Preventive Veterinary Medicine 130 (2016) 129–136

Contents lists available at ScienceDirect

Preventive Veterinary Medicine

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

Nonlinear hierarchical modeling of experimental infection data

Michael D. Singleton[a,*], Patrick J. Breheny[b]

[a] Department of Biostatistics, University of Kentucky College of Public Health, Lexington, KY 40513, United States
[b] Department of Biostatistics, University of Iowa College of Public Health, Iowa City, IA 52242, United States

ARTICLE INFO

Article history:
Received 22 July 2015
Received in revised form 21 May 2016
Accepted 10 June 2016

Keywords:
Infection response
Nonlinear hierarchical model
Viral challenge
Monte Carlo simulation
Longitudinal data analysis

ABSTRACT

In this paper, we propose a nonlinear hierarchical model (NLHM) for analyzing longitudinal experimental infection (EI) data. The NLHM offers several improvements over commonly used alternatives such as repeated measures analysis of variance (RM-ANOVA) and the linear mixed model (LMM). It enables comparison of relevant biological properties of the course of infection including peak intensity, duration and time to peak, rather than simply comparing mean responses at each observation time. We illustrate the practical benefits of this model and the insights it yields using data from experimental infection studies on equine arteritis virus. Finally, we demonstrate via simulation studies that the NLHM substantially reduces bias and improves the power to detect differences in relevant features of the infection response between two populations. For example, to detect a 20% difference in response duration between two groups (n = 15) in which the peak time and peak intensity were identical, the RM-ANOVA test had a power of just 11%, and LMM a power of just 12%. By comparison, the nonlinear model we propose had a power of 58% in the same scenario, while controlling the Type I error rate better than the other two methods.

1. Introduction

Experimental infection (EI) studies have played a central role in infectious disease research and epidemiology for more than two hundred years (Franco 2013; Academy of Medical Sciences, 2005). Often in EI studies, the aim is to determine whether the trajectory of the mean infection response differs in a treatment group relative to a control group. A straightforward approach to carrying out an EI study is to collect data at m fixed times, $t_1, t_2, \ldots, t_m$ and estimate separately the mean responses $\mu_1, \mu_2, \ldots, \mu_m$ at each of those time points, where the time of infection is t = 0. One can then test whether, say, μ_1 differs from μ_5 within a group – a "horizontal" contrast – or whether μ_3 differs between groups – a "vertical" contrast (Ludbrook, 1994).

Because the number of possible such tests is large, it is desirable to test simultaneously for whether μ_j is equal between groups across all times j=1,2,...,m. In general, this type of test is known as an ANOVA; however, because there are multiple measurements per subject, some adjustments must be made. The oldest and simplest way of making these adjustments is to introduce a single, subject-specific random effect and recalculate the various sums of squares involved in the ANOVA test. This approach is known as RM-ANOVA (Davis, 2003). One way of extending the basic RM-ANOVA model is to allow multiple subject-specific effects for each individual, one for each time point, instead of a single subject-specific effect as in the RM-ANOVA model. This type of model falls under a very broad class of models known as the linear mixed model, or LMM (Laird and Ware, 1982). RM-ANOVA and LMM are two very commonly used approaches to analyzing data from EI studies to detect group differences in the response to a pathogen (Munhoz et al., 2012; Go et al., 2012; Naylor et al., 2003).

A limitation of using these approaches for EI studies is that EI data tend to be strongly nonlinear. Following challenge there is typically a brief incubation period followed by an interval of rapid increase (the "onset" phase of the response). After the maximum or minimum value is achieved, there is a return to a stable baseline level (the "recovery" phase). This pattern has been referred to as a peaked response to distinguish it from growth curve data (Matthews et al., 1990). Peaked responses come in a variety of forms, but here, we restrict our attention to experiments in which there is exactly one challenge point, with a single-peaked response and a common pre- and post-challenge baseline level.

Many relevant research questions may be posed in terms of one or more features of a single-peaked infection response curve. For example, an investigator may be interested in whether a vaccine reduces the peak, total amount, or duration of viral shed-

* Corresponding author.
E-mail address: msingle@email.uky.edu (M.D. Singleton).

http://dx.doi.org/10.1016/j.prevetmed.2016.06.006
0167-5877/

ding; whether a certain genetic factor alters the kinetics of the immune response (e.g. maximum decrease in lymphocyte count); or whether two strains of the same virus differ in their effect on the intensity or duration of measurable disease signs such as fever (virulence studies). In transmission studies, interest centers on factors that may influence the transmission of infection. Here the timing or peak intensity of viremia, or the total amount of virus excreted, in subjects exposed to infected hosts, may be informative.

RM-ANOVA and LMM are not well-suited for addressing these types of questions, because they do not explicitly model the mean response curves and their salient features. There is also reason to expect that they will have less power to find a difference at any measurement occasion, even when differences in peak intensity and duration do, in fact, exist. This is because it is inefficient to estimate the mean at every time point when there is a clear and consistent pattern in the individual response profiles, particularly when the number of measurement occasions is large.

The objective of this paper is to present a nonlinear hierarchical model (NLHM) for single-peaked response variables. We demonstrate the application of the model and the insights it provides using data from actual EI studies on equine arteritis virus. Via simulation studies we assess the performance of the model, in terms of bias, power and type I error control. Finally, we discuss recommendations for the applying the proposed model in practice, and for designing studies to collect sufficient data for estimating the parameters of interest.

2. A nonlinear hierarchical model for EI studies

We propose a method that directly represents the quantities of biological interest (intensity, duration, time until peak response) in an EI study. To accomplish this goal we will employ a class of models known as the nonlinear hierarchical model or NLHM (Davidian and Giltinan, 1995; Vonesh and Chinchilli, 1997). NLHMs have been applied in many fields, most prominently in pharmacokinetics and pharmacodynamics research (Sheiner et al., 1972), but to the best of our knowledge, have not been used to analyze data from EI studies.

2.1. Specification of the model

To make the presentation of our NLHM more concrete, consider a hypothetical study to determine whether a vaccine reduces the intensity and/or duration of lymphopenia in horses following infection with equine influenza virus (EIV). Two groups of 10 horses are selected at random from a source population. One group is vaccinated against EIV and the other group receives a sham vaccination (normal saline solution). Two weeks after vaccination, both groups are challenged with EIV and followed for six weeks. Blood samples are drawn on all animals on days that we will refer to as $t_1, t_2,\ldots t_m$, where m represents the number of measurement occasions. We denote the observed lymphocyte counts for horse i as $y_{i1}, y_{i2}, \ldots, y_{im}$. For simplicity we are assuming that m is the same for all animals, but the NLHM does not require this to be the case.

Our NLHM consists of two models that are coupled together in a hierarchical structure. At the lower level is the model for the mean response trajectories of the individual horses in the vaccine and control groups.

Level 1 (*Individual mean response trajectories*)

$$y_{ijg} = f\left(t_j, \boldsymbol{\beta}_{ig}\right) + e_{ij},\ e_{ij} \sim N(0, \sigma^2) \tag{1}$$

Here the letter g indexes the groups (g=1 for the vaccine group and g=2 for the control/sham group), i indexes the individual horses in a group (in our example i = 1,2,. . .,10 for both groups), and j indexes the sampling times $t_1, t_2,\ldots,t_m$. Eq. (1) posits that for each horse in our study, the lymphocyte counts we actually observe consist of an unknown, systematic component, modeled as $f\left(t_j, \boldsymbol{\beta}_{ig}\right)$, plus random noise. The noise in each observed lymphocyte count is represented by the random variables e_{ij}, which are assumed to be independent and normally distributed with mean zero and common variance σ^2. Note that the trajectories share a common general form dictated by f, but that each animal has unique parameters $\boldsymbol{\beta}_{ig}$ that allow for individual differences in the features of the response. For EI studies, it is reasonable to assume that f has a single-peaked shape that rises to a maximum intensity and returns to baseline as the infection is cleared. In this paper, we use the following functional form for f:

$$f\left(t_j, \boldsymbol{\beta}_{ig}\right) = B_{ig} + I_{ig} exp\left[-\frac{\left(t_j - p_{ig}\right)^2}{2s_{ig}^2}\right] \tag{2}$$

Eq. (2) uses a Gaussian function to model the true mean lymphocyte count trajectories. The parameters of this function represent biologically meaningful features of the individual responses. The parameter p_{11}, for example, represents the estimated peak response time for horse 1 in the vaccine group. We can think of p_{11} as dividing the lymphocyte trajectory for this horse into an onset phase and a recovery phase. The symmetry of the Gaussian function implies an assumption that the onset and recovery phases of the lymphocyte response have identical durations. B_{11} and I_{11} represent the baseline lymphocyte level and maximum lymphocyte decline from baseline for the same horse. s_{11} can be considered an index of the duration of lymphopenia. It is directly related to the full width at half maximum (FWHM) of the estimated mean response curve through the formula FWHM = $2s(2\ln 2)^{1/2}$.

It will often be the case that the assumption of symmetric response trajectories is unrealistic. In such cases, the Gaussian function in Eq. (2) can be replaced with a piecewise half-Gaussian response function:

$$f\left(t_j, \boldsymbol{\beta}_{ig}\right) = \begin{cases} B_{ig} + I_{ig} exp\left[-\dfrac{\left(t_j - p_{ig}\right)^2}{2l_{ig}^2}\right], & t_j \le p_{ig} \\ B_{ig} + I_{ig} exp\left[-\dfrac{\left(t_j - p_{ig}\right)^2}{2r_{ig}^2}\right], & t_j > p_{ig} \end{cases} \tag{3}$$

Because the two halves share common baseline, time-to-peak and intensity parameters, the result is a smooth, continuous curve over time, as shown in Fig. 1. The single scale parameter s_{ig} is replaced by distinct onset and recovery scale parameters l_{ig} and r_{ig}. These can be interpreted, respectively, as indices of the duration of the onset and recovery phases of infection. Note that l_{ig} and r_{ig} are directly proportional to the half-widths of the onset and recovery curves at half of their maximum intensity (HWHM): HWHM-l_{ig} = $l_{ig}(2\ln 2)^{1/2}$and HWHM-r_{ig} = $r_{ig}(2\ln 2)^{1/2}$. The sum of these terms can be interpreted as an index of total response duration. As an example, suppose that for horse 1 in the vaccine group we estimated l_{11} = 1.5 and r_{11} = 4. Then, according to this measure, onset lasted $1.5(2\ln 2)^{1/2}$ = 1.8 days and recovery lasted $4(2\ln 2)^{1/2}$ = 4.7 days, for a total duration of 6.5 days.

It should be noted that any unimodal function could be used for f. In our experience with other unimodal functions such as Epanechnikov and tricube functions, the estimates of quantities such as intensity and time to peak are not substantially affected by the specific shape of f, nor is the computational difficulty of fitting the model affected. We used the Gaussian function here as it appeared to best match the response shapes we have observed in real data.

Level 2 specifies a model for the *population* mean response trajectory in terms of the individual mean response trajectories that were modeled at Level 1. In our hypothetical vaccine efficacy study, it is the relationship between these population mean lymphocyte trajectories that is of primary interest.