

Bayesian Nonparametric Models for Combining Heterogeneous Reliability Data

Richard L. Warr*

David H. Collins[†]

Abstract

Modern complex engineering systems often present the analyst with a mix of data types that can be used for reliability prediction: system test results, lifetime data from unit tests of components, and subsystems data, all of which may have predictive value for the system lifetime. We present a hierarchical nonparametric framework, using Dirichlet processes, in which time-to-event distributions may be estimated from sample data or derived based on physical failure mechanisms. By applying a Bayesian methodology, the framework can incorporate prior information, including expert opinion.

Key Words: Dirichlet Process, Hierarchical Modeling, Lifetime Prediction, Parallel and Series Systems

1. Introduction

Over time the complexity of the systems we design and build has expanded at an exponential rate; traditional techniques for understanding smaller systems are often inadequate for systems of a much larger scale. Modern complex systems have many components, with varying amounts of interdependence, and they are significantly affected by boundary conditions (interfaces to other systems and the general environment); this makes them difficult to understand and manage. Data acquired by observing complex systems are heterogeneous, coming from various forms of test as well as operational monitoring, acquired from different levels (component, subsystem, system), and incomplete in various ways. The challenge we face is to integrate this heterogeneous data to provide quantitative predictions, with uncertainty estimates, to aid decision-making.

Analysis of complex systems is facilitated by the fact that they are typically decomposable or “nearly decomposable”: they can be broken down into smaller units (subsystems), which in turn can be decomposed into components; i.e., they are structured in a multilevel hierarchy. This allows different levels to be treated separately via hierarchical modeling, which facilitates intellectual comprehension and enables “divide and conquer” strategies for computational tractability.

Various types of decomposition may be used:

- Based on physical structure such as series or parallel units (traditional engineering block diagrams)
- Based on event logic, e.g., fault trees or event trees
- Based on system states and stochastic transitions between states, e.g., Markov or semi-Markov processes.

In this paper we focus on the first two strategies; see Collins and Huzurbazar (2011) for more information on state-space methods. See Johnson et al. (2003), Graves et al. (2010)

*Air Force Institute of Technology, Wright-Patterson AFB, Ohio, USA. The views expressed in this article are those of the author and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the U.S. Government.

[†]Statistical Sciences Group, Los Alamos National Laboratory, Los Alamos, New Mexico, USA. This work performed under the auspices of the Los Alamos National Security, LLC, for the National Nuclear Security Administration of the U.S. Department of Energy under contract DE-AC52-06NA25396.




Figure 1: Hierarchical decomposition of an example four-component system.

and Reese et al. (2011) for additional information and general references on hierarchical reliability analysis.

The goal of this paper is to develop reliability estimates for complex systems, including estimates of uncertainty, using component, subsystem, and system data. We want to use all available data types, which may include subjective data such as expert opinion as well as data collected from various formal tests. This goal has led us to develop the hierarchical Bayesian methodology we present here.

The remainder of the paper is organized as follows: Section 2 presents background information on hierarchical decomposition, two key results in distribution theory, and brief reviews of Bayesian analysis and the Dirichlet process. In Section 3 we present a hierarchical nonparametric Bayesian framework for analyzing the kind of heterogeneous data just described. Section 4 illustrates the use of the framework in a real application, estimation of the reliability of a remotely piloted aircraft (RPA), and we conclude with a summary in Section 5.

2. Background

We use the following definitions in this paper. A *system* is a group of components that interact to function as a whole; a *component* is a part of a larger whole; a *subsystem* is component composed of interacting subcomponents; the term *unit* may refer to a system, subsystem, or component.

2.1 Hierarchical Decomposition for Reliability Analysis

Reliability data for complex systems such as aircraft typically comes from a variety of sources. Full system tests are often expensive and difficult to perform; unit tests are usually done for some components and subsystems. If physical mechanisms for failure are well-understood, we may have computer simulation results based on scientific models. Expert judgment from scientists and engineers, based on experience with this or similar systems, can also be valuable. Data types are heterogeneous: testing may provide binary pass/fail information, lifetime (time to failure) data, or degradation measurements for failure mechanisms such as corrosion and fatigue cracking. Data will have differing levels of uncertainty, depending on the source and quantity of the data.

Figure 1 is a notional view of a system composed of two subsystems, each of which has two components. The components at the lowest level of the hierarchy may have sub-components, but for practical reasons (e.g., we have no data on the subcomponents, or the component is a sealed unit) we may not wish to carry the decomposition any further. The varying shapes in Figure 2 represent the fact that we have data with different characteristics

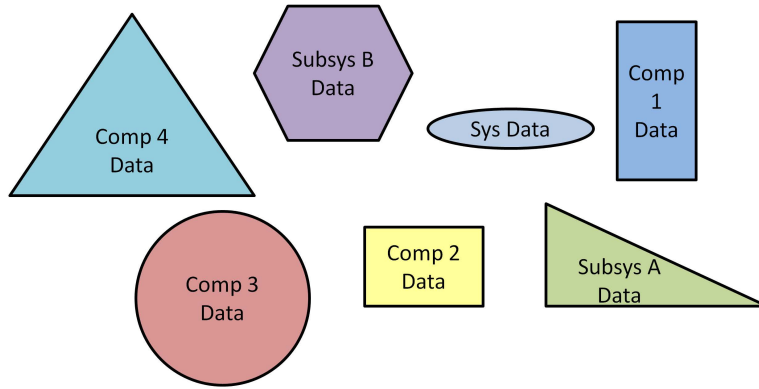


Figure 2: Heterogeneous data for the system depicted in Figure 1

for each element in the system decomposition. We wish to synthesize all this data into a reliability prediction for the whole system, possibly weighting some data more than others based on our confidence in it. Our confidence is measured on factors such as how the data were collected, the accuracy of the model, and possibly many others.

2.2 Distribution of Maxima and Minima

The following results are standard for calculating the reliability of components in series or parallel. We state these in terms of cumulative distribution functions (CDFs): if X is the random variable representing time to failure, its CDF is $F_X(x) = P(X \leq x)$; its reliability function is $R_X(x) = 1 - F_X(x) = P(X > x)$.

The failure time of a system with two independent components A and B in parallel is the *maximum* of the component failure times; if the system failure time random variable is X , and the components are independent with failure times X_A and X_B , we have $X = \max(X_A, X_B)$ and

$$P(X \leq x) = P(X_A \leq x \cap X_B \leq x) = P(X_A \leq x)P(X_B \leq x).$$

If the corresponding CDFs are F_A and F_B , then the CDF for system failure time is

$$P(X \leq x) = F_X(x) = F_A(x)F_B(x). \quad (1)$$

The failure time of a system with two components A and B in series is the *minimum* of the component failure times. With the same notation as above, using a basic formula of probability we have $X = \min(X_A, X_B)$ and

$$\begin{aligned} P(X \leq x) &= P(X_A \leq x \cup X_B \leq x) \\ &= P(X_A \leq x) + P(X_B \leq x) - P(X_A \leq x \cap X_B \leq x). \end{aligned}$$

Thus the CDF for system failure time of the series system is

$$P(X \leq x) = F_X(x) = F_A(x) + F_B(x) - F_A(x)F_B(x). \quad (2)$$

A classical way to view a series system is to use reliability functions. This provides a simple approach when dealing with many components. We have for independent random variables X_A and X_B ,

$$P(X > x) = R_X(x) = R_A(x)R_B(x). \quad (3)$$

2.3 Bayesian Analysis

The Bayesian paradigm for statistical inference is well known; here we briefly review the basic ideas and summarize the concepts of nonparametric Bayesian analysis. For a general introduction to Bayesian ideas and methods, see Christensen et al. (2010); for an in-depth study in Bayesian nonparametrics see Ghosh and Ramamoorthi (2003).

Whereas frequentists assume that distribution parameters are fixed, but unknown, quantities to be estimated, in the Bayesian framework parameters are treated as random variables with probability distributions. For example, if an exponential model $X \sim \exp(\lambda)$ is of interest, we begin by assigning a prior distribution $p(\lambda)$ which reflects an initial belief or state of knowledge about the parameter λ . After collecting data $\mathbf{x} = \{x_1, x_2, \dots, x_n\}$, where

$$X_1, X_2, \dots, X_n \mid \lambda \stackrel{iid}{\sim} \exp(\lambda)$$

we update our belief about λ to obtain a posterior distribution, using Bayes' theorem:

$$p(\lambda \mid \mathbf{x}) = \frac{p(\mathbf{x} \mid \lambda)p(\lambda)}{\int p(\mathbf{x} \mid \lambda)p(\lambda)d\lambda}. \quad (4)$$

Though parameters are treated as random, Bayesian methods normally treat the distribution model generating the data, $p(\mathbf{x} \mid \lambda)$, as fixed up to the choice of λ . In Bayesian nonparametrics this distribution is itself treated as a random quantity drawn from some function space, e.g., the space of all continuous distribution functions. The implementation of Bayes' theorem now requires not just a prior distribution over a univariate or multivariate parameter space, but a distribution over a space of distributions. In other words, such a distribution \mathcal{D} can be sampled to provide CDFs characterizing a univariate random value. For example, \mathcal{D} can be used to generate sample CDFs F_i for failure time distributions in a reliability analysis. Just as in conventional Bayesian analysis, the prior is updated, using observed failure times, to yield a posterior distribution $\mathcal{D}(F \mid \mathbf{x})$ over failure time distributions. A commonly used probability model for distribution functions is the Dirichlet Process, which is described in the next section.

2.4 Dirichlet Processes

The Dirichlet process (DP), originally introduced in Ferguson (1973), is one type of nonparametric Bayesian model. We provide a short introduction to DPs, and explain how they can be used in a reliability context to obtain an estimate of the unknown reliability function, $R(t)$, and the associated unknown CDF $F(t) = 1 - R(t)$.

DPs are defined by a mean (also called the base measure), which we call $F_0(t)$, and by a precision parameter $\alpha \geq 0$, which controls the amount of variation of the DP around $F_0(t)$. The Dirichlet process can be described as a distribution for CDFs. We denote the DP with precision α and base measure $F_0(t)$ as $\text{DP}(\alpha, F_0(t))$, or $\text{DP}(\alpha, F_0)$. The Bayes estimate for $F(t) \sim \text{DP}(\alpha, F_0)$ (with a quadratic loss function) is just $E[F(t)] = F_0(t)$; however, the DP process can also be used to quantify the uncertainty, centered around the point estimate. At any point t the $\text{DP}(\alpha, F_0(t))$ requires that $F(t) \sim \text{Beta}(\alpha F_0(t), \alpha(1 - F_0(t)))$, from which the pointwise uncertainty of $F(t)$ is easily calculated. Figure 3 illustrates what a DP looks like: the mean is a Weibull(1.8, 30) CDF, and the precision is $\alpha = 50$. The black line represents the beta distribution of $F(t)$ at some particular t , hence the "Beta domain" dimension in the plot is restricted to the interval $[0, 1]$, as we would expect for a CDF. The red line in the plot shows the mean of the DP. The areas that are more elevated in the plot represent areas more likely to contain the unknown CDF $F(t)$.




Figure 3: An example Dirichlet process

DPs can be used in a variety of ways, one of which is as a Bayesian prior for an unknown/random distribution; these are called Dirichlet process priors (DPPs). DPPs can be considered conjugate, such as the beta distribution is the conjugate prior for a binomial model; i.e., the posterior is a probability model in the same family as the prior. After collecting data and updating the DPP, $DP(\alpha, F_0(t))$ which reflects our prior belief about $F(t)$, the posterior is also a DP (this is only true for uncensored data).

To make the DP concepts more concrete, we show how a DP can be used to estimate a CDF given some data. Consider the data

$$X_1, X_2, \dots, X_n | F \stackrel{iid}{\sim} F \text{ where } F \sim DP(\alpha, F_0).$$

The prior information is contained in F_0 and α . The posterior of F (or the distribution of $F | \mathbf{X}$) is also a DP. We define

$$\hat{F}(t) = \frac{\sum_{i=1}^n I(X_i \leq t)}{n},$$

which is the empirical CDF based on the data. Then the posterior DP is

$$F | X_1, X_2, \dots, X_n \sim DP \left(\alpha + n, \frac{\alpha F_0 + n \hat{F}}{\alpha + n} \right). \quad (5)$$

Clearly, if $\alpha = 0$ then we have no prior information and F_0 has no effect on the posterior. An example of the DP posterior is shown in Figure 4. In this Figure one can see a discontinuity in the mean at each observation, corresponding to the jump in the empirical CDF at that point.

It is worthwhile to point out a few things about the posterior. The first is that the precision is increased by n , which indicates that we can roughly interpret α as the number of observations we feel the prior should account for. This is an appealing feature that allows




Figure 4: An example Dirichlet process posterior

us to understand the effect of the prior on the posterior. The next thing we observe is that the base measure is a mixture of the prior base measure and the empirical CDF. This makes sense intuitively, and helps conceptualize how the DP posterior is derived. The last thing we point out is that the DP posterior's base measure has discontinuities. Although this is not necessarily a desirable characteristic, it does not detract from our original goal.

In a reliability context, using DPs has several benefits. The DP definition allows a modeler to specify the weight and shape of the prior independently; in other words, unlike in Bayesian parametric models, the shape of the prior F_0 has no effect on the posterior precision. Therefore there is no need to perform a traditional Bayesian sensitivity analysis, because we know the effect the prior will have on the posterior. Another benefit of DPs is that if we have no prior information, α can be set to zero; this is truly an uninformative prior. If this is the case, the mean of the DP posterior is just the empirical CDF, which is a very logical result. Additionally, no effort is required to identify an adequate parametric model; parametric modeling assumptions may not be challenging when we have one distribution, but later in the paper we allow for a very large number of components, each of which may have a unique time-to-failure distribution. To properly fit a distribution to each component's failure time would most likely be time prohibitive. The last benefit we highlight is the computational ease with which DPs are handled; because of the conjugate property, typical computationally complex Bayesian methods such as Markov chain Monte Carlo (MCMC) are not required.

In this section we have briefly covered the nature of the DP and some of its properties that will be useful in the following sections. For a more in-depth treatment of Dirichlet processes, see Hjort et al. (2010) or Ghosh and Ramamoorthi (2003).

3. A Nonparametric Bayesian Framework for Hierarchical Reliability Data

Having provided some background on DPs, we now use them as building blocks in a hierarchical structure to obtain a reliability estimate for a complete system.

In concept the method is quite simple. For each component of a subsystem we construct a DP posterior and keep only the posterior base measure. Then using the structure of the subsystem and the formulas presented in Section 2.2, we find the CDF of the combined components for that subsystem. We then construct a new DPP for the subsystem using the combined component CDF as the base measure, and choose a new precision that represents the amount of data this prior information should account for. Finally, we update this DPP with any subsystem data we have. This produces another DP posterior that is an estimate of $F(t)$ for the subsystem time-to-failure. This can be iterated through any number of levels in a system's hierarchy until we attain the DP posterior for the full system we are modeling.

One possible criticism of this method is that the precision for the component DPs is lost when constructing the subsystem DPP. At first this troubled us, but the more we thought about this effect, the more it actually became appealing. The primary reason is that the observed subsystem data is usually the "gold standard", and we are only supplementing it with a prior based on component data. This assignment of α can be viewed as the weight of the DPP, that is, the amount of effect the prior should have on the posterior. Therefore it makes sense to allow the modeler to choose the weight that component's information will have in the posterior; if the amount of component data is large and the amount of subsystem data is small, the component data will not inadvertently overwhelm the subsystem data.

Ideally the base measure of the DPP from the components and the empirical CDF (ECDF) of the subsystem data will have similar shapes. If not, this could indicate one or both of two problems. One possibility is that the components are not functioning independently, in a statistical sense. Hence, there is a visible difference between the base measure of the DPP and the ECDF of the subsystem data. The other possible problem is that there is a missing component in the subsystem model. For example, if we have a two-component system and have data for each of the two components, when we connect the components the coupling mechanism (e.g., a cable or soldered connection) needs to be taken into account. These are the two main assumptions that need to be checked for this model to hold. If the base measure of the DPP from the components and the ECDF roughly agree, we should expect the model to be reasonable approximation of reality.

From an engineering perspective we would expect a shorter mean time to failure (MTTF) under the subsystem estimate than in the DPP of the components. This occurs in most systems because components do interact and new failure modes appear when they are combined. However, if the differences are negligible we can proceed using the model we have developed.

To summarize, these are the steps in this modeling technique:

1. For each component, select a DP prior, with base measure and precision (weight)
2. Collect data for each component
3. Update the priors with the data to get component DP posteriors
4. Use the base measures of the component DP posteriors to estimate the base measure for the subsystem prior, using formulas (1) and (2)
5. Construct a DP prior for each subsystem, using the base measure from the previous step and a precision that is chosen to represent the weight the component data should have on the subsystem posterior
6. Update the subsystem DP prior with observed subsystem data to obtain a DP posterior




Figure 5: Mission profile and propulsion subsystem use for a small hybrid-electric remotely piloted aircraft

7. Iterate this process through all the levels of the system, treating subsystems as components

There are several reasons why this modeling approach is effective. First, component data should not determine the influence the prior has on a subsystem posterior. In this approach the modeler can intelligently devalue the component data so that it does not overpower the subsystem data. In addition, there are relatively few assumptions built into the model. Another advantage is, the DP posterior for the full system provides an estimate of the uncertainty in the result.

4. Application


To illustrate the application of this framework to a real-world problem, we model the reliability of the propulsion system in a small hybrid-electric remotely piloted aircraft (SHERPA). SHERPA is a prototype RPA being developed at the Air Force Institute of Technology (AFIT) for a mission profile involving takeoff, ingress to a target area, loitering over the target for surveillance, egress, and landing (Ausserer (2012)). RPAs powered by gasoline engines have a high power to weight ratio and are capable of long missions, but are noisy and thus less stealthy than desired; electrically powered RPAs have a smaller acoustic signature, but battery weight is a challenge if electrical propulsion is to be used over a complete mission. A hybrid RPA using both systems provides a solution that combines enough power to complete the mission, and the ability to perform the essential part of the mission (loitering) solely on electric power.

Figure 5 shows a generic mission profile for SHERPA, along with the propulsion used. The gasoline engine (G) alone is used for ingress, egress, and landing; the electric motor is used while loitering over the target; and both are used (for maximum power) during takeoff and climb to cruising altitude. While cruising, the gas engine recharges the battery.

Table 1 lists the components in a high-level decomposition of the propulsion system. The interrelationship of components is shown by the reliability block diagram in Figure 6. The three common components on the left are in series, since a failure of any one results in loss of propulsion. The two parallel branches to the right indicate that propulsion is functional if either the gas engine or electric motor is functional (the reliability goal we consider is simply whether the aircraft can fly, though clearly its mission readiness would be

Table 1: Component List for SHERPA Propulsion System

Electric Propulsion	Gasoline Propulsion	Common Parts
Motor	Engine	Propeller
Batteries	Gas Delivery	Drive shaft
Motor controller		Gearing
Serpentine belt		

**Figure 6:** The reliability block diagram of the propulsion system of the hybrid RPA.

compromised if only one propulsion mode were functional). Each parallel branch includes the components for one propulsion mode in series.

The data we use are simulated to mimic the reliability of the components, subsystems, and system. We are not able to use the actual data due to operational issues. We also ignore at this point the fact that the actual data would be right censored, since testing typically is terminated before all units have failed. The simulated data is uncensored, which is not realistic in this setting. Our future research will enhance the framework to provide the ability to handle right-censored data. One final assumption we make is that there is no prior information for the components. This is the worst case scenario; it is straightforward to incorporate prior information if it is available.

To model $R(t)$ for the SHERPA propulsion system we start by organizing it into subsystems and components, as in Table 1 and Figure 6. (We are actually modeling $F(t)$, the time-to failure CDF, and using $R(t) = 1 - F(t)$). We have two subsystems composed of components; once we obtain a DP posterior for both subsystems we treat them as components along with the common components to obtain the DP prior for the system, which is updated using system data to get the system DP posterior. Figure 7 displays the histograms for each component, subsystem and system. (The x-axis is omitted to mask the lifetimes of the parts.)

We demonstrate this process in detail with the electrical subsystem; Figure 8 shows the flow of the modeling. First we assume no prior information is known for the Motor, MC (motor controller), Bat (battery), and Belt (serpentine belt). Next a DP posterior (labeled DP in Figure 8) is constructed for each of the four components. Then, the means of the DP posteriors are used to calculate the base measure CDF for the whole electrical subsystem. Since there are four components in series, Equation (3) finds the reliability function of the minimum time to failure (from which we obtain the CDF). Using this minimum CDF, we add a precision (weight) to construct the DPP for the electric propulsion subsystem. There are 15 observations of the electric propulsion subsystem, therefore we choose the DPP precision to be $\alpha = 0.75n = 11.25$; this ensures the prior will not overwhelm the posterior. However, other values of α could be appropriate depending on the application and the modeler's knowledge of the process. With the DPP constructed and the electric subsystem data, we use equation (5) to obtain the DP posterior for the electric propulsion

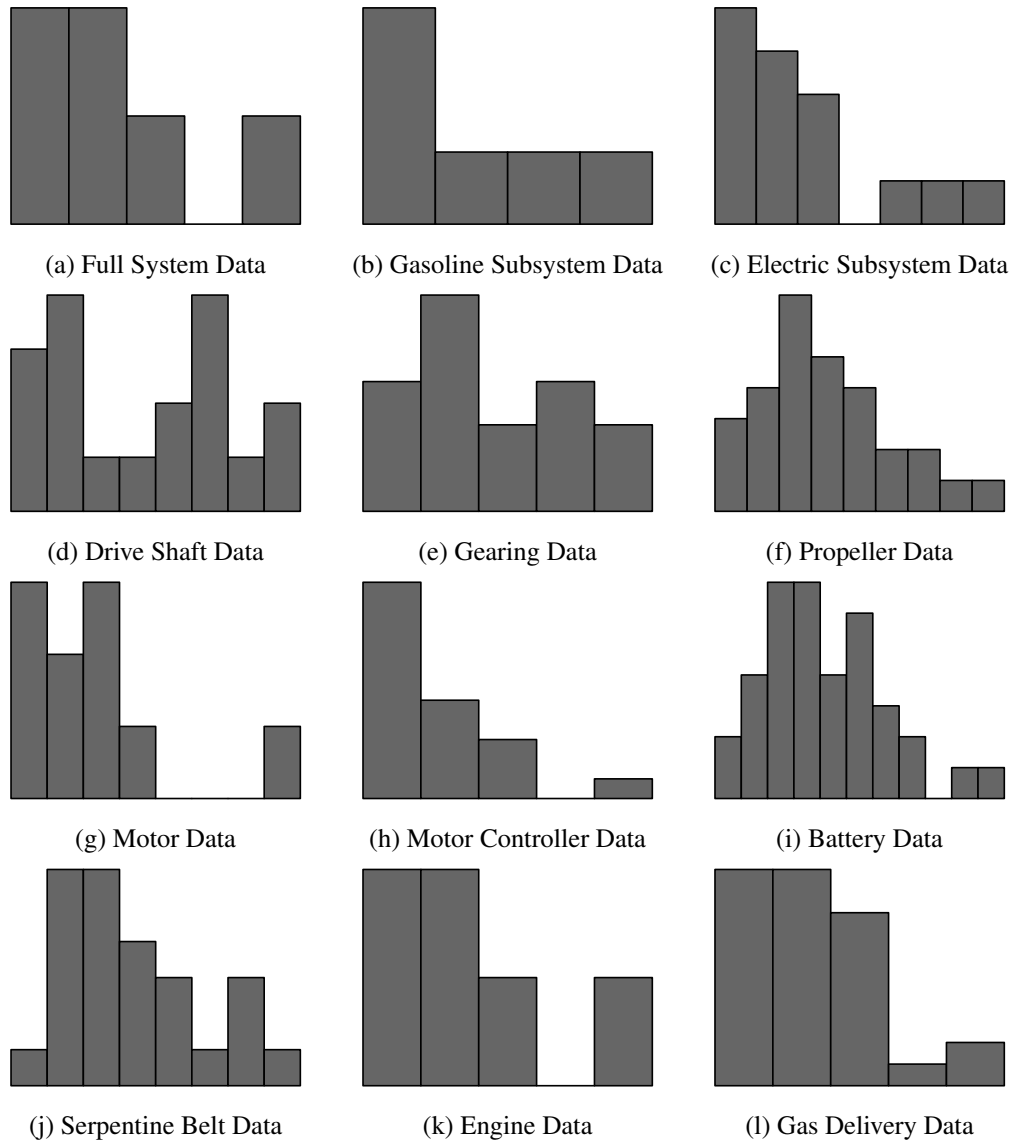


Figure 7: Histogram of System, Subsystems and Component Data

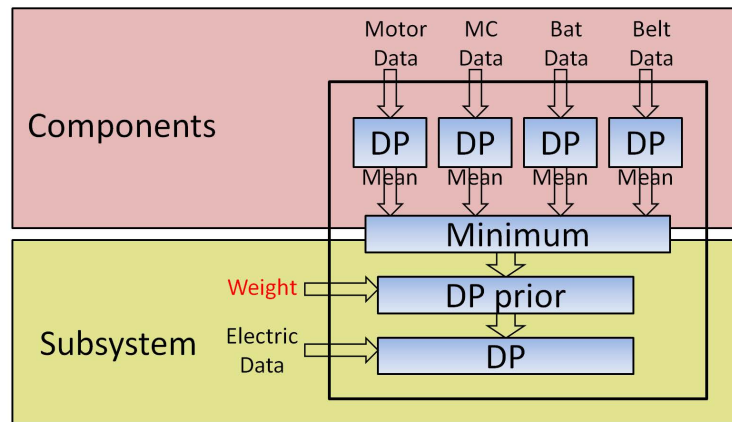


Figure 8: Modeling flow for the electric propulsion subsystem.

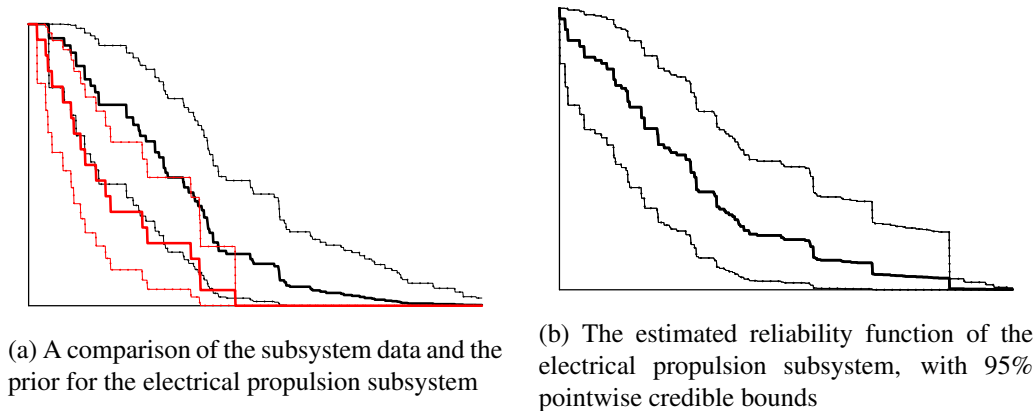


Figure 9: Plots for the electrical propulsion subsystem

subsystem.

It is usually not reasonable to obtain an analytic expression for the posterior mean of the DP. Therefore our method is to calculate these quantities on a discretized grid along the time axis; the fineness of the grid can be chosen by the modeler, with an obvious computational penalty as more points are chosen. Once these values are calculated, the posterior mean can be displayed with the pointwise 95% credible bounds of the DP posterior. These credible bounds are obtainable since the pointwise distribution of the DP is a beta distribution. Figure 9a shows two darker curves: the red one is the Kaplan-Meier estimate of $R(t)$ for the electric subsystem data, and the black one is one minus the mean of the DPP of the electric subsystem components. The lighter lines provide the 95% pointwise uncertainty associated with these curves. If the assumptions of the model are approximately correct, the two darker lines should be fairly similar, with some allowance for sample to sample variation. Figure 9b shows one minus the mean of the DP posterior with its 95% pointwise credible bounds.

The remaining calculations are similar to those for the electrical subsystem. The prior precisions are set at a conservative level with $\alpha = 0.75n$, where n is the number of test samples obtained for the subsystem or system. We show additional plots of one minus the mean of the DPPs and the system/subsystem data to verify that the assumptions are reasonable in Figures 10a and 10b. As with any model, ensuring the assumptions are

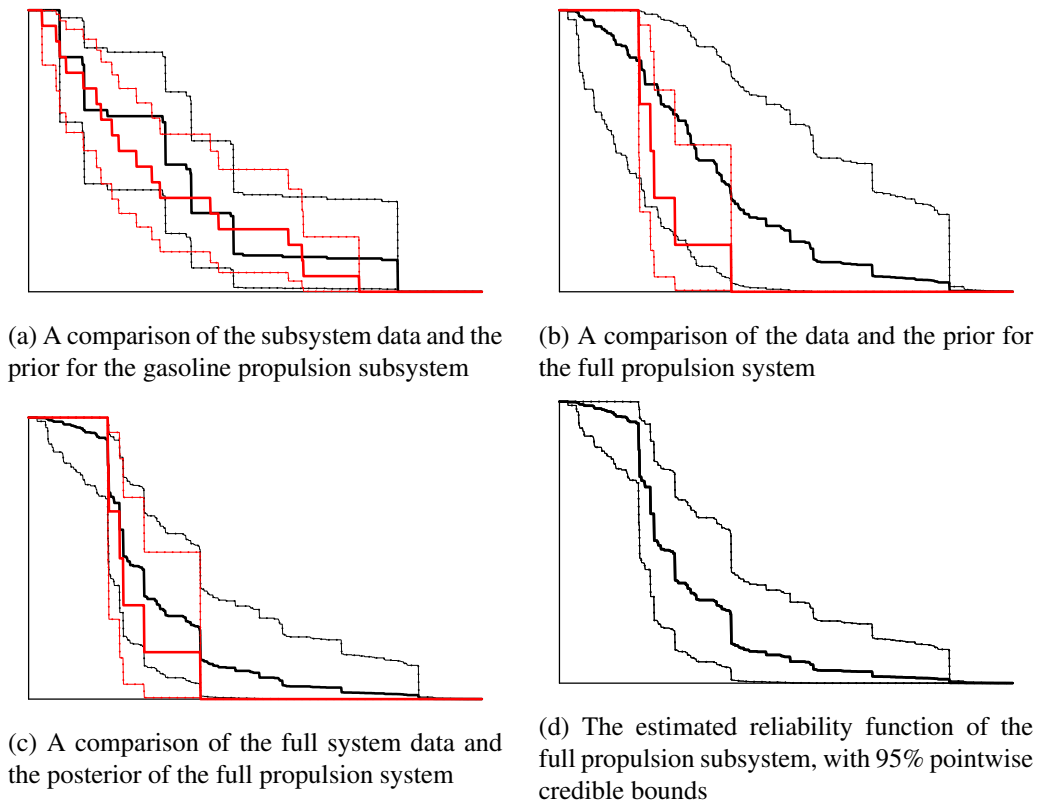


Figure 10: Plots of various estimates of reliability functions

met are vital. Figures 9a, 10a, 10b show the priors (in black) versus the subsystem or system test data (in red). If these are a fairly close match, we can proceed. Figure 10a is clearly appropriate, however the other two do show some disagreement. Since the prior accounts for 3/4 of the actual data we don't expect these deviations to dramatically impact the posterior. Since this disagreement is not too severe, we proceed on the basis that the assumptions are met.

The final output of interest is in Figure 10d. This shows the estimate of the reliability function of the full propulsion system of SHERPA with 95% pointwise credible bounds. From this we can determine quantities such as the estimated probability of failure at a particular time, the estimated 90% mission capable rate, and other measures that may be of interest to those who rely on this system. One additional result is shown in Figure 10c; the reliability estimate of the data for the six full system tests (in red) is plotted with the posterior result from our methodology. Clearly, using more data provides a better estimate, both in the fidelity of the mean and reduced uncertainty, which is shown by the narrower credible bounds.

5. Summary

In this paper we introduced a methodology that combines many levels of data from a complex hierarchical system. We applied a Bayesian nonparametric approach, which reduces the number of assumptions and provides a computationally concise solution. This method also produces a reliability estimate with uncertainty quantification.

There is still work that needs to be done to make this a truly useful method. The first major issue to be addressed is generalizing the model to incorporate randomly right

censored data. Nearly all reliability data has censored observations, so a method that is not able to handle it is not practically useful. The generalization of this model will likely use the work of Blum and Susarla (1977), which shows that a DPP updated with randomly right censored observations is a mixture of Dirichlet processes. With this information the problem will become more complex computationally, but will still be based on the same fundamental concepts. Additional work should also include the ability to set a prior on α , the precision of the DPP.

6. Acknowledgments

The authors thank Fred Harmon of AFIT and his colleagues for access to the SHERPA project, and to Aparna Huzurbazar for her suggestions and encouragement of our work.

References

- Antoniak, C. (1974). Mixtures of Dirichlet processes with applications to Bayesian nonparametric problems, *The annals of statistics* pp. 1152–1174.
- Ausserer, J. K. (2012). *Integration, testing, and validation of a small hybrid-electric remotely-piloted aircraft*, Master's thesis, Air Force Institute of Technology, Wright-Patterson AFB, OH.
- Blum, J. and Susarla, V. (1977). On the posterior distribution of a Dirichlet process given randomly right censored observations, *Stochastic Processes and their Applications* **5**(3): 207–211.
- Christensen, R., Johnson, W., Branscum, A. and Hanson, T. E. (2010). *Bayesian Ideas and Data Analysis: An Introduction for Scientists and Statisticians*, CRC Press, Boca Raton, FL.
- Collins, D. and Huzurbazar, A. (2011). Prognostic models based on statistical flowgraphs, *Applied Stochastic Models in Business and Industry* **28**(2): 141–151.
- Ferguson, T. S. (1973). A Bayesian Analysis of Some Nonparametric Problems, *The Annals of Statistics* **1**(2): 209–230.
- Ghosh, J. and Ramamoorthi, R. (2003). *Bayesian nonparametrics*, Springer Verlag.
- Graves, T., Anderson-Cook, C. and Hamada, M. (2010). Reliability models for almost-series and almost-parallel systems, *Technometrics* **52**(2): 160–171.
- Hjort, N., Holmes, C., Müller, P. and Walker, S. (2010). *Bayesian nonparametrics*, Vol. 28, Cambridge University Press.
- Johnson, V., Graves, T., Hamada, M. and Reese, C. (2003). A hierarchical model for estimating the reliability of complex systems, *Bayesian statistics* **7**: 199–213.
- Johnson, W. and Christensen, R. (1989). Nonparametric Bayesian analysis of the accelerated failure time model, *Statistics & Probability Letters* **8**(2): 179–184.
- Reese, C., Wilson, A., Guo, J., Hamada, M. and Johnson, V. (2011). A Bayesian model for integrating multiple sources of lifetime information, *Journal of quality technology* **42**(2): 127–141.