve; signal detection theory

1. Introduction

Lang and Aspelund (2001) introduce an association-marginal model to analyze bivariate ordinal rating data. The ordinal variable association is specified using standard log-linear models that typically exploit the ordinal structure of the data; see Agresti (1988). The two ordinal variable marginal distributions are assumed to follow the binormal model of classic signal detection theory (see, for example, Dorfman and Alf (1969)).

It follows that the model, called the Binormal Association-Marginal (BAM) model in Lang and Aspelund (2001), is applicable to receiver operating char-

email: thor-aspelund@uiowa.edu

Table 1
Neonatal Radiograph Rating Data

		Plain Film					Plain Film					
		1	2	3	4	5	1	2	3	4	5	
Video	1	4	1	1	0	0	6	1	0	1	2	0
	2	4	8	3	1	1	17	2	1	2	1	0
	3	0	2	2	0	0	4	3	0	2	1	1
	4	0	3	1	1	0	5	4	0	2	3	6
	5	0	0	0	1	0	1	5	1	1	3	27
		8	14	7	3	1	33	3	7	9	14	34
Normals												Abnormals

acteristic (ROC) analysis of bivariate ordinal rating data.

For more details and background, see Lang and Aspelund (2001), especially Sections 3, 5, and the Appendix.

2. Example: Neonatal Radiograph Data

The first example discussed in the paper is based on data from one radiologist in a paired-comparison study described in Franken et al. (1992).

Radiographs of 33 controls (normals) and 67 cases (abnormals) were viewed by 4 radiologists on two modalities, video image and plain film, and then given a score on a 5 point scale to indicate the suspicion of disease. The data from one of the readers is displayed in Table 1 in such a way that higher score means higher suspicion. For example, 2 controls received a score of 3 on the video and 2 on the plain film while 3 cases received a score of 5 on the video and 3 on the film.

Before the data analysis we discuss the features of the program.

3. Installing the program

The code is in a text file and can be read into S-PLUS with the source command

```
> source("sourcefile")
```

where `sourcefile` is the name of the source file. This creates the following procedures

```
"Amatrix", "bamroc", "bamroc1", "bamroc.area", "bn.area",
"bn.curves", "DA", "DirectSum", "dL", "dML", "dnum", "domega.dtheta",
"dpi.domega", "eta.ls", "gof", "gof.bam", "info.alpha", "info.omega",
"info.theta", "initial.loglinear", "L", "L2norm", "label.bam",
"local.odds", "LowerOne", "MarginalsB", "ML", "omega.theta",
"pi.omega", "roc.curve", "roc.points", "score.alpha", "score.omega",
"score.theta", "shiftX", "solve.eta", "solve.theta.efficient",
"standardM", "standardW", "summary.bam", "wald", "wald.bam",
"wls.approx"
```

so to avoid conflicts, run the source command in a new directory or workspace of S-PLUS.

4. Syntax

The program is run by

```
> bamroc(y, score.list, Amodel = 1, EqualA = F, initial.beta = 0,
           small.constant = 0, maxiter = 25, sub.maxiter = 4)
```

where

`y` are the data in a two column matrix, one for each population. See Section 5 for details.

`score.list` is a list with row and column scores, say `list(1:5,1:5)` or a vector with the number of rating categories for each margin, say `c(5,5)`.

Amodel is an integer from 1 to 8 indicating the type of association model.

The choices are; 1) Independence; 2) Quasi independence - one diagonal parameter; 3) Quasi independence - saturated diagonal 4) Linear by linear; 5) Linear by linear plus agreement - one diagonal parameter 6) Linear by linear plus agreement - saturated diagonal; 7) Quasi symmetry; 8) Saturated model.

EqualA is a flag to indicate if the populations are to have homogeneous association parameters.

initial.beta is a vector of initial values for the marginal parameters. By default initial values are generated by the program.

small.constant can be assigned a value that is added to zero cells for a few iterations (4 by default).

maxiter is the maximum number of iterations.

sub.maxiter is the maximum number of iterations to run with the value of **small.constant** added to empty cells.

5. Organization of the data

The data should be kept in a matrix with 2 columns, one for each population. By convention, S-PLUS converts a matrix into a column vector by stacking the columns one under the other. This means that the first index moves fastest. Take for example the data in Table 1. To enter these data into the program, create the matrix, say **y**, so that

```

> y
[,1] [,2]
[1,]    4    1
[2,]    4    1
[3,]    0    0
[4,]    0    0
[5,]    0    1
.
.
.
.
.
.
.
[21,]    0    0
[22,]    1    0
[23,]    0    1
[24,]    0    6
[25,]    0   27

```

This can be done with the command

```

y <- matrix(c(4, 4, 0, 0, 0, 1, 8, 2, 3, 0, 1, 3, 2, 1, 0,
            0, 1, 0, 1, 1, 0, 1, 0, 0, 0,
            1, 1, 0, 0, 1, 0, 2, 2, 2, 1, 1, 1, 1, 3, 3,
            2, 1, 1, 4, 6, 0, 0, 1, 6, 27), 25, 2)

```

to read in the data by columns. In the discussion, the row margins refer to modality 1 and the column margins to modality 2.

6. Output

The output from the program is an object with 16 components:

"theta"	"omega"	"score"	"information"
"W"	"M"	"y"	"n.ITER"
"cov.theta"	"AmodelName"	"EqualA"	"d"
"score.list"	"area"	"pi"	"cov.pi"

where

theta is the maximum likelihood estimate of the model parameter $\boldsymbol{\theta}$ if convergence is obtained.

omega is the corresponding matrix of log-linear parameters for each population.

score is the value of the score with respect to $\boldsymbol{\theta}$.

information is the value of the information with respect to $\boldsymbol{\theta}$.

W is the log-linear design matrix used.

M is the matrix to compute marginal probabilities from the data **y**.

y are the data.

n.iter is the number of iterations.

cov.theta is the estimated asymptotic covariance matrix of the $\boldsymbol{\theta}$ estimator.

AmodelName is the name of the association model. See Section 4 for a list of the available names.

EqualA is the value of the flag `EqualA`. If it is T then the association parameters are the same across population; if it is F then there are separate association parameters for the two populations.

d is the number of rating categories.

score.list is a list of the integer scores used for the rating scale.

area contains information on the areas under the curves for each modality.

pi are the fitted probability vectors.

cov.pi is the estimated asymptotic covariance matrix of the fitted probability vectors.

7. Utility functions

There are 2 utility functions available: `summary.bam` and `bn.curves`.

`summary.bam` takes the output from `bamroc` and prints; 1) the areas under the ROC curves; 2) goodness of fit statistics; 3) parameter estimates with standard errors and Wald statistics; 4) the value of the log-likelihood; and 5) the name of the association model used.

`bn.curves` takes the output from `bamroc` and plots the ROC curves.

8. Analysis of the neonatal example

The 3 models in the paper are fitted by the commands:

```
> bamroc(y,c(5,5),Amodel=1)
> bamroc(y,c(5,5),Amodel=4,EqualA=T)
> bamroc(y,c(5,5),Amodel=5,EqualA=T)
```

Suppose we keep the results from the third model by issuing

```
> ex1.baham <- bamroc(y,c(5,5),Amodel=5,EqualA=T)
```

then `summary.bam(ex1.bam)` gives

Area under the curve (AUC) estimate with standard
error for system I and II:

AI	seAI	AII	seAII	AI-AII	se(AI-AII)
0.8624	0.0367	0.8586	0.0373	0.0037	0.0384

X2	G2	df
29.888	26.737	34

	Parameter	Estimate	SE	Wald	pvalue
1	c11	-0.851	0.245	-3.468	0.001
2	c12	0.417	0.212	1.962	0.050
3	c13	0.891	0.234	3.803	0.000
4	c14	2.071	0.427	4.847	0.000
5	mu1	2.429	0.581	4.180	0.000
6	ksi1	0.688	0.255	2.700	0.007
7	c21	-0.685	0.234	-2.923	0.003
8	c22	0.406	0.210	1.936	0.053
9	c23	1.133	0.253	4.481	0.000
10	c24	1.955	0.385	5.082	0.000
11	mu2	1.983	0.446	4.450	0.000
12	ksi2	0.439	0.243	1.812	0.070
13	alpha1	0.295	0.104	2.839	0.005
14	alpha2	0.446	0.271	1.643	0.100

```
log-likelihood: -245.207
```

```
Association model: 5) Linear by linear plus  
agreement - one diagonal parameter
```

Model fitted with equal association parameters across populations.

We see that the difference between the areas under the curves is small and that the model fits adequately. However, the validity of the G2 statistic can be questioned since the data are sparse. The first 6 parameters apply to the row margins (modality 1) and the next 6 to the column margins (modality 2). The association parameters apply to both populations since the model was fitted under the assumption of homogeneous association.

By issuing `bn.curves(ex1.baham)` we get a plot of the ROC curves as in Figure 1.

9. Analysis of the simulation study data

Consider next the data from Section 7 in Lang and Aspelund (2001) shown here in Table 2.

The data can be entered table-by-table in the following way

```
y0 <- c(30,6,0,0,0,14,47,10,4,2,4,8,20,  
       11,12,2,1,3,12,10,0,0,0,0,4)  
  
y1 <- c(3,0,1,0,0,6,12,0,1,1,5,2,15,11,  
       23,0,2,4,15,40,0,0,1,4,54)  
  
sy <- matrix(c(y0,y1),ncol=2)
```

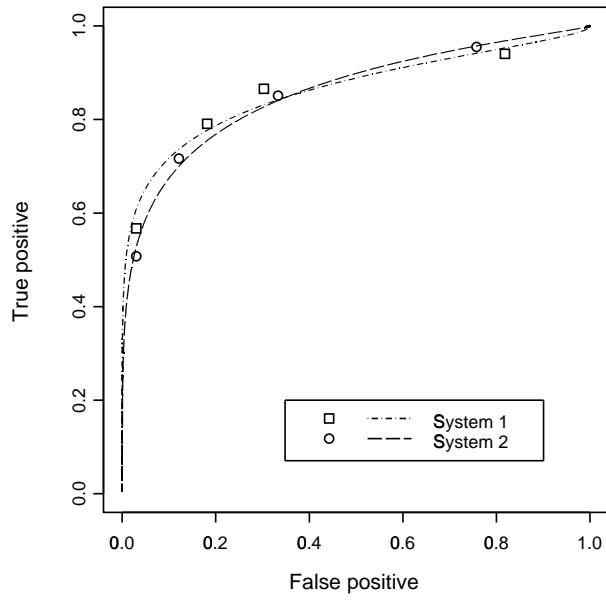


Figure 1. Fitted ROC curves for the 2 modalities.

Table 2
Simulated BAM data. D 1 and D 2 is short for Diagnostic 1 and 2.

D 2						D 2						
	1	2	3	4	5		1	2	3	4	5	
D 1	1	30	14	4	2	0	50	1	3	6	5	0
	2	6	47	8	1	0	62	2	0	12	2	2
	3	0	10	20	3	0	33	3	1	0	15	4
	4	0	4	11	12	0	27	4	0	1	11	15
	5	0	2	12	10	4	28	5	0	1	23	40
						200	4	20	56	61	59	200
						Signal-Absent						
						Signal-Present						

Notice how the contingency tables are read into S-PLUS by columns. Again, margin 1 refers to modality 1 and margin 2 to modality 2.

The models are fitted and the results stored by submitting, for example,

```
ex2.baim <- bamroc(sy,c(5,5),Amodel=1)
ex2.bahum <- bamroc(sy,c(5,5),Amodel=4,EqualA=T)
ex2.baham <- bamroc(sy,c(5,5),Amodel=5,EqualA=T)
```

To view the results, submit

```
summary.bam(ex2.baim)
summary.bam(ex2.bahum)
summary.bam(ex2.baham)
```

and perhaps bn.curves(ex2.baham) to view the ROC curves.

REFERENCES

- Agresti, A. (1988). A model for agreement between ratings on an ordinal scale. *Biometrics* **44**, 539–548.
- Dorfman, D. D. and Alf, Jr., E. (1969). Maximum-likelihood estimation of parameters of signal-detection theory and determination of confidence intervals—rating method data. *Journal of Mathematical Psychology* **6**, 487–496.
- Franken, Jr., E. A., Berbaum, K. S., Marley, S. M., Smith, W. L., Sato, Y., Kao, C. S. and Milam, S. G. (1992). Evaluation of a digital workstation for interpreting neonatal examinations: A receiver operating characteristic study. *Investigative Radiology* **27**, 732–737.

Lang, J. B. and Aspelund, T. (2001). Binormal association-marginal models for empirically evaluating and comparing diagnostics. *Statistical Modelling Forthcoming*.

S-PLUS (2000). *User's Guide*. Data Analysis Products Division, MathSoft, Seattle, WA.