THE ELICITATION AND USE OF EXPERT JUDGMENT FOR MAINTENANCE OPTIMIZATION

T. A. Mazzuchi

The George Washington University, Washington, D. C., USA

R. M. Cooke, J. R. van Dorp, and J. M. van Noortwijk

Delft University of Technology, Delft, The Netherlands

R. Dekker

Koninklijke/Shell-Laboratorium, Amsterdam (Shell Research B.V.), The Netherlands

ABSTRACT

Herein, we present an overview of the elicitation and use of expert opinion in developing optimal maintenance policies. The procedure developed is based on restrictions found in practice. That is, where the "expert" has little statistical training and the elicitation process most be performed in a clear and quick manner. Due to these restrictions, a histogram is elicited from the expert and feedback and analysis is based on combining the elicited histogram with a fitted right tail to form a continuous distribution. Expressions for the the pdf, failure rate, percentile life, and mean life are developed and used to calculate the optimal maintenance interval for given cost data.

INTRODUCTION

The topic of maintenance optimization has been a focus of research interest for some time, however, few actual applications exist in practice (see for example Dekker [3]). Often the main bottleneck in the implementation of maintenance optimization procedures is the determination of the life length distributions of the system components. The problem is that, due to the scarcity of good component failure data, the determination of such distributions via known statistical estimation procedures is, in many cases, impossible. Scarcity of failure data is inherent to a preventive maintenance environment primarily because preventively maintaining components implies that the complete component life cycle will rarely be observed. The occurrence of many failures will in fact lead to equipment modification, making past data obsolete. In addition, the data that is available is often incomplete (with the cause of failure often unspecified) or of poor quality (with only an ambiguous description of the maintenance and failure of components).

One approach to overcome this scarcity of good data, is to determine the lifetime distributions based on the use expert opinion. However, defining a proper elicitation scheme for obtaining the expert opinion is also not an easy task. The use of expert opinion is not a new topic (see e. g., Cooke [2], Mosleh, Bier, and Apostolakis [6], Singpurwalla [10], Spetzler and Staël von Holstein [5], and Wallsten and Budescu [13]), but the application of expert opinion in a maintenance environment is new and presents a different set of constraints governing its application. Herein we present an overview of an elicitation procedure presented in van Noortwijk, Dekker, Cooke, and Mazzuchi [7] and show how it may be implemented and used to provide statistical quantities of interest for feedback and analysis purposes. We present the analysis for the case of a single expert and single component. The problem of combining and updating the expert opinion is developed in [7].

THE ELICITATION PROCEDURE

The selection of an appropriate procedure for eliciting expert subjective probabilities must take into account the nature of the information required and the background and training of the experts. For the problem at hand, the required information is the failure distribution of a component which displays aging (nonconstant) and revealed failure behavior. It is only by attempting to model this aging behavior that a time-based preventive maintenance policy can be justified. Most experts have little experience with failure models and so requiring them to select an appropriate parametric family as the best representation of failure behavior is impossible. Furthermore, given the constraints on the experts' time and powers of concentration, the elicitation process should be kept as short as possible. In one elicitation session several component life time distributions (in addition to other pertinent questions relating to the effects of preventive maintenance and the consequences of failure etc.) must be assessed.

An extensive review of elicitation techniques can be found in Wallsten and Budescu [13] and Spetzler and Staël von Holstein [11]. However, none of the techniques discussed in these sources was specifically designed for the present problem characteristics, and thus none is directly applicable. Due to the lack of statistical training among the experts, it was felt that a discretised version of the continuous density function (i. e. a histogram) would be easier for the experts to comprehend since the concept of probability density would be replaced by the concept of the probability that a lifetime lies in a fixed time interval. A continuous distribution could then be fit to the discrete probability function obtained. Kabus [5] also reported success with using a histogram technique in predicting the interest rate of certificates of deposit.

The Histogram Technique

Let the domain of lifetimes $(0,\infty)$ be divided into m disjoint time intervals $(t_{i-1}, t_i]$, $i = 1, \dots, m-1$, and (t_{m-1}, t_m) where $0 \equiv t_0 < t_1 < \cdots < t_{m-1} < t_m \equiv \infty$. The experts need only specify their subjective probability of failure for the component in each interval. For convenience we denote time interval $(t_{i-1}, t_i]$ by time interval i, $i = 1, \dots, m-1$. Defining interval m, (t_{m-1}, ∞), as an open interval is motivated by the fact that maintenance engineers only have experiences only with the first part of a component life cycle since most components will be replaced before failing. A number of implementational issues still must be addressed to make the procedure more understandable for the expert. Specifically the representation of the interval failure probabilities, and selection of the t_i's and m.

Spetzler and Staël von Holstein [6], for example, have determined that indication of a probability of say 0.001 as "one in one thousand" is more clear for non-technical experts. Applying this concept to assess a histogram of lifetimes, the expert is asked to imagine that there are n components of the same type installed at time t_0 and requested to provide his expected number, n_i, of components which will fail within time interval i. From this, the expert's (subjective) probability of failure within time interval i is easily calculated as $p_{ie} = n_{ie}/n$, $i = 1, \dots, m$. From an accuracy view point it is recommended to choose n as large as possible, however, from a realistic point of view it is better to avoid large n values, so as not to require an accuracy from the expert which cannot be achieved. Furthermore, it would be intuitively appealing to define n=100 and thus the expert is dealing with percentages. Thus we choose n = 100.

For the choice of the number of time intervals, again there is a conflict between the desired and the attainable accuracy. We note that if we eventually wish to fit the elicited distribution to a parametric form with two unknown parameters, we require at least three time intervals. Except for this restriction then, we allow the expert to specify as many intervals as he feels comfortable.

The choice of $t_1, \dots, t_{m-1}$ under the ordering restriction does not influence the developed mathematical model; however, for better perception it is recommended to let these values be equidistant. If the t_i, $i = 1, \dots, m-1$, are not equidistant, the histogram would not clearly display the failure behavior of a particular component type and this may in fact confuse the experts and introduce extra bias. We can take for example $t_i = i*x'$, $i = 0, \dots, m-1$, where x' (or 2x') may represent a maintenance interval used in the past. This is advantageous in that the maintenance expert's familiarity with x' enables the expert to assess the probabilities by a comparison of known maintenance intervals. A disadvantage could be that if m is small, the elicited distribution may have a large tail.

In order to enhance the speed and accuracy of the elicitation process, an interactive PC-based program was developed for histogram elicitation. Using only the cursor controls, the expert generates a histogram display by iteratively increasing (or decreasing) the number of time intervals and defining the (subjective) number of expected failures in each interval. The expert receives on-line visual feedback on his assessment. Providing (mathematically) redundant feedback often prompts an expert to reassess his subjective distribution and helps reduce the experts' ambiguity as to the consequences of his assessment. Various forms of feedback are considered in the next section. In addition to providing feedback, these quantities may be of use/interest in their own right for providing analytical results.

ANALYSIS AND FEEDBACK

Feedback in general can be divided in two categories, outcome feedback and task feedback. Outcome feedback is concerned with the outcome of the previously predicted event, whereas task feedback emphasizes the relationships between the cues in the environment and the variable to be predicted. Within a maintenance context, examples of outcome feedback are the histogram of lifetimes, estimated failure density, percentile life, mean time to failure and the component failure rate function. An example of task feedback is the optimal maintenance interval. Other types of feedback are possible as well.

Obtaining a Continuous Distribution for Component Life Length

Providing the histogram for the elicited expert distributions is, of course, straightforward. However, the other forms of feedback require the specification of the right tail of the failure density. Van Noortwijk [8] suggested fitting a specified continuous parametric distribution with cdf $F(t \mid \underset{\sim}{\theta})$ to the discrete lifetime distribution using the idea of minimizing the relative information between the elicited and fitted histogram probabilities (see Cooke [2]) to obtain the unknown parameter values $\underset{\sim}{\theta}$. The choice of the parametric family would be based on the appropriateness of that family for the component failure distribution. A usual assumption in reliability is that the life lengths are described adequately by the Weibull distribution.

Another approach is to use as much information as possible from the discretised distribution. In considering the appropriateness of any continuous distribution, a particular problem that may arise is that the failure rate of most of the known parametric distributions used in reliability is monotonic, while the estimated discrete failure rate may not be monotonic. Thus we consider fitting a continuous parametric distribution using only the part of the discrete distribution which exhibits a monotonic failure rate (denoted MFR). To determine the point of MFR we use the *interval failure rate*,

$$h_i = \frac{p_i}{(t_i - t_{i-1}) \sum_{j=i}^{m} p_j} \qquad i = 1, \ldots, m \tag{1}$$

which represents the probability of failure in the ith interval given survival to time t_{i-1} normalized by the interval length. If a value ℓ exists such that $\{h_i,\ i = \ell, \ldots, m\}$ is a nondecreasing (or nonincreasing) sequence then the interval endpoint $t_{\ell-1}$ is assumed to be the start of the MFR part of the continuous distribution. Note that ℓ must be less than $m-k$ to fit a continuous distribution with k parameters, i.e. if $\mathrm{DIM}(\underset{\sim}{\theta}) = k$. If the MFR part of the discrete distribution begins in the ℓth interval, the continuous distribution is fit using intervals $\ell, \ldots, m$ again using the idea of minimum relative information. To preserve as much of the elicited information as possible, the continuous distribution will only be used for time $t \geq t_{m-1}$. Thus in determining $\underset{\sim}{\theta}$, a constraint is added in order to guarantee continuity of the cdf at $t = t_{m-1}$ and $\underset{\sim}{\theta}$ is obtained as

$$\begin{aligned} &\underset{\underset{\sim}{\theta}}{\mathrm{MAX}} \left\{ \sum_{i=\ell}^{m} p_{iD} \, \mathrm{Ln}\left[F(t_i \mid \underset{\sim}{\theta}) - F(t_{i-1} \mid \underset{\sim}{\theta})\right] \right\} \\ &\text{s.t.} \quad F(t_{m-1} \mid \underset{\sim}{\theta}) = \sum_{j=1}^{m-1} p_j \end{aligned} \tag{2}$$

Note that, using the constraint, this is easily converted into a single variable optimization problem for the usual case of $\mathrm{DIM}(\underset{\sim}{\theta}) = 2$.

To make the distribution completely continuous we assume that within each time interval the probability of failure is uniformly distributed and thus the "continuized" distribution $f_c(t)$ is given as

$$f_c(t) = \begin{cases} 0 & t = 0 \\ \dfrac{p_i}{t_i - t_{i-1}} & t_{i-1} < t \le t_i,\ i = 1,\ldots, m-1 \\ f(t|\underset{\sim}{\theta}) & t \ge t_{m-1} \end{cases} \tag{3}$$

Mean Component Life

The expected lifetime can be calculated directly using (3) as

$$E[T] = \int_0^\infty f_c(t)dt = \sum_{i=1}^{m-1} p_i \frac{t_i + t_{i-1}}{2} + \int_{t_{m-1}}^{\infty} t\, f(t|\underset{\sim}{\theta})dt \tag{4}$$

where, for most cases, the integral must be evaluated numerically.

Component Percentile Life

The component percentile life (defined as t_p such that $F_c(t_p) = p$) can be obtained once the continuized cdf, $F_c(t)$ is defined,

$$F_c(t) = \begin{cases} 0 & t = 0 \\ \sum_{j=1}^{i-1} p_j + p_i \cdot \dfrac{t - t_{i-1}}{t_i - t_{i-1}} & t_{i-1} < t \le t_i,\ i = 1,\ldots, m-1 \\ F(t|\underset{\sim}{\theta}) & t \ge t_{m-1} \end{cases} \tag{5}$$

Two cases are possible:

Case i: $p > \sum_{j=1}^{m-1} p_j \Rightarrow t_p$ is given as the solution to: $F(t_p|\underset{\sim}{\theta}) = p$ (6)

Case ii: $\sum_{j=1}^{i-1} p_j < p \le \sum_{j=1}^{i} p_j \Rightarrow t_p$ is given by: $t_p = \dfrac{p - \sum_{j=1}^{i-1} p_j}{p_i}(t_i - t_{i-1}) + t_{i-1}$ (7)

Component Failure Rate Function

For a continuous pdf f(t) and cdf F(t), the failure rate function r(t) is defined as:

$$r(t) = \frac{f(t)}{1 - F(t)} \tag{8}$$

Note that the discontinuous nature of the interval failure rate makes it's use for visual feedback or analytical use doubtful. In addition, because $f_c(t)$ is not continuous for $t \le t_{m-1}$, discontinuity will also occur in the calculated continuized failure rate. A continuous approximation for the failure rate for the continuized distribution continuous failure rate of

$t < t_{m-1}$, we use the well known relationship (see for example Barlow and Proschan [1]) between the reliability and failure rate

$$1 - F(t) = \exp\left\{ - \int_0^t r(u)du \right\} \tag{9}$$

and approximating $\int_t^{t+\delta} r(u)du \approx r(t)\delta$ for suitable small value δ, we define

$$r(t) = \frac{\mathrm{Ln}\left(\dfrac{1 - F_c(t)}{1 - F_c(t+\delta)}\right)}{\delta} \tag{10}$$

and use this as an approximation to the continuous failure rate with $r(0) \equiv 0$.

The Optimal Maintenance Interval

A basic model within the context of maintenance optimization is the age replacement model (see Barlow and Proschan [1] for an extensive discussion). In the age replacement model, a maintenance activity is carried out at a prespecified age of the component and both a preventive maintenance action and a failure replacement serve to renew the life of the component. A "renewal cycle", is then defined as the time between two consecutive renewals (either preventive maintenance or a failure) of the component. Denoting the prespecified replacement age by x, the long term average cost using this age x, g(x), can be calculated from the expected costs during a renewal cycle (see the renewal reward theorem, Tijms [12]). This is given by

$$g(x) = \frac{(C_f - C_p)F(x) + C_p}{\int_0^x (1 - F(u))\, du} \tag{11}$$

where F is the cdf of the life length distribution and C_f and C_p are the expected costs associated with a failure of the component and a preventive replacement of the component respectively. The optimal maintenance interval will then be the time x for which g(x) is minimal. For the subsequent analysis it is assumed that the cost figures are given and furthermore that $C_f > C_p$, otherwise the optimal maintenance strategy would be "replace at failures only". Using the continuized cdf (5) we may obtain the optimal maintenance interval by solving a univariate optimization problem (a solution algorithm and associated theory is presented in van Dorp [4]).

EXAMPLE

Consider the following hypothetical example. We elicit the (subjective) number of failures from two experts based on a previous maintenance interval x' = 6 months, i.e. $t_i = i*6$, $i = 0, 1, \ldots, 4$. The results are given in Table I. We compare the associated feedback for the two experts. and fit a Weibull tail to both elicited histograms using (2) with $F(t \mid \underset{\sim}{\theta} = (\alpha,\beta)) = 1 - \exp\{-(t/\alpha)^{\beta}\}$. Both tails are fit from $\ell = 0$ since the interval failure rates (given in Table II) are increasing. The estimated cdf's (5) and continuous failure rates (10) are presented for comparison in Figures 1 and 2 respectively. The percentile life for the elicited distributions are presented in Table III. The calculated mean life (4) is 50.46 for expert 1 and 27.01 for expert 2.

TABLE I
Elicited Values for Expert Histograms

Expert	n_{1e}	n_{2e}	n_{3e}	n_{4e}	n_{5e}
1	2	2	4	8	84
2	4	4	12	16	64

TABLE II
Calculated Expert Interval Failure Rates

Expert	h_1	h_2	h_3	h_4	h_5
1	.0033	.0034	.0069	.0145	.1667
2	.0067	.0069	.0217	.0333	.1667

TABLE III
Calculated Expert Percent Life

Expert	p = .1	.2	.3	.4	.5	.6	.7	.8	.9
	19.50	27.30	34.86	42.04	49.30	57.02	65.75	76.49	92.20
	13.00	18.00	21.75	25.75	30.19	34.93	40.27	46.85	56.47

It is clear that expert 1 is more optimistic than expert 2 and this optimism is reflected in the optimal maintenance interval. Figure 3 displays plots of the long term average cost using (11) with $F(\cdot) = F_c(\cdot)$, $C_f = 100$ and $C_p = 10$. Note that the optimal maintenance interval is approximately 50% longer for expert 1 (that is, $x^* \approx 12$ months for expert 2 and $x^* \approx 18$ months for expert 1). In addition the estimated long term average cost is approximately 50% less for expert 1.

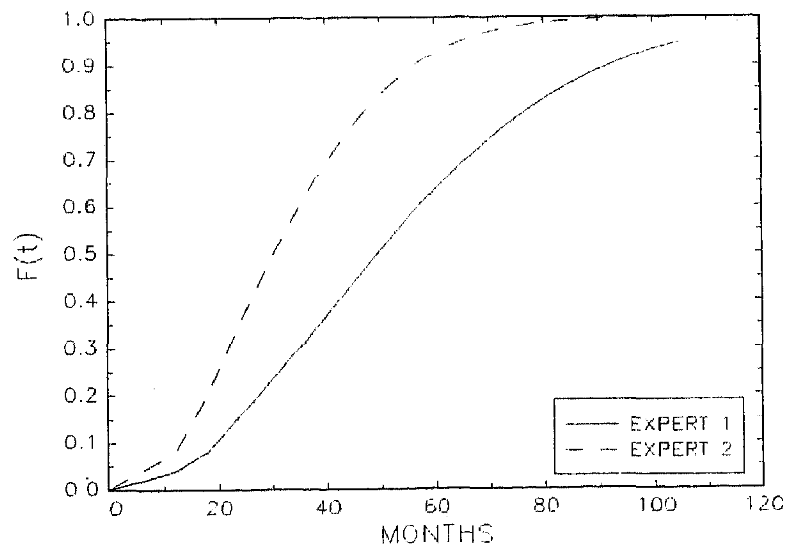

FIGURE 1. Comparison of Estimated Cumulative Distribution Functions.

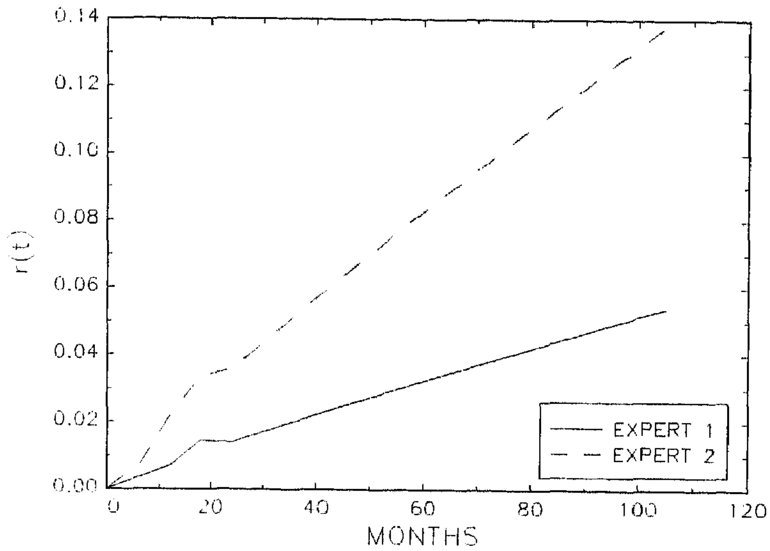

FIGURE 2. Comparison of Estimated Failure Rates.

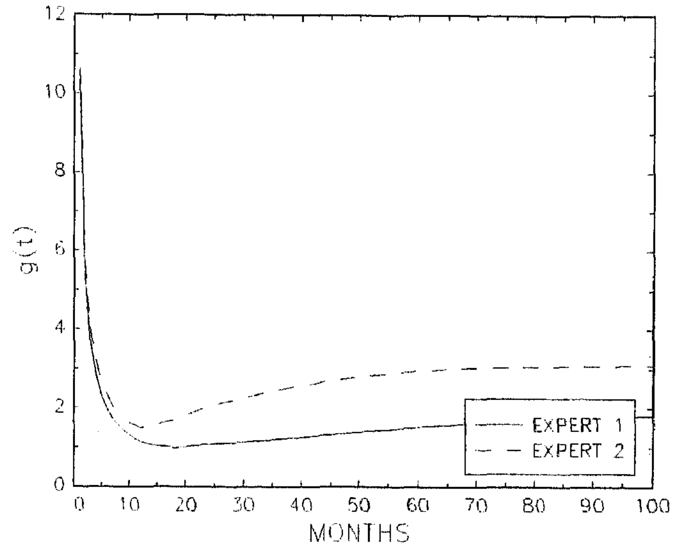

FIGURE 3. Comparison of Estimated Long Term Average Cost.

CLOSING COMMENTS

The most likely comment on the above approach is "why not just assume a continuous distribution or fit a continuous distribution to the histogram and use the fitted distribution for the analysis"? The answer is simple. We are dealing with experts with very little statistical experience and only partial information on the failure process. The selection of a "best" distribution is difficult and even if one was assumed, we cannot ask questions about the distribution parameters from experts who have little idea about their meaning. Even more common parameters like mean, median, percentile are commonly confused and miscommunicated by experts. We fit a right tail to our distribution because it is required for analysis. Yet the goal is to use as much of the elicited information as possible so that we are dealing with the expert opinion and not some assumed best distribution where best is arbitrarily described. The difference between our approach and just using a fitted distribution can best be seen from expert 2's results in the above example. The calculated mean value is 27.01 and the mean value for the fitted distribution is 31.25. Thus using the fitted distribution would provide a more optimistic view than presented by the expert. This would also be reflected in the determination of the optimal maintenance interval. Note also that for expert 1 and 2 the optimal maintenance interval was not in the fitted part of the distribution.

REFERENCES

[1] Barlow, R. E. and Proschan, F. (1981). Statistical theory of reliability and life testing, (2^{nd} Ed). To Begin With.

[2] Cooke, R. M. (1991). Experts in uncertainty: expert opinion and subjective probability in science. to be published, Oxford University Press.

[3] Dekker, R. (1991). Applications of maintenance optimization models. Koninklijke/Shell Laboratorium, Amsterdam, external report to be issued.

[4] van Dorp, J. R. (1989). Expert opinion and maintenance data to determine lifetime distributions. Master's Thesis, Delft University of Technology.

[5] Kabus, I. (1976). You can bank on uncertainty. Harvard Business Review, 56, 95-105.

[6] Mosleh, A., Bier, V. M., and Apostolakis, G. (1988). A critique of current practice for the use of expert opinions in probabilistic risk assessment. Reliability Engineering and System Safety, 20, 63-85.

[7] van Noortwijk J., Dekker, R., Cooke, R. M. and Mazzuchi, T. A. (1991). Expert judgment in maintenance optimization. to appear IEEE Transactions in Reliability.

[8] van Noortwijk, J. (1989). Use of expert opinion for maintenance optimization. Master's thesis, Delft University of Technology.

[10] Singpurwalla, N. D. (1988). An interactive P.C.-based procedure for reliability assessment incorporating expert opinion and survival data. Journal of the American Statistical Association, 83, 44-51.

[11] Spetzler, C. S., and Staël von Holstein, C. S. (1975). Probability encoding in decision analysis. Management Science, 22, 340-358.

[12] Tijms, H. C. (1986). Stochastic modelling and analysis: A computational approach. Wiley Series in Probability and Mathematical Statistics.

[13] Wallsten, T. S. and Budescu, D. V. (1983). Encoding subjective probabilities: a psychological and psychometric review. Management Science, 29, 151-173.