

DISSERTATION

STATISTICAL METHODS FOR THE DETECTION AND ANALYSIS OF RADIOACTIVE  
SOURCES

Submitted by

John Klumpp

Department of Environmental and Radiological Health Sciences

In partial fulfillment of the requirements

For the Degree of Doctor of Philosophy

Colorado State University

Fort Collins, Colorado

Fall 2014

Doctoral Committee:

Advisor: Alexander Brandl

Thomas Johnson  
Georg Steinhauser  
Susan LaRue  
Geof Givens

Copyright by John Allan Klumpp 2014

All Rights Reserved

## ABSTRACT

### STATISTICAL METHODS FOR THE DETECTION AND ANALYSIS OF RADIOACTIVE SOURCES

We consider four topics from areas of radioactive statistical analysis in the present study: Bayesian methods for the analysis of count rate data, analysis of energy data, a model for non-constant background count rate distributions, and a zero-inflated model of the sample count rate. The study begins with a review of Bayesian statistics and techniques for analyzing count rate data. Next, we consider a novel system for incorporating energy information into count rate measurements which searches for elevated count rates in multiple energy regions simultaneously. The system analyzes time-interval data in real time to sequentially update a probability distribution for the sample count rate. We then consider a “moving target” model of background radiation in which the instantaneous background count rate is a function of time, rather than being fixed. Unlike the sequential update system, this model assumes a large body of pre-existing data which can be analyzed retrospectively. Finally, we propose a novel Bayesian technique which allows for simultaneous source detection and count rate analysis. This technique is fully compatible with, but independent of, the sequential update system and moving target model.

## ACKNOWLEDGEMENTS

The author was supported in part by the CDC/NIOSH Mountain & Plains Educational and Research Center. The contents herein, however, are solely the responsibility of the author and do not represent the official views of the CDC, NIOSH and MAP ERC.

I would like to respectfully acknowledge Guthrie Miller, Alexander Brandl, and Thomas Johnson for their significant contributions to this work, as well as Dawn Lewis, then at Los Alamos National Laboratory, for providing the  $^{238}\text{U}$  method blank data set.

Finally, I would like to thank my wife Teresa for loving me and believing in me even when I was most frustrated.

## TABLE OF CONTENTS

ABSTRACT.....	ii
ACKNOWLEDGEMENTS.....	iii
TABLE OF CONTENTS.....	iv
EXECUTIVE SUMMARY .....	1
Chapter 1.....	6
BAYESIAN TECHNIQUES FOR THE ANALYSIS OF COUNT RATE DATA .....	6
1.1 Introduction.....	6
1.2 The Likelihood Function.....	7
1.3 The Prior .....	8
1.4 The Predictive Distributions .....	11
1.5 Credible Intervals.....	13
1.6 Count Rate Analysis .....	15
1.7 Bayesian Hypothesis Testing.....	19
Chapter 2.....	27
SEQUENTIALLY UPDATED BAYESIAN DETECTION SYSTEM UTILIZING ENERGY AND COUNT RATE DATA FOR THE DETECTION OF RADIOACTIVE SOURCES .....	27
2.1 Introduction.....	27
2.3 Materials and Methods.....	31
2.3.1 Theoretical Considerations .....	31
2.3.2 Instrumentation .....	35
2.4 Results and Discussion .....	36
2.4.1 Fixed source detection systems:.....	36
2.4.2 Moving prior detection systems:.....	39
2.5 Conclusion .....	41
Chapter 3.....	43
CHARACTERIZATION OF BACKGROUND RADIATION WITHOUT THE ASSUMPTION OF IDENTICALLY DISTRIBUTED COUNT RATE MEASUREMENTS.....	43

3.1 Introduction.....	43
3.3 Materials and Methods.....	45
3.3.1 Background.....	45
3.3.2 Maximum Likelihood Estimation.....	48
3.3.3 Bayesian Probabilistic Analysis.....	50
3.3.4 Experimental Analysis.....	51
3.4 Results.....	51
3.4.1 <sup>234</sup> U Bioassay Data Set.....	51
3.4.2 Simulated Data Sets.....	55
3.5 Conclusion.....	56
Chapter 4.....	58
<b>SIMULTANEOUS SOURCE DETECTION AND ANALYSIS USING A ZERO-INFLATED COUNT RATE.....</b>	<b>58</b>
4.1 Synopsis.....	58
4.2 Introduction.....	58
4.3 Zero-Inflated Model.....	61
4.4 Examples.....	68
4.4.1 Example 1: Negative Observed Source Count Rate.....	69
4.4.2 Example 2: Strong Signal, Rare Source.....	73
4.5 Discussion.....	76
4.6 Conclusion.....	77
Chapter 5.....	79
<b>SUMMARY AND CONCLUSIONS.....</b>	<b>79</b>
5.1 Bayesian Analysis.....	79
5.2 Sequentially-Updated Source Detection Systems.....	80
5.3 Moving Target Model of Background.....	81
5.4 Zero-Inflated Priors.....	82
References.....	84
Appendix.....	87

## EXECUTIVE SUMMARY

A common application of radiation counting measurements is to estimate the count rate from a source which is present in a sample of material. This application, which is referred to as “count rate analysis,” is typically carried out by taking counts from the sample of interest and from background. The net sample count rate is then estimated by subtracting the background count rate from the gross count rate of the sample. Source detection is a related application whose aim is to estimate the probability that any radioactive material is present in the sample (e.g., that the sample count rate is non-zero). Because samples cannot typically be measured independently of background, source detection and analysis depend on the ability to accurately model both the background and sample distributions. Although source detection is a special case of count rate analysis, in practice the two applications are often approached very differently. The ability to detect and identify radioactive sources has a number of important applications. It plays a key role in monitoring vehicles and cargo for unauthorized nuclear materials, localization of lost or stolen sources, facility decontamination and decommissioning, and in personal and environmental monitoring for radiation protection purposes.

In a classical radiation monitoring set-up, one assumes a known fixed background, decides on a decision threshold, and compares the results of a measurement to that decision threshold. If the number of observed counts in a given time interval is greater than the decision threshold, the observer concludes that a source has been detected. Traditionally, the false-positive rate is controlled to 5% by setting the decision threshold to mean background +  $1.645\sigma$ . This set-up is easy to understand and implement, but has the disadvantage of unnecessarily

discarding past measurement data. Similarly, on-line detectors often employ techniques, such as moving averages and recursive digital filters, which are designed to discard and de-emphasize old data (Fehlau 1993). For that reason, alternative techniques have been developed to more completely make use of the available data (e.g., Brandl 2013). Specifically, there has been increasing interest in Bayesian techniques, which allow for the formal incorporation of prior belief, data, and knowledge, and which produce probabilities for the underlying parameter (i.e., mean count rate). This is in contrast to classical techniques, which do not have the ability to incorporate prior data, and which can only make statements about whether an observation was drawn from a particular distribution.

Chapter 1 reviews the standard Bayesian techniques for analyzing count rate data. Bayesian analysis has a number of advantages over classical analysis of data. The most obvious is that Bayesian analysis allows the investigator to make direct statements about the quantities of interest. For example, after a measurement of a possibly radioactive source, classical analysis can only estimate the probability of having observed  $N$  counts for a given source count rate. This is very different from a probability distribution of the source count rate, which can only be obtained by incorporating a prior, but it is easy to see how the two probabilities could be confused. In contrast, Bayesian analysis can directly answer the question being asked – in this case, “is there a source, and, if so, how strong is it likely to be?” For that reason, Bayesian techniques are uniquely suited for the analysis of radiation count rate data.

Chapter 2 presents a novel system for incorporating energy information into count rate measurements which searches for elevated count rates in multiple energy regions simultaneously. The system analyzes time-interval data (e.g., time between counts), as this was shown to be a more sensitive technique for detecting low count rate sources compared to analyzing counts per



unit interval (Luo et al. 2013). Two distinct versions of the detection system are developed. The first is intended for situations in which the sample is fixed and can be measured for an unlimited amount of time. While developing this system, it was hypothesized that the multiple-energy fixed sample detection systems would detect sources in at least 20% less time than the single-energy detector. The second version of the detection system is intended to detect transient sources – i.e., sources which are physically moving relative to the detector, such as a truck moving past a fixed roadside detector, or a waste storage facility under an airplane. The performance of the transient source detectors will be evaluated in the following scenario: The moving target detection system is run for a total of 15 s. It measures background for 10 s, then the source for 5 s. The mean background count rate is 2 counts per second, and the source count rate varies from 0 to 4 cps. The energy of the source is 1200 keV (corresponding to  $^{60}\text{Co}$ ). We predict that, if the decision thresholds are chosen to control the false positive rate at 5%, the multiple-energy detection system will significantly out-perform the single-energy detector in terms of failed detections. Similarly, we predict that, if the decision thresholds are chosen to control the failed detection rates at 5%, the multiple-energy detection system will significantly out-perform the single-energy detector in terms of false positives. Sequentially updating Bayesian radiation detectors are a recent development, and have not previously been designed to incorporate energy information. Therefore, the detection systems proposed in this document represent a new development in the field of health physics. In addition, the improved performance of these detection systems could represent a significant advance in the ability of health physicists to rapidly detect low count rate sources.

Chapter 3 proposes a “moving target” model of background radiation in which the instantaneous background count rate is a function of time, rather than being fixed. Specifically,

the moving target model treats the average Poisson mean in a given time interval as a random variable drawn from a gamma distribution. Because of the nature of radioactive decay, the number of counts in a given time interval from any fixed source of radiation will follow a Poisson distribution which, for high enough counts, can be well approximated by a Normal distribution. Virtually all classical decision thresholds are based on the assumption that the number of counts from background is also sampled from a Normal distribution. However, it is well known that background varies with time of day, location, detector performance, meteorological conditions, and many other factors (Hemic 1988). In cases where background is varying, this results in a broader inferred distribution of background count rate than a “fixed-target” background. Left unaccounted for, this leads to inaccurate posterior distributions, higher than expected false positive rates in measurements with decision thresholds, and faulty performance estimates of detectors. We predict that 95% decision thresholds based on the  $^{238}\text{U}$  background data will be at least 10% higher under the moving target analysis compared to when the background is assumed to be stationary and that decision thresholds will be proportional to the standard deviation of the distribution of the mean count rate. This is the first time that gamma distributions have been used to model actual variation as well as uncertainty in the mean count rate of background radiation. This approach represents a significant advance in the ability of health physicists to accurately characterize background radiation, and is a significant step towards solving the problem of detecting sources in the presence of a non-constant background.

Chapter 4 proposes a novel Bayesian technique which allows for simultaneous source detection and count rate analysis. The technique involves using priors which include a finite probability that the source count rate is exactly zero. Such priors are called “zero-inflated.” Solving the posterior distribution of a zero-inflated count rate model provides the probability that

the sample contains a source and a probability distribution for the source count rate if the source exists, without the need to perform redundant computations. Sampling from zero-inflated distributions is straightforward, and can be accomplished with easily accessible open source software. In addition, zero-inflated priors lead to finite posterior probabilities of “no source,” which is an easy to understand and satisfying result. While zero-inflated likelihood functions have become relatively common, especially in Bayesian statistics, zero-inflated priors are extremely rare in the statistics literature, and their potential for simultaneous signal detection and analysis does not appear to have been recognized. In particular, models with zero-inflated priors solve the problem for health physicists of how to characterize the count rate of samples which may or may not contain sources. By presenting a compact solution to the two major problems of radiation counting measurements, this approach may also ease the transition of health physicists to Bayesian statistics. In spite of their great utility, models of radioactive sources which include zero-inflated priors are entirely new to the field of health physics.

## Chapter 1

### BAYESIAN TECHNIQUES FOR THE ANALYSIS OF COUNT RATE DATA

#### 1.1 Introduction

Bayes's theorem produces a probability distribution of the parameter(s) of interest by combining experimental data with the investigator's pre-experiment (i.e., prior) knowledge or intuition. In the case of radiation sample measurements, the pre-experiment knowledge about the probability of possible values of the mean source count rate,  $\lambda$ , is expressed as the prior probability distribution,  $p(\lambda)$ . This  $p(\lambda)$  is then multiplied by the likelihood function, which gives the likelihood of having observed  $n$  counts given the possible range of mean source count rates,  $\lambda$  (e.g., the probability of having observed a specific number of counts as a function of the mean decay rate – for counting measurements, the Poisson distribution). The product of these probability distributions is then normalized to give an updated posterior probability distribution of the mean count rate  $\lambda$  – the probability distribution of  $\lambda$  given the observed data. Stated mathematically, given the result of a measurement,  $n$ , taken on a radioactive source with mean decay rate,  $\lambda$ , Bayes's theorem states:

$$p(\lambda|n) = \frac{p(n|\lambda)p(\lambda)}{p(n)}, \quad (1.1)$$

where  $p(\lambda|n)$  (read “the probability of  $\lambda$  given  $n$ ) is the posterior probability distribution,  $p(n|\lambda)$  (read: “the probability of  $n$  given  $\lambda$ ) is the likelihood function of the observed data,  $p(\lambda)$  is the prior, and  $p(n)$  is the total probability of having observed the data under the prior and likelihood distributions (Kruschke 2010). The denominator,  $p(n)$ , serves as a normalizing

constant to the numerator. It is therefore given by the integral of the numerator over the parameter space, so that Bayes's theorem may be re-written,

$$p(\lambda|n) = \frac{p(n|\lambda)p(\lambda)}{\int_0^{\infty} p(n|\lambda)p(\lambda)d\lambda}. \quad (1.2)$$

Because the prior and likelihood function together constitute a 'model' of the mechanisms which led to the observed data, it follows that  $p(n)$  can be interpreted as the posterior likelihood of the model.

## 1.2 The Likelihood Function

Radioactive decay is a Poisson process, meaning that given a radioactive source with mean count rate,  $\lambda$ , the probability of the source producing  $n$  counts in a given interval  $t$  (provided the interval is very short compared to the half-life) is given by the Poisson distribution:

$$p(n|\lambda t) = \frac{(\lambda t)^n}{n!} e^{-\lambda t}. \quad (1.3)$$

When expressed as a function of the number of counts,  $n$ , the Poisson distribution is a normalized probability distribution for the number of observed counts. However, if the Poisson distribution is expressed as a function of mean count rate  $\lambda$ , holding the value of  $n$  constant, it provides the relative likelihood of different values of  $\lambda$  for a given observed count rate,  $n$ . This is the likelihood function of the measurement,  $p(n|\lambda)$ . It is worth noting that the integral of the Poisson distribution over the mean count rate,  $\lambda$ , does not evaluate to one, and is therefore not a proper probability distribution (thus the term "likelihood function" as opposed to "likelihood

distribution”). The likelihood function, like the prior, is a part of the investigator’s model describing radioactive decay. In other words, when specifying a Poisson likelihood function, the investigator is postulating that the number of observed counts is indeed generated by a Poisson process. This is equivalent to postulating that the mean-count rate is fixed, that each count occurs independently of the time since the last event, that the mean of the observed counts is equal to the variance of the observed counts, and that the number of observed counts is centered on only one “peak.” In Chapter 3 we consider a situation in which some of those postulates fail, and propose an alternative likelihood function.

### 1.3 The Prior

The selection of a prior probability distribution is the most controversial and perhaps the most important aspect of applying Bayes’s theorem. For count rate measurements, the prior represents the investigator’s pre-existing knowledge or belief about the possible values of the mean count rate,  $\lambda$ . Because it is not possible to have a negative mean count rate, the prior distribution should exclude negative values.

For the purposes of rapid computation, it is desirable to select a prior which is mathematically tractable. A prior which will produce a posterior of the same parametric form as itself for a given likelihood function is called the conjugate prior for that likelihood. The conjugate prior for the Poisson distribution is the gamma distribution. That means that if the prior is a gamma distribution and the likelihood is a Poisson distribution, the posterior will also be a gamma distribution. (Ntzoufras 2009) The gamma distribution is given by:

$$\text{gamma}(\alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} \lambda^{\alpha-1} e^{-\beta\lambda}, \quad (1.4)$$

where  $\lambda$  is the mean count rate. The mean of the gamma distribution is given by  $\alpha/\beta$ , and the variance by  $\alpha/\beta^2$ . Therefore, the absolute width (e.g. variance) is small when  $\alpha$  is small compared to  $\beta^2$ . The relative width of the distribution can be expressed as the standard deviation divided by the mean, which evaluates to  $\alpha^{-1/2}$ . Therefore, the relative width is small when  $\alpha$  is large, and is not a function of  $\beta$ . For the case of measuring count rate data, the parameter  $\alpha$  can be interpreted as the total number of observed counts, while  $\beta$  is the total time in which those counts occurred. Multiplying the gamma distribution by the likelihood function, it follows that the posterior distribution of  $\lambda$  is given by

$$p(\lambda | n) \propto \text{gamma}(\alpha + n, \beta + t) \quad (1.5)$$

One advantage of choosing a conjugate prior is that the posterior can be used as a prior for subsequent measurements. The likelihood of observing a series of events is equal to the product of the likelihood of each event individually. Therefore, the likelihood of an observation of a specific combination of counts,  $N = (n_1, n_2, \dots, n_N)$ , is given by

$$p(n|\lambda) = \prod_i \frac{\lambda^{n_i} e^{-\lambda}}{n_i!}, \quad (1.6)$$

where  $n_i$  is the number of counts per time observation (Luo et al. 2013). From this, it follows that the posterior distribution of  $\lambda$  is given by

$$p(\lambda|t) \propto \text{gamma}(\alpha + N, \beta + \Sigma t_i), \quad (1.7)$$

where  $t_i$  is the duration of the  $i^{\text{th}}$  measurement.

It follows that after  $N$  total counts spanning  $t$  time intervals (e.g., seconds, hours, etc.) the parameters of the posterior distribution will be  $\alpha + N$  and  $\beta + t$ . If little or no information about the mean count rate is available before the measurement, a prior should be chosen which has minimal effect on the posterior distribution. Such a prior is often called “uninformative.” It is also important to select a prior which assigns some non-zero probability to all possible values of  $\lambda$ . From inspection of Eq. 1.7, this is accomplished by selecting parameter values which are very small. If  $\alpha = 1$  and  $\beta$  is very small, the gamma distribution becomes uniform, meaning that it is constant for all positive values of  $\lambda$ . This is a common choice, but in reality may be biased towards large values of  $\lambda$ . For example, given a uniform prior, the odds ratio of  $\lambda$  being within 10% of 10 vs. being within 10% of 1000 is

$$OR = \frac{P(9 < \lambda < 11)}{P(900 < \lambda < 1100)} = \frac{2}{200} = 1:100. \quad (1.8)$$

A somewhat more objective prior is obtained for  $\alpha = \frac{1}{2}$  and  $\beta \rightarrow 0$ , such that

$$P(\lambda) \propto \lambda^{-1/2}, \quad (1.9)$$

which is the Jeffreys prior for the Poisson mean (Jeffreys 1946). The Jeffreys prior is invariant under reparameterization of  $\lambda$ , making it preferable to the uniform prior for scale-invariant parameters such as the Poisson mean. Either choice of prior is usually acceptable in practice, as both will be quickly drowned out by the observed data.

When a conjugate prior has been chosen, the numerator is already a normalized probability distribution, and there is no need to compute the integral. However, as we shall see,



this is not usually the case for most real-world scenarios. In fact, it is often the case that the integral cannot be evaluated in closed form. In such cases, Markov Chain Monte Carlo techniques must be called upon to evaluate the posterior.

#### 1.4 The Predictive Distributions

The distribution which gives the probability of all possible observations (e.g., every possible number of observed counts,  $n$ ) is called the predictive distribution. In classical statistics, the predictive distribution only exists after a measurement, and is typically given by the likelihood function whose parameters are fixed at the maximum likelihood parameter estimates. For example, after an observation of  $n_{first}$  counts generated by a Poisson process, the maximum likelihood value for the Poisson mean is  $\lambda = n_{first}$ , and the predictive distribution for the next observation,  $n_{next}$ , is given by,  $p(n_{next}|n_{first}) = \text{Poisson}(n_{next}|\lambda = n_{first})$ .

In contrast, Bayesian statistics typically treats the parameters themselves as random variables drawn from a probability distribution (e.g., the prior or the posterior). Because there is some pre-existing knowledge about the possible parameter values, it is possible to generate a predictive distribution for the next observation before any prior observations have been made. This distribution is called the prior predictive. After a measurement, the probability distribution of the parameters is updated, and is now called the posterior distribution. Combining the posterior distribution with the likelihood function gives the posterior predictive distribution. Note that if the prior is conjugate to the likelihood function then the posterior will be of the same form as the prior, and the posterior predictive will therefore be of the same form as the prior predictive.

The probability of a particular observation under a given likelihood-prior combination is given by the probability of that observation under a given parameter value multiplied by the probability of that parameter value, summed over all possible parameter values. For example, the probability of observing a particular number of counts,  $n$ , following a Poisson distribution whose mean is drawn from a gamma distribution with parameters  $\alpha$  and  $\beta$  can be determined by integrating the probability of observing  $n$  counts for a given count rate  $\lambda$ ,  $\text{Poisson}(n|\lambda)$ , over the normalized probability distribution of the count rate,  $\text{gamma}(\lambda|\alpha, \beta)$ :

$$p(n|\alpha, \beta) = \int_{\lambda=0}^{\infty} \text{Poisson}(n|\lambda) * \text{gamma}(\lambda|\alpha, \beta) d\lambda = \int \frac{(\lambda t)^n}{n!} e^{-\lambda t} \times \frac{\beta^\alpha}{\Gamma(\alpha)} \lambda^{\alpha-1} e^{-\beta\lambda} d\lambda. \quad (1.10)$$

This process is called marginalizing over, or integrating out,  $\lambda$ . For that reason, the predictive distribution is also sometimes called the marginal likelihood. The integral evaluates to a normalized parametric probability distribution which, through a simple re-parameterization, takes the form of a negative binomial distribution (Ntzoufras 2009):

$$p(n|\alpha, \beta) = \text{N. B.} \left( r = \alpha, p = \frac{1}{1 + \beta} \right) = \binom{n + \alpha - 1}{n} \left( \frac{\beta}{\beta + 1} \right)^\alpha \left( \frac{1}{\beta + 1} \right)^n. \quad (1.11)$$

The negative binomial corresponding to the gamma prior is therefore the prior predictive, or marginal likelihood, of the observed data under the Poisson likelihood/gamma prior model, which is the normalizing constant in Eq. 1.2. If, after the measurement,  $\alpha$  and  $\beta$  are the parameters of the posterior gamma distribution, Eq. 1.11 provides the posterior predictive distribution for future measurements. The posterior predictive results also hold if the gamma

distribution represents actual variations in the value of  $\lambda$  between measurements. In that case the likelihood function is a negative binomial distribution, and priors may be put on  $\alpha$  and/or  $\beta$ . This situation is considered in Chapter 3.

The mean of the negative binomial distribution is given by  $pr/(1-p) = \alpha/\beta$ , which is the same result one would obtain using the stationary target (Poisson) model. In contrast, the variance of the negative binomial distribution is given by:

$$\frac{pr}{(1-p)^2} = \frac{\alpha}{\beta} + \frac{\alpha}{\beta^2}. \tag{1.12}$$

The variance of the moving target sampling distribution is therefore equal to the sum of the variance of the stationary target sampling distribution ( $\alpha/\beta$ ) and the variance of the gamma distribution from which the Poisson mean is sampled.

## 1.5 Credible Intervals

Credible intervals are the Bayesian analogue of classical confidence interval (Kruschke 2010). They are intervals on the prior or posterior distributions which contain a likely range of values. For example, if there is a 95% chance that  $5 \text{ cps} < \lambda < 10 \text{ cps}$ , then (5, 10) is a 95% credible interval on  $p(\lambda)$ . Note that there are multiple intervals which contain 95% of the probability mass, and therefore multiple valid techniques for computing credible intervals. The preferred technique used in this document is to compute the highest density interval (HDI), that is, the interval of highest average density (equivalently, the narrowest possible interval). Another

common technique is to compute the interval between the 2.5% quantile and the 97.5% quantile. The quantile interval is guaranteed to be continuous and is more straightforward to compute, but does not necessarily contain the most probable values of the mean count rate.

Although some statistical analysis tools have built-in algorithms to compute HDIs, it is useful for understanding to derive the algorithms by hand. Because every point in the 95% HDI is more likely than every point outside of the 95% HDI, the points of a discrete distribution contained in the HDI can be identified by sorting the points from most to least likely, and adding up the probability of these points until the collective probability reaches 95%. This is illustrated in the following algorithm:

#### Discrete HDI Algorithm

Given:

X – A list containing all points in the probability distribution

PX – A list containing the probability of the respective points in list X

Limit – The probability contained by the HDI

Algorithm:

```

PX = [n/sum(PX) for n in PX]           #Normalize the distribution
pairs = [[px,x] for px,x in zip(PX,X)] #Keep track of which probability is associated with which x
sortedpairs = reversed(sorted(pairs))  #Copy the pairs to a list sorted by probability
pairsInHDI = []
sum = 0
for p in sortedpairs:                  #iterates through every probability/value pair
    sum += p[0]                          #integrates the total probability
    pairsInHDI.append(p)                 #Add to the list of pairs in the HDI
    if sum>(HDI*sum(PX)):                 #Stops when the probability is greater than the HDI limit
        break;
#Add each pair in the original list which is in the HDI (Preserves the original ordering of the lists)
HDI = [[x,px] for px,x in pairs if pairsInHDI.contains([px,x])]

```

To compute the HDI of continuous distributions, it is helpful to think of the HDI as the region of land which rests above a rising water line (Kruschke 2010). In this analogy, the 95% HDI is the region of land above water when the water covers 5% of the land. It follows that the

boundaries of the HDI of a continuous distribution,  $dist$ , are given when the integral between the roots (zero's) of  $dist - p$  is equal to 0.95. This technique can involve evaluating complex integrals many times, making it prohibitively slow. However, HDI's of continuous distributions can be estimated quickly and efficiently by discretizing the distribution and using the discrete HDI Algorithm shown above.

## 1.6 Count Rate Analysis

The classic two-count problem (Little 1982) consists of a background measurement of time  $t_B$  and a sample measurement of time  $t_G$ . The number of counts observed in the background measurement,  $N_B$ , derives from a Poisson distribution centered at  $\lambda_B t_B$ , where  $\lambda_B$  is the mean background count rate per unit time. Similarly, the number of observed counts from a sample measurement,  $N_G$ , derives from a Poisson distribution centered at  $(\lambda_b + \lambda_s)t_G$ , where  $\lambda_s$  is the mean sample count rate. The goal is to derive a probability distribution for the mean sample count rate,  $\lambda_s$ .

The first step is to determine the appropriate prior probability distributions to describe  $\lambda_B$  and  $\lambda_s$  (Chalmers and LaBone 2013). Assuming we have no prior knowledge about the background and sample count rates (other than that they are positive), it is traditional to choose priors which are expected to have minimal impact on the posterior distribution, such as the Jeffreys prior:

$$p(\lambda_B) = \text{gamma}\left(\frac{1}{2}, 0.01\right) = \frac{0.01^{0.5}}{\Gamma(0.5)} \lambda_B^{(0.5-1)} e^{-0.01\lambda_B}, \quad (1.13)$$

$$p(\lambda_S) = \text{gamma}\left(\frac{1}{2}, 0.01\right) = \frac{0.01^{0.5}}{\Gamma(0.5)} \lambda_S^{(0.5-1)} e^{-0.01\lambda_S}.$$

Before proceeding, it is worth examining the prior to determine if it really expresses our beliefs about the likely values of the background and sample count rates. The gamma distribution has a 95% credible interval of (0.49, 2512). This very broad credible interval means that a wide range of values are possible under this prior, which may not accurately reflect our belief about the relative likelihood of very large count rates. This prior also assumes the sample count rate is positive, meaning that it rules out the possibility of there being no source. Nonetheless, very different priors will result in very similar posteriors, provided that the prior parameters,  $\alpha_B$ ,  $\beta_B$ ,  $\alpha_G$ , and  $\beta_G$ , are small compared to the observed data and measurement times,  $N_B$ ,  $t_B$ ,  $N_G$ , and  $t_G$ . However, this is not the case when count-rate analysis is combined with hypothesis testing.

The next step in solving the two-count problem is to write down the likelihood function. As discussed above, radiation count data are typically assumed to be Poisson distributed. Therefore, the likelihood function of the background measurement is

$$p(N_B | \lambda_B) = \frac{(\lambda_B t_B)^{N_B}}{N_B!} e^{-\lambda_B t_B}. \quad (1.14)$$

Similarly, the likelihood function for the gross counts from the sample measurement is

$$p(N_G | \lambda_B, \lambda_S) = \frac{((\lambda_B + \lambda_S) t_G)^{N_G}}{N_G!} e^{-(\lambda_B + \lambda_S) t_G}. \quad (1.15)$$

The likelihood function of the combined measurements, also known as the joint likelihood, is the product of the likelihoods of the two individual measurements:

$$p(N_B, N_G | \lambda_B, \lambda_S) = p(N_G | ((\lambda_B + \lambda_S)t_B)) \times p(N_B | \lambda_B t_B). \quad (1.16)$$

Bayes's Theorem (Eq. 1.2) requires that the posterior probability distribution is proportional to the product of the prior distribution and the likelihood function. Therefore,

$$p(\lambda_B, \lambda_S | N_B, N_G) \propto p(N_G | ((\lambda_B + \lambda_S)t_B)) \times p(N_B | \lambda_B t_B) \times p(\lambda_B) \times p(\lambda_S), \quad (1.17)$$

Where  $\propto$  is read 'is proportional to.' Eq. 1.17 is the un-normalized joint posterior of the background and sample count rate. A normalized probability distribution is equal to 1 when integrated over the parameter space. It follows that the normalizing constant is equal to

$$p(N_B, N_G) = \int_0^\infty \int_0^\infty p(N_G | ((\lambda_B + \lambda_S)t_B)) \times p(N_B | \lambda_B t_B) \times p(\lambda_B) \times p(\lambda_S) d\lambda_B d\lambda_S, \quad (1.18)$$

and the normalized posterior distribution is

$$p(\lambda_B, \lambda_S | N_B, N_G) = \frac{p(N_G | ((\lambda_B + \lambda_S)t_B)) \times p(N_B | \lambda_B t_B) \times p(\lambda_B) \times p(\lambda_S)}{p(N_B, N_G)}. \quad (1.19)$$

Given the normalized posterior, it is now possible to compute the expectation value of any function of  $\lambda_B$  or  $\lambda_G$ . For example, the expectation value of the source count rate (i.e., the mean) is given by

$$\bar{\lambda}_S = \int_0^\infty \lambda_S \int_0^\infty p(\lambda_B, \lambda_S | N_B, N_G) d\lambda_B d\lambda_S. \quad (1.20)$$

Similarly, the variance of the sample count rate is

$$\text{Var}(\lambda_S) = \int_0^\infty (\lambda_S - \bar{\lambda}_S)^2 \int_0^\infty p(\lambda_B, \lambda_S | N_B, N_G) d\lambda_B d\lambda_S . \quad (1.21)$$

The posterior can also be used to calculate credible intervals for the background and sample count rates. Note that, unlike classical confidence intervals, these cannot include negative values, as these are excluded by the prior. However, the prior also rules out the possibility that the sample contains no source at all (i.e., that the sample count rate is exactly zero). Therefore, the posterior can be interpreted as the probability distribution of the background and sample count rates given that the sample contains a source.

It is immediately obvious that the most challenging aspect of calculating the posterior is evaluating the normalization constant. In this idealized example, which neglects uncertainty on, for example, detector calibration and energy response, it is possible to evaluate the integral using standard numerical integration techniques. However, these techniques are inadequate for more complex models, and must be replaced with Monte Carlo integration methods. These have the advantage of allowing the analysis of a model by simply writing it down (i.e., writing down the prior and likelihood functions) and entering the observed data, bypassing complex calculations altogether. However, Markov Chain Monte Carlo (MCMC) techniques can only produce estimates of the posterior, and care must be taken to ensure that these estimates have converged. A common, freely available software package for performing MCMC is called OpenBUGS (Lunn et al., 2000). Other than the observed data, MCMC software such as OpenBUGS only requires input of the “generative model” of the data.

A generative model is a representation of the processes which are believed to have led to the observed data. For the above example, the mean count rates are random variables drawn from



gamma distributions, and the observed data are random variables drawn from Poisson distributions centered at those means. Therefore, the generative model for the data would read:

Classic Two-Count Model:

$$\begin{array}{ll}
 \lambda_B \sim \text{gamma}(\alpha = 0.5, \beta = 0.01), & \# \text{Mean background count rate} \\
 \lambda_S \sim \text{gamma}(\alpha = 0.5, \beta = 0.01), & \# \text{Mean source count rate} \\
 \lambda_G = \lambda_B + \lambda_S, & \# \text{Mean gross count rate} \\
 N_B \sim \text{Poisson}(\lambda_B t_B), & \# \text{Counts observed in background measurement} \\
 N_G \sim \text{Poisson}(\lambda_G t_G). & \# \text{Counts observed in sample measurement.}
 \end{array}$$

Where the ‘ $\sim$ ’ symbol represents a stochastic relationship, while the ‘=’ symbol represents a deterministic relationship. This model indicates that  $\lambda_B$  and  $\lambda_S$  are random variables drawn from gamma distributions. In contrast, the gross count rate,  $\lambda_G$ , can be precisely determined given the background and sample count rates. The observed values,  $N_B$  and  $N_G$  are also known to be random samples drawn from Poisson distributions. In addition to being necessary for OpenBUGS, generative models are powerful tools for understanding and communicating complex statistical models. A detailed example of how to analyze data using generative models and OpenBUGS is given in the Appendix.

**1.7 Bayesian Hypothesis Testing**

The above analysis provides a probability distribution for the sample and background count rates given that the sample contains a source. It can be shown that if a continuous, differentiable prior is used, analysis of the posterior distribution will always result in the probability of there being no source as strictly equal to zero. The traditional Bayesian approach to hypothesis testing is called Bayesian model comparison (Kruschke 2010). In this approach,

the hypotheses themselves are treated as parameters. In order to calculate the likelihood function of the observed data, the hypotheses must first be defined mathematically. In the case of source detection, the natural choice of hypotheses is, *Hypothesis 1 = No Source*, *Hypothesis 2 = Source*. These will be referred to as the *No Source* hypothesis and the *Source* hypothesis, respectively. These can be expressed mathematically as different prior probability distributions on the mean source count rate. If *Hypothesis 1* is correct then there is no source and the mean count rate is zero. If *Hypothesis 2* is correct then there is a source but (assuming we don't know what we're looking for) the mean source count rate could assume any value permitted by the prior.

For a given hypothesis, Bayes's Theorem (Eq. 1.1) states that, after an observation of  $n$  counts, the posterior probability of that hypothesis being correct is

$$p(\text{Hypothesis}|n) = \frac{p(n|\text{Hypothesis})p(\text{Hypothesis})}{p(n)}. \quad (1.22)$$

Given an observation of  $n$  counts, Eq. 1.22 establishes that the probability of a given hypothesis (e.g., source or no source) is equal to the probability that the hypothesis was correct before the observation,  $p(\text{Hypothesis})$ , multiplied by the likelihood of observing  $n$  counts if that hypothesis was correct,  $p(n|\text{Hypothesis})$ , divided by the probability of observing  $n$  counts under any of the hypotheses,  $p(n)$ . If we are only interested in the relative probability of two hypotheses, the normalizing constant,  $p(n)$ , drops out, so that the Probability Ratio,  $PR$ , of the two hypotheses given the data is the product of the marginal likelihood ratio, or Bayes Factor ( $BF$ ), with the prior odds ratio:

$$PR = \frac{p(\text{Source}|N_B, N_G)}{p(\text{No Source}|N_B, N_G)} = \frac{p(N_B, N_G|\text{Source})}{p(N_B, N_G|\text{No Source})} \times \frac{p_{\text{Source}}}{1 - p_{\text{Source}}}. \quad (1.23)$$

In this case, there are only two hypotheses, so the probability of the two hypotheses must add up to 1. It follows that

$$\begin{aligned} p(\text{Source}|N_B, N_G) &= \frac{PR}{1 + PR} \\ p(\text{No Source}|N_B, N_G) &= \frac{1}{1 + PR}. \end{aligned} \tag{1.24}$$

Therefore, computing the probability ratio is equivalent to computing the posterior probabilities of the two hypotheses. The posterior odds of the two hypotheses is a direct function of the prior odds. This means that the prior will always have a strong impact on the posterior, regardless of the amount of data collected. It follows that there is no such thing as an “uninformative” prior on the likelihood of a given hypothesis. However, note that an order of magnitude difference in posterior probabilities obtained from different priors updated with a large amount of data,  $p(\text{Source}|N_B, N_G) = 0.99$  and  $p(\text{Source}|N_B, N_G) = 0.999$ , may not be significant in terms of decision making. Regardless, the investigator must use professional judgment and past experience to determine realistic priors. For example in the case of internal dosimetry, where the hypotheses might be *Intake* and *No Intake*, the investigator could use historical data about the fraction of incidents which have resulted in actual intakes of radionuclides.

Bayesian hypothesis testing as formulated in Eq. 1.23 only provides a point estimate for  $PR$ . This is due to the fact that the prior odds ratio was assumed to be precisely known. In some cases it is reasonable to assume that the prior probability ratio,  $\frac{p_{\text{Source}}}{1 - p_{\text{Source}}}$ , is exactly known – for example, if the labels were misplaced on two samples, only one of which was known to contain a source. Note that choosing  $p_{\text{Source}} = 0.5$  is equivalent to putting a uniform prior from 0 to 1 on  $p_{\text{Source}}$ , in which case, from Eq. 1.23, the Posterior probability ratio is equal to the Bayes

Factor. However, in many cases it may only be possible to estimate the probability ratio based on empirical observations or professional judgment, and in some cases the probability ratio may be completely unknown. In such cases, the prior uncertainty on the probability ratio can be expressed as a probability distribution,  $p(PR)$ , in which case Bayes's theorem gives:

$$p(PR|N_B, N_G) = \frac{p(N_B, N_G|PR) \times p(PR)}{p(N_B, N_G)}, \quad (1.25)$$

where, from Eq 1.23, the likelihood function is

$$p(N_B, N_G|PR) = \frac{p(N_B, N_G|Source)}{p(N_B, N_G|No Source)} \times \frac{p(Source|PR)}{p(No Source|PR)} = \frac{p(N_B, N_G|Source)}{p(N_B, N_G|No Source)} \times PR, \quad (1.26)$$

And, from the denominator of Eq. 1.2, the normalizing constant is calculated by integrating the numerator over the possible values of  $PR$ :

$$p(N_B, N_G) = \int_0^1 \frac{p(N_B, N_G|Source)}{p(N_B, N_G|No Source)} \times PR \times p(PR) dPR, \quad (1.27)$$

Note that putting a uniform prior on  $p_{Source}$  is not equivalent to putting a uniform prior on the probability ratio.

Both the *Source* and *No Source* hypotheses correspond to prior probability distributions for the sample count rate. Because  $N_B$  does not depend on the hypotheses, the background measurement can be used to update the prior on  $\lambda_B$ . From Eq. 1.7, the posterior distribution of  $\lambda_B$  after the background measurement becomes

$$p(\lambda_B|N_B) = \text{gamma}(0.5 + N_B, 0.01 + t_B). \quad (1.28)$$

The likelihood of the sample measurement under a given hypothesis (a.k.a. the marginal likelihood) is equal to the sum of the likelihood for a given value of  $\lambda_G$ , multiplied by the probability of  $\lambda_G$  under the given hypothesis. In the case of the *No Source* hypothesis, this is

$$p(N_B, N_G|No\ Source) = \int_0^\infty \text{Poisson}(N_G|\lambda_B)p(\lambda_B|N_B)d\lambda_B. \quad (1.29)$$

This is exactly the same as the denominator (normalizing constant) from Eq. 1.2, except that the prior of  $\lambda_B$  has been updated by the background measurement.

The *Source* hypothesis corresponds to a prior probability distribution representing the investigator's belief about the likely values of  $\lambda_S$  if the sample contains a source. The likelihood function is thus

$$\begin{aligned} p(N_B, N_G|Source) \\ = \int_0^\infty \int_0^\infty p(N_G|\lambda_G = \lambda_B + \lambda_S) \times p(\lambda_S) \times p(\lambda_B|N_B)d\lambda_B d\lambda_S. \end{aligned} \quad (1.30)$$

Note the similarity between this and the normalizing coefficient given by Eq. 1.18. As in the count rate analysis, these equations result in a rather unwieldy formula for the posterior odds which needs to be evaluated numerically or using MCMC.

An important characteristic of Bayesian hypothesis testing is that the Bayes Factor naturally favors the simpler hypothesis (Ntzoufras 2009), and is thus exquisitely sensitive to the priors used by the two models (e.g.,  $p(\lambda_S)$ ). In particular, *BF* is highly sensitive to the magnitude of the variance of the priors (Bartlett 1957). Assuming nothing is known about the count rate of

the source if it exists, the traditional choice would be to represent it with the Jeffreys prior, such that,

$$p(\lambda_S = 0|No\ Source) = 1, \quad (1.31)$$

$$p(\lambda_S|Source) = \text{gamma}\left(\alpha = \frac{1}{2}, \beta = 0.01\right).$$

However, as we saw earlier, this formulation of the Jeffreys prior corresponds to a 95% credible interval for the source count rate of (0.49, 2512), with a mean (expected) count rate of 50. This model is clearly unrealistic, and will therefore be prone to fail, even if there seems to be strong evidence in favor of a source. Because the Jeffreys prior implies taking the limit as  $\beta \rightarrow 0$ , and the variance of the gamma distribution is inversely proportional to the square of  $\beta$ , decreasing  $\beta$  by a factor of ten, such that  $p(\lambda_S|Source) = \text{gamma}\left(\alpha = \frac{1}{2}, \beta = 0.001\right)$ , greatly increases the prior probability of very large source count rates, and thus greatly reduces the Probability Ratio in favor of the *No Source* hypothesis. Using a uniform prior  $p(\lambda_S|Source)$  is even worse, as it is equivalent to stating that there is a nearly equal probability of the source activity being 1 Bq and 100 MBq. This illustrates that, rather than expressing the investigators' lack of knowledge, uniform priors actually imply strong and irrational assumptions on the part of the investigator. It follows that models which utilize uninformative priors for the source count rate will be prone to fail in Bayesian hypothesis testing. Crucially, different "uninformative" priors on the source count rate will result in different posteriors, regardless of the magnitude of the observed data. Furthermore, models which utilize improper priors (i.e., priors which do not integrate to 1) such as the uniform prior and the Jeffreys prior will always fail, regardless of the observed data. As with the priors on the different hypotheses, priors on the sample count rate are

best obtained using professional judgment and historical data. These can be updated formally using Bayes's theorem on the measurement of multiple different samples.

The prohibition of improper priors may be considered a feature, as hypothesis testing prohibits unrealistic models, but it also raises the question of how to identify a prior which is realistic but at the same time adequately represents the uncertainty on the source count rate. Ideally, priors should be based on empirical data from past observations or prior knowledge (e.g., we know we are looking for a 10 cps source). If that is not possible, a simple option is to use the Posterior Bayes Factor (*PBF*), which compares the likelihood of the data given the posterior distributions,

$$PBF = \frac{p(N_B, N_G | N_B, N_G, Source)}{p(N_B, N_G | N_B, N_G, No Source)} \quad (1.32)$$

The *PBF* allows for the use of uninformative priors, but the practice of evaluating the probability of the data given the data is problematic and has been criticized as theoretically unsound (Ntzoufras 2009).

A number of alternatives have been suggested to the *PBF* (e.g., O'Hagan 1995, Berger and Pericchi 1996). The most straightforward approach is to establish an informative prior on the source count rate by using a short independent measurement. In other words, take a short sample measurement, collecting  $N_{G0}$  counts in time  $t_{G0}$ , and, using Eq. 1.19, calculate an intermediate posterior distribution of the sample count rate. This distribution can then be used as the prior for the posterior odds ratio. This measurement should be as short as possible relative to the main sample measurement so that the intermediate posterior will have minimal impact on the final posterior distribution.

If two sample measurements are taken, the ‘Intermediate’ Bayes Factor is given by

$$IBF = \frac{p(N_B, N_{G1} | N_B, N_{G0}, Source)}{p(N_B, N_{G1} | N_B, N_{G0}, No Source)}. \quad (1.34)$$

This technique may work particularly well if a “half-Cauchy” prior (that is, a Student’s  $t$  distribution with one degree of freedom truncated at zero) is put on  $\lambda_S$  before the initial sample measurement (Gelman 2006). However, a detailed examination of such techniques is beyond the scope of this research. There is no universally accepted method for selecting a prior in the absence of empirical data. When empirical data are not available, priors should be based on scientific judgment, excluding values which are ruled out by theory (such as negative count rates), and reflecting as accurately as possible the beliefs of the investigator. Uninformative priors are intended to have as little impact as possible on the posterior distribution, with the goal of obtaining an “objective” posterior. However, objective priors are often not believable or even reasonable models of the parameter space. The rejection of models utilizing uninformative priors is a reflection of this fact.



## Chapter 2

# SEQUENTIALLY UPDATED BAYESIAN DETECTION SYSTEM UTILIZING ENERGY AND COUNT RATE DATA FOR THE DETECTION OF RADIOACTIVE SOURCES<sup>1</sup>

### 2.1 Introduction

In this chapter, we consider a particle counting and detection system which searches for elevated count rates in multiple energy regions simultaneously. The system analyzes time-interval data (e.g., time between counts), as this was shown to be a more sensitive technique for detecting low count rate sources compared to analyzing counts per unit time (Luo et al. 2013). Two distinct versions of the detection system are developed. The first is intended for situations in which the sample is fixed and can be measured for an unlimited amount of time. The second version is intended to detect sources which are physically moving relative to the detector, such as a truck moving past a fixed roadside detector, or a waste storage facility under an airplane. In both cases, the detection system is expected to be active indefinitely – i.e., it is an on-line detection system. Both versions of the multi-energy detection system are compared to their respective gross count rate detection systems in terms of Type I and Type II error rates and sensitivity.

By characterizing and accounting for variations in the background, general purpose radiation detectors can be improved with little or no increase in cost. The statistical and computational techniques to perform this kind of analysis have been developed. The necessary signal analysis can be accomplished using existing Bayesian algorithms which account for

---

<sup>1</sup> This chapter is based on a manuscript which has been accepted for publication by the journal ‘Health Physics.’

multiple channels, multiple detectors, and multiple time intervals. Furthermore, Bayesian machine-learning techniques have already been developed which, with trivial modifications, can generate appropriate decision thresholds based on the comparison of new measurements against a non-parametric sampling distribution.

A detection system which is always on – e.g., an on-line detection system – and which has any chance of making a detection decision in a given time interval will always eventually produce an alarm, regardless of whether or not a source is present. Therefore, false-positive and false-negative rates of on-line detectors can only be evaluated over fixed measurement intervals, and are not the most useful metrics for evaluating the system. An effective on-line radiation detection system should 1) quickly detect and identify changes in radiation levels with a small number of failed detections, and 2) be capable of running continuously with a very small number of false positive identifications. These qualities can be quantified using the average run length (ARL), which is defined as the average time interval between detection decisions at a given source count rate. The ARL can be determined empirically by activating the detector and measuring the time interval between detection decisions (with the detection system being reset after each detection decision). Therefore, the background ARL can be understood as the expected frequency of alarms for an online detector measuring background (e.g. false alarms). Similarly, the ARL in the presence of a source represents the expected amount of time it will take to detect the source. The fundamental challenge in the detection and identification of low count rate sources is the natural variation in count rates from both the background and the source.

Because of the nature of radioactive decay, the number of counts in a given time interval from any fixed source of radiation will follow a Poisson distribution which, for high enough

counts, can be well approximated by a Normal distribution. Virtually all classical decision thresholds are based on the assumption that the number of counts from background is also sampled from a Normal distribution. However, it is well known that background varies with time of day, location, detector performance, meteorological conditions, and many other factors (Hemic 1988). Regardless, the natural variation of counts from the background and source results in some probability of error in any detection decision. In any detection system, there will be a trade-off between false positive and false negative rates. The appropriate balance between the two types of errors depends on the application and preferences of the user, whose decision is embodied in the decision threshold.

In a classical radiation monitoring set-up, one assumes a known fixed background, decides on a decision threshold, and compares the results of a measurement to that decision threshold. If the observed counts in a given time interval is greater than the decision threshold, the observer concludes that a source has been detected. Traditionally, the false-positive rate is controlled to 5% by setting the decision threshold to mean background +  $1.645\sigma$ . This set-up is easy to understand and implement, but has the disadvantage of unnecessarily discarding past measurement data. Similarly, on-line detectors often employ techniques, such as moving averages and recursive digital filters, which are designed to discard and de-emphasize old data (Fehlau 1993). For that reason, alternative techniques have been developed to more completely make use of the available data (Brandl 2013). Specifically, there has been increasing interest in Bayesian techniques, which allow for the formal incorporation of prior belief, data, and knowledge, and which produce probabilities for the underlying parameter (i.e., mean count rate). This is in contrast to classical techniques, which do not have the ability to incorporate prior data,

and which can only make statements about whether an observation was drawn from a particular distribution.

The Health Physics literature has discussed Bayesian techniques for the analysis of low count rate radiation since 1982, when Bayesian analysis was applied in situations where classical analysis produced negative estimates of total activity (Little 1982). In these situations, a prior distribution was used with zero probability of negative values. This resulted in meaningful estimates of net count rates. This work was continued for applications in environmental monitoring and internal dosimetry (Miller et al 1993, 2001, 2002), and in 2003 it was suggested that ISO adopt new methods for the evaluation of uncertainty of measurements to provide consistency from a Bayesian viewpoint (Kaeker and Jones 2003; ISO 2000). Subsequently, Bayesian techniques were used to more accurately measure uncertainty in radiation measurements, and further revision of some parts of the ISO 11929 standard was suggested (Weise et al. 2006). Bayesian techniques have also been developed for reducing false-positive rates in evaluation of low count rate data (Strom et al. 2001). False positive and false negative rates have been favorably compared between classical control chart methods (the Shewart and cumulative sum charts (Berger 1986)) and a Bayesian Shiryaev-Roberts (Shiryaev 1961; Roberts 1966) control chart, and the Bayesian method was shown to be the best way to control false positives compared to both control charts (DeVol et al. 2009). Recently, techniques have been developed for the Bayesian analysis of time-interval data (Luo et al. 2013). These techniques provide meaningful, non-negative estimates of net radioactivity, and compare favorably to classical methods in terms of sensitivity and false positive rate. The techniques explored in this chapter are largely an extension of those developed by Luo et al.

Radiation data are delivered one count at a time by the detector. Most detection systems record the number of counts detected in a given time interval. This method is technically easier to handle, and the variation in counts per time interval is typically well characterized by a Poisson distribution. An alternative strategy for recording count data is to deal with the time-intervals between individual counts. Several authors have discussed this technique (e.g., Dowdy et al. 1978; Arandjelovic et al 2002), which was shown to be superior in cases where source information is limited (Luo et al. 2013). The reason for this is that a statistically significant number of pulses may often be gathered in a time shorter than any arbitrarily chosen fixed count time.

An important challenge in the use of Bayesian analysis of radiation count data is the incorporation of energy information. The detection system developed by Luo et al. considers the gross count rate. However, background radiation is derived largely from known sources, and thus has a specific energy distribution, or “fingerprint.” Therefore, there clearly is much to be gained by characterizing the count rate distribution independently in multiple energy regions. A high observed count rate in an energy region in which background is low would represent strong evidence of a source. Conversely, a high count rate in an energy region typically inhabited by background is much weaker evidence. This suggests that a detection system which analyzes energy data will reduce false negatives and false positives.

## **2.3 Materials and Methods**

### **2.3.1 Theoretical Considerations**

The proposed on-line detection system begins with a prior probability distribution for the sample count rate centered at the mean background count rate, which is known with minimal

uncertainty. The prior is then updated sequentially as new information is available, and makes a detection decision (i.e., an alarm) if the posterior probability that the sample count rate is less than background drops below some pre-determined threshold. It is therefore important to select a prior which assigns some non-zero probability to all possible values of  $\lambda_S$ , but which will also quickly take on the characteristics of the observed data. Without loss of generality, suppose the mean background count rate is equal to 2 counts per time interval. In that case, an initial prior of  $\text{gamma}(2,1)$ , which is equivalent to 2 counts observed in one time interval, is a reasonable approximation for background, and can be considered to be a ‘low information’ prior which allows the observed data to dominate the posterior distribution (Luo et al 2013).

After each observation, the probability that the true count rate,  $\lambda_S$ , is above a given background count rate,  $\lambda_B$ , is inferred from the posterior distribution. This allows the investigator to set a decision threshold. If the probability of  $\lambda_S > \lambda_B$  is higher than the decision threshold, a detection decision is made. The detector will continue to collect new observations until a detection decision is made or the user ends the measurement. For the purposes of detecting transient sources, it is important that the prior never obtain so much information as to drown out new signal. This can be solved via the adoption of a ‘moving prior’ technique in which a maximum amount of observation time is permitted to be used in the posterior distribution. As new data are gathered, the oldest observations are discarded so that the total observation time incorporated in the posterior is constant. Unlike previously suggested moving prior techniques (Luo et al 2013), which limit the number of observations included in the posterior, this technique guarantees that the posterior will only incorporate data collected during a fixed time interval. Detection systems which refresh their priors in this way will be referred to as a “moving prior detection systems,” while detection systems which incorporate all of the gathered data will be

referred to as “fixed source detection systems.” The sequential update technique can be understood as an adaptation of the Kalman Filter for change-point analysis of count rate data (Kalman 1960, Welch and Bishop 1995, Taylor 2000).

In extrapolating this technique to multiple energy regions, separate background count rates and prior count rate distributions are assigned to each region. A simple “proof of concept” algorithm is defined for a gamma ray detection system operating in a 2 cps background. The algorithm considers two energy regions: 0 – 511 keV (Low), and 511+ keV (High). Given that most background counts will fall into the low energy region, initial backgrounds of  $\lambda_{SLow} = 1.5$  counts/interval and  $\lambda_{BHigh} = 0.5$  counts/interval are assigned. Accordingly, the initial priors are assigned

$$\begin{aligned} p(\lambda_{SLow}) &= \text{gamma}(1.5, 0.5) \\ p(\lambda_{SHigh}) &= \text{gamma}(0.5, 0.5). \end{aligned} \tag{2.1}$$

These priors are then updated independently with new data. A detection decision is made if the probability of  $\lambda_S > \lambda_B$  is higher than the decision threshold in either energy region. The choice of energy regions will naturally depend on the type of radiation being measured. We have chosen energy regions which correspond to a gamma-ray detector, but in principal this technique could be applied equally well to alpha, beta, or even neutron detectors.

Finally, the algorithm is extended to consider eleven separate energy regions. For simulated data, the relative activities in each region are adjusted to correspond to published average background fluence rates (Hemic 1988). For real-world data, fluence rates can be

estimated using observed background count rates. In this case, it is important to remember that such measurements are specific to the location and detector.

In order to maintain a constant false positive rate in a fixed source detection system, the decision threshold in each energy bin must be scaled by the total number of bins. Applying the Dunn-Sidak correction (Sidak 1967), a detector with  $n$  bins must have the following decision threshold in each bin in order to maintain a constant false positive rate:

$$(1 - \alpha_{bin}) = (1 - \alpha_{tot})^{1/n} . \quad (2.2)$$

Therefore, a classical two-energy detector where the decision threshold in each bin is 97.5% is guaranteed to have a total false-positive rate less than or equal to a gross count rate detector with a decision threshold of 95% (Sidak 1967). However, the false-positive rate in a Bayesian detector is dependent on the average count rate in each energy bin, as well as the choice of prior. In a moving-prior detector, the false positive rate is also a function of how much information is stored in the prior at a given time (i.e., the moving prior limit), and how long the detector is run. Because moving prior detection systems are designed to ‘reset’ themselves after a certain amount of time, an indefinite run time is equivalent to performing an infinite number of independent measurements. Therefore, there is some probability of a false detection during each time interval, and the false positive rate of any moving prior will increase with run time. The exact relationship between decision thresholds and false positives has yet to be thoroughly examined, and in practice decision thresholds are chosen empirically to provide the desired false-positive rates.



### 2.3.2 Instrumentation

All statistical algorithms and simulations were developed using the Python programming language (Van Rossum and Drake 2001), with heavy reliance upon the scientific plug-in SciPy (Jones et al. 2001). Three distinct gamma-ray detector models were tested. The first was a replica of the Luo gross count rate detector (Luo et al. 2013). The second and third were the two-energy and eleven-energy detectors described above. Because the huge quantity of radiation data required to comprehensively test these would be prohibitive in terms of acquisition time and computer memory, the algorithms were initially tested using computer generated data. A Monte Carlo method was used to simulate background radiation, as well as the radioactive decay of  $^{60}\text{Co}$ . In order to facilitate study of the relationship between decision thresholds in each bin with average run length, these simulations use an idealized detector with complete energy absorption and perfect resolution. This causes all of the decays from  $^{60}\text{Co}$  fall into the 511+ keV bin for the two-energy detector, and the 1 MeV bin for the eleven-energy detector. However, it is worth noting that, due to incomplete energy absorption,  $^{60}\text{Co}$  would cause an increased count rate in lower energy bins in a real-world detector.

The distribution of the background data and the performance of the algorithms were validated using measurement data from a NaI scintillation detector. The detector was attached to a Canberra LYNX<sup>TM</sup> amplifier/multi-channel analyzer, which was controlled using the ProSPECT workstation environment. The background was measured at 1 ms precision for approximately 3600 ms, with an average count rate of 0.4 counts  $\text{ms}^{-1}$  (420 cps). These data confirmed both the assumed energy distribution of the background, as well as the predicted performance of the detectors in the absence of any source.

The experimental and simulated data were analyzed using the Luo time-interval method (Luo et al. 2013) and its corresponding energy-region analysis method. The performance of fixed source detection systems was evaluated using average run length (ARL). A long ARL when measuring background is indicative of a low false positive rate, while a short ARL when measuring a source indicates a low false negative rate. In contrast, moving prior detection systems were evaluated in terms of detection probability (which represents sensitivity when measuring a source and false positive rate when measuring background) during fixed interval measurements. Detection probability and false-positive rates were also tested as a function of moving prior length (i.e., how much data was collected in the moving prior).

## **2.4 Results and Discussion**

Gross count rate detectors were compared to two-energy and eleven-energy detectors in fixed source and moving-prior modes using simulated and laboratory data. Simulated data were used to validate previous findings about the performance of gross count rate detectors by replicating the exact conditions reported in those studies (Luo et al. 2013), as well as to investigate the performance of multiple-energy detectors under a broad range of carefully controlled conditions. Experimental data were used to confirm predictions about detector performance derived from the simulated data.

### **2.4.1 Fixed source detection systems:**

Average run lengths (ARLs) of fixed source detection systems were calculated as a function of source count rate using simulated and experimental data. The detectors performed

similarly on simulated and laboratory data. The decision threshold for the gross count rate detector was set at 95%, meaning that a detection decision was made when the true mean count rate from the posterior distribution had a 5% or less probability of being less than or equal to the mean observed background count rate. The decision threshold in each bin of the multi-energy detectors was set according to the Sidak correction. Therefore, the two-energy detector decision threshold was 97.5%, and the eleven-energy detector decision threshold was 99.5% in each bin. The mean count rate in each energy region was given a prior of  $\text{gamma}(\lambda_{Be}, 1)$ , where  $\lambda_{Be}$  is the observed background count rate in that energy region.

Fig. 2.1 shows plots of gross count rate vs. ARL for the three detectors on simulated data.

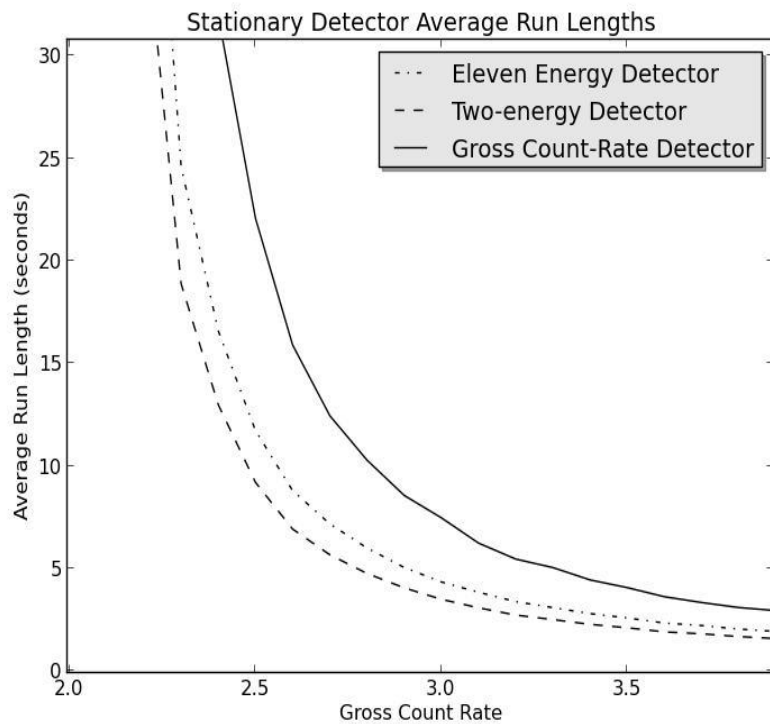


Figure 2.1: Average run lengths of fixed source detection systems on simulated data plotted as a function of gross count rate (background = 2 cps). The detectors performed similarly on simulated and laboratory data.

When the mean gross count rate approaches background (2 cps), the ARLs of all three detectors approach infinity. This is due to the fact that as the posterior collects more data, it becomes less sensitive to random fluctuations in the observed counts, and the probability of a detection decision in any given time interval rapidly approaches zero as the duration of the measurement increases. When the mean count rate is exactly equal to background, the probability of a detection decision (i.e., a false positive) approaches zero so rapidly that there are instances where a detection decision is never made. This results in an infinite ARL regardless of the fraction of runs which actually result in false positives (if one run is infinitely long, the ARL will be infinite, regardless of the outcome of the other runs). Therefore, the ARL is not a useful statistic for evaluating the false-positive rate of fixed source detection systems. This is also indicative of the reason that fixed source detection systems are not adequate for applications in which the gross count rate is likely to change. When the gross count rate increases, the ARLs of all three detectors drop very rapidly. As expected, the ARLs of the energy-spectrum detectors drop more quickly and are shorter for all source count rates compared to the gross count rate detector, with the two-energy detector performing somewhat better than the eleven-energy detector. This suggests that the Sidak correction may be too conservative for the Bayesian multi-energy detectors. At 3 cps, which corresponds to a source count rate of 1 cps, the ARL of the two-energy detector is almost 50% less than that of the gross count rate detector, which confirms the first research hypothesis regarding the multi-energy detector method. The shorter ARL's of the multi-energy detectors indicate that they would, in a fixed time interval, be able to reliably detect lower count rate sources than the gross count rate detector.

#### **2.4.2 Moving prior detection systems:**

The detection probabilities of moving prior detection systems on transient radioactive sources were calculated using experimental data. To simulate a transient source, each detector was repeatedly exposed to 5 s of background data, followed by 5 s of source data, followed by 5 additional seconds of background data, with the source strength varied from 0 to 4 cps, and background fixed at 2 cps. When the source count rate is 0, the detection probability is equivalent to the false-positive rate. For all detectors, the moving prior limit was set to 5 s.

Although there is no universally acceptable failed-detection and false-positive rate, the relative performance of different detectors can be evaluated by controlling the detectors to produce equal detection probabilities at a given count rate and comparing the detection probabilities at a different count rate. For example, a common metric for comparing detectors is the lower limit of detection (LLD), defined as the weakest source strength which will be detected 95% of the time when the false-positive rate is controlled to 5%.

The LLD of each detector was determined by controlling the false positive rate to 5%, and reading off the source strengths at which the detection probability for each detector was 95%. This is illustrated in Fig. 2.2, which plots the average detection probability of each detector as a function of gross count rate (background + source) when the false-positive rate is limited to 5%. By this metric, the LLD of the multi-energy detectors is approximately 1.5 cps. In contrast, the gross count rate detector only detected approximately 60% of the 1.5 cps sources, and did not obtain a 5% failed-detection (e.g., non-detection) rate until the source strength reached 3 cps. In this scenario, the eleven-energy detector has the highest detection probability at all source count rates, closely followed by the two-energy detector, with both significantly out-performing the gross count rate detector. The relatively small performance gain in the eleven-energy detector

may be due to the unequal detection probability in each energy bin owing to the different flux rates in each region.

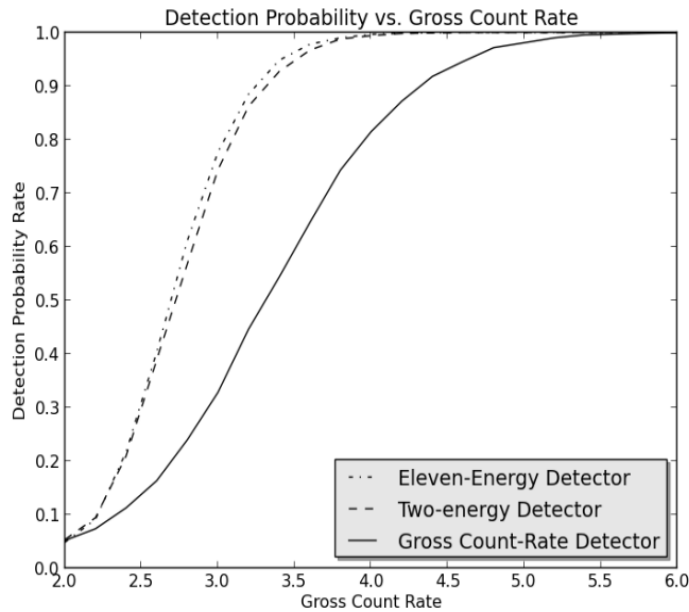


Figure 2.2: Detection probabilities for the three moving prior detection systems as a function of gross count rate with decision thresholds chosen to control the false positive rate to 5%. The background count rate was 2 cps.

Conversely, when the failed-detection (e.g., non-detection) rate for 1.5 cps sources is controlled to 5%, the eleven-energy detector has the lowest false-positive rate, with both multi-energy detectors again significantly out-performing the gross count rate detector. In this case, the false positive rates of the multi-energy detectors are relatively unchanged at approximately 5%, while the false-positive rate of the gross count rate detector dropped to approximately 45%. This is illustrated in Fig. 2.3. As predicted, the multiple-energy detectors significantly out-perform the single-energy detectors in terms of false positives and failed detections.




Figure 2.3: Detection probabilities for the three moving prior detection systems as a function of gross count rate with decision thresholds chosen to control the failed detection rate of 1.5 cps sources to 5%. The background count rate was 2 cps.

## 2.5 Conclusion

Time-interval analysis for particle detectors compares favorably in terms of sensitivity and false positive rate when compared to classical methods. Incorporation of energy data decreases the ARL of fixed source detection systems, effectively allowing them to detect much weaker sources in a given time interval than gross count rate detectors. Incorporation of energy data also significantly improves performance when monitoring for transient sources. The lower limits of detection (LLDs) of the multi-energy detectors were approximately half that of the gross count rate detector. Similarly, when the sensitivity of the gross count rate detector is adjusted to equal those of the multi-energy detectors at their LLDs, its false-positive rate jumps to 60%.

Utilization of multi-energy moving prior detection systems for the detection of transient radioactive sources is the subject of ongoing research. Optimal settings for relative bin size, prior parameters, and moving prior limits have yet to be determined. As yet there is also no theoretical model to predict false-positive rates at different decision thresholds for the multi-energy moving prior detection system. Finally, the performance of the multi-energy detection systems has yet to be evaluated on sources with different (or multiple) energies. Although it seems reasonable to assume that the eleven-energy detector will have a more constant performance across different source energies, this has yet to be confirmed experimentally.

By assuming that the background count rate comes from a Poisson distribution with a perfectly-known mean, the multi-energy detection systems described above are able to make detection decisions in real time. However, it is often the case that the background count rate varies with time, and therefore could benefit from a more detailed analysis. The next chapter considers a method for analyzing the background count rate under a more sophisticated model. Although the analysis is likely too complex to be carried out in real time, we show that it can significantly improve control over the false-positive and false-negative error rates.



## Chapter 3

# CHARACTERIZATION OF BACKGROUND RADIATION WITHOUT THE ASSUMPTION OF IDENTICALLY DISTRIBUTED COUNT RATE MEASUREMENTS<sup>2</sup>

### 3.1 Introduction

In the previous chapter, we considered a case in which the mean background count rate was assumed to be fixed and perfectly well known – or at least known with minimal uncertainty relative to the uncertainty on the sample count rate. This chapter considers the opposite extreme – a situation in which the mean of the radiation background count rate (i.e., the Poisson mean) is a function of time rather than being fixed. In this “moving target” model, the Poisson mean for the background is assumed to change at specified time intervals. As we shall see, it is convenient to assume that each new Poisson mean is drawn from a gamma distribution with parameters  $\alpha$  and  $\beta$ . Given a dataset consisting of replicate background count measurements, we determine the values of  $\alpha$  and  $\beta$  using maximum likelihood and Bayesian probabilistic analyses using Markov Chain Monte Carlo (MCMC). The  $\alpha$  and  $\beta$  so determined replace the usual fixed- $\lambda$  determination of background. In the Bayesian analysis, the parameters are assumed to have log-uniform priors, and the parameter space is explored using Markov Chain Monte Carlo. The probabilistic analysis results in probability distributions for the parameters, which are approximately centered at the maximum likelihood (M.L.) values. We confirm the efficacy of these techniques using Monte Carlo generated data, and apply them to data sets generated during bioassay monitoring of radiation workers.

---

<sup>2</sup> This chapter is based on a manuscript which has been accepted for publication by the journal, ‘Radiation Protection Dosimetry.’

The ability to accurately characterize background radiation is of interest to a wide range of fields, including health physics, radiation detection, medical and industrial radiology, and astrophysics. Specifically, it is essential in applications such as determining decision thresholds in radiation monitoring, determining source count rates, and verifying the performance of detectors. Left unaccounted for, a varying background can lead to higher than expected false positive rates in measurements to determine whether a source is present, inaccurate confidence intervals on source count rates, and inaccurate performance estimates of detectors. Nonetheless, virtually all widely used statistical models - classical and Bayesian - operate on the assumption that the background count rate is also sampled from a Gaussian distribution (compare ISO 2000). While both models acknowledge uncertainty as to the true mean background count rate, both also assume that some fixed true mean count rate does exist, and that uncertainty on that count rate can therefore be reduced indefinitely by repeated observations. As a result, it is common practice to attempt to obtain very accurate estimates of true mean count rates by conducting long background measurements (NCRP 1985; Turner 2007).

The classical model compiles background data into “sufficient statistics” (e.g., total number of counts,  $N_b$ , and total measurement time,  $t_b$ ), and in the process discards a great deal of the available information. In particular, the complete measurement data contain information about the time dependence of background which is lost when it is compiled into sufficient statistics. This simplification was necessary when routine statistical calculations were performed by hand, but as computational power has increased, there has been a corresponding increase of interest in developing models which can make use of the great wealth of information obtained in long background measurements (Brandl 2013). We explore two techniques which take advantage

of long background measurements to make inference on the variability of background and discuss the relative merits of each.

### 3.3 Materials and Methods

#### 3.3.1 Background

As in the traditional model, we assume that the number of counts observed in a given time interval is a random variable drawn from a probability distribution, which we will call the sampling distribution. Furthermore, we assume that the count rate is approximately constant during any sufficiently small time interval,  $\Delta t$ . The number of counts observed during  $\Delta t$  is then drawn from a Poisson sampling distribution centered at the true mean count rate during that time interval,  $\lambda_{\Delta t}$  (hereafter termed the Poisson mean):

$$P(n; \lambda_{\Delta t} \Delta t) = \frac{(\lambda_{\Delta t} \Delta t)^n}{n!} e^{-\lambda_{\Delta t} \Delta t}. \quad (3.1)$$

However, as the Poisson mean is generally not constant, it is more appropriate to treat it as a random variable sampled from a probability distribution. We will use a gamma distribution with parameters  $\alpha$  and  $\beta$  (alternatively, the dimensionless  $\beta' = \beta/t$ , where  $t$  is the counting period, can be used). The mean of the gamma distribution is given by  $\alpha/\beta$ , and variance by  $\alpha/\beta^2$ . Therefore, the absolute width (e.g. variance) is small when  $\alpha$  is small compared with  $\beta^2$  (the variance is equal to  $\alpha/\beta^2$ ). The relative variation of the mean count rate can be represented by standard deviation divided by the mean, which evaluates to  $\alpha^{-1/2}$ . Therefore, the relative width is small when  $\alpha$  is large, and is not a function of  $\beta$ .

Recall from Chapter 1 that the probability of observing a particular number of counts,  $n$ , following a Poisson distribution whose mean is drawn from a gamma distribution with parameters  $\alpha$  and  $\beta'$  can be determined by integrating the probability of observing  $n$  counts for a given count rate  $\lambda$ ,  $\text{Poisson}(n|\lambda)$ , over the normalized probability distribution of the count rate,  $\text{gamma}(\lambda|\alpha, \beta')$ :

$$\begin{aligned} p(n|\alpha, \beta) &= \int_{\lambda=0}^{\infty} \text{Poisson}(n|\lambda) * \text{gamma}(\lambda|\alpha, \beta') d\lambda \\ &= \int \frac{(\lambda t)^n}{n!} e^{-\lambda t} * \frac{\beta'^{\alpha}}{\Gamma(\alpha)} \lambda^{\alpha-1} e^{-\beta' \lambda} d\lambda \end{aligned} \quad (3.2)$$

This integral evaluates to a normalized parametric probability distribution which, through a simple re-parameterization, takes the form of a negative binomial distribution (Ntzoufras 2009):

$$\begin{aligned} P(n|\alpha, \beta) &= \text{N. B.} \left( r = \alpha, p = \frac{1}{1 + \beta'} \right) \\ &= \binom{n + \alpha - 1}{n} \left( \frac{\beta'}{\beta' + 1} \right)^{\alpha} \left( \frac{1}{\beta' + 1} \right)^n \end{aligned} \quad (3.3)$$

Therefore, we propose that the observed number of counts from a changing background cannot be modeled as a Poisson process, and suggest instead using a negative binomial likelihood function to model background data. The mean of the negative binomial distribution is given by  $pr/(1 - p) = \alpha/\beta$ , which is the same result one would obtain using the stationary target (Poisson) model. In contrast, the variance of the negative binomial distribution is given by:

$$\frac{pr}{(1-p)^2} = \frac{\alpha}{\beta} + \frac{\alpha}{\beta^2} = \frac{\alpha}{\beta^2}(1 + \beta) \quad (3.4)$$

The variance of the moving target sampling distribution is therefore equal to the sum of the variance of the stationary target sampling distribution,  $\alpha/\beta$ , and the variance of the gamma distribution from which the Poisson mean is sampled,  $\alpha/\beta^2$ . Equivalently, the standard deviation of the negative binomial distribution increases linearly with the standard deviation of the mean count rate.

From Eq. 3.4, we see that the standard deviation of the negative binomial distribution is proportional to the standard deviation of the mean count rate,  $\sqrt{\alpha}/\beta$ . It follows that, because classical decision thresholds for a given sampling distribution are proportional to the standard deviation, classical decision thresholds will also be proportional to the standard deviation of the mean count rate. Similarly, in the presence of a perfectly known background a Bayesian decision threshold might be set when 95% of the posterior lies above the background count rate, which is described by a gamma distribution. In that case, the decision threshold would again be proportional to the standard deviation of the gamma distribution is  $\sqrt{\alpha}/\beta$ . This confirms the second research hypothesis regarding the moving target model.

Given this sampling distribution, a given data set can be analyzed to determine estimates of the parameters  $\alpha$  and  $\beta$ . We employ classical and Bayesian approaches to produce maximum likelihood estimates and probability distributions, respectively, for these parameters.

### 3.3.2 Maximum Likelihood Estimation

The maximum likelihood (M.L.) method attempts to find the parameters  $\alpha$  and  $\beta'$  which maximize the combined likelihood of a set of observations, which is given by the product of the probability of each observation.

$$p(Data) = \prod_{i=1}^N p(x_i|\alpha, \beta') \quad (3.5)$$

In this case,  $x_i$  can represent the number of counts observed in the  $i^{\text{th}}$  observation, or the time interval between counts  $i-1$  and  $i$ . Maximizing the likelihood is equivalent to minimizing the negative logarithm of this function. We therefore define an error function, to be minimized, for which each data point is defined as:

$$\chi_i^2 \equiv -2 \ln(p(x_i|\alpha, \beta')). \quad (3.6)$$

For purposes of interpretation, it is useful to scale the likelihood in the error function by the maximum value of the stationary target likelihood of  $x_i$ , such that

$$p(x_i|\alpha, \beta) \rightarrow \frac{p(x_i|\alpha, \beta)}{\max_{\lambda} p(x_i|\lambda)}. \quad (3.7)$$

Under this transformation, the average optimum value of  $\langle \chi^2 \rangle / N$  for a large number of repetitions is somewhat greater than, but approximately equal to, 1. Therefore,  $\chi^2$  becomes a useful metric for interpreting how well the model fits the observed data.

When the data are expressed as a sequence of counts, i.e.,  $Data = [n_1, n_2, \dots, n_N]$ ,

$\chi_i^2$  evaluates to:

$$\begin{aligned}\chi_i^2 &= -2 \ln \left( \frac{\text{N. B.} \left( n_i | \alpha, \frac{1}{1 + \beta'} \right)}{\text{Poisson}(n_i | \lambda = n_i)} \right) \\ &= -2 \ln \left( \frac{\binom{n_i + \alpha - 1}{n_i} \left( \frac{\beta'}{\beta' + 1} \right)^\alpha \left( \frac{1}{\beta' + 1} \right)^{n_i}}{\frac{n_i^{n_i} e^{-n_i}}{n_i!}} \right).\end{aligned}\tag{3.8}$$

In this case, an analytic solution to the optimization problem can be found by taking the partial derivative with respect to  $\beta'$ :

$$0 = \frac{\partial}{\partial \beta'} \sum_{i=1}^N \ln(p(n_i | \alpha, \beta')) = \sum_{i=1}^N \frac{1}{\beta' + 1} \left( \frac{\alpha}{\beta'} - n_i \right).\tag{3.9}$$

This implies that  $\beta' = \alpha / \bar{N}$  at the minimum, where  $\bar{N} = (\sum_i n_i) / N$ . Given that the mean of the Gamma distribution is given by  $\alpha / \beta$ , this leads to the intuitive result that the mean of the M.L. distribution of the true mean count rate will be equal to the observed mean count rate.

Incorporating these results, the error function becomes:

$$\chi_i^2 = -2 \ln \left[ \left( \frac{\alpha}{\alpha + \bar{N}} \right)^\alpha \left( \frac{e \bar{N}}{n_i (\alpha + \bar{N})} \right)^{n_i} \frac{(n_i + \alpha - 1)!}{(\alpha - 1)!} \right]\tag{3.10}$$

For a single data point, the minimum value of  $\chi^2$  is zero, which occurs for  $\alpha \rightarrow \infty$ . For multiple observations that are not identical, the minimum of the  $\chi^2$  function occurs for some finite value of  $\alpha$ , which can be easily determined by a variety of straightforward numerical optimization algorithms.

### 3.3.3 Bayesian Probabilistic Analysis

In the probabilistic approach, prior probability distributions need to be assumed for the parameters  $\alpha$  and  $\beta'$ . Although  $\alpha$  and  $\beta'$  are needed for the sampling distribution, the two parameters are highly correlated, with the most likely values being those whose ratio is equal to the observed mean count rate. For that reason, the analysis is done using parameters  $\mu = \alpha/\beta'$  and relative precision  $\alpha$ . Both parameters are assigned broad log-uniform priors. The parameter space is then explored using Markov Chain Monte Carlo. This was accomplished using General Markov Chain Monte Carlo (GMCMC) routines, which sample the joint posterior distribution via either the Metropolis-Rosenbluth-Teller algorithm, the Barker algorithm, or the Multiple Candidate algorithm (depending on user preference) (Miller 2013). By calculating the probability of observing the data under varying values of  $\alpha$  and  $\beta'$ , the GMCMC routines produce empirical probability distributions for the parameters. The modes of these distributions are almost identical to the M.L. estimates, indicating that the chosen priors are not significantly influencing the outcome. However, in contrast to the maximum-likelihood approach, the probabilistic analysis also produces reliable estimates for the uncertainty on the distribution parameters.



### 3.3.4 Experimental Analysis

Both methods were applied to a collection of bioassay data, which, like all radiation measurements on living organisms, exhibit significant time dependence. We examined 252 replicate  $^{234}\text{U}$ -in-urine “method blank” background measurements taken over a two-year period. These are measurements of samples known to not contain  $^{234}\text{U}$ , carried through the entire chemical processing and alpha-spectrometry counting procedure, using the exact analysis procedures routinely employed for actual samples. The counting period was 42 hours for one measurement. The total number of detected counts for all 252 replicate measurements was about 2200 (only about 0.2 counts per hour).

As a check, both methods were also tested on Monte Carlo simulated data. For each simulated data set, the Poisson mean was drawn from a Gamma distribution with known parameter values. In this case, because the true result was known, we could look at the agreement between calculated and true values. To evaluate the efficacy of the algorithms in general, they were tested on a variety of different sized simulated data sets, which were generated using mean count rates drawn from a wide range of parameter values.

## 3.4 Results

### 3.4.1 $^{234}\text{U}$ Bioassay Data Set

The M.L. estimates of the background variation parameters for the  $^{234}\text{U}$  bioassay data are  $\alpha = 8.7$  and  $\beta' = 1.0$ , which implies that the Poisson mean  $\lambda_{\Delta t}$  is drawn from

$$\lambda_{\Delta t} \sim \text{Gamma}(\alpha = 8.7, \beta = 1.0) \quad (3.11)$$

The sampling distribution for the expected counts is given by

$$p(n|\alpha, \beta') = \text{N. B. } (\alpha = 8.7, \beta = 1) \quad (3.12)$$

The mean and standard deviation of the sampling distribution are given by

$$\begin{aligned} \text{Mean}[\text{N. B. } (\alpha, \beta')] &= \frac{\alpha}{\beta'} = 8.7 \\ \text{S. D. } [\text{N. B. } (\alpha, \beta')] &= \frac{\alpha}{\beta'} + \frac{\alpha}{\beta'^2} = 4.17 \end{aligned} \quad (3.13)$$

These parameters are the M.L. estimates of the predicted observed counts and expected standard deviation of multiple count measurements. This sampling distribution is analogous to the Poisson distribution which would be computed using the traditional fixed target model. Under this model, the expected standard deviation of multiple count measurements is equal to the square root of the mean counts, 2.95. Both distributions have the same mean counts but, as expected, the moving target model predicts greater variation in the observed counts. Both sampling distributions are plotted in Figure 3.1.

If these distributions were used in a classical null-hypothesis test, the 95% decision threshold for the stationary target model would be 14 counts. In other words, one would expect more than 14 counts to be observed in less than 5% of background measurements. In contrast, the 95% decision threshold corresponding to the moving target model is 16 counts.

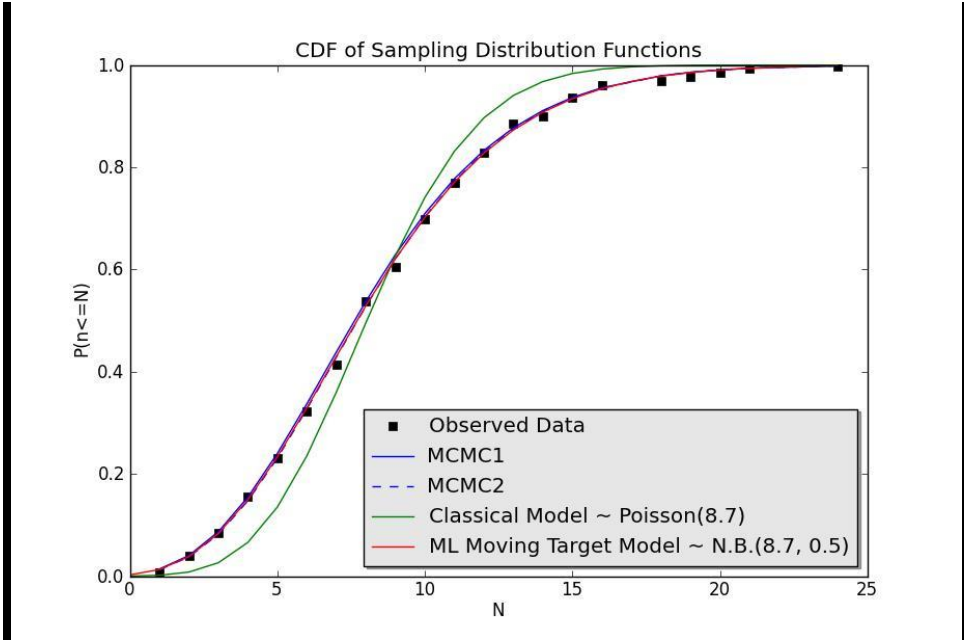


Figure 3.1: The cumulative distribution functions of the observed counts under the stationary target and moving target models.

The classical 95% decision threshold is therefore 14% higher under the moving target model (16) than it would be if the background were assumed to be stationary (14). This confirms the first research hypothesis regarding the moving target model.

The probabilistic analysis provides probability distributions rather than point estimates for the parameters. Probability distributions from analysis of the  $^{234}\text{U}$  data are shown in Figures 3.2a and 3.2b. These are clearly well-aligned with the M.L. point estimates. The probabilistic analysis also generates its own non-parametric estimate of the sampling distribution. The cumulative distribution function of the Poisson mean  $\lambda_{\Delta t}$  is plotted in Figure 3.3 (from two Monte Carlo runs with different initial values), along with cumulative distributions from the M.L. estimate and fixed target estimate for comparison. The probabilistic sampling distribution is very similar to the M.L. sampling distribution when applied to large data sets.

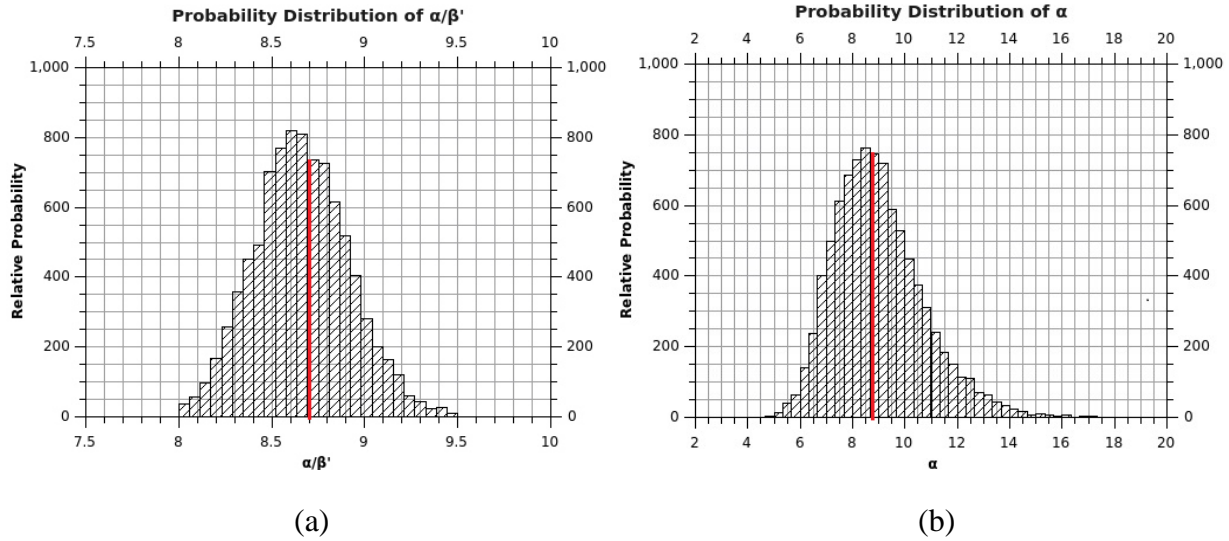


Figure 3.2: Probability distributions of the parameters of the Poisson mean. These figures illustrate the uncertainty on the values of the parameters of the gamma distribution from which the Poisson mean is drawn.

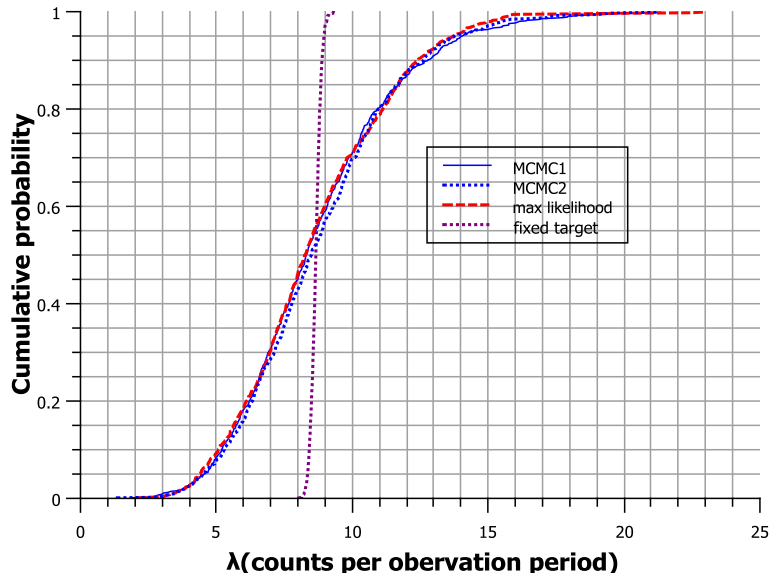


Figure 3.3: The posterior cumulative distribution function of the true background count rate, calculated using the probabilistic method, the maximum likelihood method, and the fixed target method.

### 3.4.2 Simulated Data Sets

When applied to large datasets, Bayesian analysis produces narrow parameter distributions which are well represented by the M.L. point estimates. In this case, both methods accurately estimate the true parameter values of the simulated data sets. As the size of the data set decreases, the uncertainty on the parameters increases, and the maximum likelihood point estimates become extremely noisy, and the uncertainties on the parameter values become essential to accurately interpreting the data. As a general rule, analysis of measurements on the order of one thousand independent observations or less tend to result in significant uncertainty on the relative width of the Poisson mean,  $\alpha^{-1/2}$ . For that reason, it is essential that only probabilistic analysis be used to analyze small datasets.

When  $\lambda_{\Delta t}$  is drawn from an extremely narrow probability distribution ( $\alpha$  greater than 100), the sampling distribution effectively reduces to a Poisson distribution. In this case, the likelihood of the data is essentially equal for a wide range of parameter values, and the probability distribution of  $\alpha$  may appear to be very broad. However, when  $\alpha$  is large, the relative width ( $\sigma/\mu = 1/\sqrt{\alpha}$ ) becomes very small, so that large differences in  $\alpha$  correspond to very small differences in relative width. For example, if the 95% credibility interval on  $\alpha$  is given by  $\alpha \in (6000, 8000)$ , the corresponding credibility interval on the relative width will be  $\alpha^{-1/2} \in (8000^{-1/2}, 6000^{-1/2}) = (0.011, 0.013)$ . In other words, great uncertainty and variation on  $\alpha$  does not necessarily imply great variation on the distribution of  $\lambda_{\Delta t}$ .

### 3.5 Conclusion

We have examined probabilistic and maximum likelihood methods for estimating background radiation, and have shown that the two produce very similar results when applied to large datasets. The choice of which method to employ depends on the goals of the user, and on the size of the dataset. Moving target analysis can always be applied, even to small datasets, using the probabilistic approach. The probabilistic approach produces uncertainty estimates on the parameters which are essential for understanding the degree of uncertainty on the variation of the Poisson mean in all but very large data sets. Conversely, the M.L. method is easier to implement, but can only be safely applied to large datasets.

Although moving target analysis does not assume that the background is fixed, it also does not require that it be moving. Moving target analysis on a sufficiently large dataset will correctly identify cases where the background is fixed. However, it is not generally possible to conclude that the background is fixed on the basis of a relatively small data set.

Maximum Likelihood analysis should be used with caution when the relative width of the Poisson mean distribution is greater than about 5%, which typically occurs for data sets with fewer than one thousand independent observations. Failure to account for the uncertainty on the parameters can lead to poor control of the false positive and false negative rates. For that reason, if smaller data sets are to be analyzed under the moving target model, it is essential that the probabilistic approach be used.

Given the need for large background datasets, it may be tempting to reuse background measurements for multiple sample measurements. However, this temptation must be avoided, as

it will result in correlated sample measurements which cannot be interpreted in a straightforward way.

Variation of the true mean background count rate can occur for a number of reasons, and should be considered the norm in counts occurring over long time intervals or on living organisms. Failure to account for this variation may produce a dramatic misrepresentation of the true background distribution. This will adversely affect a variety of critical endeavors, including attempts to control false positive rates, estimate lower limits of detection (LLD's), estimate confidence intervals for source count rates, and assess instrumentation performance. Moving target analysis results in a much higher probability of observing counts far above the mean, meaning that such observations may no longer strongly imply the presence of a source. Given the ease with which modern computers can incorporate count rate variation into the background model, it is no longer necessary to analyze background data under the assumption that the count rate is fixed.

In general, the purpose of a detailed model of the background count rate is to better detect and characterize radioactive sources. In the next chapter, we discuss a technique for performing both of those tasks simultaneously.

## Chapter 4

# SIMULTANEOUS SOURCE DETECTION AND ANALYSIS USING A ZERO-INFLATED COUNT RATE<sup>3</sup>

### 4.1 Synopsis

This chapter proposes a novel Bayesian technique which allows for simultaneous source detection and count rate analysis. The technique involves using priors which include a finite probability that the source count rate is exactly zero. Such priors are called “zero-inflated.” Solving the posterior distribution of a zero-inflated count rate model provides the probability that the sample contains a source and a probability distribution for the source count rate if the source exists, without the need to perform redundant computations. Sampling from zero-inflated distributions is straightforward, and can be accomplished with easily accessible open source software. In addition, zero-inflated priors lead to finite posterior probabilities of “no source,” which is an easy to understand and satisfying result.

### 4.2 Introduction

A common application of radiation counting measurements is to estimate the count rate from a source which is present in a sample of material. This application, which is referred to as “count rate analysis,” is typically carried out by taking counts from the sample of interest and from background. The net sample count rate is then estimated by subtracting the background count rate from the gross count rate of the sample. Source detection is a related application whose aim is to estimate the probability that any radioactive material is present in the sample

---

<sup>3</sup> This chapter is based on a manuscript which has been submitted for publication in the journal, ‘Health Physics.’



(i.e., that the sample count rate is non-zero). Although source detection is a special case of count rate analysis, in practice the two applications are often approached very differently.

The classical approach to count rate analysis involves making inference on the sampling distribution – in this case the probability distribution of the observed data as a function of the true sample count rate. Source detection is accomplished using the “null-hypothesis” test, which considers the probability of a given measurement if the true sample count rate is zero. However, because classical tests cannot produce probability distributions for the sample count rate, which is the goal of count rate analysis, they have been largely rejected in favor of the Bayesian approach (Little 1982, Potter 1999, Strom and MacLellan 2001).

The Bayesian approach to source detection involves posterior probability ratios (Jeffreys 1961). In this test, a prior probability is assigned to the existence of the source. Probability distributions are also assigned to the source (if it exists) and background count rates, but these are typically assumed to be known and are therefore not updated with new measurements. Following a measurement, the posterior probability ratio is given by the prior probability of the source existing multiplied by the likelihood ratio of the model in which the sample contains a source and the model in which it does not.

Bayesian count rate analysis is very different from Bayesian source detection. In Bayesian count rate analysis, prior probability distributions are assigned to the source and background count rates. The priors are then updated with subsequent measurements, giving posterior probability distributions for the count rates. The choice of priors will depend on how much prior knowledge is available, as well as the specific test chosen. However, even uninformative priors typically assume a priori that the sample count rate is greater than zero.

Virtually all commonly used probability density functions are finite, continuous, and differentiable. This may be due to the fact that such functions are mathematically tractable, or because they adequately describe most naturally occurring phenomena. Because finite continuous probability density functions only give non-zero probabilities to ranges of values, the probability of any precise value will be precisely zero. It follows that the prior, and therefore posterior, probability that the source count rate is exactly zero is equal to zero. Because source detection is equivalent to determining the probability that the source count rate is exactly zero, and finite continuous priors rule out this possibility, such priors cannot be used for source detection.

These problems can be partially overcome by the implementation of an “alpha” prior on the source count rate (Miller et al. 2001). The alpha distribution is a limit on the gamma distribution in which the shape parameters,  $\alpha$  and  $\beta$ , are very small and the first parameter,  $\alpha$ , is related to the prior probability of a sample. The distribution is such that it approximates a Delta function at  $p(\lambda = 0)$ , and becomes log-uniform where  $p(\lambda > 0)$ . Using the alpha prior, the posterior will include a finite probability that the sample count rate is extremely close to zero. In principal, this allows one to avoid hypothesis testing altogether.

However, a number of practical issues arise in the implementation of an alpha prior. The first issue is that it cannot be sampled by commonly available software, meaning that the user needs to develop code from scratch in order to sample from it. This is exacerbated by the fact that the Delta function at zero makes this sampler particularly difficult to write. Another issue is that the prior probability that the source count rate is extremely low is encoded into the first parameter of the distribution, making it somewhat difficult to read and modify. Finally, because the Delta function can only be approached by a gamma function in a limit situation, the alpha

prior still leads to a zero posterior probability of “no source.” This issue might seem pedantic when the posterior gives a high probability that the source count rate is less than (for example)  $10^{-25}$ , but it can be confusing for the uninitiated user.

This chapter proposes using zero-inflated priors as an alternative to the alpha prior. Like the alpha prior, zero-inflated priors allow for simultaneous source detection and count rate analysis. However, sampling from zero-inflated distributions is straightforward compared to sampling from an alpha distribution, and can be accomplished with easily accessible open source software. In addition, zero-inflated priors lead to finite posterior probabilities of “no source,” which is a less confusing and more satisfying result.

### 4.3 Zero-Inflated Model

Recall from the discussion of the two-count problem in Chapter 1 that the joint posterior distribution is obtained by evaluating

$$p(\lambda_B, \lambda_S | N_B, N_G) = \frac{p(N_G | ((\lambda_B + \lambda_S)t_B)) \times p(N_B | \lambda_B t_B) \times p(\lambda_B) \times p(\lambda_S)}{p(N_B, N_G)}, \quad (4.1)$$

where

$$p(N_B, N_G) = \int_0^\infty \int_0^\infty p(N_G | ((\lambda_B + \lambda_S)t_G)) \times p(N_B | \lambda_B t_B) \times p(\lambda_B) \times p(\lambda_S) d\lambda_B d\lambda_S, \quad (4.2)$$

and that calculating the Probability Ratio,  $PR$ , between two hypotheses requires performing the most difficult part of the count rate analysis problem – evaluating the double integral of Eq. 4.2. Having evaluated that integral for both models, it becomes trivial to calculate the joint posterior.

All that is required is to extract the posterior probabilities,  $p(\text{Source}|N_B, N_G)$  and  $p(\text{No Source}|N_B, N_G)$ , from the Probability Ratio using the following relationship:

$$\begin{aligned} p(\text{Source}|N_B, N_G) &= \frac{PR}{1 + PR} \\ p(\text{No Source}|N_B, N_G) &= \frac{1}{1 + PR}, \end{aligned} \tag{4.3}$$

and to combine them with the joint posterior given that there is a source to obtain the true joint posterior. The problem with this approach is that it requires the combination of two complex analyses after they have been carried out. This not only involves redundant computations, but is conceptually rather unclear. Putting a zero-inflated prior distribution on the source count can lead to an intuitive model which unifies the two approaches.

A zero-inflated probability distribution is a probability density function (e.g., Gaussian, Poisson, gamma, etc.) with an increased, or “inflated,” probability mass at zero. A classic example is the distribution of cigarettes smoked per day by a given person (Neelon et. al 2010). Amongst smokers, the number of cigarettes smoked per day might come from a Poisson distribution centered at  $\lambda_{\text{smokers}}$ . However, non-smokers typically smoke precisely zero cigarettes per day. Therefore, the distribution of cigarettes per day across the whole population would be a Poisson distribution centered at  $\lambda_{\text{smokers}}$  with a large spike at zero.

Similarly, the number of decays in a given time interval from a sample containing a source will come from a Poisson distribution which is likely to be centered well above zero. However, the number of decays occurring in a sample containing no source will be precisely zero. Therefore, an accurate representation of our prior belief about the sample count rate would be a distribution which assigns a finite probability that the mean count rate is exactly zero,

followed by an uninformative prior for count greater than zero – in other words, a zero-inflated prior. If, as before, we choose a gamma distribution to represent the probability of the source count rate if it is greater than zero, then the complete prior is a zero-inflated gamma (ZIG) distribution. If the prior probability that a given sample contains a source is  $p_S$ , the prior probability density function of  $\lambda_S$  can be written:

$$p(\lambda_S) = (1 - p_S)\delta(0) + p_S * \text{gamma}(\alpha, \beta) \equiv \text{ZIG}(p_S, \alpha, \beta) \quad (4.4)$$

where  $\delta(0)$  is a Delta function located at zero. The Delta function simply indicates that this is really a probability mass function at zero – in other words, there is a finite probability that  $\lambda_S$  is exactly equal to zero.

Rather than expressing the ZIG distribution as a single probability density function containing a Delta function, it may be clearer to express it as the product of two independent probability distributions – a probability mass function for the existence of a source, and a probability density function for the strength of the source if it exists. If we represent the existence of a source with a binary variable, *Source*, where *Source* = 1 if the source exists, and *Source* = 0 if the source does not exist, then the probability density of  $\lambda_S$  given *Source* is

$$p(\lambda_S | \textit{Source}) = \textit{Source} * \text{gamma}(\alpha, \beta) \equiv \text{ZIG}(p_S, \alpha, \beta), \quad (4.5)$$

where *Source* is a random variable with a Bernoulli probability mass function,

$$\textit{Source} \sim \text{Bernoulli}(p_S). \quad (4.6)$$

Or, equivalently,

$$p(\textit{Source}) = \begin{cases} p_S & \textit{Source} = 1 \\ 1 - p_S & \textit{Source} = 0. \end{cases} \quad (4.7)$$

As discussed in Chapter 1, it will often be the case that  $p_S$  is not precisely known. In that case, a prior distribution may be put on  $p_S$  to represent the uncertainty on its possible values, such that

$$p_S \sim p(p_S) \quad 0 \leq p_S \leq 1. \quad (4.8)$$

Models in which priors are put on the parameters of other priors are called hierarchical models, and such priors are called hyperpriors, and their parameters are called hyperparameters (Ntzoufras 2009). If the prior distribution on  $p_S$  is uniform, then

$$p(p_S) = \frac{1}{2} \quad 0 \leq p_S \leq 1. \quad (4.9)$$

It can be easily shown that this leads to identical results as assuming  $p_S = 0.5$ . Including this  $p_S$  and *Source*, the joint posterior now becomes  $p(\lambda_B, \lambda_S, p_S, \textit{Source} | N_B, N_G)$ . The marginal posterior of *Source*,  $p(\textit{Source} = 0 | N_B, N_G)$ , is the probability that the sample does not contain a source or, equivalently, that the sample count rate is exactly zero.

The marginal posterior of  $p_S$  can be thought of as the probability distribution that the next sample will contain a source. The hyperparameter,  $p_S$ , is not of interest for interpreting the given sample, and may be integrated out of the joint posterior as follows,

$$p(\lambda_B, \lambda_S, \textit{Source} | N_B, N_G) = \int_0^1 p(\lambda_B, \lambda_S, p_S, \textit{Source} | N_B, N_G) \times p(p_S) dp_S. \quad (4.10)$$

Because *Source* is a binary variable,  $p(\lambda_B, \lambda_S, \text{Source} | p_S, N_B, N_G)$  is equal to the combination of the two marginal posteriors,  $p(\lambda_B, \lambda_S, \text{Source} = 1 | N_B, N_G)$ , and  $p(\lambda_B, \lambda_S, \text{Source} = 0 | N_B, N_G)$ .

Because *Source* is really a hypothesis about  $\lambda_S$ , it follows that,<sup>4</sup>

$$p(\lambda_B, \lambda_S | N_B, N_G) = \begin{cases} p(\lambda_B, \lambda_S | \text{Source} = 1, N_B, N_G) \times p(\text{Source} = 1 | N_B, N_G) \times p(p_S | N_B, N_G) & \lambda_S > 0 \\ p(\lambda_B, \lambda_S | \text{Source} = 0, N_B, N_G) \times p(\text{Source} = 0 | N_B, N_G) \times p(p_S | N_B, N_G) & \lambda_S = 0, \end{cases} \quad (4.11)$$

where  $p(p_S | N_B, N_G)$  can be derived from the posterior of the odds ratio,  $p(PR | N_B, N_G)$ , given by Eq.

1.25. All of the expressions on the right hand sides of these equations can be calculated using Eq. 4.1, Eq. 4.2, and Eq. 4.3. In particular, if the source does not exist then  $\lambda_S = 0$ , meaning that both measurements were of background only. In that case, we can combine our observations such that we have observed  $N_B + N_G$  counts from background in  $t_B + t_G$  time intervals. It follows that the marginal posterior of the mean count rates if the source does not exist reduces to

$$p(\lambda_B, \lambda_S | \text{No Source}, N_B, N_G) = p(\lambda_B | N = N_B + N_G) \quad \lambda_S = 0, \quad (4.12)$$

Where, from Eq. 1.7,  $p(\lambda_B | N = N_B + N_G)$  is given by

$$p(\lambda_B | N = N_B + N_G) = \text{gamma}(\alpha + N, \beta + t_B + t_G). \quad (4.13)$$

Alternatively, if multiple background measurements have been carried out, it is possible to determine the background distribution by performing a moving target background analysis as described in Chapter 3. Because the moving target analysis models background as a gamma

---

<sup>4</sup> This follows from the product rule, which states that  $p(A, B) = p(A|B) \times p(B)$ , or that the probability of two observations, *A* and *B*, occurring together is equal to the probability of *A* given *B* multiplied by the probability of *B*.

distribution, this will have no impact on the subsequent calculations except to use different parameters for the background distribution.

Regardless of how the parameters of background distribution are obtained, the marginal posterior of the mean count rates if the source exists,  $p(\lambda_B, \lambda_S | Source = 1, N_B, N_G)$ , is given by Eq. 4.1, which is the result of the original sample count rate analysis.

Finally,  $p(Source = 1 | N_B, N_G)$  and  $p(Source = 0 | N_B, N_G)$  are functions of the Probability Ratio given by Eq. 4.3. Making these substitutions results in the following expression for the joint posterior:

$$\begin{aligned}
 & p(\lambda_B, \lambda_S | N_B, N_G) \\
 &= \begin{cases} \frac{PR}{1 + PR} \times \frac{p(N_G | ((\lambda_B + \lambda_S)t_B)) \times p(N_B | \lambda_B t_B) \times p(\lambda_B) \times p(\lambda_S)}{p(N_B, N_G | Source)} & \lambda_S > 0 \\ \frac{1}{1 + PR} \times \text{gamma}(\alpha + N_B + N_G, \beta + t_B + t_G) & \lambda_S = 0. \end{cases} \quad (4.14)
 \end{aligned}$$

Equation 4.14 shows that, after integrating out  $p_S$ , solving for the joint posterior under a zero-inflated prior is equivalent to performing traditional count rate analysis and hypothesis testing. From the joint posterior, the marginal posterior of the source count rate can be obtained by integrating over the background count rate:

$$p(\lambda_S | N_B, N_G) = \int_0^\infty p(\lambda_B, \lambda_S | N_B, N_G) d\lambda_B. \quad (4.15)$$

Although the formula for the posterior under a zero-inflated prior is complex and difficult to solve directly, zero-inflated distributions are very straightforward to sample from, which makes them ideal for MCMC analysis. The generative model corresponding to the two-count problem with a zero-inflated prior is as follows:



### Two-Count Model with Zero-Inflated Gamma Prior:

$\lambda_B \sim \text{gamma}(\alpha = 0.5, \beta = 0.01)$ , #Mean background count rate  
 $p_S \sim p(p_S)$ , #The uncertainty on the odds of a source,  $p_S$   
 $Source \sim \text{Bernoulli}(p_S)$ , #*Source* is 1 with odds  $p_S$ , 0 with odds  $(1 - p_S)$   
 $\hat{\lambda}_S \sim \text{gamma}(\alpha = 0.5, \beta = 0.01)$ , #Sample count rate if there is a source  
 $\lambda_S = Source \times \hat{\lambda}_S$  #True sample count rate  
 $\lambda_G = \lambda_B + \lambda_S$ , #Gross count rate  
 $N_B \sim \text{Poisson}(\lambda_B t_B)$ , #Counts observed in background measurement  
 $N_G \sim \text{Poisson}(\lambda_G t_G)$ . #Counts observed in sample measurement.

This technique has been used to describe the generative model for zero-inflated Poisson likelihood functions (Ghosh 2006), but does not appear to have ever been applied to zero-inflated prior distributions. As in the classic two-count model, this model indicates that  $\lambda_B$  and  $\hat{\lambda}_S$  are random variables drawn from gamma distributions. In addition, the existence of a source is encapsulated in the random variable *Source*, which is equal to 1 with odds  $p_S$  and 0 with odds  $1 - p_S$ . The true sample count rate is therefore equal to 0 if *Source* = 0, and  $\hat{\lambda}_S$  if *Source* = 1. Breaking up the zero-inflated prior (with its Delta term) into the product of two common probability functions makes it possible to sample from the prior, and therefore analyze the data, using easily accessible open source software such as OpenBUGS. The result of this analysis will include a finite posterior probability of “no source,” which has obvious advantages for communicating with non-technical co-workers and the public.

Markov Chain Monte Carlo techniques produce a chain of values which, after a certain number of iterations, can be considered unbiased random samples from the posterior distribution. In order for this to occur, it must be possible for the chain to reach any value of the posterior from every other possible value of the posterior. In other words, it must never be possible for the chain to include any value from which it cannot then sample from the complete posterior. More formally, the Markov Chain only converges to the posterior distribution if it is irreducible,

positive recurrent, and aperiodic (Ntzoufras 2009). In general, it is important to understand the MCMC algorithm well enough to determine if one's model will satisfy these conditions. Given the simplicity of the model and widespread use of similar models, it is expected that the two count radioactive sample problem will satisfy these conditions for the Metropolis Hastings and Gibbs sampling algorithms. In addition, the posteriors for the following examples are calculated numerically as well as with MCMC. The strong agreement between MCMC and deterministic techniques provides further evidence that the Markov Chains can be expected to converge to the posterior. Nonetheless, care should always be taken to ensure convergence when relying on MCMC (see Appendix for details).

#### **4.4 Examples**

This section presents some simple examples to illustrate the application and usefulness of zero-inflated models compared to classic hypothesis testing and count rate analysis. Although the ZIG model presented above represents an idealized model, it is easily extended to account for real-world considerations such as detector efficiency, geometric factors, energy-dependence, etc., and their respective uncertainties. However, for the purposes of illustration it is helpful to use the idealized model. This also makes it possible to use the same data, and therefore to compare results, with previous publications.

#### 4.4.1 Example 1: Negative Observed Source Count Rate

In the classic analysis of the two-count problem (Little 1982), the measurement data are as follows:

$$\begin{aligned}N_B &= 84, \\t_B &= 18.56 \text{ s}, \\N_G &= 16, \\t_G &= 4.64 \text{ s}.\end{aligned}$$

Note that the average count rate during the sample measurement,  $16/4.64 = 3.5$  cps is lower than the average count rate during the background measurement,  $84/18.56 = 4.5$  cps. This situation is problematic if the strength of the possible source is unknown, as there is no obvious way to extract an empirical prior for the source count rate from the measurement. In this case, there is no choice but to use a “subjective” prior – a prior which represents the investigator’s belief about how strong the source could be. It seems reasonable to assume that the investigators would have known in advance that they were looking for a relatively weak source. Therefore, a prior that is uniform between zero and five, and zero beyond five would seem reasonable. Note also that the maximum mean count rate permitted by this prior, 5 cps, is approximately twice the background, at which point experienced persons can typically make decisions as to the presence of radioactive sources without the need for statistical methods. Such a prior can be represented as a uniform gamma distribution truncated at five:

$$p(\lambda_S | \text{Source}) = \text{Uniform}(0,5) \propto \lim_{\beta \rightarrow 0} \text{gamma}(\lambda_S | \alpha = 1, \beta); \quad 0 < \lambda_S < 5. \quad (4.16)$$

Because it is truncated, the gamma distribution must be renormalized so it will integrate to 1. If the investigators knew that a source existed, entering the observed data into Eq. 4.1 with a

uniform prior would result in the joint posterior obtained by Little. In particular, the mean sample count rate and 95% credible interval found by Little are,

$$\begin{aligned}\bar{\lambda}_S &= 0.63, \\ 95\%CI &= (0, 1.77).\end{aligned}$$

However, the evidence clearly suggests that there is no source. Therefore, the posterior on the source count rate must have significant mass at  $\lambda_S = 0$ , which will tend to pull the confidence interval towards zero. As shown in Eq. 4.14, the posterior distribution of the source count rate is a function of the posterior given that the source exists,  $p(\lambda_S|N_B, N_G, Source)$ , and the Probability Ratio that the source exists ( $PR$ ).

In order to compute  $PR$ , it first is necessary to explicitly state the probability before the measurement that the sample would contain a source. In practice, the investigator will often possess empirical data and/or relevant experience regarding this prior. However, in order to simplify the numerical integration, let us select a uniform prior for  $p_S$ , which is equivalent to  $p_S = 0.5$ . Note that the choice of  $p_S$  necessarily has a strong influence on  $PR$ . Using this choice, the ratio of priors equals one, and the posterior Probability ratio reduces to

$$PR = \frac{p(Source|N_B, N_G)}{p(No Source|N_B, N_G)} = \frac{p(N_B, N_G|Source)}{p(N_B, N_G|No Source)}. \quad (4.17)$$

Using uniform priors for  $p(\lambda_B)$  and  $p(\lambda_S)$  and inserting Eq. 4.2, this becomes,

$$PR = \frac{\int_0^\infty \int_0^5 \text{Poisson}(16|\lambda = 4.64(\lambda_B + \lambda_S)) \times \text{Poisson}(84|\lambda = 18.56\lambda_B) d\lambda_B d\lambda_S}{\int_0^\infty \text{Poisson}(100|\lambda = 23\lambda_B) d\lambda_B}, \quad (4.18)$$

where the integral in the denominator was simplified by recognizing that the *No Source* hypothesis requires that  $\lambda_S$  be exactly 0. Solving these integrals numerically yields,  $PR = 0.155$ . These integrals are challenging to solve with numerical algorithms, and it is easy to see how a more complex model would produce equations which could only be solved using Monte Carlo techniques. In addition, once  $PR$  is calculated, it is still necessary to calculate the posterior of the sample count rate for  $\lambda_S > 0$ . From Eq. 4.6, the joint posterior is,

$$p(\lambda_B, \lambda_S | N_B, N_G) = \begin{cases} \frac{PR}{1 + PR} \frac{\text{Poisson}(16 | \lambda = 4.64(\lambda_B + \lambda_S)) \times \text{Poisson}(84 | \lambda = 18.56\lambda_B)}{\int_0^\infty \int_0^5 \text{Poisson}(16 | \lambda = 4.64(\lambda_B + \lambda_S)) \times \text{Poisson}(84 | \lambda = 18.56\lambda_B) d\lambda_B d\lambda_S} & \lambda_S > 0 \\ \frac{1}{1 + PR} \times \text{gamma}(\lambda_B | \alpha = 101, \beta = 23) & \lambda_S = 0. \end{cases} \quad (4.19)$$

When working by hand, computing the posterior of  $\lambda_S$  from a zero-inflated prior is equivalent to performing the source-detection computation followed by the count rate analysis computation. Using Markov Chain Monte Carlo (MCMC), however, analyzing the zero-inflated prior model is almost identical to analyzing the traditional count rate analysis. Inserting the priors and observed data yields the following generative model:

#### ZIG Generative Model for Example 1:

$\lambda_B \sim \text{gamma}(\alpha = 1, \beta = 0.01),$	#Mean background count rate
$p_S \sim \text{Uniform}(0,1),$	#The odds of a source, $p_S$ , is equally likely to assume any value between 0 and 1
$Source \sim \text{Bernoulli}(p_S),$	#50% chance <i>Source</i> exists
$\hat{\lambda}_S \sim \text{Uniform}(0,5),$	#Sample count rate if there is a source
$\lambda_S = Source \times \hat{\lambda}_S$	#True sample count rate
$\lambda_G = \lambda_B + \lambda_S,$	#Gross count rate
$84 \sim \text{Poisson}(18.56\lambda_B),$	#Counts observed in background measurement
$16 \sim \text{Poisson}(4.64\lambda_G).$	#Counts observed in sample measurement.

The MCMC analysis produces marginal posterior distributions for the source count rate, and for the existence of the source. From the analysis, the probability that a source is present is,  $p(\lambda_S > 0 | N_B, N_G) = 0.132$ , or about 13%. This corresponds to  $PR = 0.152$ , which agrees with the numerical calculation to within less than 3%. It follows that  $p(\lambda_S = 0 | N_B, N_G) = 0.868$ , or about 87%. Equivalently, there is an 87% probability that  $\lambda_S$  is precisely zero, and that there is, by definition, no source. The mean and 95% credible interval of the posterior of the source count rate are

$$\bar{\lambda}_S = 0.083,$$

$$95\% \text{CI} = \begin{cases} [0] & 87\% \\ (0, 0.68) & 13\% \end{cases}$$

This means that there is a 95% chance the sample count rate is between 0 and 0.64 cps. Because  $p(\lambda_S)$  is a probability mass, rather than a probability density, it will always be a mode of the posterior distribution, and will often cause very unlikely values of the source count rate to be included by default in quantile confidence intervals. For that reason, HDI confidence intervals are strongly preferred for zero-inflated count rate distributions.

Figure 4.1 shows a plot of the marginal posterior distribution of the sample count rate. From this plot, it is clear that the most likely value of the source count rate is zero, and that beyond zero the probability is inversely proportional to the count rate.

The MCMC analysis also produces a marginal posterior distribution for  $p_S$ , which can be used as a prior for a measurement of the next source. This technique can be applied retroactively to historical measurements to develop an empirical prior for  $p_S$ . In this approach, the posterior of each measurement,  $p(p_S: N_B, t_B, N_G, t_G)$ , serving as the prior for the next measurement. By measuring multiple different samples, this technique provides a quantitative approach for developing empirical priors for future measurements.




Figure 4.1: The marginal posterior distribution of  $\lambda_S$  for Example 1. The large spike at  $\lambda_S = 0$  is actually a probability mass representing the probability that the sample contains no source. The 95% HDI, represented by the dashed line, is  $\lambda_S \in (0, 0.68)$ .

#### 4.4.2 Example 2: Strong Signal, Rare Source

The previous example considered a sample which, before measurement, was believed to have a 50% chance of containing a source with mean count rate between 0 and 5. The sample measurement produced a weak (negative) signal, which was used to update the probability of the sample containing a source, as well as the likely values of the source count rate. In practice, the investigators usually have a much less vague notion (perhaps from past experience) of the

probability that a given sample will contain a source, and of the count rate which sources are likely to exhibit. Consider a scenario in which precisely one in twenty samples contains a source, and where sources typically have mean count rates of about 2.5 cps. This might be represented by  $\text{gamma}(\lambda_S | \alpha = 2.5, \beta = 1)$ , which corresponds to a mean source count rate of  $\bar{\lambda}_S = \frac{\alpha}{\beta} = 2.5$  with variance  $\text{var}(\lambda_S) = \frac{\alpha}{\beta^2} = 2.5$ . The results of a background and sample measurement are as

follows:

$$\begin{aligned} N_B &= 100, \\ t_B &= 25 \text{ s}, \\ N_G &= 90, \\ t_G &= 15 \text{ s}. \end{aligned}$$

In this case, there is a clear signal for a source with a count rate of approximately 2 cps, which is well within the expected range of the source count rate. However, that signal must be balanced against the prior knowledge that only about 5% of samples actually contain a source. As before, it is useful to explicitly state the generative model for the observed data:

ZIG Generative Model for Example 2:

$\lambda_B \sim \text{gamma}(\alpha = 1, \beta = 0.01),$	#Mean background count rate
$Source \sim \text{Bernoulli}(0.05),$	#5% chance <i>Source</i> exists
$\hat{\lambda}_S \sim \text{gamma}(2.5, 1),$	#Sample count rate if there is a source
$\lambda_S = Source \times \hat{\lambda}_S$	#True sample count rate
$\lambda_G = \lambda_B + \lambda_S,$	#Gross count rate
$100 \sim \text{Poisson}(25\lambda_B),$	#Counts observed in background measurement
$90 \sim \text{Poisson}(15\lambda_G).$	#Counts observed in sample measurement.

This model can once again be analyzed either numerically or using MCMC. The probability ratio can be expressed analytically as



$$\begin{aligned}
PR = & \frac{\int_0^\infty \int_0^\infty \text{Poisson}(90|15(\lambda_B + \lambda_S)) \times \text{Poisson}(100|23\lambda_B) \times \text{gamma}(\lambda_S|2.5,1) d\lambda_B d\lambda_S}{\int_0^\infty \text{Poisson}(100|\lambda = 23\lambda_B) d\lambda_B} \\
& \times \frac{p(\text{Source}) = 0.05}{p(\text{No Source}) = 0.95}.
\end{aligned} \tag{4.20}$$

This integral evaluates to  $PR = 1.16$ , which agrees with the Monte Carlo analysis to within less than 1%. Analyzing the model with MCMC is much more straightforward (See Appendix), and results in the probability that a source is present as,  $p(\lambda_S > 0|N_B, N_G) = 0.539$ , or about 54%. It follows that  $p(\lambda_S = 0|N_B, N_G) = 0.461$ , or about 46%. This corresponds to  $PR = \frac{0.539}{0.461} = 1.17$ .

The mean and 95% highest density credible interval (HDI) of the posterior of the sample count rate are

$$\begin{aligned}
\bar{\lambda}_S &= 1.06, \\
95\% \text{ HDI} &= \begin{cases} [0] & 46\% \\ (0.9, 3.1) & 54\% \end{cases}
\end{aligned}$$

An interesting feature of the posterior of  $\lambda_S$  is that the 95% HDI is discontinuous. This means that there is a 95% chance that the sample count rate is either precisely zero (no source) or between 0.9 and 3.1 cps. This is indicative of a bimodal distribution – i.e., a distribution with two distinct peaks. The 95% quantile confidence interval for Example 2 would include, in addition to higher values, the interval between 0 and 0.9, obscuring the bimodal nature of the posterior. The marginal posterior distribution and 95% HDI of  $\lambda_S$  are shown in Figure 4.2.

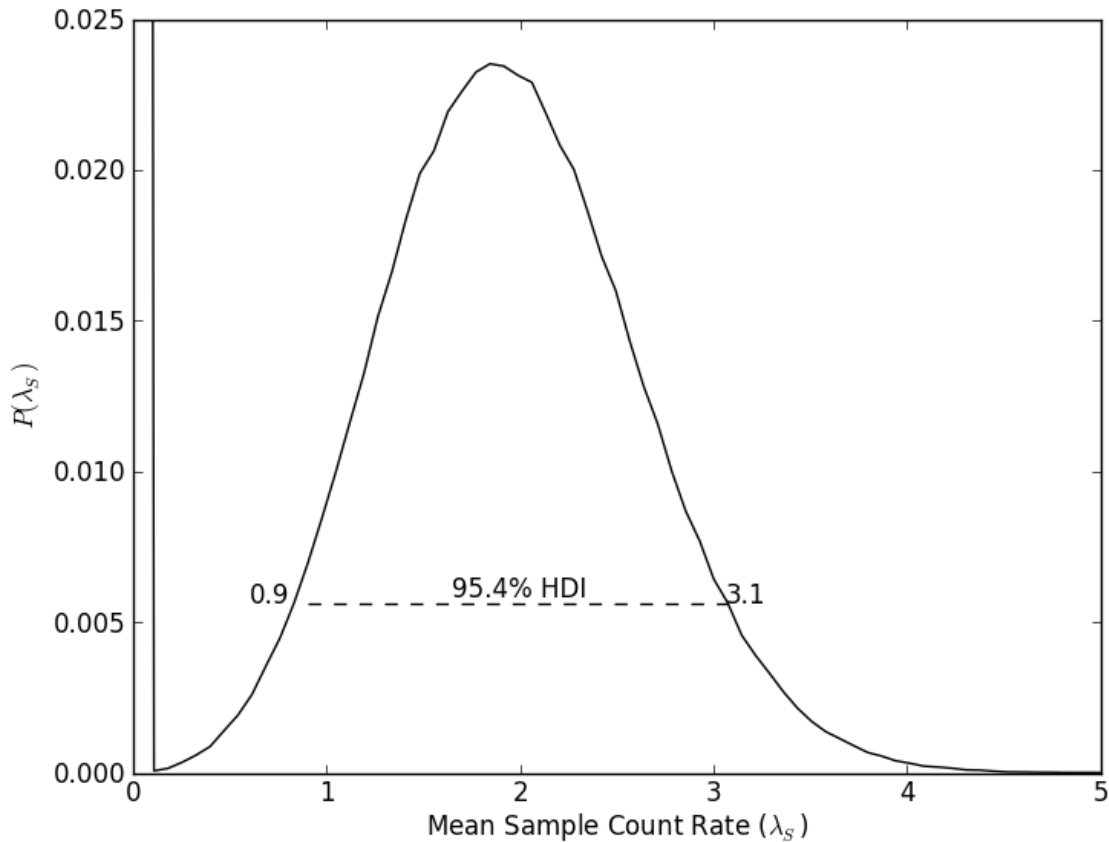


Figure 4.2: The marginal posterior distribution of  $\lambda_S$  for Example 2. The large spike at  $\lambda_S = 0$  represents the probability that the sample contains no source. The 95% HDI, represented by the dashed line, includes  $\lambda_S = 0$  and  $\lambda_S \in (0.9, 3.1)$ .

#### 4.5 Discussion

Bayesian hypothesis testing (and thus analysis with zero-inflated priors) also naturally discourages unreasonable priors. For example, putting a uniform prior on the source count rate is equivalent to stating that there is an equal probability of the source activity being 1 Bq and 100 MBq. As shown above, hypotheses which correspond to uniform (i.e., unreasonable) priors inevitably fail when compared to informative priors in Bayesian hypothesis testing. Although uniform priors are clearly unrealistic, they nonetheless do not favor any one parameter value over any other, and therefore often have little impact on the posterior. For that reason, they are

commonly used to represent the investigator's lack of knowledge about the true values of the parameters, and are perceived as producing "objective" posteriors. The unique sensitivity of hypothesis testing to the variance of the priors makes it difficult to identify priors which will not unduly influence the results. Ideally, there will be empirical knowledge about the likely values of the source count rate. If there is not, it is possible to generate empirical data prospectively by performing multiple observations. However, the ideal technique for generating empirical priors is a topic for further research and discussion.

#### **4.6 Conclusion**

Two of the most common applications of counting measurements are source detection and sample count rate analysis. The traditional Bayesian approaches to these two problems are very different, and are therefore difficult to combine. In any situation where a sample may or may not contain a source, traditional priors cannot adequately represent the possibility that the sample count rate is precisely zero. One approach to this problem is to use a uniform prior to obtain a probability distribution for the sample count rate if the sample contains a source, and then to perform hypothesis testing. By performing two different techniques without attempting to unify them, the investigator is essentially forced to conduct twice as much work. Another disadvantage to this approach is that it allows the investigator to 'get away' with using an unreasonable uniform prior on the source count rate.

Another approach to dealing with samples which may contain no source is to use a prior which compresses some fraction of the probability density as close as possible to zero – e.g., an 'alpha' prior. A disadvantage of this approach is that such distributions are extremely difficult to formulate and to sample from for the purposes of MCMC. In addition, alpha priors

may lead to posterior distributions which are difficult to communicate to the general public, for whom there is a very big difference between a “very small source” and “no source.”

Given the importance of priors which honestly represent the investigator’s beliefs, the advantage of zero-inflated priors is clear. Zero-inflated priors can compactly encapsulate the beliefs of the investigator – specifically, how likely it is that the sample contains a source and how strong that source is likely to be if it exists. By incorporating hypothesis testing into count rate analysis, they eliminate the need for redundant calculations, while at the same time forcing investigators to postulate reasonable priors. Zero-inflated priors are also relatively easy to sample from, meaning that data can be analyzed using open-source MCMC software such as OpenBUGS (See Appendix). Finally, zero-inflated priors result in zero-inflated posteriors, meaning that it will always be possible to report a finite probability that the sample count rate is precisely zero, and thus does not contain a source. It is hoped that the use and interpretation of models utilizing zero-inflated priors will prove so simple and straightforward that the technique will become widely adopted for statistical decision making in health physics.

## Chapter 5

### SUMMARY AND CONCLUSIONS

#### 5.1 Bayesian Analysis

The use of Bayesian techniques for analyzing radiation count rate data throughout this paper reveals a number of advantages over classical analysis. The most obvious advantage is that Bayesian analysis allows the investigator to make direct statements about the quantities of interest. For example, after a measurement of a possibly radioactive source, classical analysis can only estimate the probability of having observed  $N$  counts for a given source count rate. This is very different from a probability distribution of the source count rate, which can only be obtained by incorporating a prior, but it is easy to see how the two probabilities could be confused. This document shows how Bayesian analysis can directly answer the question being asked – in this case, “is there a source, and, if so, how strong is it likely to be?”

The prior probability distribution is the most important and the most controversial aspect of Bayesian analysis. Although requiring the investigator to explicitly state their beliefs in advance may seem inconvenient and potentially hard to justify, it also avoids the problems of a classical analysis, in which assumptions are often implicit and not well understood, even by the investigator (e.g., the assumption that a sample is Normally distributed, or that the samples are independent). Hypothetically, the hidden assumptions of classical significance tests makes it possible for an investigator to analyze data using a number of different classical tests until one produces the desired outcome. In contrast, the requirement in Bayesian analysis to explicitly state a prior forces the investigator to think seriously about the assumptions necessary for an

analysis, knowing that those assumptions will be questioned and considered by others. This leads to more realistic assumptions and, therefore, more realistic results. For example, the use of a strictly positive prior makes it impossible to obtain a confidence interval for the sample count rate which contains negative numbers. The results of a Bayesian analysis are therefore more reasonable, easier to understand, and more easily discussed by the scientific community.

## **5.2 Sequentially-Updated Source Detection Systems**

Two of the most common applications of counting measurements are source detection and sample count rate analysis. In Chapter 2, a method of source detection is proposed which operates under the idealized assumption that the background distribution is fixed and perfectly well known. This assumption greatly simplifies computations, enabling the detection system to update the posterior distribution in real-time and allowing for near instantaneous detection decisions. By considering multiple energy-regions simultaneously, the detection system is able to detect sources more quickly and reliably with a decreased incidence of false positives.

In particular, incorporation of energy data decreases the Average Run Length (ARL) of fixed source detection systems, effectively allowing them to detect much weaker sources in a given time interval than gross count rate detectors. At 5 cps, which corresponds to a source count rate of 3 cps, the ARL of the two-energy detector is almost 50% less than that of the gross count rate detector. This confirms the first research hypothesis regarding the multi-energy detector method.

Incorporation of energy data also significantly improves performance when monitoring for transient sources. The LLD of each detector is determined by controlling the false positive rate to 5%, and reading off the source strengths at which the detection probability for each

detector was 95%. By this metric, the LLD of the multi-energy detectors is found to be approximately 1.5 cps. This exceeds the performance gain predicted in the second research hypothesis regarding the multi-energy detector. Conversely, when the failed-detection (e.g., non-detection) rate for 1.5 cps sources is controlled to 5%, the false positive rates of the multi-energy detectors are relatively unchanged at approximately 5%, while the false-positive rate of the gross count rate detector drop to approximately 45%. As predicted, the multiple-energy transient source detectors significantly out-perform their single-energy counterparts in terms of false positives and failed detections.

### **5.3 Moving Target Model of Background**

Source detection and source count rate analysis both depend on an accurate model of radiation background. The third chapter shows that, unlike fixed radioactive sources, the mean of the radiation background count rate (i.e. the Poisson mean) may sometimes be a function of time. Non-constant background radiation is characterized with a “moving target” model, in which the Poisson mean for the background is assumed to change at specified time intervals, with the new mean being drawn from a gamma distribution with parameters  $\alpha$  and  $\beta$ . Given a dataset consisting of replicate background count measurements, we determine the values of  $\alpha$  and  $\beta$  using maximum likelihood and Bayesian probabilistic analysis using Markov Chain Monte Carlo (MCMC).

When a new Poisson mean is drawn from a gamma distribution between each observation, the resulting sampling distribution is a negative binomial distribution. As a result, decision thresholds are different under the moving target model than they are under the usual assumption of constant background. In particular, we find that classical decision thresholds

increase with the standard deviation of the mean count rate,  $\sqrt{\alpha}/\beta$ . This confirms the second research hypothesis regarding the moving target model. We also find that Bayesian decision thresholds would be proportional the standard deviation of the mean count rate. This confirms the second research hypothesis regarding the moving target model.

If negative binomial sampling distributions derived from the observed data were used in a classical null-hypothesis test, the 95% decision threshold for the stationary target model would be 15 counts. In other words, one would expect more than 15 counts to be observed in less than 5% of background measurements. In contrast, the 95% decision threshold corresponding to the moving target model is 17 counts. The classical 95% decision threshold is therefore 13% higher under the moving target model than it would be if the background were assumed to be stationary. This confirms the first research hypothesis regarding the moving target model.

#### **5.4 Zero-Inflated Priors**

In any situation where a sample may or may not contain a source, traditional priors cannot adequately represent the possibility that the sample count rate is precisely zero. Therefore, Chapter 4 proposes using non-traditional, “zero-inflated,” priors to model the mean count rate of a possibly radioactive sample. Zero-inflated priors compactly encapsulate the beliefs of the investigator – specifically, how likely it is that the sample contains a source and how strong that source is likely to be if it exists. By incorporating hypothesis testing into count rate analysis, they eliminate the need for redundant calculations, while at the same time forcing investigators to postulate reasonable priors. Zero-inflated priors are also relatively easy to sample from, meaning that data can be analyzed using open-source MCMC software such as



OpenBUGS (See Appendix). Finally, zero-inflated priors result in zero-inflated posteriors, meaning that it will always be possible to report a finite probability that the sample count rate is precisely zero, and thus does not contain a source. It is hoped that the use and interpretation of models utilizing zero-inflated priors will prove so simple and straightforward that the technique will become widely adopted for statistical decision making in health physics.

## References

- Arandjelovic V, Koturovic A, Vukanovic R. A software method for suppressing statistical fluctuations in preset count digital-rate meter algorithms. *IEEE Trans Nucl Sci* 49:2561-2566; 2002.
- Bartlett P., Jordan M, and Mcauliffe J. A comment on D. V. Lindley's statistical paradox. *Biometrika* 44, 533-534; 1957.
- Berger, James O., and Luis R. Pericchi. "The intrinsic Bayes factor for model selection and prediction." *Journal of the American Statistical Association* 91.433: 109-122; 1996.
- Brandl, A. "Statistical Considerations for Improved Signal Identification from Repeated Measurements at Low Signal-to-background Ratios." *Health physics* 104.3: 256-263; 2013.
- Chalmers N., LaBone T. Random Thoughts: Bayesian Statistics. Unpublished manuscript from personal correspondence; 2013.
- DeVol TA, Gohres AA, Williams CL. Application of classical versus Bayesian statistical control charts to on-line radiological monitoring. *J Radioanalytical Nucl Chem* 282: 933 – 938; 2009.
- Dowdy EJ, Henry CN, Hastings RD, France SW. Neutron detector suitcase for the nuclear emergency search team. Los Alamos, NM: Los Alamos Scientific Lab; Report LA-1708; 1978.
- Fehlau PE. Comparing a recursive digital-filter with the moving-averaging and sequential probability-ratio detection methods for SNM portal monitors. *IEEE Trans Nucl Sci* 40:143 – 146; 1993.
- Gelman, A. Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper). *Bayesian analysis*, 1(3), 515-534; 2006.
- Ghosh S. K., Mukhopadhyay P., and Lu, J. C. J. Bayesian analysis of zero-inflated regression models. *Journal of Statistical Planning and Inference* 136(4), 1360-1375; 2006.
- Hemic, G. Environmental radiation monitoring in the context of regulations on dose limits to the public. *Journal Office de la Republique Française* 4:88-521; 1988.
- Huiming Z, Lili C, Yu D. Some Properties of the Generalized Stuttering Poisson Distribution and its Applications. *Studies in Mathematical Sciences* 5 (1): 11–26; 2012.

- ISO, 2000f. Standard ISO 11929-1 through –4 Determination of the detection limit and decision threshold for ionizing radiation measurements, Geneva.
- Jeffreys H. An invariant form for the prior probability in estimation problems. Proc. Roy. Soc. A 186, 453-461; 1946
- Jeffreys H. The theory of probability. Oxford University Press. 1998.
- Jones E, Oliphant T, Peterson P, et al. Scipy: Open Source Scientific Tools for Python; 2001. Available at <http://scipy.org>.
- Kaeker R, Jones A. On use of Bayesian statistics to make the guide to the expression of uncertainty in measurement consistent. Metrologia 40: 235 -248; 2003.
- Kalman R. E. A New Approach to Linear Filtering and Prediction Problems. Transaction of the ASME – Journal of Basic Engineering. 35 -45; 1960.
- Kruschke J. Doing Bayesian data analysis: A tutorial introduction with R. Academic Press. 88-90; 2010.
- Little RJA. The statistical-analysis of low-level radioactivity in the presence of background counts. Health Phys 43: 693-709; 1982.
- Lunn D.J., Thomas, A., Best, N., and Spiegelhalter, D. WinBUGS -- a Bayesian modelling framework: concepts, structure, and extensibility. Statistics and Computing 10:325-337; 2000
- Luo P, Sharp J, Devol T. Bayesian Analyses of Time-Interval Data for Environmental Radiation Monitoring. Health Phys. 104:15-25; 2013.
- Miller G, Inkret WC, Martz HF. Bayesian detection analysis for radiation exposure. Radiat Protect Dosim 48:251 – 256; 1993.
- Miller G, Inkret WC, Little TT, Martz HF, Schillaci ME. Bayesian prior probability distributions for internal dosimetry. Radiat. Protect. Dosim. 94:347 -352; 2001.
- Miller, G., Martz, H. F., Little, T. T., & Guilmette, R. "Using exact Poisson likelihood functions in Bayesian interpretation of counting measurements." Health physics 83.4: 512-518; 2002.
- Miller, G. Probabilistic Interpretation of Data. <http://www.lulu.com/content/paperback-book/probabilistic-interpretation-of-data/13620353>: Guthrie Miller, 2013.
- National Council on Radiation Protection and Measurements. A handbook of radioactivity measurements procedures (2nd edition)NCRP, Bethesda, MD (1985) NCRP Report Number 58

- Nayak, Tapan Kumar. Multivariate Lomax distribution: properties and usefulness in reliability theory. *Journal of Applied Probability*. 170-177; 1987.
- Neelon B. H., O'Malley A. J., and Normand S. L. T. A Bayesian model for repeated measures zero-inflated count data with application to outpatient psychiatric service use. *Statistical Modelling* 10(4), 421-439; 2010.
- Ntzoufras, I. *Bayesian Modeling Using WinBUGS*. New Jersey: John Wiley & Sons, Inc; 2009.
- O'Hagan A. Fractional Bayes factors for model comparison. *Journal of the Royal Statistical Society. Series B (Methodological)* 99-138; 1995
- Papoulis, A, Unnikrishna P. *Probability, random variables, and stochastic processes*. Tata McGraw-Hill Education; 2002.
- Potter WE. Neyman-Pearson confidence intervals for extreme low-level paired counting. *Health Physics* 76:186–187; 1999.
- Roberts, S. W. A comparison of some control chart procedures. *Technometrics* 8: 411-430; 1966
- Shiryayev, A. N. The problem of the most rapid detection of a disturbance in a stationary process. *Soviet Mathematics-Doklady* 2: 795-799; 1961.
- Šidák, Z.K. Rectangular Confidence Regions for the Means of Multivariate Normal Distribution. *Journal of the American Statistical Association* 62 (318): 626–633; 1967.
- Strom DJ, MacLellan JA. Evaluation of eight decision rules for low-level radioactivity counting. *Health Phys.* 81:27 – 34; 2001.
- Taylor, W. A. Change-point analysis: a powerful new tool for detecting changes. *http://www.variation.com/cpa/tech/changepoint.html*; 2000.
- Turner, James E. *Atoms, radiation, and radiation protection*. John Wiley & Sons; 2008.
- Van Rossum G, Drake F.L. *Python Reference Manual*, Python Labs, Virginia, USA; 2001. Available at <http://www.python.org>.
- Weise K, Huebel K, Rose E, Schlaeger M, Schrammel D, Taeschner M, Michel R. Bayesian decision threshold, detection limit and confidence limits in ionizing radiation measurement. *Radiat Protect Dosim* 121:52 – 63; 2006.
- Welch G, and Bishop G. *An introduction to the Kalman filter*. (1995).

## Appendix

This section provides a step-by-step description of how to analyze count measurement data using OpenBUGS. OpenBUGS is a freely available, open-source software package designed for “Bayesian Analysis Using Gibbs Sampling” (BUGS). It runs on all major operating systems, including Windows, Macintosh, and Linux. After downloading OpenBUGS, open a new project by clicking ‘File -> New.’ This is shown in Figure 4.A.1.




Figure 4.A.1: Open a new project by clicking ‘File -> New.’

A window will appear in which the model, data, and initial values may be entered. To enter the model, type “model { ... }”, with the generative model between the brackets. The nomenclature is very similar to that used in this paper, except that deterministic relationships are expressed using backwards arrows, ‘<-’, instead of ‘=’. Text following the pound sign, ‘#’, is ignored by the compiler. OpenBUGS has a number of built-in probability distributions, such as ‘dgamma,’ ‘dpois’, ‘dbern’, ‘dunif’, and ‘dnorm’, corresponding to the gamma, Poisson, Bernoulli, uniform, and Normal distributions, respectively. Finally, algebra is not permitted in lines indicating stochastic relationships (with a few rare exceptions). For example, the line,

```
N ~ dpois(lambda*t)
```

would become

```
mu <- lambda*t  
N~dpois(mus)
```

The generative model used in Example 1 would read as follows:

```
model {  
#Counts from background in the detector (b0)  
lambdaB ~ dgamma(1,.01)           #Mean background count rate  
muB <- lambdaB*tB                 #Mean background counts in  
interval tb  
NB ~ dpois (muB)                  #Observed counts in background  
measurement  
  
#Counts from Source  
source ~ dbern(p_source)          #0 if no source, 1 if there  
is a source  
lambdaSource ~ dunif(0,5)         #Source Activity (Bq) assuming  
there is a source)  
lambdaSample <- source*lambdaSource #True source activity  
  
#Measured Counts  
muG <- (lambdaB+lambdaSample)*tg  #Mean gross counts in  
interval tg  
NG ~ dpois(muG)                   #Observed counts in sample  
measurement  
}
```

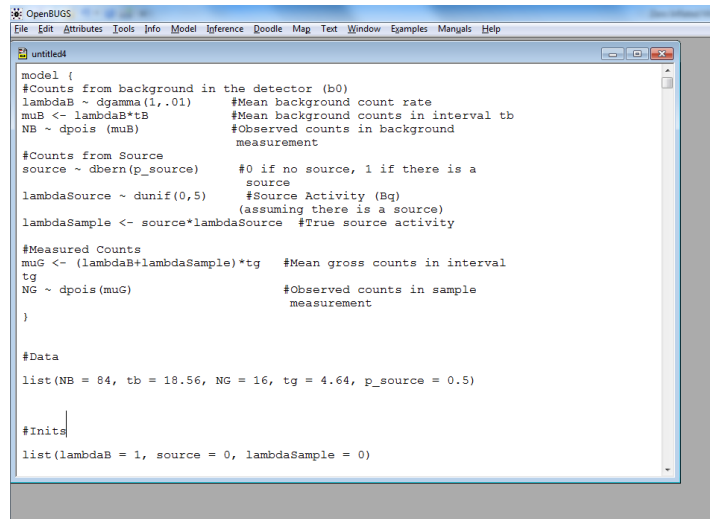
After specifying the model, the next step is to specify the observed data. These are stored in a single list, denoted 'list(...)'. For Example 1, this would read,

```
list(NB = 84, tb = 18.56, NG = 16, tg = 4.64, p_source = 0.5)
```

Optionally, initial values of the undefined stochastic variables (i.e., those which have not been measured) can be specified in another list. For example,

```
list(lambdaB = 1, source = 0, lambdaSample = 0)
```

The final result should look something like Figure 4.A.2.



```
untitled4
model {
#Counts from background in the detector (b0)
lambdaB ~ dgamma(1,.01) #Mean background count rate
muB <- lambdaB*tb #Mean background counts in interval tb
NB ~ dpois(muB) #Observed counts in background
measurement

#Counts from Source
source ~ dbern(p_source) #0 if no source, 1 if there is a
source
lambdaSource ~ dunif(0,5) #Source Activity (Bq)
#(assuming there is a source)
lambdaSample <- source*lambdaSource #True source activity

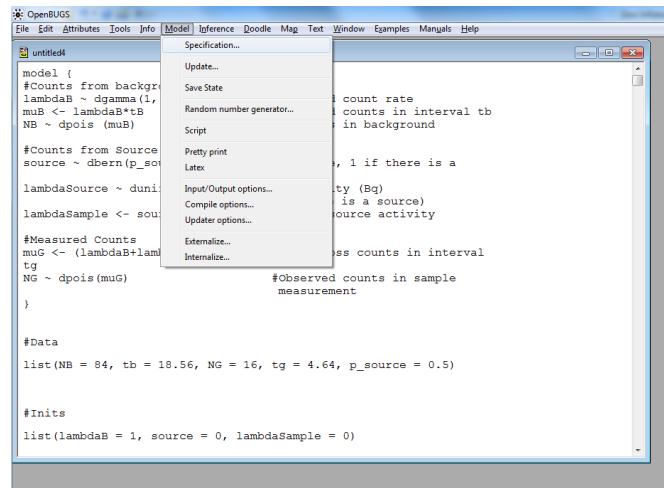
#Measured Counts
muG <- (lambdaB+lambdaSample)*tg #Mean gross counts in interval
tg
NG ~ dpois(muG) #Observed counts in sample
measurement
}

#Data
list(NB = 84, tb = 18.56, NG = 16, tg = 4.64, p_source = 0.5)

#Inits
list(lambdaB = 1, source = 0, lambdaSample = 0)
```

Figure 4.A.2: The complete model specification with data and initial values.

The next step is to load the model, data, and initial values into the program. To do that, select ‘Specification...’ from the ‘Model’ drop-down menu (Figure 4.A.3).



```
untitled4
model {
#Counts from backgr
lambdaB ~ dgamma(1,
muB <- lambdaB*tb
NB ~ dpois(muB)

#Counts from Source
source ~ dbern(p_sou

lambdaSource ~ duni
lambdaSample <- sou

#Measured Counts
muG <- (lambdaB+lamb
tg
NG ~ dpois(muG)
measurement

#Data
list(NB = 84, tb = 18.56, NG = 16, tg = 4.64, p_source = 0.5)

#Inits
list(lambdaB = 1, source = 0, lambdaSample = 0)
```

Figure 4.A.3: Open the model specification window by clicking ‘File -> New.’

Highlight the word ‘model’ (right before the brackets), and click ‘check model.’ If the model is not syntactically correct, an error message will appear in the lower-left corner of the screen. This is shown in Figure 4.A.4. If there are no errors, load the data by highlighting the word ‘list’ before the list of observed data and click ‘load data.’ Again, if there are any errors, a message will appear in the lower left corner. Next, compile the model by clicking ‘compile.’ Finally, load the initial values in the same way. If you do not wish to specify initial values, you can have the program generate them for you by clicking ‘gen inits.’

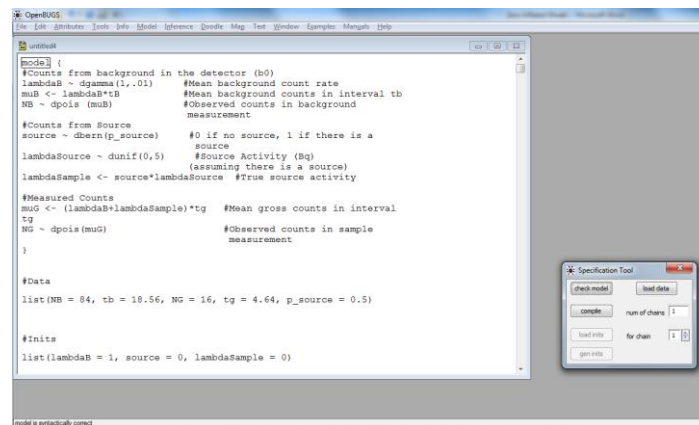


Figure 4.A.4: The screen after clicking ‘check model’ in the ‘Specification Tool’ window. Note that the word ‘model’ is highlighted, and that there is a message in the lower left corner saying ‘model is syntactically correct.’

The next step is to ‘burn in’ the Markov Chain by having BUGS run several thousand untracked MCMC iterations, which allows it to converge on the posterior distribution. From the ‘Model’ drop down menu, select ‘Update...’ Enter the desired burn-in iterations (20,000 is usually enough), and click ‘update.’ Once the burn-in is complete, a message will appear in the lower left corner saying “x updates took n seconds.” This is shown in Figure 4.A.5.



The next step is to specify the parameters whose posterior distribution you want to calculate. To do this, select ‘Samples...’ from the ‘Inference’ drop-down menu. In the text box/drop down menu next to the word ‘node,’ type in the name of each parameter you would like to follow one at a time, then press ‘set’ after typing in each parameter. In this case, the parameters of interest are *lambdaSample*, *source*, and (maybe) *lambdaB*. This is shown in Figure 4.A.6.

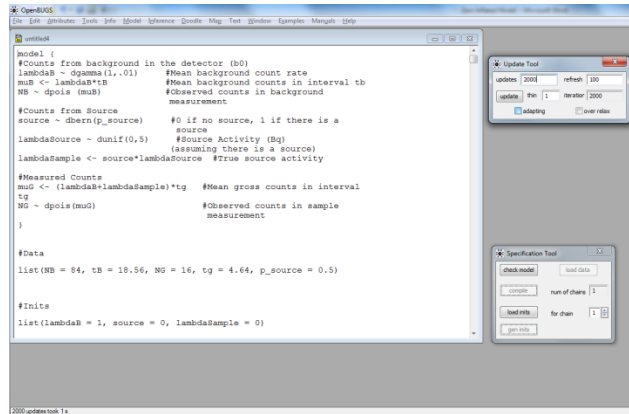


Figure 4.A.5: The screen after running 2000 ‘burn-in’ iterations

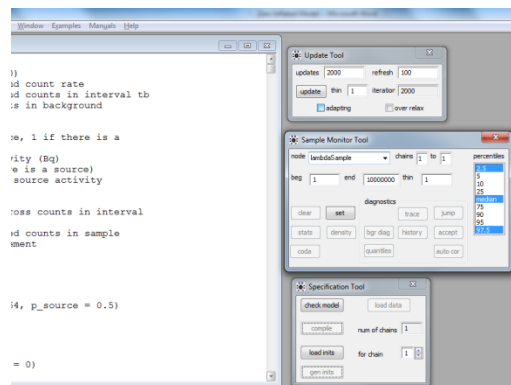


Figure 4.A.6: Entering the parameters we wish to track. Note that the ‘set’ button darkens when you have typed in a parameter name which is the model. The next step is to click that button.

To run the analysis, go back to the ‘Update Tool’ window and select the number of iterations you would like to run. Every iteration produces one independent sample from the joint posterior

distribution. The more parameters in the model, the more iterations it takes to explore the sample space. A reasonable number to start with is 10,000, but it is not uncommon to use 100,000 or even 1,000,000.

Once the simulation is finished, return to the ‘Sample Monitor Tool’ window to review the results. To view results for one parameter at a time, select the desired parameter from the ‘node’ text box/drop-down menu. To view all parameters at once, type ‘\*’ into the node text box/drop-down menu. To view summary statistics, click the ‘stat’ button. Plots of the distributions can be viewed by pressing the ‘density’ button. To view the raw data, press the ‘coda’ button. The result of pressing all of these buttons is shown in Figure 4.A.7.

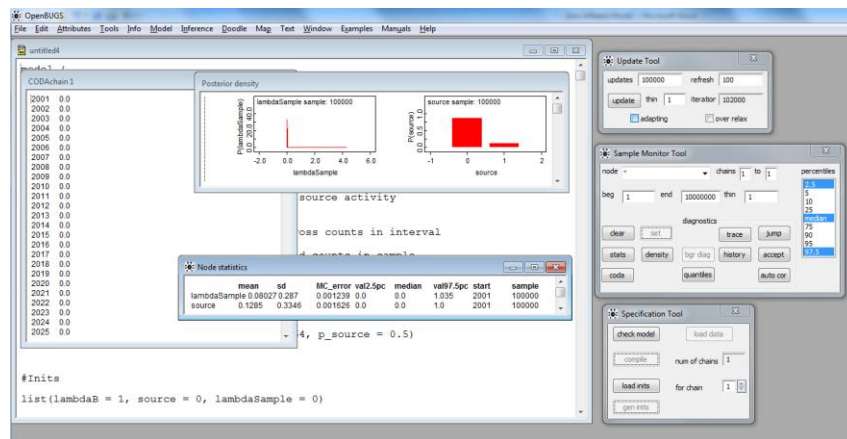


Figure 4.A.7: The summary statistics, density plots, and raw data for the specified parameters after 100,000 iterations.

These data can then be imported and analyzed using a program of choice. To ensure convergence of the chain, press the ‘history’ button. The plot should look like white noise. If it looks like a flat line with occasional jumps, or if it trends up or down, the chain has not converged, and the model needs to be modified. This can happen if the initial values are far from the posterior mean, if the priors are inconsistent with the observed data, or if not enough iterations have been run.

Figure 4.A.8 shows a history plot for *lambdaSample*. The plot shows no obvious trends and appears like white noise, indicating that the chain has likely converged.




Figure 4.A.8: history plot for the lambdaSample chain. Note that the plot looks like white noise, which is indicative of convergence.

This is obviously only a brief introduction to OpenBUGS. Users should read the instruction manual, and educate themselves on the details of Markov Chain Monte Carlo, especially those related to chain convergence.