Chapter 8

Applied Statistics and Exposition

Karl W. Broman

I particularly admire Terry for the breadth of his thinking, the depth of his understanding, the strength and openness with which he expresses his opinions, and the clarity of his writing. All of these qualities are on display in the manuscripts in this section.

But before I comment on those articles, I want to first mention that the best illustrations of Terry's admirable qualities are his regular commentaries for the IMS Bulletin, begun when he was President of the Institute of Mathematical Statistics and continued by popular demand. (To date, he has written over 70 such commentaries! They are available online at <http://bulletin.imstat.org>.) Among my favorites: In Praise of Postdocs [10], Keep Gender on the Agenda [8], A Toast to Posters! [11], Books Worth Reading [7] (his five favorite books are Feller Volume I [2], editions 1, 2 and 3, and Feller Volume II [3], editions 1 and 2), and It's Job-hunting Time! [9]. Terry has strong opinions on a wide range of topics, and we all benefit from his willingness to share them.

Terry's technical report on probabilities related to the reliability of nuclear reactors [14], commenting on probability statements in the Reactor Safety Study [4], demonstrates Terry's willingness to delve deeply and thoroughly into a problem and the clarity and insight he then gives us. (The manuscript is especially relevant today, given the problems with nuclear reactors following the tsunami in Japan.) Misuse of the addition and multiplication rules of probability is well known to statisticians. His further point, that probability statements are meaningless without an understanding of how they were derived, is obvious in retrospect yet often overlooked. It is related to a point that Terry has repeatedly emphasized: an estimate is of no value without some statement on its uncertainty (such as a standard error). I can imagine no better (or more amusing) illustration of this issue than the statement from the Chairman of the U.S. National Transportation Safety Board regarding the chance of two jets colliding on the ground, and the response to Terry's aerogramme. (The word "aero-

K.W. Broman

Department of Biostatistics & Medical Informatics, University of Wisconsin–Madison
e-mail: kbroman@biostat.wisc.edu

gramme” itself makes me smile.) Also amusing is Terry’s comment [14, p. 24], “The whole exercise can now be recognized as being totally without relevance to the real world.” Ouch!

Terry’s work on salmon populations [16] and bioassays [1] were among his last projects before he began to focus almost exclusively on genetics and genomics. The salmon work is a particularly good illustration of Terry’s approach: know the science and the scientific questions and goals, and let the methods follow. I saw this in action in the bioassay project; our (or really *his*) normal-Poisson mixture model, which now seems so natural, was the last of a long sequence of preliminary models. My most vivid memory of that effort was our initial error (or really *my* initial error) in implementing an EM algorithm for the normal-Poisson mixture. The sufficient statistics include not just $\sum x_i$ but also $\sum x_i^2$, and so the E-step requires not just the expected values of the x_i , but also of the x_i^2 . That one should follow the log likelihood across iterations, as a diagnostic for the correct implementation of an EM algorithm, was another important lesson.

Terry’s comment [12] on Robinson’s Statistical Science article on BLUP [5] is masterly. The diverse applications of BLUP that Terry provides add great emphasis to the importance of Robinson’s paper. Terry’s comment has been widely cited as being the first articulation of the connection between smoothing splines and BLUPs.

Finally, let me comment on the two encyclopedia entries reprinted in this section, on restricted maximum likelihood [15] and on iterative proportion fitting (IPF) [13], though I can think of little to say except that they are superb examples of the clarity of Terry’s writing (and to note that Terry has an IMS Bulletin commentary [6] on IPF, too!). The breadth and depth of Terry’s thinking often fills me with envy, but the effort he devotes to sharing his ideas allows me to get past the envy and simply appreciate his insights.

References

- [1] K. Broman, T. Speed, and M. Tigges. Estimation of antigen-responsive T cell frequencies in PBMC from human subjects. *J. Immunol. Methods*, 198: 119–132, 1996.
- [2] W. Feller. *An Introduction to Probability Theory and its Applications*, volume I. Wiley, New York, 3rd edition, 1968.
- [3] W. Feller. *An Introduction to Probability Theory and its Applications*, volume II. Wiley, New York, 2nd edition, 1971.
- [4] N. C. Rasmussen et al. Reactor Safety Study: An assessment of accident risk in U.S. commercial nuclear power plants. Technical report, Nuclear Regulatory Commission, Washington, D.C. (USA), October 1975. Report Number(s): WASH-1400-MR; NUREG-75/014-MR.
- [5] G. K. Robinson. That BLUP is a good thing: The estimation of random effects. *Stat. Sci.*, 6:15–51, 1991.

- [6] T. Speed. Terence's stuff: My favourite algorithm. *IMS Bulletin*, 37 (9):14, 2008.
- [7] T. Speed. Terence's stuff: Books worth reading. *IMS Bulletin*, 34 (2):6, 2005.
- [8] T. Speed. President's column: Keep gender on the agenda. *IMS Bulletin*, 33 (2):4, 2004.
- [9] T. Speed. Terence's stuff: It's job-hunting time! *IMS Bulletin*, 38 (1):14, 2009.
- [10] T. Speed. President's column: In praise of postdocs. *IMS Bulletin*, 32 (6):4, 2003.
- [11] T. Speed. Terence's stuff: A toast to posters! *IMS Bulletin*, 34 (6):6, 2005.
- [12] T. P. Speed. Comment on G. K. Robinson, "That BLUP is a good thing: The estimation of random effects". *Stat. Sci.*, 6(1):42–44, 1991.
- [13] T. P. Speed. Iterative proportional fitting. In P. Armitage and T. Colton, editors, *Encyclopedia of Biostatistics*. John Wiley & Sons, New York, 1999.
- [14] T. P. Speed. Negligible probabilities and nuclear reactor safety: Another misuse of probability? Technical report, Mathematics Department, University of Western Australia, 1977.
- [15] T. P. Speed. Restricted maximum likelihood (REML). In S. Kotz, editor, *Encyclopedia of Statistical Sciences*, volume 2, pages 472–481. John Wiley & Sons, New York, 1997.
- [16] T. P. Speed. Modelling and managing a salmon population. In V. Barnett and K. T. Turkman, editors, *Statistics for the Environment: Water Related Issues*, volume 1, pages 267–292. John Wiley & Sons, 1993.

Journal of Immunological Methods 198 (1996) 119–132

**JOURNAL OF
IMMUNOLOGICAL
METHODS**

Estimation of antigen-responsive T cell frequencies in PBMC from human subjects

Karl Broman^a, Terry Speed^{a,*}, Michael Tigges^b^a Department of Statistics, University of California, Berkeley, CA 94720, USA^b The Biocine Company, Chiron Corporation, Emeryville, CA 94608, USA

Received 5 April 1996; revised 10 June 1996; accepted 10 June 1996

Abstract

A new method for estimating the frequency of antigen-responsive T cells, using a cell proliferation assay, is described. In this assay, the uptake of tritiated thymidine by peripheral blood mononuclear cells which have been exposed to antigen, is measured for each well on a microtiter plate. Whereas this assay is generally used as part of a limiting dilution assay, here we estimate the frequency of responding cells using a single, carefully chosen cell density. The traditional analysis of such data uses a cut-off to separate wells which contain no responding cells and wells which contain at least one responding cell. The new method uses the scintillation count to estimate the number of responding cells for each well on the plate. We do this by fitting a two-stage model, the first stage being a Poisson model with antigen-specific frequency parameters, and the second stage a linear model with plate-specific parameters.

Keywords: Antigen-responsive T cell frequency; Limiting dilution assay; Maximum likelihood; EM algorithm; Plate-specific parameter

1. Introduction

In this paper we describe a method for estimating the frequency of antigen-responsive T cells among peripheral blood mononuclear cells (PBMC) from human subjects. It was developed in a particular context, but we believe that the approach may have wider applicability. The proliferation assay which motivates our method seeks to quantitate the re-

sponse of human subjects to a herpes simplex vaccine consisting of HSV type 2 glycoproteins D and B, expressed as recombinant proteins in Chinese hamster ovary cells and administered as a vaccine combined with alum (Parr et al., 1991; Straus et al., 1993) or an oil-in-water emulsion adjuvant MF59 (Langenberg et al., 1995). A standard proliferation assay (James, 1991) utilizing the responses in triplicate wells was found to be inadequate in this context, while a full limiting dilution assay (LDA) was not feasible because of the need for more PBMC than were available from the vaccine recipients. Accordingly, our analysis was developed for estimating the number of antigen-responsive T cells based on a single carefully chosen dilution, and so we sought to make greater use of the data than is usually the case.

Abbreviations: PBMC, peripheral blood mononuclear cells; LDA, limiting dilution assay; SHPM, single-hit Poisson model; PM, Poisson model; MLE, maximum likelihood estimate; SD, standard deviation; HSV, herpes simplex virus; PHS, pooled human serum; PHA, phytohemagglutinin.

* Corresponding author. Tel.: (510) 642-2781, Fax: (510) 642-7892.

Two other factors prompted our approach, which we now describe. In the present context, a standard LDA begins with a classification of the wells in all or part of a microtiter plate into the categories positive (contains at least one responding cell) and negative (contains no responding cells). Use is then made of a statistical model, typically the single-hit Poisson model (SHPM), which relates the frequency of positive wells to the frequency of responding cells in the wells. The frequency of responding cells is then estimated from the data on wells using the method of maximum likelihood or minimum χ^2 . All such analyses have to use some method of determining which wells contain responding cells, that is, of classifying wells into positives and negatives. The data with which this classification is done is generally a scintillation count, and it is usual for the assay to contain some wells which should all be negative to provide an estimate of the size of the background count. A common approach (see e.g. Langhorne and Fischer-Lindahl, 1981) is to select a threshold, often the mean plus two or three standard deviations of the background counts, and consider a well positive if its count exceeds this threshold. This approach clearly works well much of the time, as indicated by the straight lines frequently obtained when plotting the log of the frequency of negative wells against the number of cells per well. However, the background counts are usually not normally distributed about their mean, the more common situation being where there is considerable skewness, with the left-hand tail of the distribution being much shorter than its right-hand tail. Under these circumstances, reliance upon a threshold defined in terms of the mean and standard deviation of such counts can be problematic. This was the case with the data we consider in the present paper: determination of a threshold was not at all straightforward. A comparison of the counts corresponding to the wells in which no responding cells were expected with those for wells to which antigen was added revealed no clear cut-off in many cases. Efforts to develop more elaborate methods of determining the threshold for positive wells were not completely satisfactory. This led us to seek an analysis of the data which made direct use of the scintillation counts, and which did not reduce them to positive and negative well frequencies. There are good statistical reasons for avoiding the use of thresholds

in situations where they are not entirely clear. In such cases, the actual threshold used can have a very great impact, indeed be the dominant contributor to the final frequency estimates, yet the very real uncertainty in the determination of the threshold is typically not reflected in the standard errors or confidence limits given for these frequency estimates. Doing so presents formidable technical statistical problems, yet ignoring this important source of uncertainty can lead to quite unrealistic impressions concerning the precision of the frequency estimates.

There was a second, independent reason why it was desirable to avoid classifying wells as positive or negative: in many cases the entire set of wells for a given antigen would be positive. This arose whenever the density of cells chosen for the assay was a poor guess, something that could not always be avoided. In such cases the standard analysis of the data, given as a frequency of 100% of wells positive, does not yield a point estimate of the frequency of responding cells, but only a lower confidence limit. This causes difficulties later, when such results are to be compared or combined with other results. Since up to 20% of our assays would yield all positive wells, however the threshold was defined, we had a strong incentive to develop a method of estimating the frequency of responding cells another way. Similar remarks apply to the less frequent cases in which all wells would have been scored negative.

As with most studies of this kind, the approach we adopt below makes use of the Poisson model (PM) for the distribution of responding cells in a well. But instead of relating frequencies of responding cells to proportions of positive wells, we relate them to averages of suitably transformed scintillation counts. Our model involves plate-specific parameters which need to be estimated, but the evidence so far suggests that this can be done well enough to permit useful estimates of the frequencies of responding T cells to be obtained from replicate pairs of microtiter plates involving cells at a single density. Details are given in Section 2.

In order to demonstrate the validity of the new method, we analyse four LDAs on cells from three subjects, each involving three antigens and a control, carried out on replicate pairs of plates at six, five and four dilutions. We then compare the results obtained from a single dilution with those obtained using the

entire LDA. We also analyse a number of assays run at a single dilution. As a further demonstration of the usefulness of this approach, we reanalyse the data from a quite different type of LDA, see Langhorne and Fischer-Lindhahl (1981), namely a ^{51}Cr -release assay designed to estimate the frequency of cytotoxic T lymphocyte precursor cells in mixed lymphocyte cultures. Finally, we analyse two additional proliferation assays of different designs to demonstrate the flexibility of this method: one a single density assay applied to samples pre- and post-immunization (S. Rodda, Chiron Mimotopes, Melbourne, Australia) and a set of three assays designed as limiting dilution assays (D. Koelle, University of Washington, Seattle, WA).

2. Materials and methods

2.1. Vaccine, antigens and subjects

Two subjects (#711 and #713) in a clinical trial of an HSV vaccine provided informed consent for the collection of additional blood for development of the assay. These subjects had never been infected with either HSV-1 or HSV-2 prior to vaccination with a vaccine consisting of 30 μg each of two HSV glycoproteins (D and B) expressed as a recombinant product in Chinese hamster ovary (CHO) cells. The proteins were combined with an oil-in-water emulsion adjuvant (MF59, Chiron Corporation) and administered by intramuscular injection in the deltoid muscle at 0, 1 and 6 months (Langenberg et al., 1995). A third subject (NIH 1394) who was subject to frequent recurrences of genital herpes and who is HSV-2 positive was also recruited (Kost et al., 1993). The gD and gB proteins used in the assay were from the vaccine lots or comparable lots manufactured by Chiron Biocine (Emeryville, CA). Tetanus toxoid was a gift from Wyeth Laboratories (Nutley, NJ). Phytohemagglutinin was purchased from Sigma (St. Louis, MO).

2.2. Preparation of cells

Peripheral blood mononuclear cells (PBMC) were collected from the two vaccine study subjects 60–69 days after the third immunization by three pass

leukapheresis. PBMC were prepared from the third subject by one pass leukapheresis. The recovered cells were further purified by density centrifugation onto Ficoll-Hypaque (Pharmacia), washed free of separation medium and prepared for cryopreservation in RPMI 1640 medium containing 20% pooled human serum and 7.5–10% DMSO. The cells were stored in multiple aliquots in vapor-phase liquid nitrogen until assay.

2.3. Description of the assays

Six-point LDAs of PBMC from subjects #713 and #711 and a second five-point LDA of PBMC from subject #713 were set up using information obtained from a frequency analysis assay (Reece et al., 1994a; M. Tigges, unpublished observations). PBMC from #713 contain a high frequency of gD2 and gB2 specific T cells while #711 PBMC responded poorly to these two antigens. The expected frequencies were 100 and 200 responders per 10^6 PBMC to gD2 and gB2 respectively from #713 and five and two responders per 10^6 PBMC to gD2 and gB2 respectively from #711. The frequency of gB2 and gD2 responsive T cells in the NIH 1394 sample was expected to be on the order of 20–40 per 10^6 PBMC based on the results of assays with PBMC from HSV-2 seropositive subjects in other trials.

The PBMC were thawed and washed free of preservative with LGM-1 (RPMI 1640, JRH Biosciences, supplemented with 1 mM Na pyruvate, JRH Biosciences, 5 mM Hepes pH 7.2, Gibco, 2 mM glutamine, JRH Biosciences, 50 $\mu\text{g}/\text{ml}$ gentamicin, Gibco, and 1% pooled human serum). The pooled human serum (PHS) was prepared from screened units of recovered plasma. The washed cells were resuspended in lymphocyte growth medium containing 10% PHS (LGM-10), the cell concentration was adjusted, and the appropriate volume of the cell suspension was placed in the culture plates. The test antigens consisted of gD2 and gB2 at a concentration of 1 $\mu\text{g}/\text{ml}$ in 48 wells each distributed between two plates at each concentration. The control antigen consisted of tetanus toxoid at 2 Lf/ml in 44 wells and the control mitogen, PHA, was added to two wells on each plate. The PBMC from subject #711 were seeded into U-bottom plates and those from subjects #713 and NIH 1394 were seeded into V-

bottom plates. For U-bottom plates, the test and control antigens were diluted to 10 $\mu\text{g}/\text{ml}$ in LGM-10 then 20 μl were transferred to the appropriate wells. The #711 PBMC were diluted into three cell suspensions with densities of 10^6 , 5×10^5 and 2×10^5 cells/ml then either 200 or 150 μl of the suspension were transferred to two plates. For PBMC from subject #713, the washed cells were adjusted to 2×10^5 cells/ml in LGM-10 then mixed with an equal volume of LGM-10 containing the test and control antigens at 2 $\mu\text{g}/\text{ml}$. The cells were then seeded into 48 replicate V-bottom wells in volumes of 150, 100, 75, 50, 35 and 25 μl . In a second assay the #713 cells were prepared similarly. The cells were resuspended at 5×10^5 cells/ml and seeded into replicate V-bottom wells in volumes of 100, 80 or 60 μl or resuspended at 2×10^5 cells/ml and seeded in volumes of 100 or 50 μl . The assay design for the NIH 1394 PBMC differed slightly in that the antigens and cells alone were seeded into 36 wells distributed over three plates (12 wells/plate); the tetanus toxoid was included in only six wells and the PHA in three wells. After the cells were diluted for dispensing into the wells, a sample was taken to determine the actual number of cells seeded per well. All of the microwell plates were cultured for four days in humidified boxes at 37°C in 7% CO_2 before being labeled for 6 h with 0.5 $\mu\text{Ci}/\text{well}$ of [^3H]thymidine. PHA was added to the appropriate wells on day 2. The plates were harvested using a Cambridge Technologies automated harvester and the filters counted in a Wallac/Pharmacia β plate scintillation counter.

The assay design for the data from S. Rodda was similar to that described above, except that 64 replicates were plated for each antigen tested and cells were plated at 200 000 cells/well. The PBMC were obtained from a single individual before and 3 weeks after immunization with tetanus toxoid (Reece et al., 1994a). The assays included preimmune cells, post-immunization cells, and a 1:1 mixture of the two. Test antigens included influenza ribonucleoprotein and a peptide that includes an epitope from tetanus toxoid that is recognized by CD4^+ T cells from many individuals (Reece et al., 1994b). The data from D. Koelle were taken from an LDA design that consisted of 10 000 irradiated autologous PBMC/well, a graded number of between 50 000

and 780 PBMC and HSV-2 viral antigen in 24 replicates. The cells alone controls were seeded in replicates of 12 or 24 on plates separate from the antigen containing wells. The cells were cultured for 5 days before being labeled with [^3H]thymidine and harvested on day 6.

2.4. Description of the data analysis

Our analysis of the data from a single microtiter plate makes use of all but two of the counts for the 96 wells: we do not use the counts for the two wells with PHA, which simply serve as a positive control for helping select usable data. We will describe the analysis of data from a single plate first, although our final analysis involves replicate pairs of plates.

The analysis is based upon the same PM that underlies most LDAs, but we need an extra relationship connecting the scintillation count to the number of responding cells in a well. Specifically, we suppose that there are plate-specific parameters a , b and σ , and a widely applicable power parameter p such that the p th power of the scintillation count in a well with k responding cells is approximately normally distributed with mean $a + bk$ and standard deviation σ . As with other analyses using the PM, we suppose that the number of responding cells in a well with c PBMC is Poisson distributed with mean $fc \times 10^{-6}$, where f is the frequency of responding cells per million PBMC. Since we have 24 wells with cells alone, 24 with gD2, 24 with gB2 and 22 with tetanus toxoid in any given plate, there will be four frequencies and three additional parameters for each plate. When we analyse replicate pairs of plates, there will be ten parameters: the frequencies of each of the four classes of responding cells, assumed the same in each plate, and a set of three plate-specific parameters a , b and σ for each plate.

Algebraically, our assumptions are as follows. Let y_{ij} denote the transformed scintillation count for well j of class i , and let k_{ij} denote the corresponding number of responding cells. Here $i = c, d, b, t$ corresponds to the cells alone, gD2, gB2 and tetanus toxoid classes, respectively. We assume that the (y_{ij}, k_{ij}) are mutually independent, that k_{ij} follows a Poisson distribution with mean λ_i , and that, given k_{ij} , y_{ij} follows a normal distribution with mean $a + bk_{ij}$ and standard deviation σ . The aim of our

analysis is to estimate the parameters λ_i , and hence, using the estimated numbers of cells per well, the frequencies f_i of responding cells per million PBMC.

The power parameter was selected by maximum likelihood (Box and Cox, 1964) from the values 1, 1/2, 1/4 and 0 (corresponding to log). The model itself was fitted by the method of maximum likelihood, specifically, using a form of the so-called EM algorithm (Dempster et al., 1977), although we also carried out a number of confirmatory analyses using the fully calculated likelihood. Standard errors for the parameter estimates were computed using the SEM algorithm (Meng and Rubin, 1991). A more detailed description of the statistical methods can be found in Broman et al. (1996).

A by-product of the EM algorithm, which regards the 'complete' data for any well as a pair (y, k) , where y is the observed scintillation count, and k the unobserved number of responding cells, is an estimate of k for each well. These can be plotted against the count, and provide an informative diagnostic for the analysis of a plate. In particular, one can usually see the effective threshold distinguishing wells with no estimated responding cells from those with one or more.

The foregoing analysis provides estimates of the frequencies λ_c , λ_d , λ_h and λ_i of responding cells *per well* for each of the four classes: cells alone, gD2, gB2 and tetanus toxoid. We next obtained MLEs of the frequencies of responding cells *per well above background*, i.e. of $\lambda_d - \lambda_c$, $\lambda_h - \lambda_c$ and $\lambda_i - \lambda_c$, and then converted the resulting figures to frequencies per 10^6 cells. This last step involved scaling by the estimated number of cells/well in each plate. We note that as long as $\hat{\lambda}_d$, $\hat{\lambda}_h$ and $\hat{\lambda}_i$ are all $\geq \hat{\lambda}_c$, the MLEs of the differences $\lambda_d - \lambda_c$, etc.,

are just the differences $\hat{\lambda}_d - \hat{\lambda}_c$, etc., of the MLEs. In the rare cases where one or more of $\hat{\lambda}_d$, $\hat{\lambda}_h$ and $\hat{\lambda}_i$ was $\leq \hat{\lambda}_c$, a slightly modified analysis was necessary, involving combining sets of counts. The details are straightforward, see e.g. Barlow et al. (1972), and will be omitted.

In order to analyse the ^{51}Cr -release assay, which has a more standard structure, with 24 replicate wells at each of eight concentrations, we used the same basic method, modified to correspond to only a single unknown frequency of responder cells at each concentration.

The single density assay using data from S. Rodda consisted of six plates with cells from one subject taken before and after tetanus immunization: a pair of plates with 200 000 preimmune cells per well, a pair with 200 000 post-immunization cells per well, and a pair containing a mixture of 100 000 preimmune cells per well and 100 000 post-immunization cells per well. Each plate consisted of three groups of 32 wells containing cells alone, cells treated with influenza RNP antigen, and cells treated with an epitope from tetanus toxoid, respectively. For each pair of plates, we estimate the frequency of responding cells per well for the three groups of wells (denoted λ_c , λ_i and λ_t , corresponding to cells alone, tetanus, and influenza RNP, respectively), and two sets of plate-specific parameters (a, b, σ) .

The data from D. Koelle consist of three seven-point LDAs, corresponding to three different subjects. For each LDA, four plates were used. On the first plate, 24 wells were dedicated to each of four cell densities: 50 000, 25 000, 12 500 and 6250 cells per well; antigen was added to each well. On the second plate, 24 wells were dedicated to each of three cell densities: 3125, 1563 and 781 cells per

Table 1
Maximum likelihood estimates and estimated standard deviations of model parameters for the results of the first assay of subject #713 at density 11 400 cells/well

	λ_c	λ_d	λ_h	λ_i	a	b	σ
<i>Joint</i>							
Plate 1	0.4 (0.1)	3.5 (0.3)	3.3 (0.3)	4.7 (0.3)	16.4 (0.9)	10.3 (0.3)	3.6 (0.5)
Plate 2	0.4 (0.1)	3.5 (0.3)	3.3 (0.3)	4.7 (0.3)	14.8 (0.8)	9.4 (0.2)	2.9 (0.4)
<i>Separate</i>							
Plate 1	0.3 (0.1)	3.0 (0.4)	2.8 (0.4)	4.4 (0.5)	16.7 (0.9)	10.3 (0.3)	3.5 (0.4)
Plate 2	0.5 (0.1)	3.9 (0.4)	3.9 (0.4)	5.0 (0.5)	14.5 (0.7)	9.3 (0.2)	2.8 (0.3)

well; again, antigen was added to each well. The third and fourth plates in each LDA were like the first two, but with no antigen added. We analysed these data using the first and third plates together, with the parameters (a , b , σ) constrained to be equal for the two plates. The second and fourth plates were analysed similarly.

3. Results

Maximum likelihood estimates (MLEs) under the model (Finney, 1978) of the mean number of responding *cells/well* for a pair of plates from LDA #713, with 11 400 cells per well, are presented in Table 1, together with MLEs of the parameters a , b and σ and estimated standard deviations (SDs) for

each estimate. We used the square root of the scintillation count, as indicated by a Box-Cox analysis. Estimates were obtained by treating each plate *separately*, and also for the *joint* analysis of the pair of plates, where the mean numbers of responding cells/well were constrained to be equal. The estimated SDs given in parentheses after each parameter estimate take into account only *within-plate* variation. No attempt has been made to include a component of variation between duplicates within duplicate pairs, and thus the SDs for the *joint* estimates are somewhat understated. The results for different members of a duplicate pair are usually quite close, and so this underestimation is unlikely to be a problem. When the two sets of results are quite different, it is usually the case that one of the pair is simply a bad plate, and the results are discarded. In any event,

Fig. 1. Estimated number of responding cells vs. square root of scintillation count. b and d display the lower left regions (marked by dashed lines) of a and c , respectively.

the variability between duplicates within duplicate pairs is usually very much smaller than the between-assay variability, as we shall see.

After obtaining estimates of the numbers of responding cells/well we calculate the frequency of responding cells/ 10^6 cells for each of the three antigens. The results for the first assay of subject #713 at density 11 400 cells/well are displayed in Table 2. In this table, these frequencies have been corrected for background using the estimated number of cells/well responding in the cells alone category. Furthermore, the estimated SDs presented incorporate variation due to cell counting and dilution errors, though they still do not reflect the between-assay variability. It can be seen that at this stage our single dilution frequency estimates have coefficients of variation around 25%.

Table 2

Maximum likelihood estimates and estimated standard deviations of frequencies of responding cells per 10^6 cells for the results of the first assay of subject #713 at density 11 400 cells/well

	f_d	f_h	f_t
Joint	271 (67)	257 (64)	378 (92)
Separate			
Plate 1	238 (60)	219 (56)	356 (87)
Plate 2	306 (75)	299 (73)	400 (97)

A by-product of the (EM) algorithm we use for maximum likelihood estimation is an estimate, for each well, of the number of responding cells in that well. These estimates are plotted against the square root of the scintillation count and exhibited in Fig. 1

Fig. 2. Subject #713 (one six-point LDA, one five-point LDA, plus single assays). Maximum likelihood estimates of frequencies ($\times 10^6$) of responding cells using two plates at each dilution: estimates plotted against # cells/well. Error bars correspond to ± 1 SD. Dotted line corresponds to estimated frequency of responding cells ($\times 10^6$) obtained using all the data.

for two plates from LDA #713. Fig. 1b and Fig. 1d display the lower left regions (marked by dashed lines) of Fig. 1a and Fig. 1c, respectively. In Fig. 1a and Fig. 1b (#713: 11 400 cells/well, plate 1), we see that wells with counts less than 400 (20 on the square root scale) have been assigned 0 responding cells, while wells with counts greater than 625 (25 on the square root scale) have been assigned ≥ 1 responding cell. In between 400 and 625 is a grey area: there is no clear cut-off for this data. The mean and SD of the untransformed counts are 461 and 401 respectively, and so use of mean $+2$ SD as cut-off would lead to a figure of 1263. It is evident that at least two, perhaps three wells in the cells alone group seem to contain a responding cell. Their scin-

tillation counts inflate the SD of the cells alone counts, which in turn can lead to an unduly large cut-off under the traditional analysis of such data. Fig. 1c and Fig. 1d (#713: 1900 cells/well, plate 2) exhibit a much sharper cut-off.

In Figs. 2–4, estimates of the frequencies of responding cells per 10^6 cells are plotted against cell density, with error bars corresponding to ± 1 SD. Here the SD incorporates both within plate variation and errors involved in counting the number of cells/well and dilution errors. The dotted line in each plot corresponds to the estimated frequency of responding cells/ 10^6 cells obtained by carrying out a maximum likelihood analysis using all the data from both LDAs, but not the single assays. It is

Fig. 3. Subject #711 (one five-point LDA, plus single assays). Maximum likelihood estimates of frequencies ($\times 10^6$) of responding cells using two plates at each dilution: estimates plotted against # cells/well. Error bars correspond to ± 1 SD. Dotted line corresponds to estimated frequency of responding cells ($\times 10^6$) obtained using all the data.

immediately clear from Figs. 2–4 that the error bars we have calculated understate the variability exhibited by the estimates. Roughly speaking, we would expect about 68% of these ± 1 SD intervals to contain the true (but unknown) frequency, if they incorporated *all* sources of variability, and it seems evident that this is unlikely to be the case. In particular, the intervals given in the second assay for #713 at the three highest cell densities, and those for the antigens gD2 and gB2 at the four highest cell densities for #711, seem too small by perhaps a factor of 2. Despite these difficulties with the estimation of error, it is apparent that the use of these assays at a single carefully chosen density of cells will yield estimates of the frequency of responding cells with a coefficient of variation of the order of 20–25% or better. The longer error bars in the results from NIH

Table 3

Maximum likelihood estimates of CTL-precursor frequencies and plate parameters at each density for the data from Langhorne and Fischer-Lindhahl (1981)

Cell density	λ	a	b	σ
0	0.0 (0.0)	5.9 (0.1)	0.0 (0.0)	0.4 (0.1)
100	0.1 (0.1)	5.7 (0.1)	1.9 (0.2)	0.4 (0.1)
300	0.5 (0.1)	6.0 (0.1)	2.2 (0.1)	0.4 (0.1)
500	1.1 (0.2)	6.3 (0.2)	2.1 (0.1)	0.5 (0.1)
750	1.7 (0.3)	6.5 (0.1)	1.8 (0.1)	0.4 (0.1)
1000	2.5 (0.3)	6.4 (0.3)	1.7 (0.1)	0.4 (0.1)
3000	3.4 (0.4)	8.2 (0.2)	1.4 (0.1)	0.3 (0.1)
10000	3.2 (0.4)	7.6 (0.5)	1.6 (0.2)	0.4 (0.1)

1394 (Fig. 4) reflect the smaller number of wells used for each antigen. The data from this assay were particularly difficult to analyse by the traditional

Fig. 4. Subject NIH 1394 (one four-point LDA). Maximum likelihood estimates of frequencies ($\times 10^6$) of responding cells using two plates at each dilution: estimates plotted against # cells/well. Error bars correspond to ± 1 SD. Dotted line corresponds to estimated frequency of responding cells ($\times 10^6$) obtained using all the data.

K. Broman et al. / Journal of Immunological Methods 198 (1996) 119–132

Table 4

Maximum likelihood estimates of responder frequencies and plate parameters for each experimental group for the data from S. Rodda

Group	λ_c	λ_r	λ_i	a	b	σ
Post immunization	0.0 (0.0)	0.4 (0.1)	1.3 (0.2)	14.3 (0.2)	6.9 (0.2)	1.8 (0.2)
				13.9 (0.2)	4.8 (0.3)	1.5 (0.1)
1:1 mixture	0.0 (0.0)	0.2 (0.1)	2.9 (0.2)	16.0 (0.3)	7.1 (0.2)	2.1 (0.2)
				16.1 (0.2)	6.0 (0.2)	1.6 (0.1)
Preimmune	0.0 (0.0)	0.2 (0.1)	5.1 (0.3)	18.1 (0.4)	5.3 (0.2)	2.3 (0.3)
				19.1 (0.3)	8.2 (0.2)	2.5 (0.2)

cut-off method, as five out of 12 observed responses were all negative.

In order to re-analyse the data from Langhorne

and Fischer-Lindahl (1981), it was first necessary to read the counts per minute ($\times 10^{-2}$) from their Fig.

2. In Table 3 we exhibit the MLEs of the CTL-pre-

Table 5

Maximum likelihood estimates of responder frequencies and plate parameters for the three LDAs from D. Koelle: DK2, KD and EL

(a)

# cells/well	λ_c	λ_a				
50 000	1.1 (0.3)	15.1 (1.0)				
25 000	0.0 (0.0)	9.3 (0.7)				
12 500	0.3 (0.2)	5.4 (0.5)				
6 250	0.2 (0.1)	2.4 (0.4)				
3 125	1.4 (1.1)	3.5 (2.4)				
1 563	0.0 (0.0)	1.6 (1.3)	Plates	a	b	σ
781	0.2 (0.6)	1.6 (1.2)	1, 3	19.7 (0.6)	5.6 (0.2)	3.3 (0.4)
			2, 4	20.1 (0.7)	1.7 (1.0)	3.5 (0.4)

(b)

# cells/well	λ_c	λ_a				
50 000	0.2 (0.1)	16.3 (0.9)				
25 000	0.1 (0.1)	11.8 (0.7)				
12 500	0.0 (0.0)	8.4 (0.6)				
6 250	0.3 (0.1)	5.2 (0.5)				
3 125	0.0 (0.0)	7.3 (1.6)				
1 563	0.1 (0.2)	4.9 (1.1)	Plates	a	b	σ
781	0.5 (0.3)	2.5 (0.7)	1, 3	17.9 (0.4)	7.7 (0.1)	2.8 (0.2)
			2, 4	21.1 (0.7)	4.8 (0.9)	4.2 (0.4)

(c)

# cells/well	λ_c	λ_a				
50 000	0.3 (0.2)	14.5 (0.9)				
25 000	0.0 (0.0)	8.9 (0.7)				
12 500	0.1 (0.1)	5.3 (0.5)				
6 250	0.0 (0.0)	2.1 (0.3)				
3 125	0.4 (0.2)	2.3 (0.3)				
1 563	0.5 (0.2)	1.7 (0.3)	Plates	a	b	σ
781	0.8 (0.2)	0.8 (0.2)	1, 3	19.5 (0.5)	6.3 (0.2)	3.5 (0.3)
			2, 4	17.4 (0.4)	4.6 (0.3)	1.8 (0.2)

K. Broman et al. / *Journal of Immunological Methods* 198 (1996) 119–132

Fig. 5. Maximum likelihood estimates of frequencies of CTL-Ps plotted against number of responding cells.

cursor frequencies for each cell density, and also the associated estimates of plate-specific parameters. As before, the SDs reflect only within-plate variation. Fig. 5 shows these estimates plotted against the number of responders, where the estimates from the traditional analysis are also presented. Note that for these data, the untransformed scintillation counts were used, as indicated by a Box-Cox analysis.

The effectiveness of our analysis of the CTL-precursor assay data from Langhorne and Fischer-Lindahl (1981) is evident from Fig. 5. Not only do our estimates of the CTL-precursor frequencies give a slightly better linear relation than the traditional ones, based upon the first six frequencies, our estimates of b and σ were remarkably stable over this range, while the estimates of a increase with cell density. It is worth pointing out that we have analysed these

Fig. 6. Maximum likelihood estimates of frequencies of responders plotted against number of cells per well for the three LDAs from D. Koelle.

data as though the different sets of 25 counts at each precursor cell density were obtained from different microtiter plates, estimating a new a , b and σ for each set. It is not clear from the paper whether this was the case, but if not, then an analysis constraining some of the plate-specific parameters to be equal would be both more appropriate and more efficient.

Table 4 presents the results of our analysis of the data from S. Rodda. Here λ_c , λ_t and λ_i denote the estimated frequencies of responding cells per well for the cells alone, tetanus and influenza RNP groups, respectively. For these data, the square roots of the scintillation counts were used. Note that the estimated frequency for the 50:50 mixture is approximately the average of the estimated frequencies of its components, within the estimated error.

Table 5 presents the results of our analysis of the data from three LDAs from D. Koelle, while Fig. 6 gives a plot of the estimated frequencies of responding cells per well (above background) against cell density. For these data, the log scintillation counts were used. We notice that although there appears to be a reasonable linear relationship between estimated frequency of responders and number of cells per well, the estimated ratio of b to σ for two pairs of plates (DK2, #2,4; KD, #2,4) suggests difficulty fitting the model, while those for two other pairs (DK2, #1,3; EL, #1,3) are only marginally satisfactory.

4. Discussion

The main objective of our analysis was to obtain estimates of frequencies of responding cells based upon a single dilution. It is clear from the results in Figs. 2–4 that we can do this with a coefficient of variation of about 30% or lower, apart from a component of assay-to-assay variation which we discuss shortly. But before doing so, we recall the stated aim of our single dilution assay: it was intended to be a significantly more sensitive version of the standard proliferation assay, which obtains a 'stimulation index' as the ratio of the mean scintillation count from three test wells to the mean count from three wells with cells alone. Our assay was not intended to be a replacement for a standard LDA, whether analysed in the usual way by reducing the well counts to

quantal responses, or under the model presented in this paper, although it does seem that our model and analysis will provide an alternative, possibly more efficient analysis of such LDAs, under certain circumstances. We have presented analyses of a variety of full LDAs simply to enable us to assess the extent to which our single dilution estimates can be relied upon.

Between assay variation is clearly an important issue, and its extent can be gauged from Figs. 2–4. Apparently it can be quite substantial, with the results for tetanus toxoid being of particular concern. It might therefore be thought essential that the SDs assigned to frequency estimates should incorporate a component of between-assay variation, so that we get a realistic impression of the true imprecision in these estimates. We feel that this topic is best studied within the context of a larger ongoing trial, in which some samples are routinely analysed in two or more different assays. For this reason we will not discuss the matter any further here, apart from noting that on the basis of the evidence presented here, it might not be unreasonable to double the estimated SDs of frequency estimates if between-assay variation is to be incorporated.

The method of analysis we have presented in this paper highlights certain issues relating to the design of assays of the kind we discuss. The most important concerns the relation of the negative control wells, our cells alone, to wells containing antigen. We have examined and analysed a number of assays in which the negative control wells and the wells containing antigen were located on different microtiter plates. Although this practice is not necessarily injudicious, there are at least three reasons why it should be avoided, particularly if our method of analysis is to be used, but even more generally. As is made explicit by our estimation of the plate-specific parameters a , b and σ , the responses of cells from the same source can differ from plate to plate, even in well-conducted assays. It is good general practice to carry out comparative analyses with the greatest possible degree of control over the conditions which could cause differences, in this case, between negative control and antigen wells, for the responses in wells with antigen will be adjusted to the extent that the negative control wells respond. If it is not possible to make such intra-plate adjustments, it becomes

necessary to assume that the plate containing the negative control wells has the same parameters as that containing the wells with antigen, for otherwise the two classes of wells will not be treated similarly in their analysis.

A few comments on the interpretation of our parameters and the range of reasonable values for them are in order. The parameter a is most easily interpreted as the average or median value of the counts or transformed counts observed in the wells in the cells alone category. Of course it will also have the same relation to the counts from those wells with antigen in which there are no responding cells, but we will not usually know definitely which these are, whereas we can generally be confident that the overwhelming majority of counts from wells in the cells alone category are simply background. The parameter b is the slope of the line relating average count or transformed count to frequency of responding cells within a plate. Finally, the parameter σ corresponds to the spread of counts or transformed counts about their mean, for wells with the same number of responding cells. Under our normality assumption, about 95% of such counts would be within 2σ of the mean, which has the form $a + kb$. Thus satisfactory discrimination between adjacent values of k is only possible when σ is $b/2$ or smaller, and is really good if $\sigma \leq b/4$. Our algorithm can converge and give reasonable frequency estimates with values of σ as large as b , or even larger than b . However, we are inclined to regard such situations as failures of the model to fit the data, and include them only when there is no other way to get a frequency estimate, and a rough one is desired.

Finally, we remark that our model and analysis, originally designed for the particular assay described in Section 2, does appear to have a wider usefulness. With only relatively minor adaptations, it could be applied to the CTL-precursor assay of Langhorne and Fischer-Lindahl (1981), as well as to data from LDAs (D. Koelle), and the single density assay of S. Rodda, all of which had a structure quite different from that of the original assays. In each case analysis gave satisfactory results, and in a way which avoided the arbitrary choice of cut-offs. Assuming that the model we used is appropriate, our analysis will also be more efficient.

5. Note added in proof

A computer program that incorporates the method described here has been written for the Windows and Macintosh platforms.

Acknowledgements

Many thanks are due to David Freedman and Peter Hall for their helpful comments on the material in this paper. We also wish to thank Mario Geysen and Stuart Rodda for the initial application of the single cell density and for many helpful discussions in the early development of the frequency analysis.

References

- Barlow, R.E., Bartholomew, D.J., Bremner, J.M. and Brunk, H.D. (1972) *Statistical Inference under Order Restrictions: The Theory and Application of Isotonic Regression*. Wiley, New York.
- Box, G.E.P. and Cox D.R. (1964) The analysis of transformations. *J.R. Stat. Soc. Ser. B* 26, 211.
- Broman, K.W., Speed, T.P. and Tigges, M. (1996) Estimation of antigen-responsive T-cell frequencies in PBMC from human subjects. Technical Report #454. Department of Statistics, University of California, Berkeley, CA.
- Dempster, A., Laird, N. and Rubin, D. (1977) Maximum likelihood estimation from incomplete data via the EM algorithm. *J.R. Stat. Soc. Ser. B* 39, 1.
- Finney, D.J. (1978) *Statistical Method in Biological Assay*. Academic Press, New York.
- James, S.P. (1991) Measurement of basic immunologic characteristics of human mononuclear cells. In: J.E. Coligan, A.M. Kruisbeek, D.H. Margulies, E.M. Shevach and W. Stroger (Eds.), *Current Protocols in Immunology*. Green Publishing Company and Wiley-Interscience, New York, section 7.10.
- Kost, R.G., Hill, E.L., Tigges, M. and Straus, S.E. (1993) Brief report: Recurrent acyclovir-resistant genital herpes in an immunocompetent patient. *New Engl. J. Med.* 329, 1777.
- Langenberg, A.G.M., Burke, R.L., Adair, S.F., Sekulovich, R., Tigges, M., Dekker, C.L. and Corey, L. (1995) A recombinant glycoprotein vaccine for herpes simplex type 2: Safety and efficacy. *Ann. Intern. Med.* 122, 889.
- Langhorne, J. and Fischer-Lindahl, K. (1981) Limiting dilution analysis of precursors of cytotoxic T-lymphocytes. In: I. Lefkovits and B. Pernis (Eds.), *Immunological Methods*, Vol. II. Academic Press, New York, section 12.
- Meng, X.-L. and Rubin, D.B. (1991) Using EM to obtain asymptotic variance-covariance matrices: the SEM algorithm. *J. Am. Stat. Assoc.* 86, 899.

K. Broman et al. / Journal of Immunological Methods 198 (1996) 119–132

- Parr, D., Savarese, B., Burke, R.L., Margolis, D., Meier, J., Markoff, L., Ashley, R., Corey, L., Adair, S., Dekker, C. and Straus, S.E. (1991) Ability of a recombinant herpes simplex virus type 2 glycoprotein D vaccine to induce antibody titers comparable to those following genital herpes. *Clin. Res.* 39, 216A.
- Reece, J.C., McGregor, D.L., Geysen, H.M. and Rodda, S.J. (1994a) Scanning for T helper epitopes with human PBMC using pools of short synthetic peptides. *J. Immunol. Methods* 172, 241.
- Reece, J.C., Geysen, H.M. and Rodda, S.J. (1994b) Mapping the major human T helper epitopes of tetanus toxoid. The emerging picture. *J. Immunol.* 151, 6175.
- Straus, S.E., Savarese, B., Tigges, M., Freifeld, A.G., Krause, P.R., Margolis, D.M., Meier, J.L., Paar, D.P., Adair, S.F., Dina, D., Dekker, C. and Burke, R.L. (1993) Induction and enhancement of immune responses to herpes simplex virus type 2 in humans by use of a recombinant glycoprotein D vaccine. *J. Infect. Dis.* 167, 1045.

G. K. ROBINSON

Comment

Terry Speed

Geoff Robinson is to be congratulated for writing this paper. It is lucidly written, it bridges a number of gulfs that have developed in our subject, and it is provocative. That he wrote it is clearly a Good Thing! I welcome the opportunity to say this and to make a few remarks that he might have made. I believe that these remarks will strengthen his already strong case for a much more explicit recognition of the role of BLUPs in our subject.

1. THE BAYESIAN DERIVATION

In Section 4.2 Robinson describes a Bayesian derivation, stating that the posterior mode is given by the BLUP estimates when β is regarded "as a parameter with a uniform, improper prior distribution and u as a parameter which has a prior distribution which has mean zero and variance $G\sigma^2$, independent of β ." All this is certainly true, but it may be helpful to add that if β is given a proper prior (normal) distribution with mean zero and variance $B\sigma^2$, say, with u as before, then all of the results one could possibly want (posterior means, posterior variances, etc.) can be derived straightforwardly by the standard Bayesian formulae. Then all one has to do to derive the corresponding BLUP formulae is let $B^{-1} \rightarrow 0$. An identity which I have found useful, perhaps even indispensable, for carrying out this last step, is discussed in de Hoog, Speed and Williams (1990). Note that the approach just described is essentially that adopted in Dempster, Rubin and Tsutakawa (1981).

2. FORMULAE FOR $\hat{u}$

The only actual formulae given in the paper for $\hat{u}$ in the general case is the rather complicated one in Section 4.3. This is a pity, because there is an obvious "plug-in" expression, namely

$$(1) \quad \hat{u} = GZ^T V^{-1}(y - X\hat{\beta}),$$

where $V = ZGZ^T + R$. This may be viewed as the result of regressing u on y , with the mean $X\hat{\beta}$ of y replaced by its obvious linear estimator.

A variant of (1) is

$$(1') \quad \hat{u} = (Z^T R^{-1} Z + G^{-1})^{-1} Z^T R^{-1} (y - X\hat{\beta}).$$

Terry Speed is Professor, Department of Statistics, University of California, Berkeley, California 94720.

The simpler formulae (5.3) and (5.4) arising when there are no fixed effects also have more general analogues, namely

$$(2) \quad (Z^T A Z + G^{-1}) \hat{u} = Z^T A y,$$

where $A = R^{-1}(I - S)$, $S = P_{\mathcal{R}(X)}^R$ being the projector onto $\mathcal{R}(X)$ orthogonal with respect to $(a, b) = a^T R^{-1} b$, and for the variance-covariance matrix of $\hat{u}$:

$$(3) \quad \{G^{-1} - (Z^T A Z + G^{-1})^{-1}\} \sigma^2.$$

These expressions can be derived readily using the Bayesian approach outlined in (1) above, together with the matrix identity already referred to. I note in passing that Robinson's formulae (5.4) is in fact the variance-covariance matrix of $\hat{u} - u$, not, as stated, of $\hat{u}$.

3. SOLVING THE BLUP EQUATIONS

Perhaps in order to avoid messy algebra, Robinson has said little about the actual solution of the BLUP equations. I know that he has worked on this problem with some enormous data sets, and so I am hesitant to comment here. However, it does seem worthwhile to make one easy point, in order to connect this topic with another, closely related one. The obvious rearrangement of the first equation in (1.2),

$$(4) \quad X^T R^{-1} X \hat{\beta} = X^T R^{-1} (y - Z\hat{u}),$$

can be combined with either (1) or (1') above, to form the basis of an iterative solution of the BLUP equations, provided, of course, that the separate problems are readily solved. Just such a strategy is recommended more generally in Green (1985) in the context of smoothing, a topic to which I shall return.

It is also worth pointing out that (1') or (2) is to be preferred when G^{-1} has simple structure, whereas if G is simple and V is readily inverted, (1) is more useful. In many animal breeding problems it is G^{-1} which has the simpler structure, as it also does in the Kalman filter case.

4. REML AND BLUP

In Section 5.4 Robinson states that "REML is the method of estimating variance components that seems to have the best credentials from a Classical

THE ESTIMATION OF RANDOM EFFECTS

viewpoint." What he does not say, which should be of interest to readers of his paper, is that REML and BLUP are intimately connected. Indeed one view—certainly not the only one—of the REML equations for variance components is that they are simply equating *observed* with *expected* sums of squares of BLUPs. This observation goes back to the original paper by Patterson and Thompson (1971; see also Harville, 1977) and can be concisely stated within Robinson's framework as follows.

Suppose that $Z = [Z_1 : \cdots : Z_c]$ is blocked, corresponding to c random effects, with Z_i being $n \times q_i$, $i = 1, \dots, c$, $u = (u_1 : \cdots : u_c)$ is similarly blocked into q sets of random effects, and finally $G = \text{diag}(G_1, \dots, G_c)$ is diagonally blocked with $G_i = \gamma_i I_{q_i}$, where $\sigma_i^2 = \gamma_i \sigma^2$ is the variance of each independent component of the i th random effect u_i . It is also convenient to denote e by u_0 , put $Z_0 = I_n$ and $\gamma_0 = 1$.

With this notation the REML equations take the form

$$(5) \quad y^T \left\{ V^{-1} \bar{Q} \frac{\partial V}{\partial \sigma_i^2} V^{-1} \bar{Q} \right\} y = \text{tr} \left\{ \frac{\partial V}{\partial \sigma_i^2} V^{-1} \bar{Q} \right\}$$

$i = 0, \dots, c$, where $\bar{Q}y = X\hat{\beta}$ and $\bar{Q} = I - Q$. (By contrast, the ML equations have no $\bar{Q}$ term in the right-hand expression.)

Turning now to BLUPs in this context, they are (in the form (1) above)

$$(6) \quad \hat{u}_i = \gamma_i Z_i^T V^{-1} \bar{Q} y$$

$i = 0, \dots, c$, and

$$(7a) \quad \text{var}(\hat{u}_i) = (G_i - U_i) \sigma^2$$

$i = 1, \dots, c$, where U_i is the i th diagonal block of the matrix $(Z^T A Z + G^{-1})^{-1}$. Furthermore,

$$(7b) \quad \text{var}(\hat{e}) = V^{-1} \bar{Q} \sigma^2 = (A - A Z U Z^T A) \sigma^2$$

where $A = R^{-1} \bar{S}$ was defined earlier, and $U = (Z^T A Z + G^{-1})^{-1}$. If we write $p_i = \gamma_i^{-1} \text{tr}(U_i)$, $i = 1, \dots, c$, then it follows that for $i = 1, \dots, c$

$$(8a) \quad \mathbb{E} |\hat{u}_i|^2 = (q_i - p_i) \sigma_i^2$$

and

$$(8b) \quad \mathbb{E} |\hat{e}|^2 = \left[(n - p) - \sum_{i=1}^c (q_i - p_i) \right] \sigma^2.$$

Now the striking thing is this: the REML equations (5) can rather easily be manipulated into a form just like (8a) and (8b), with the expectation symbol $\mathbb{E}$ omitted. Although this is not necessarily the best way to solve these equations, the repeated calculation of BLUPs and then updating the vari-

ance components is one simple iterative scheme which works quite well.

5. PENALIZED LEAST SQUARES

Suppose that we regard (1.1) as an ordinary ("fixed effects") linear model, and that we wished to estimate β and u by R -weighted least squares with a "penalty" $u^T G^{-1} u$ being added to the sum of squares term being minimized. Then we would obtain just the expression given in Section 4.1, which Henderson minimized. Such penalties are added for many reasons: to smooth, to improve the condition of the matrix to be inverted, and so on, and it has long been recognized that this is a way of making one's linear model "quasi-Bayesian." More precisely, it turns the standard least squares problem into a case of BLUP. This practice has a long history, dating back at least to Whittaker (1923).

6. SMOOTHING SPLINES ARE BLUPS

Continuing with the theme of the previous remark, let us see how the smoothing splines popularized by G. Wahba (see her 1990 monograph for a comprehensive exposition) are just BLUPs. This observation corrects the terminology which has been used in the spline literature for over a decade, for the Bayesian interpretation of the smoothing spline—with a partially improper prior—is just the statement heading this section.

It is simplest to deal with cubic smoothing splines on the interval $[0, 1]$. If the observations are taken at $0 \leq t_1 < \cdots < t_n \leq 1$, and are

$$y_i = g(t_i) + \epsilon_i,$$

$i = 1, \dots, n$, where g is an unknown smooth function, then the function g_λ which minimizes

$$n^{-1} \sum_{i=1}^n (y_i - g(t_i))^2 + \lambda \int_0^1 \{g''(u)\}^2 du$$

over a suitable class $G_0 \oplus G_1$ of functions, has the values

$$\hat{y} = (g_\lambda(t_i))$$

$$(9) \quad = X(X^T V^{-1} X)^{-1} X^T V^{-1} y + Q_n V^{-1} (I - X(X^T V^{-1} X)^{-1} X^T V^{-1}) y,$$

where

$$X = (t_i^{k-1}), \quad i = 1, \dots, n; k = 1, 2;$$

$$Q_n(i, j) = Q(t_i, t_j), \quad 1 \leq i, j \leq n;$$

and

$$Q(s, t) = \int_0^1 (s - w)_+ (t - w)_+ dw, \quad 0 \leq s, t \leq 1.$$

G. K. ROBINSON

It is easy to check that (9) is just the fitted value

$$\hat{y} = X\hat{\beta} + Z\hat{u},$$

where, in Robinson's notation, $\hat{\beta}$ and $\hat{u}$ are the BLUPs, X is as given above, $Z = I_n = R$, and $G = (n\lambda)^{-1}Q_n$.

Certainly there is more to smoothing splines than BLUPs; for example, estimates of the value of the function g at values of t other than those observed, but in many applications (9) and related expressions are all that is required.

By now it should come as no surprise to hear that the technique termed generalized maximum likelihood (GML) for estimating the smoothing parameter λ is no other than REML in this BLUP problem. This is readily checked by comparing formulae in this paper with ones in Wahba (1990).

With only very few changes, the identification just made to show that smoothing splines are BLUPs shows that the model robust response surface designs of Steinberg (1985) are also BLUPs. In this case the u term corresponds to sums of tensor products of orthogonal polynomials.

7. LINEAR SMOOTHERS ARE ALMOST BLUPS

There is a sense in which all linear smoothers (see Buja, Hastie and Tibshirani, 1989) are intimately related to BLUPs. A typical linear smoother S satisfies $S^n \rightarrow T$ as $n \rightarrow \infty$, where T is idempotent. This corresponds to a projector onto the subspace $\mathcal{A}(X)$ in Robinson's model (1.1), and so Ty corresponds to $X\hat{\beta}$. Thus $(S - T)y$ corresponds to $Z\hat{u}$, and in some situations it is even possible to construct a covariance matrix V such that this correspondence is precise. Furthermore, many smoothers S have form $S(\lambda)$, where λ is a parameter (bandwidth, variance ratio, smoothness penalty, etc.) that defines a family of similar smoothers. In such cases $S(\infty)$ often has the form $T + W$, where

W is another projector, while $S(0) = T$. Many of the problems and the formulae in the theory of linear smoothers are analogues of ones arising in the theory of BLUPs.

8. INTERVAL ESTIMATES INVOLVING BLUPS

In Section 5.6 Robinson briefly alludes to work done on estimating the precision of BLUP estimates when uncertainty in the dispersion parameter is taken into account. This general problem, and in particular the assignment of interval estimates, has attracted a lot of attention in the literature on smoothing splines (see, e.g., Nychka, 1988, for a recent review). Much concern has been given to the question of what, if any, coverage properties can be expected of a "Bayesian" posterior interval. Making interval statements about an object which is an estimate of the sum of fixed and random effects is bound to cause problems of interpretation to many people, and I would be interested to hear Geoff Robinson's comments on this point. I know that he has studied these matters closely in the past.

9. SUMMARY

In closing these few remarks, I cannot resist paraphrasing I. J. Good's memorable aphorism: "To a Bayesian, all things are Bayesian." How does "To a non-Bayesian, all things are BLUPs" sound as a summary of this fine paper?

ACKNOWLEDGMENT

The author gratefully acknowledges partial support from NSF Grant DMS-88-02378, and the hospitality of the Department of Statistics, IAS, Australian National University whilst preparing this comment.

Iterative Proportional Fitting

Iterative proportional fitting (IPF), also known as iterative proportional scaling, is an **algorithm** for constructing tables of numbers satisfying certain constraints. In its simplest form, the algorithm enables one to construct two-way **contingency tables** with specified marginal totals and a prescribed degree of association; from a more general perspective, it may be viewed as a cyclic ascent algorithm which maximizes a specific objective function. The algorithm can also be used to construct **maximum likelihood** estimators for table entries based upon hierarchical **log-linear models** for **Poisson**, **multinomial**, or product multinomial models. We will illustrate these aspects of the algorithm and its applications by describing some simple cases.

Suppose that we are given two pairs $\mathbf{u} = (u_1, u_2)$ and $\mathbf{v} = (v_1, v_2)$ of positive numbers satisfying $u_1 + u_2 = v_1 + v_2$, and a further positive number ψ . The IPF algorithm will enable us to construct the *unique two-by-two table* $\mathbf{b} = (b_{ij})$ such that, for all i and j ,

$$b_{i+} = u_i, \quad b_{+j} = v_j, \quad \frac{b_{11}b_{22}}{b_{12}b_{21}} = \psi,$$

where the subscript $+$ denotes the result of summing over the subscript it replaces. The algorithm goes like this. Begin with the 2×2 table $\mathbf{a} = (a_{ij})$ defined by $a_{11} = \psi$, $a_{12} = a_{21} = a_{22} = 1$, noting that the cross-ratio $a_{11}a_{22}/a_{12}a_{21} = \psi$ (see **Odds Ratio**). Next, scale the rows of $\mathbf{a}$ to form the table $\mathbf{a}' = (a'_{ij})$:

$$a'_{ij} = a_{ij} \times \frac{u_i}{a_{i+}}, \quad (1)$$

for $i = 1, 2$ and $j = 1, 2$. It is easy to check that $\mathbf{a}'$ has the desired row sums, as well as having cross-ratio ψ . We now scale the columns of $\mathbf{a}'$ to form the table $\mathbf{a}'' = (a''_{ij})$:

$$a''_{ij} = a'_{ij} \times \frac{v_j}{a'_{+j}}. \quad (2)$$

One can check that $\mathbf{a}''$ has the desired column sums and cross-ratio, although the row sums are no longer (u_i) . This completes one cycle of the IPF algorithm, beginning with the table $\mathbf{a}$.

The algorithm continues by repeatedly scaling the rows, as in (1), and then the columns, as in (2), to have the desired totals. After a number of cycles, the row totals are closer to (u_i) than they were initially, the column totals are exactly (v_j) , and the cross-ratio is exactly ψ . The sequence of tables so defined converges pointwise to a 2×2 table $\mathbf{b}$ with all the desired properties; uniqueness also follows.

It is instructive to examine why these assertions are true, for in doing so we obtain further insights into the IPF algorithm. To do this, we introduce the notion of **information** (or I -) divergence between two tables $\mathbf{c} = (c_{ij})$ and $\mathbf{d} = (d_{ij})$, satisfying $c_{++} = d_{++}$, defined as follows:

$$I(\mathbf{c}|\mathbf{d}) = \sum_{ij} c_{ij} \log \left(\frac{c_{ij}}{d_{ij}} \right).$$

(A similar definition applies to singly indexed arrays.) It can be proved that $I(\mathbf{c}|\mathbf{d}) \geq 0$, and that $I(\mathbf{c}|\mathbf{d}) = 0$ if and only if $\mathbf{c} = \mathbf{d}$. Although not a symmetric function of its arguments, I behaves in many ways like a metric on tables, and it provides the basis of a proof of convergence of the IPF algorithm. We return to our construction of a table $\mathbf{b}$ having row totals $\mathbf{u}$, column totals $\mathbf{v}$, and cross-ratio ψ . First define the table $\mathbf{c} = (c_{ij})$ as follows:

$$c_{ij} = \frac{u_i v_j}{w},$$

where $w = u_+ = v_+$. The tables $\mathbf{a}$, $\mathbf{a}'$, $\mathbf{a}''$, ... become closer to $\mathbf{c}$ as the iterations continue, closeness here being in the sense of I -divergence. More precisely, we can check that

$$I(\mathbf{c}|\mathbf{a}) = I(\mathbf{c}|\mathbf{a}'') + I(\mathbf{v}|\mathbf{a}'_2) + I(\mathbf{u}|\mathbf{a}_1), \quad (3)$$

where $\mathbf{a}_1 = (a_{i+})$ and $\mathbf{a}'_2 = (a'_{+j})$. The convergence and uniqueness assertions above all follow from repeated use of this expansion and the stated properties of I . As long as there exists at least one table $\mathbf{c}$ with the desired marginal totals, we can begin the IPF algorithm with any table having the desired cross-ratio, and expect to converge to the stated limit. The repeated scaling gives tables closer and closer in the sense of I -divergence to the table $\mathbf{c}$, all the while retaining the original cross-ratio, and the row and column totals converge to their desired values.

All of the discussion so far applies with minimal changes to $r \times s$ tables; in the more general case, there are further cross-ratios to take into account.

Iterative Proportional Fitting

Whereas in a 2×2 table there is only one cross-ratio whose value can be fixed, in an $r \times s$ table, there are $(r-1)(s-1)$ multiplicatively independent cross-ratios. A convenient set (cf. [12]) is the following:

$$\psi_{ij} = \frac{b_{ij}b_{rs}}{b_{is}b_{rj}}, \quad i = 1, \dots, r-1; j = 1, \dots, s-1.$$

Here we constructed our cross-ratios in relation to the index values r and s . Other choices give equivalent results; indeed there are quite different ways of defining the quantities which are preserved. This issue is addressed in the theory of **loglinear models**; see [2, 11] and [12]. Given an arbitrary set of $(r-1)(s-1)$ positive numbers (ψ_{ij}) , and positive numbers $\mathbf{u} = (u_i)$ and $\mathbf{v} = (v_j)$ satisfying $u_+ = v_+$, the IPF algorithm may be initiated with the table $\mathbf{a} = (a_{ij})$ given by $a_{ij} = \psi_{ij}$, $i = 1, \dots, r-1$; $j = 1, \dots, s-1$, and $a_{rj} = 1 = a_{is}$, $i = 1, \dots, r$; $j = 1, \dots, s$. With this initial table, the steps are just as before, and the resulting sequence of tables converges to the unique table having row totals (u_i) , column totals (v_j) , and cross-ratios (ψ_{ij}) .

We turn now to reasons for constructing such tables. One is simply to demonstrate the fact that the row totals, column totals, and cross-ratios of two-way tables may be specified independently, and to show how to obtain tables with arbitrarily specified (but consistent) values of these quantities. Historically, the algorithm was first used to adjust sample frequencies to expected marginal totals. In the examples in Deming [5], we have a table $\mathbf{n} = (n_{ij})$ based upon a **sample survey**, and marginal totals (N_{i+}) and (N_{-j}) , but *not* the individual cell frequencies $\mathbf{N} = (N_{ij})$, from a census of the population. The result of applying the IPF algorithm with initial table $\mathbf{n}$, and desired marginal totals (N_{i+}) and (N_{-j}) , can then be regarded as an estimate of what would have been obtained by cross-tabulating the entire population, instead of only a sample thereof. A modern treatment of these ideas can be found in [2], where the procedure is known as *raking* the table $\mathbf{n}$. The third application of the algorithm we note is to the construction of maximum likelihood estimates of table entries under loglinear models. We simply describe the results here; the reader may consult standard references such as [2, 11], or [1] for fuller details. Suppose that $\mathbf{n} = (n_{ij})$ is a two-way table of independent Poisson counts with parameters $\lambda = (\lambda_{ij})$. Then the maximum likelihood estimate $\hat{\lambda}$

of λ under the *multiplicative model* for the (λ_{ij}) , has the same row and column totals as $\mathbf{n}$, and all cross-ratios equal to 1. In this case, the IPF algorithm begins with a table all of whose entries are 1, and scales the row and column totals to match those of the data $\mathbf{n}$. The algorithm converges after a single cycle to the unique maximum likelihood estimator $\hat{\lambda}$.

Three- and Higher-Way Tables

There are a number of ways in which the IPF algorithm may be used with three-way tables. We illustrate two of these. Suppose that we have an $r \times s$ table $\mathbf{u} = (u_{ij})$ and an $s \times t$ table $\mathbf{v} = (v_{jk})$ of positive numbers satisfying $u_{+j} = v_{j+}$ for $j = 1, \dots, s$. By analogy with our earlier construction, we might be interested in obtaining an $r \times s \times t$ table $\mathbf{b} = (b_{ijk})$ having

$$b_{ij+} = u_{ij}, \quad b_{+jk} = v_{jk}.$$

This can be solved rather straightforwardly. For example, the table $\mathbf{c} = (c_{ijk})$ given by

$$c_{ijk} = \frac{u_{ij}v_{jk}}{w_j},$$

where $w_j = u_{+j} = v_{j+}$, $j = 1, \dots, s$, is readily checked to have ij -margin $\mathbf{u}$ and jk -margin $\mathbf{v}$.

Of course, this is not the end of the story. We may also be interested in any further structure concerning the table $\mathbf{b}$ which may be specified, in addition to these marginal totals. It turns out that we may also ask that the table has predetermined values of certain cross-ratios. In this example, and more generally, we need rules to tell us which marginal totals and which cross-ratios can be specified independently. The issue is best discussed in the language of **hierarchical loglinear models** for multiway tables, where these are commonly described in terms of the marginal subtables which constitute the **sufficient statistics** for the models (under either independent Poisson, multinomial, or independent multinomial sampling). We refer to [1, 2], and [11] for details concerning these models. In this language, the cross-ratios that we have been specifying are the antilogarithms of elements of subspaces **orthogonal** to those that define the hierarchical loglinear model corresponding to the specified marginal totals. For example, by specifying margins corresponding to the indices

Iterative Proportional Fitting

ij and jk , as we did in our example, we are also able to specify independently cross-ratios corresponding to the pair ik and the triple ijk – that is, all interactions other than those involved in the log-linear model defined by the prescribed marginal totals.

Now let us suppose that, in addition to $\mathbf{u}$ and $\mathbf{v}$ as above, we are given a $t \times r$ table $\mathbf{w} = (w_{ki})$ of positive numbers satisfying $w_{k+} = v_{+k}$ and $w_{+i} = u_{i+}$ for all k and i . Can we use IPF to construct a table $\mathbf{b} = (b_{ijk})$ satisfying

$$b_{ij+} = u_{ij}, \quad b_{+jk} = v_{jk}, \quad b_{i+k} = w_{ki},$$

and having prescribed values for the ijk cross-ratios? One might think that this would be quite straightforward. Begin with a suitable initial table $\mathbf{a}$. Then scale to achieve the ij , jk , and ki marginal totals $\mathbf{u}$, $\mathbf{v}$, and $\mathbf{w}$, respectively. One cycle of the algorithm would be three such scalings, and after a few cycles, we might expect to have a table with the specified cross-ratios, and essentially the desired marginal totals.

How can this version of IPF go wrong? A clue is provided by our indication of the method used to prove that IPF converges. We made use of the existence of a table $\mathbf{c}$ satisfying the marginal constraints, and then everything followed. However, in the case of three-way tables, it is not hard to specify three consistent, positive two-way tables, for which *no* three-way table exists having positive entries, and the three specified tables as two-way marginal totals. A simple example is given by three 2×2 tables each having 1 in the diagonal cells and 2 in the off-diagonal cells. Although they are clearly consistent, it is easy to check that no $2 \times 2 \times 2$ table can exist with positive entries and these margins. Use of the IPF algorithm with an initial table whose entries are all 1, and these three marginal tables, results in a cycle through the same three tables. The tables constructed do not converge. Summarizing this discussion, we can say that only if there exists a three-way table with the given two-way tables as marginal totals is the IPF algorithm guaranteed to converge to a limiting table with the desired marginal totals and three-way cross-ratios. When it does, this table is uniquely specified by these properties.

We note that in the application of this result to maximum likelihood estimation with loglinear models, the assumption of the existence of *some* table with the given marginal totals is trivially satisfied

as long as the observed table $\mathbf{n} = (n_{ijk})$ has positive entries, for in this case $\mathbf{n}$ itself suffices. If the observed table has some zero entries, but positive two-way marginal totals, the IPF algorithm still converges, but to a table with some zero entries. In a sense, this is an extended maximum likelihood estimator: one on the boundary of the natural parameter space.

The foregoing discussion applies without change to higher-way tables. For example, suppose that we have an initial four-way table $\mathbf{a} = (a_{ijkl})$, and we wish to scale it to have prescribed ij , jk , kl , and li marginal totals. What cross-ratios (equivalently, what loglinear structure) of this initial table will be preserved throughout the iterations, and could therefore be specified independently of the marginal totals? The answer is: all interactions other than those involved in the loglinear model defined by the prescribed marginal totals, that is, the ik , jl , ijk , ijl , ikl , jkl , and $ijkl$ interactions. Note that we still need to know that there exists a table with the specified marginal totals before the algorithm is guaranteed to converge to a limiting table with all the desired properties.

Finite Termination: Decomposable Models

Decomposable models are a class of loglinear models for complete multiway tables which possess closed-form expressions for their MLEs under the standard sampling models; see [11] and [2]. It turns out that the IPF algorithm behaves rather well for this class of models. Suppose that a set of marginal totals to be fitted via IPF defines a decomposable loglinear model. If the initial table is constant, and the margins to be fitted are taken in a suitable order, the algorithm converges after just one cycle. Furthermore, there *always* exists a table with the given set of tables as marginal subtables, when the corresponding model is decomposable. Finally, as long as the specified tables are all positive, the table whose existence has just been described has positive entries.

History

Fienberg [7] presents a discussion of the history of the IPF algorithm. Some additional references can be found in [8]. The most important early papers are [6] and [14].

Iterative Proportional Fitting

Numerical Aspects

Haberman [11] proves that tables constructed by the IPF algorithm converge to their limit at a geometric (also called first-order) rate. This means that, asymptotically, the difference between the n th iterate and the limit is bounded above by ρ^n for some ρ between zero and unity. (This compares unfavorably with the behavior of Newton or modified Newton algorithms, which typically exhibit what is known as quadratic convergence.) In many cases, ρ may be quite close to unity, and so convergence may be rather slow, giving rise to a literature concerning speeding up of the algorithm. However, at that point, the algorithm ceases to be the one we are discussing.

The great advantage of the IPF algorithm is its simplicity, stability, and economy of space. When a table is large, and the number of iterations is not a limiting factor, it is the method of choice for the problems we have discussed. For other problems, such as the calculation of MLEs under loglinear models, Newton-type methods are preferred, because of their speed of convergence and the fact that variance–covariance matrices are an automatic byproduct. FORTRAN IV versions of the IPF algorithm can be found in [9] and [10].

Variants and Generalizations

It is implicit in the foregoing discussion that the tables being considered are all *complete*, that is, are fully rectangular, or rectangular parallelepipeds, etc., and have no so-called **structural zeros**. This was because the algorithm is mostly used, and its properties are most easily discussed, in that context. However, variant forms of the algorithm are used successfully with tables having a variety of other structures, and preserving features corresponding to models other than hierarchical loglinear models; see [11].

For generalizations of a different kind, see [3, 4, 13]. In these papers, applications of the algorithm

beyond contingency tables are given, and its connections to the information measure I and entropy are more fully explored.

References

- [1] Agresti, A. (1990). *Categorical Data Analysis*. Wiley, New York.
- [2] Bishop, Y.M.M., Fienberg, S.E. & Holland, P.W. (1975). *Discrete Multivariate Analysis*. MIT Press, Cambridge, Mass.
- [3] Csiszar, I. (1975). I -divergence geometry of probability distributions and minimization problems, *Annals of Probability* **3**, 146–158.
- [4] Darroch, J.N. & Ratcliff, D. (1972). Generalized iterative scaling for loglinear models, *Annals of Mathematical Statistics* **43**, 1470–1480.
- [5] Deming, W.E. (1964). *Statistical Adjustment of Data*. Dover, New York.
- [6] Deming, W.E. & Stephan, F.F. (1940). On a least squares adjustment of a sampled frequency table when the expected marginal totals are known, *Annals of Mathematical Statistics* **11**, 427–444.
- [7] Fienberg, S.E. (1970). An iterative procedure for estimation in contingency tables, *Annals of Mathematical Statistics* **41**, 907–917.
- [8] Fienberg, S.E. & Meyer, M.M. (1983). *Encyclopedia of Statistical Sciences*, Vol. 4, S. Kotz & N.L. Johnson, eds. Wiley, New York, p. 2275.
- [9] Haberman, S.J. (1972). Loglinear fit for contingency tables, *Applied Statistics* **21**, 218–225.
- [10] Haberman, S.J. (1973). Printing multidimensional tables, *Applied Statistics* **22**, 118–126.
- [11] Haberman, S.J. (1974). *The Analysis of Frequency Data*. University of Chicago Press, Chicago.
- [12] Plackett, R.L. (1981). *The Analysis of Categorical Data*, 2nd Ed. Griffin, London.
- [13] Ruschendorf, L. (1995). Convergence of the iterative proportional fitting procedure, *Annals of Statistics* **23**, 1160–1174.
- [14] Stephan, F.F. (1942). An iterative method of adjusting sample frequency tables when the expected marginal totals are known, *Annals of Mathematical Statistics* **13**, 166–178.

(See also **Categorical Data Analysis**)

TERRY P. SPEED