Computers in Biology and Medicine 82 (2017) 3–11

Contents lists available at ScienceDirect

Computers in Biology and Medicine

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

Bayesian approach to decompression sickness model parameter estimation

L.E. Howle[a,b,*], P.W. Weber[a,b], J.M. Nichols[c]

[a] Mechanical Engineering and Materials Science Department, Duke University, 144 Hudson Hall, Durham, NC 27708-0300, United States
[b] BelleQuant Engineering, PLLC, Mebane, NC 27302-9281, United States
[c] Naval Research Laboratory, Maritime Sensing Section (Code 5673), 4555 Overlook Ave. SW, Washington, DC 20375, United States

ARTICLE INFO

Keywords:
Bayesian
Markov-Chain-Monte-Carlo
Decompression sickness
Decompression model
Decompression illness

ABSTRACT

We examine both maximum likelihood and Bayesian approaches for estimating probabilistic decompression sickness model parameters. Maximum likelihood estimation treats parameters as fixed values and determines the best estimate through repeated trials, whereas the Bayesian approach treats parameters as random variables and determines the parameter probability distributions. We would ultimately like to know the probability that a parameter lies in a certain range rather than simply make statements about the repeatability of our estimator. Although both represent powerful methods of inference, for models with complex or multi-peaked likelihoods, maximum likelihood parameter estimates can prove more difficult to interpret than the estimates of the parameter distributions provided by the Bayesian approach. For models of decompression sickness, we show that while these two estimation methods are complementary, the credible intervals generated by the Bayesian approach are more naturally suited to quantifying uncertainty in the model parameters.

1. Introduction

Decompression sickness (DCS) is a concern for astronauts, aviators, and scuba divers (military, commercial and recreational). The widely-accepted cause of DCS is inert gas bubble formation in the body following one or more exposure(s) to decompression. The symptoms of DCS can be severe (paralysis, death) to mild (itching, joint pain), depending heavily upon factors such as the pressure profile to which the subject was exposed as well as factors related to the individual [1]. Besides its ability to cause human injury, DCS can also be responsible for mission failure in the military and lost work time in commercial endeavors, so avoidance of DCS is an area of great interest.

Much emphasis has been placed on DCS modeling in the past. The Haldane theory was the first to postulate that DCS was caused by bubbles forming due to nitrogen supersaturation in the body tissues [2]. The Haldane model prescribed that a diver should ascend at a rate that does not exceed a 2:1 pressure ratio between the calculated nitrogen tension in any of five hypothetical body compartments and the ambient pressure. This deterministic procedure, commonly referred to as stage decompression, is still in use today.

Berghage et al. [3] and Weathersby et al. [4] were the first to introduce the concept of probabilistic modeling into the DCS field. Prior to this, DCS prediction algorithms were deterministic. That is, the output from the algorithm generated a sharp delineation between a 'safe' and 'unsafe' dive and, therefore, generated a two state (binary) prediction of dive safety. Probabilistic modeling assumes that the binary (yes/no) DCS event can be described by a model that assigns a continuous probability that the dive profile will result in DCS. Also, the same diver can execute the same dive on different occasions yet only sometimes experience symptoms. This approach is more in line with observed empirical data, which has shown that multiple divers can execute the exact same dive profile but not all of them will experience symptoms. Empirical dive data were used with these probabilistic models to estimate model parameters using the method of maximum likelihood. As the name implies, the method of maximum likelihood seeks the parameter values that maximize the probability of having observed the empirical data.

All estimation problems require that we first construct a model[1] that we believe predicts the physical phenomenon of interest, in this case DCS. The job of the estimator is to produce parameter values, $\vec{\chi}$, that are close to the 'true' parameter values that define our model. A good estimator is one that performs well in the face of the inevitable uncertainty we have about the underlying process (e.g. noise, model error, etc.). Additionally, a good model is expected to perform reasonably well when applied to data, in this case dive profiles, that are outside of the training set used in optimizing the parameter values.

There are two basic schools of thought regarding estimation problems, stemming from two different interpretations of probability.

* Corresponding author at: Mechanical Engineering and Materials Science Department, Duke University, 144 Hudson Hall, Durham, NC 27708-0300, United States.
E-mail address: laurens.howle@duke.edu (L.E. Howle).

[1] Note that we could also prescribe multiple models. The growing field of multi-model inference addresses this topic and is not discussed here.

http://dx.doi.org/10.1016/j.compbiomed.2017.01.006
Received 20 January 2016; Received in revised form 13 January 2017; Accepted 16 January 2017

The first, sometimes referred to as the frequentist interpretation, defines probability in terms of the relative frequencies with which an event occurs. Using this approach to estimation, one treats the model parameters as fixed quantities, finds the 'best' estimate (to be defined in the next section), and generates measures of confidence in the result by repeating the estimation procedure multiple times, or by making asymptotic approximations about the estimator variance.

The alternative is to adhere to the axiomatic definition of probability and treat the parameters as random variables, each obeying its own probability distribution. The procedure by which this parameter probability distribution is estimated is described by Bayes' Rule, hence is referred to as Bayesian estimation. The final parameter estimate is often taken as the value that maximizes this distribution (i.e. most probable); a measure of confidence in the result is obtained by quantifying the spread of this distribution.

The goal of this work is to compare these two approaches to estimation in the context of predicting the probability of experiencing DCS. We are interested both in generating model parameter estimates that accurately predict observed instances of DCS and in developing intervals of confidence for those parameter estimates. The DCS model used in this study was investigated previously by the lead authors, and was aimed towards predicting the probability of DCS for Navy divers, as the empirical data to which the model was fitted contained the results of military dive trials. While the general model is described here the reader is referred to [5] for a more detailed presentation.

2. Probabilistic DCS model

For the purposes of demonstrating the Bayesian parameter estimation method as applied to probabilistic DCS modeling, we will use the simple three-compartment exponential model, the EE1(nt) model, that has been extensively studied elsewhere [5–7]. The model begins with a dive profile containing a series of 'dive legs' during which the ambient pressure is either constant or changes with a constant slope in time. For example, the nitrogen partial pressure, P_{N_2}, for a given dive leg can be written as

$$P_{N_2} = P_{N_2}^0 + R_{N_2}(t - t_0) \tag{1}$$

where $P_{N_2}^0$ is the nitrogen partial pressure at the beginning of the dive leg, the constant-slope rate of change of the nitrogen partial pressure is given by R_{N_2}, and t_0 represents the time at which the particular dive leg begins. An entire dive profile is thus created by connecting a series of either constant depth or constant-slope dive legs. This creates dive profiles with continuous ambient pressure traces. The nitrogen partial pressure, however, need not be continuous in time. A discontinuous nitrogen partial pressure can result if abrupt changes in the breathing gas take place. An example of an abrupt breathing gas change occurs if a diver undergoes decompression procedures involving changes in breathing gases. In this case, the diver might breathe a specific gas while at maximum depth, such as 21/35 normoxic trimix (21% O_2, 35% He, remainder N_2), and then switch to oxygen or nitrox during a portion of the decompression phase of the dive. However, even with gas switching, the change between the two gasses is assumed to occur over a short duration, typically, 0.1 min.

For each of the three parallel, well-perfused tissue compartments in the EE1(nt) model, the arterial nitrogen partial pressure – or nitrogen tension – is assumed to be equal to the alveolar nitrogen tension. Therefore, using a well stirred analogy to a mixing problem, we may write a change in the nitrogen tissue tension for the i^{th} compartment P_{T_i}, as

$$\frac{dP_{T_i}}{dt} = k_i(P_{T_{in}} - P_{T_{out}}) \tag{2}$$

where each compartment has a unique tissue rate, k_i. For the present work, $i = 1..3$. The tissue inlet and outlet inert gas tensions are shown in Eq. (2) by the respective symbols $P_{T_{in}}$ and $P_{T_{out}}$. When we make the assumption that $P_{T_{in}}$ changes according to Eq. (1) and that $P_{T_i,out} = P_{T_i}$, we arrive at a simple system of uncoupled linear, first-order differential equations for the change in the compartment gas tensions

$$\frac{dP_{T_i}}{dt} + k_i P_{T_i} = k_i(P_{N_2}^0 + R_{N_2} t) \tag{3}$$

which admit the solutions

$$P_{T_i}(t) = \alpha_i e^{-k_i t} + R_{N_2} t + \beta_i \tag{4}$$

where the constants

$$\alpha_i = P_{T_i}^0 - P_{N_2}^0 + k_i^{-1} R_{N_2}$$
$$\beta_i = P_{N_2}^0 - k_i^{-1} R_{N_2} \tag{5}$$

are fixed for each leg of the dive profile.

Following Weathersby et al. [4,8–10] and Thalmann et al. [6,7] the probability that a diver may experience DCS during a particular dive profile can be approximated using the failure model

$$P(DCS) = 1 - e^{-\sum_i g_i \int r_i \, dt} \tag{6}$$

for a suitable risk (or hazard) function, r_i. We choose the supersaturation ratio risk function

$$r_i = \max\left(\frac{P_{T_i} - P_{amb} + P_{FVG}}{P_{amb}}, 0\right) \tag{7}$$

where the subscript i was previously defined, $P_{FVG} = 6.34 fsw$ (feet of sea water, $33.066\, fsw = 1\ atm$) is the partial pressure of the fixed venous gases, and P_{amb} is the ambient pressure. The integrated risk for each tissue compartment is scaled by a compartment-specific gain parameter, g_i, as shown in Eq. (6).

In order to fit the model to empirical data, many previous studies have used log likelihood maximization as originally proposed for decompression models by Weathersby et al. [4]. A number of techniques may be used to maximize the log likelihood such as various gradient-based methods, simulated annealing, simplex, or other methods. All of these methods attempt to adjust the free parameter vector, here, $\vec{\chi} = (g_1, g_2, g_3, k_1, k_2, k_3)$, until the best fit is achieved between the model and the empirical data. For the present problem, the log likelihood, LL, is defined by

$$LL = \sum_{d=1}^{ND} log\,[P_d(DCS)^{\delta}(1 - P_d(DCS))^{1-\delta}] \tag{8}$$

where the summation occurs over the ND dive profiles comprising the data set. Additionally, in Eq. (8), $P_d(DCS)$ is the probability that a diver develops DCS for the d^{th} dive profile and δ is the binary outcome indicator. That is, $\delta = 1$ if DCS occurred in the data for that dive profile and $\delta = 0$ otherwise. When viewed from the frequentist approach, the log likelihood (Eq. (8)) gives us the log probability of observing the outcomes of the empirical data, $\vec{y}$, given the parameter set $\vec{\chi}$. We will denote this as $P_L(\vec{y}|\vec{\chi})$.

3. Brief review of maximum likelihood estimation and Bayesian estimation

Central to both maximum likelihood and Bayesian approaches to inference is the likelihood [11] which describes the joint probability of our sequence of observations $\vec{y}$ given model parameters $\vec{\chi}$. It would therefore make sense to define the 'best' estimate as one that maximizes this probability. Indeed, maximizing Eq. (8) over all parameters produces maximum likelihood parameter estimates (MLPEs), denoted $\hat{\chi}$. The method of maximum likelihood is a powerful approach to estimation that comes with some very important guarantees regarding the quality of the estimator. There are (at least) two properties we would like in an estimator: 1) it be un-biased and 2) it possesses a small variance. That is to say, if we were to repeatedly