Reliability Engineering and System Safety 123 (2014) 47–56

Contents lists available at ScienceDirect

Reliability Engineering and System Safety

journal homepage: www.elsevier.com/locate/ress

Improving estimates of critical lower percentiles by induced censoring

David J. Edwards [a,*], Frank M. Guess [b], Ramón V. León [b], Timothy M. Young [c], Kevin A. Crookston [c]

[a] Department of Statistical Sciences & Operations Research, Virginia Commonwealth University, Richmond, VA 23284, USA
[b] Department of Statistics, Operations, & Management Science, University of Tennessee, Knoxville, TN 37996, USA
[c] Center for Renewable Carbon, University of Tennessee, Knoxville, TN 37996, USA

ARTICLE INFO

Article history:
Received 9 December 2011
Received in revised form
1 October 2013
Accepted 7 October 2013
Available online 23 October 2013

Keywords:
Aging period
Bathtub curve
Lower percentile estimation
Model misspecification
Reliability
Weibull

ABSTRACT

In this article, we present an approach based on induced censoring for improving the estimation of critical lower percentiles. We validate this technique via simulation results and practical industrial insights. Data from product components that have at least two aging periods (e.g., bathtub failure rate) is investigated. When such data are improperly fit by certain reliability distributions, estimates of lower percentiles are impacted by longer-lasting failures, resulting in larger root mean square errors (RMSE) and bias. In lieu of utilizing a more complex bathtub model, we propose induced right censoring of data at various points to substantially reduce RMSE and bias of lower percentile estimates. A technique for finding optimal or near optimal censoring points is discussed and two real world examples illustrate how this works in practice.

1. Introduction

As product complexity increases, customers require reassurance that products will perform adequately over time [1]. For such customer concerns, warranty issues, specification improvements, and/or liability, practitioners often need to know when to expect early product failures to occur by estimating critical lower percentiles. Reliability studies (or life tests) are conducted by placing samples of products on test and measuring the time (or strength) to failure for each unit, often under accelerated life tests (ALT). See, for example, [2] for details. Determining product reliability via ALTs can be costly. However, inadequate product reliability can result in even greater costs [3]. The Weibull distribution, among others, plays an important role in the modeling of ALT's (and other settings, such as strength until failure). Indeed, the researcher Weibull himself studied such data and subsequently developed the distribution that carries his namesake (see [4,5]). Compare also [6,7]. The probability density function of the Weibull distribution is given by

$$f(t) = \frac{\beta}{\eta}\left(\frac{t}{\eta}\right)^{\beta-1} e^{-(t/\eta)^{\beta}}, \quad (1)$$

where $t \geq 0$, $\beta > 0$ is the shape parameter, and $\eta > 0$ is the scale parameter.

Recall that the cumulative distribution function (cdf) gives the probability that a unit will fail before time *t*. In reliability studies, the derivation of the cdf is often determined through the use of the hazard function (or failure rate), $h(t) = f(t)/S(t)$, where the denominator $S(t)$ is the survival function (see the Appendix for more details). For the Weibull distribution, $h(t)$ is a decreasing function for $0 < \beta < 1$ and increasing for $\beta > 1$. In many applications, the hazard function consists of two or three intervals. The first time interval consists of early failures (infant mortality) with a decreasing hazard rate. During the second stage of usable life of a product, components will often fail at an approximately constant rate. Many integrated circuits have only these first two aging stages during screening, etc. Other components (e.g., those subject to loads) that are still in use by the end of the usable life interval will begin to wear out, and the rate of failures will increase again. This type of hazard function with all three aging stages is well known as the bathtub curve (see Fig. 1).

The flexibility of the Weibull distribution has allowed for it to be used to model any of the parts of the bathtub curve. However, a single Weibull model cannot be used to model all three phases of a bathtub curve at the same time. As a result, alternative modeling approaches for bathtub-shaped failure rates, based on generalizations of standard distributions, have been proposed. For instance, Hjorth [8] proposed the bathtub-shaped hazard rate function, $h(t) = \alpha(1+\beta t)^{-1} + 2\gamma t$, $0 \leq \gamma \leq \alpha\beta/2$, which is based

* Corresponding author. Tel.: +1 804 828 2936; fax: +1 804 828 8785.
E-mail addresses: dedwards7@vcu.edu (D.J. Edwards), fguess@utk.edu (F.M. Guess), rleon@utk.edu (R.V. León), tmyoung1@utk.edu (T.M. Young), kevincrookston@gmail.com (K.A. Crookston).

0951-8320/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.ress.2013.10.001

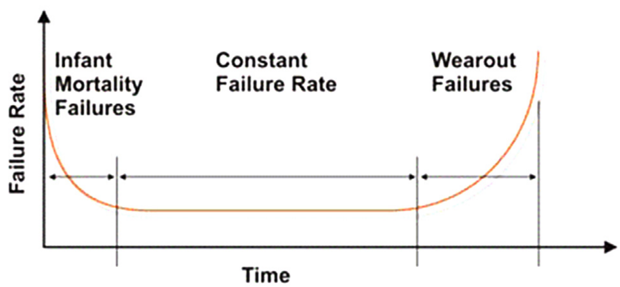

Fig. 1. Bathtub curve.

on the generalized Rayleigh family. Although it possesses a rather complex hazard function, Mudolkar and Srivastava [9] suggested the exponentiated Weibull distribution with cdf $F(t) = (1 - \exp[-(\lambda t)^{\beta}])^{\nu}$ to model bathtub-shaped failure data. For other proposed bathtub-shaped failure rate models, see [10–12], among others.

As stressed by [13], however, models for bathtub-shaped failure rates such as those mentioned above are often not practical for use by reliability engineers. As a consequence, [13] suggests using an additive (piecewise) Weibull model. In this case, the failure rate function is expressed as the sum of two failure rate functions of the Weibull form. That is,

$$h(t) = ac(at)^{b-1} + cd(ct)^{d-1} \tag{2}$$

with parameters $a, c \geq 0, \; b > 1, \; d < 1$. This approach offers a practical advantage to the aforementioned failure rate models given the simplicity of combining two Weibull distributions. See also [14,15] for other examples of additive models for bathtub-shaped failure rates.

The use of models such as those described previously for bathtub-shaped failure rates may indeed be necessary when the complete life cycle of a system is to be modeled. In practice (e.g., in real time manufacturing settings), however, one might consider only a portion of the bathtub curve to be relevant when interest lies with just a part of the lifetime (e.g., when considering early failure warranties). Such is the focus of this paper. In particular, this article will investigate the practice of induced *right* censoring of reliability data in order to reduce the effect of the long-term failure mechanism on estimates of critical lower percentiles. That is, if the problem of interest is the estimation of lower percentiles (e.g., 1st, 5th and 10th), this paper proposes to transform larger failure times into right-censored observations and subsequently fit a simple lifetime distribution to this transformed dataset. Doing so offers a practical alternative to the fitting of complex models with bathtub-shaped failure rates. Simulation results reveal practical improvements due to this technique as the lower tail of the data is more heavily weighted.

The rest of our paper is organized as follows. In Section 2, we describe the practice of induced censoring for the estimation of critical lower percentiles. Section 3 discusses in detail a simulation study, its results, and implications. Section 4 offers practical advice to the user as to when and how induced censoring should (or should not) be implemented. Section 5 illustrates the practice of induced censoring for estimation of lower percentiles using two real world data examples. We conclude the article in Section 6 with a brief discussion of our findings.

2. Lower percentiles and induced censoring

Lower percentiles are helpful for understanding early failures, specification limit improvement, warranty analysis, etc. Thus, in reliability studies, it is generally of high interest to obtain good estimates of these key parameters. See, for example, [16] which emphasizes the role of percentiles in system life optimization. Percentiles can be estimated using either parametric or nonparametric techniques. If a distribution can be identified that satisfactorily fits a given data set, parametric estimation is often recommended it is produces more precise percentile estimates versus those obtained via nonparametric methods [17]. As a consequence, in this article, we make use of the inverse cdf (or maximum likelihood) approach [2] for percentile estimation. It should be noted, however, that the technique of induced censoring could be fruitfully applied regardless of the method utilized for percentile estimation. For instance, if suitable prior information is available, Bayesian methods for percentile estimation [2] could likewise be employed alongside induced censoring.

It is known that when data involving components or devices with multiple aging periods are improperly fit by reliability distributions, the estimates of lower percentiles are impacted by the longer lasting failures and can be biased. In Fig. 2, the solid curve represents a Weibull distribution with constant parameters ($\beta = 0.5$, $\eta = 400$) over the entire time interval ($0 \leq t \leq 1000$), and the dashed curve represents a two-part Weibull distribution with a decreasing then constant hazard rate. That is, $\beta = 0.5$ for $0 \leq t < 100$ and $\beta = 1$ for $t \geq 100$ (η remains fixed at 400 over the entire time interval). While the choice of $\eta = 400$ is somewhat arbitrary, this value is representative of what one might expect to obtain with a 1000 h ALT in which a push is made to obtain a large number of failures by the end of 400 h (or approximately two weeks). Recall that η is also the 0.632 quantile.

Our overarching motivation for induced censoring is attributed to the observation that when the correct reliability model is a simple Weibull distribution with one aging behavior and interest lies in the estimation of the lower percentiles, induced right censoring will produce percentile estimates with less bias than estimation without induced censoring. For example, suppose that the solid curve in Fig. 2 represents the data we are modeling. A small-scale simulation study (results not shown for brevity) based on assuming a Weibull model with $\beta = 0.5$ for $0 \leq t < 1000$ indicates that the percentile estimates with minimum bias are actually obtained when approximately 6% of the data is right censored. That is, right censoring the most extreme 6% of the data will result in a more helpful fit for estimating the extreme lower percentiles.

Now, assuming the two-part Weibull distribution in Fig. 2 is representative of the data we are modeling, an unhelpful fit will clearly occur when the *full* set of data is modeled inappropriately by a reliability distribution that specifies only one type of aging behavior. The result is higher root mean square errors and bias of lower percentile estimates. To correct for this model misspecification and, as a result, allow the practitioner to make use of a simple reliability distribution in lieu of a more complex bathtub curve models, we propose a strategy based on induced censoring.

Induced censoring is a technique that restructures a reliability data set by censoring a portion of the raw data (i.e., the data is weighted rather than “thrown away”). In doing so, the data involving the lifetime period of interest is given more emphasis without violating the integrity of the data. Thus, the fundamental idea of induced censoring is to extract the information of interest from the weighted data while also maintaining as much information as possible from the raw data.

The problem of determining the amount of data that should be right censored can be linked with the failure rate change point. In the case of lifetime distributions, the change point of $h(t)$ is considered a critical time [18]. Thus, for the purpose of estimation of lower percentiles, one would consider induced right censoring of observations beyond the change point from the ‘infant mortality’ phase to the ‘useful life’ phase. Naturally, though, such a change point is unknown in practice. Estimation of change points