EXAMPLES OF PARAMETRIC EMPIRICAL BAYES METHODS FOR THE ESTIMATION OF FAILURE PROCESSES FOR REPAIRABLE SYSTEMS

Asit P. Basu
Department of Statistics
University of Missouri
Columbia, Missouri 65211

Steven E. Rigdon
Mound Facility*
P. O. Box 32
Miamisburgh, Ohio 45342

* Mound Facility is operated by Monsanto Research Corporation for the U. S. Department of Energy under Contract No. DE-AC04-76-DP00053.

Abstract

Suppose that several systems, e.g. airplanes, computers, etc., are operating independently and are subject to failure and repair. The failure times are assumed to have been generated by a stochastic point process, and the repair times are assumed to be negligible. In this paper the homogeneous Poisson process is assumed as the stochastic point process. If the systems are similar, for example, they might have been produced in the same factory in the same time period, then it seems appropriate that information from all systems could be used to estimate the failure intensities of one particular system. Here, we treat the failure intensities of the systems as random quantities which were generated from a gamma or a log-normal prior distribution whose parameters are considered fixed but unknown. The marginal maximum likelihood estimators of the parameters of the prior distribution are found by using numerical techniques. Point estimates for the failure intensities are then obtained by finding the mean of the posterior distribution given the estimated values of the prior distribution.

Key Words: Empirical Bayes, Repairable system, EM algorithm, Poisson process.

1. Introduction

Suppose that N similar systems, e.g. airplanes, computers, etc., are operating independently and are subject to failure at random points in time. Repair time is assumed to be negligible. In most cases, the parameter(s) of the stochastic point process are unknown and must be estimated from the past history of the systems. The classical approach would be to treat this problem as N separate problems. That is, the parameters of the stochastic point process for system i, say, are estimated using only the data gathered from that system. However, if the systems are similar, manufactured in the same factory in the same time period, say, then it seems appropriate that data from all the systems could be used to estimate the parameters of system i. To this end, we assume that the parameters of the N stochastic point processes make up a random sample from some prior distribution. If this prior distribution is known completely, then a Bayesian analysis is appropriate and straightforward. If the prior distrubution is not known, one of three approaches may be taken. First, a parametric form may be chosen for the prior distribution, with the parameters of this distribution considered fixed but unknown quantities. This leads to a parametric empirical Bayes (PEB) model. Most work on PEB models involves the normal distribution; see Morris (1983). Second, there may be no specific parametric form chosen for the prior. This leads to a nonparametric empirical Bayes model similar to the models introduced by Robbins (1955). For a third approach, a parametric form may be chosen for the prior, and the parameters of this distribution are themselves considered random quantities with a completely specified prior distribution. In this paper, the first approach, the PEB approach, will be the one taken. For simplicity, the homogeneous Poisson process will be used to model the failure times. Martz (1975) gives a nonparametric empirical Bayes analysis for this same type of problem. Here, the prior distribution is taken to be the gamma distribution or the log-normal distribution. The values of the parameters which maximize the marginal likelihood function, called the marginal maximum likelihood estimates (MMLE's), are found by using numerical methods. For the gamma prior, the MMLE's are obtained by applying the Newton-Raphson algorithm. For the log-normal prior, the MMLE's are found by applying the EM algorithm of Dempster, Laird and Rubin (1977). Once the MMLE's are obtained, point estimates for the failure intensities are obtained by taking the posterior expectations given the MMLE's. Approximate posterior probability intervals can also be calculated.

In Section 2, the details are shown for the case of the gamma prior distribution. The details for the log-normal prior distribution are shown in Section 3. In Section 4, two sets of data are discussed. One involves the times between failures of the airconditioning equipment in thirteen Boeing 720 aircraft. The other set of data involves the number of entries between errors for five bookkeepers. This set of data was first discussed by Davis (1952).

2. Gamma Prior Distribution

In this section, the details of the PEB model are presented when the stochastic point process is assumed to be the homogeneous Poisson process and the prior distribution of the failure intensities is the gamma distribution. Let $\lambda_1, \lambda_2, \ldots, \lambda_N$ be the failure intensities of the N systems. The density of the gamma prior is then

$$p(\lambda|\alpha,\theta) = (\theta^\alpha/\Gamma(\alpha))\, \lambda^{\alpha-1} \exp(-\theta\lambda).$$

Let $t_{i,1} < t_{i,2} < \ldots < t_{i,n(i)}$ be the observed failure times of system i. We assume that it was predetermined that system i would be observed until n(i) failures occurred. This is called failure truncation. It is analagous to failure, or Type II censoring when referring to nonrepairable systems. Let $\underline{t}_i = (t_{i,1}, t_{i,2}, \ldots, t_{i,n(i)})$ and $\underline{t} = (\underline{t}_1, \underline{t}_2, \ldots, \underline{t}_i)$. Since the distributions of times between failures for each system are independent and exponentially distributed, and since the systems operate independently, the density of $\underline{t}$ given λ is

$$p(\underline{t}|\underline{\lambda}) = \prod_{i=1}^{N} \lambda_i^{n(i)} \exp(-\lambda_i t_{i,n(i)}). \tag{1}$$

The density of $\underline{\lambda}$ given α and θ is then

$$p(\underline{\lambda}|\alpha,\theta) = (\theta^\alpha/\Gamma(\alpha))^N \prod_{i=1}^{N} \lambda_i \quad \exp(-\lambda\theta_i). \tag{2}$$

At this point it should be made clear that the λ's are unobservable random variables. Equations (1) and (2) were derived so that the density of $\underline{t}$ given α and θ may be obtained, since inference for α and θ must be based solely on $\underline{t}$. This yields

$$p(\underline{t}|\alpha,\theta) = \int p(\underline{t},\underline{\lambda}|\alpha,\theta) \; d\underline{\lambda}$$

$$= \int p(\underline{t}|\underline{\lambda}) \; p(\underline{\lambda}|\alpha,\theta) \; d\underline{\lambda}$$

$$= (\theta^{\alpha}/\Gamma(\alpha))^{N} \prod_{i=1}^{N} \int_{0}^{\infty} \lambda_i^{n(i)+\alpha-1} \; \exp[\, -(\theta + t_{i,n(i)}) \; \lambda_i] \; d\lambda_i \; . \qquad (3)$$

Once estimates for α and θ are obtained, point estimates for the λ's can be based on the posterior distribution of λ_i given the estimated values of α and θ. We propose that the marginal maximum likelihood estimates (MMLE's) be used as point estimates for the failure intensities. The logarithm of the marginal likelihood function, which is the function that needs to be maximized, is then found to be

$$\log p(\underline{t}|\alpha,\theta) = N \alpha \log \theta - N \log \Gamma(\alpha) + \sum_{i=1}^{N} \log \Gamma(n(i) + \alpha) - \sum_{i=1}^{N} (n(i) + \alpha) \log (\theta + t_{i,n(i)}). \qquad (4)$$

Given observed data $\underline{t}$, (4) can be maximized by applying the Newton-Raphson algorithm. Algorithms for evaluating the digamma (phi) function $\psi(x) = (d/dx) \log(\Gamma(x))$ and the trigamma function $\psi'(x)$ can be found in Bernardo (1976) and Schneider (1978), respectively.

Once the MMLE's $\hat{\alpha}$ and $\hat{\theta}$ are obtained, point estimates for the λ's can be obtained by taking the posterior expectations. The posterior distribution of λ_i given $\underline{t}$, $\hat{\alpha}$ and $\hat{\theta}$ is the gamma distribution with density

$$p(\lambda_i|\underline{t},\hat{\alpha},\hat{\theta})$$

$$= [\ \Gamma(n(i) + \hat{\alpha}) \quad (\hat{\theta} + t_{i,n(i)})^{n(i)+\hat{\alpha}}\ ]^{-1}\ \lambda_i^{n(i)+\hat{\alpha}-1} \quad \exp[-(\hat{\theta} + t_{i,n(i)})\ \lambda_i].$$

Thus all posterior moments have closed form expressions. In particular, the posterior mean, which is used as the point estimate for λ_i is

$$\hat{\lambda}_i = \ (n(i) + \hat{\alpha}) \quad (\hat{\theta} + t_{i,n(i)}).$$

Estimated posterior probability intervals can be evaluated since $2\ \lambda_i\ (t_{i,n(i)}+\hat{\theta})$ has a chi-square distribution with $2(n(i)+\hat{\alpha})$ degrees of freedom. Many computer routines that compute the tabled values of the chi-square distribution will work correctly even for noninteger degrees of freedom.

3. Log-Normal Prior Distribution

In this section the details of the PEB model are presented when the prior distribution is the log-normal distribution. The reason that this section is included is so that the effect of the particular parametric form of the prior can be determined. In section 4, where two examples are discussed, it will be seen that this effect is almost negligible. We will continue to use the same notation as in the previous section.

The density of the prior is now

$$p(\lambda|\mu,\sigma) = (2\pi\lambda^2\sigma^2)^{-1/2} \quad \exp[-(\log \lambda_i - \mu)^2/2\sigma^2]. \tag{5}$$

The density of $\underline{t}$ given μ and σ is then found to be

$$p(\underline{t}|\mu,\sigma) = (2\pi\sigma^2)^{-N/2} \tag{6}$$

$$\times \prod_{i=1}^{N} \int_0^\infty \lambda_i^{n(i)-1} \exp[-(\log \lambda_i - \mu)^2/2\sigma^2 - \lambda_i t_{i,n(i)}] \, d\lambda_i .$$

This expression, or the logarithm of this expression, must be maximized. Here, however, evaluation of the log marginal likelihood function invloves evaluating a sum of improper integrals. This makes direct maximization of the log marginal likelihood function more difficult. If we consider $\underline{\lambda}$ to be missing data from the complete data $(\underline{t},\underline{\lambda})$, then the EM algorithm of Dempster, Laird and Rubin (1977) may be applied. In this case, we must use the extended form of the EM algorithm since the distribution of the complete data does not belong to the exponential family of distributions. The EM algorithm consists of alternating between the E (expectation) step and the M (maximization) step. Initially, a guess (μ_0,σ_0) for the MMLE of (μ,σ) is chosen, and k is set equal to 1. The E step consists of evaluating

$$E[\ \log p(\underline{t},\underline{\lambda}\ |\ \mu,\sigma\ )\ |\ \mu_{k-1},\sigma_{k-1},\underline{t}\ ] \tag{7}$$

presumably, for all values of (μ,σ) . The M step involves finding that value of (μ,σ) which maximizes (7) . (In practice (7) is not evaluated at all values of (μ,σ) ; it is maximized directly, in the M step. Thus, the extended form of the EM algorithm involves, in essense, only an M step.) The value of (μ,σ) which maximizes (7) then becomes (μ_k,σ_k) , k is incremented, and this process is repeated until convergence is attained.

Because of an error in Theorem 2 in Dempster, Laird and Rubin (1977), it is not guaranteed that this iterative process converges, nor that it converges to the global maximum if indeed it does converge. For a further discussion of this problem, see Wu (1983) and Boyles (1983). Despite this, the estimates of (μ,σ) and the point estimates of the λ's, obtained from applying the EM algorithm have been reasonable in the cases examined so far.

Improper integrals still need to be evaluated, even though the EM algorithm is applied. These can be numerically evaluated using Gauss-Laguerre quadrature formulas; see Stroud and Secrest (1966). Whereas the posterior moments of the λ's, given the MMLE's of the prior parameters, have closed form expressions when the prior is the gamma distribution, such is not the case here. The posterior moments must be evaluated using numerical integration techniques such as Gauss-Laguerre quadrature. Approximate posterior probability intervals for the λ's can be found by finding the posterior standard deviation and treating the posterior distribution as if it were a normal distribution.

4. Examples

To illustrate the PEB approaches to estimating failure intensities of several repairable systems, we will use two sets of data. The first is the well known set of data on the operating times between failures of the airconditioning equipment in thirteen Boeing 720 aircraft. This data was presented and first discussed by Proschan (1963). The second set of data contains the number of entries between errors for five bookkeepers. This set of data was given in a paper by Davis (1952). While the random variable which governs the number of entries between errors is actually discrete (the geometric distribution), the numbers are large enough so that the geometric distribution can be approximated by the exponential distribution. The goal of this section is to illustrate the PEB approaches to estimating failure intensities and to compare the results obtained by choosing two different forms for the prior. The goal is not to reanalyze these particular sets of data.

4.1 Aircraft Airconditioning Equipment

The collection of times between failures of the airconditioning equipment of the thirteen aircraft, is given in Table 1. As mentioned previously, the Newton-Raphson algorithm was used to find the MMLE's of α and θ (the parameters of the gamma prior distribution). This yielded the following estimates:

$$\hat{\alpha} = 18.41$$

$$\hat{\theta} = 1733.10 .$$

The EM algorithm was used to find the MMLE's of μ and σ^2 when the prior distribution was the log-normal distribution. The EM algorithm yielded the following estimates:

$$\hat{\mu} = -4.57$$

$$\hat{\sigma}^2 = 0.0522 .$$

Classical MLE's and PEB point estimates of failures per 1000 hours (1000 λ_i) are shown in Table 2. Also in Table 2 are the classical confidence intervals, the posterior probability intervals when the gamma distribution is the prior, and the approximate posterior probability intervals when the log-normal distribution is the prior. It can be seen that the PEB point and interval estimates obtained from the gamma prior and the log-normal prior are nearly identical. It can also be seen that the PEB point estimates of the failure intensities are considerably less disperse than the classical MLE's. The most extreme case of this "shrinkage" effect is for aircraft number 11 which experienced only two failures. Here, there is little information available from aircraft 11. Thus more information is "borrowed" from the collection of aircraft. In general, the more data available for a particular system, the less the discrepancy between the classical and the PEB estimates. The interval estimates using the PEB method are somewhat narrower than the classical interval estimates. This is the result of the added assumption concerning the prior distribution for the λ's.

4.2 Entries Between Errors for Bookkeepers

Now we discuss the data given by Davis (1952) on the number of entries between errors for five bookkeepers. The entire set of data is somewhat voluminous, so it will not be reproduced here. It is however summarized in Table 3. The MMLE's of α and θ were found to be

$$\hat{\alpha} = 4.074$$

$$\hat{\theta} = 1597.44$$

and the MMLE's of μ and σ^2 were found to be

$$\hat{\mu} = -6.10$$

$$\hat{\sigma}^2 = 0.240.$$

The classical and PEB point and interval estimates of the "error" intensities are shown in Table 3. Again, the PEB estimates obtained from the gamma prior and the log-normal prior are nearly identical. For this set of data, the classical and the PEB point and interval estimates are approximately the same. This appears to be due to there being so few systems (bookkeepers) and so many observations for each.

5. Conclusions

We have discussed the problem of estimating the parameters of several stochastic processes simultaneously. The homogeneous Poisson process was assumed throughout much of the paper. The failure intensities of these processes were assumed to make up a random sample from some parametric prior distribution whose parameters are condidered fixed but unknown quantities. This is commonly called a parametric empirical Bayes (PEB) model. The details were presented for the cases of the gamma prior and the log-normal prior. Within this PEB framework, we have obtained point and approximate interval estimates for the failure intensities. The point and interval estimates obtained from the gamma prior and the log-normal prior distributions were found to be nearly identical in the examples considered. We believe that the results would be similar for most two parameter prior distributions. In general, the PEB point estimates are less disperse than the classical MLE's. When there are few systems and many observed failure times per system, the point estimates are only slightly less disperse. PEB posterior probability interval estimates are generally narrower than the classical confidence intervals.

The problem addressed in this paper appears to be somewhat common in statistics. That is, it is clear that the parametrers of interest were generated from some prior distribution. This prior distribution is often unknown however. In this paper we have described how to find point estimates and approximate interval estimates of the parameters of interest in one particular statistical problem, the problem of simultaneously estimating the failure intensities of several systems.

REFERENCES

:rnardo, J. M. (1976). Algorithm AS 103, Psi (digamma) function, Journal of the Royal Statistical Society, Series C, 25, 315-317.

›yles, R. A. (1983). On the convergence of the EM algorithm, Journal of the Royal Statistical Society, Series B, 45, 47-50.

avis, D. (1952). An analysis of some failure data, Journal of the American Statistical Association, 47, 113-134.

empster, A. P., N. Laird, and D. B. Rubin (1977). Maximum likelihood estimation from incomplete data via the EM Algorithm, with discussion, Journal of the Royal Statistical Society, Series B, 39, 1-38.

iartz, H. F. (1975). Pooling life test data by means of the empirical Bayes method, IEEE Transactions on Reliability, R-24, 27-30.

lorris, C. N. (1983). Parametric empirical Bayes inference: theory and applications, with discusion, Journal of the American Statistical Association, 78, 47-65.

'roschan, F. (1963). Theoretical explanation of observed decreasing failure rate, Technometrics, 5, 375-383.

lobbins, H. (1955). An empirical Bayes approach to statitics, Proceedings of the third Berkeley symposium on probability and mathematical statistics, 1, 157-163.

Schneider, B. E. (1978). Algorithm AS 121, Trigamma function, Journal of the Royal Statistical Society, Series C, 27, 97-99.

Stroud, A. and D. Secrest (1966). Gaussian quadrature formulas, Prentice-Hall: Englewood Cliffs.

Wu, C. F. (1983). On the convergence properties of the EM algorithm. Annals of Statistics, 11, 95-103.

TABLE 1

OPERATING TIME BETWEEN FAILURES OF AIRCONDITIONING

EQUIPMENT IN THIRTEEN BOEING 720 AIRCRAFT

Plane Number

1	2	3	4	5	6	7	8	9	10	11	12	13
194	413	90	74	55	23	97	50	359	50	130	487	102
15	14	10	57	320	261	51	44	9	254	493	18	209
41	58	60	48	56	87	11	102	12	5		100	14
29	37	186	29	104	7	4	72	270	283		7	57
33	100	61	502	220	120	141	22	603	35		98	54
181	65	49	12	239	14	18	39	3	12		5	32
	9	14	70	47	62	142	3	104			85	67
	169	24	21	246	47	68	15	2			91	59
	447	56	29	176	225	77	197	438			43	134
	184	20	386	182	71	80	188				230	152
	36	79	59	33	246	1	79				3	27
	201	84	27	15	21	16	88				130	14
	118	44	153	104	42	106	46					230
	34	59	26	35	20	206	5					66
	31	29	326		5	82	5					61
	18	118			12	54	36					34
	18	25			120	31	22					
	67	156			11	216	139					
	57	310			3	46	210					
	62	76			14	111	97					
	7	26			71	39	30					
	22	44			11	63	23					
	34	23			14	18	13					
		62			11	191	14					
		130			16	18						
		208			90	168						
		70			1	24						
		101			16							
		208			52							
					95							

TABLE 2

CLASSICAL AND PEB POINT AND INTERVAL

ESTIMATES OF FAILURES PER 1000 HOURS

FOR AIRCRAFT AIRCONDITIONING DATA

Plane No.	No. of Fail.	Total Time	Classical MLE	Classical 95% C.I.	Gamma Pt.Est.	Gamma 95% P.I.	Log-Normal Pt.Est.	Log-Normal 95% P.I.
1	6	493	12.17	(4.46,23.67)	10.97	(7.06,15.72)	10.92	(6.58,15.27)
2	23	2201	10.45	(6.62,15.07)	10.53	(7.57,13.97)	10.49	(7.33,13.65)
3	29	2422	11.97	(8.02,16.71)	11.41	(8.40,15.12)	11.33	(8.11,14.55)
4	15	1819	8.25	(4.62,12.91)	9.41	(6.49,12.85)	9.41	(6.34,12.49)
5	14	1832	7.64	(4.18,12.13)	9.09	(6.23,12.48)	9.12	(6.13,12.12)
6	30	1788	16.78	(11.32,23.29)	13.75	(10.15,17.88)	13.84	(9.78,17.90)
7	27	2074	13.02	(8.58,18.37)	11.93	(8.71,15.64)	11.89	(8.40,15.38)
8	24	1539	15.59	(9.99,22.42)	12.96	(9.36,17.14)	12.99	(8.89,17.09)
9	9	1800	5.00	(2.29, 8.76)	7.76	(5.13,10.92)	7.96	(5.26,10.67)
10	6	639	9.39	(3.44,18.26)	10.29	(6.62,14.76)	10.27	(6.28,14.25)
11	2	623	3.21	(0.39, 8.94)	8.66	(5.32,12.81)	8.81	(5.31,12.32)
12	12	1297	9.25	(4.78,15.17)	10.04	(6.79,13.90)	10.01	(6.54,13.47)
13	16	1312	12.20	(6.97,18.86)	11.30	(7.84,15.38)	11.24	(7.44,15.03)

TABLE 3

CLASSICAL AND PEB POINT AND

INTERVAL ESTIMATES OF ERRORS

PER 1000 ENTRIES FOR FIVE BOOKKEEPERS

Bkpr	No. of	Total No.	Classical		Gamma		Log-Normal	
No.	Errors	Entries	MLE	95% C.I.	Pt.Est.	95% P.I.	Pt.Est.	95% P.I.
1	31	17,991	1.72	(1.17,2.38)	1.79	(1.25,2.43)	1.78	(1.20,2.36)
2	26	17,533	1.48	(0.97,2.10)	1.57	(1.06,2.18)	1.57	(1.03,2.11)
3	26	18,742	1.39	(0.91,1.97)	1.48	(1.00,2.05)	1.48	(0.97,1.99)
4	54	18,273	2.96	(2.22,3.79)	2.92	(2.22,3.72)	2.90	(2.15,3.65)
5	81	15,446	5.24	(4.16,6.45)	5.00	(3.99,6.11)	5.03	(3.94,6.12)