A note on OSL age estimates in the presence of dose rate heterogeneity

Rex Galbraith,[1*]

[1] Department of Statistical Science, University College London, Gower Street, London, WC1E 6BT, UK

[*]E-mail: r.galbraith@ucl.ac.uk

1. Introduction

There has been some discussion on the estimation of OSL ages in the presence of beta dose rate heterogeneity. For example, Jacobs et al. (2008) used an 'adjusted dose rate' method to account for observed equivalent doses that appeared to follow a finite mixture distribution, while Guérin et al. (2013) argued that an estimate based on an average or central age was more appropriate. I recently commented on the latter article to say, among other things, that the mathematical argument presented there did not justify that conclusion and that further statistical analysis was needed (Galbraith, 2015). In this note I will consider dose rate heterogeneity from a statistical point of view.

2. OSL age estimates

An OSL age is usually estimated as a ratio: an equivalent dose divided by a dose rate, each of which is estimated separately. The numerator of that ratio, the equivalent dose, is usually some sort of average or representative value obtained from a sample of mineral grains, that is intended to represent a radiation dose of interest — such as the radiation energy absorbed by the mineral grains in the sample since they were last exposed to sunlight. A variety of methods can be used, depending on the context, to estimate that numerator, including the use of common age, central age, minimum age and finite mixture models. Such models may be expressed either in terms of observed equivalent doses or in terms of their logarithms, depending on whether the dominant sources of variation are additive or multiplicative. Details and rationale of these models can be found in Galbraith & Roberts (2012).

There is a separate industry devoted to estimating the denominator of that ratio, i.e., the relevant 'environmental' dose rate. This is typically a weighted sum of contributions from several sources, including alpha, beta and gamma radiation, and cosmic rays, all of which are estimated or measured separately in the field and in the laboratory, using a variety of emission counting techniques and techniques that directly measure elemental concentrations of U, Th and K. In practice, a single dose rate is usually obtained that represents an 'average' value pertaining to the sample location and its near environment. It is recognised that individual grains in the sample might experience different dose rates, especially from beta sources which can vary across small spatial distances, but dose rates experienced by individual grains are not measured by the current standard techniques. Nevertheless such variation in dose rates will be reflected to some extent in the observed single-grain equivalent doses. What implications might this have for OSL age estimation?

3. A common age model with additive errors

Suppose that we have observed equivalent doses for n mineral grains along with their empirically determined standard errors. Denote these by y_i and s_i, respectively, for $i = 1, 2, \ldots, n$. Let us consider the simplest case where we think that every grain has the same true age t. If the dose rates vary between grains, then the observed equivalent doses will vary partly because of this and partly for other reasons, including natural variation and measurement error. We can express this as

$$y_i = t\xi_i + e_i \qquad (1)$$

where ξ_i is the unobserved dose rate experienced by grain i and e_i is an unobserved random 'error' drawn from a distribution with mean zero and standard deviation σ_i, say. I have used the Greek notation ξ_i, rather than a friendlier symbol such as x_i, to remind us that the individual dose rates are not observed. The error standard deviation σ_i will include variation from all sources other than varying dose rates — including natural variation between grains and measurement error — and will typically be larger than s_i. We should remember too that the assumption that the error distribution has mean 0 is not trivial.

There is little we can do in practice with equation (1),

as ξ_i is not observed or measured. Some further theoretical analysis might be done by making assumptions about the statistical distribution of the dose rates, but before doing so, it is worth noting that if we average both sides of equation (1) we get

$$\bar{y} = t\bar{\xi} + \bar{e} \qquad (2)$$

where $\bar{y}$, $\bar{\xi}$ and $\bar{e}$ are the average observed equivalent dose, unobserved dose rate and unobserved error, respectively, for the n grains. A similar equation could be obtained by taking a weighted average: that is, we could regard $\bar{y}$, $\bar{\xi}$ and $\bar{e}$ in (2) as corresponding weighted, rather than unweighted (or straight), averages. In either case, $\bar{e}$ is also a random 'error' drawn from a distribution with mean zero. It follows that if we could estimate $\bar{\xi}$, we could then estimate the age t by dividing $\bar{y}$ by that estimate of $\bar{\xi}$; that is, by dividing an (unweighted or weighted) average equivalent dose by an estimate of the corresponding (unweighted or weighted) average dose rate experienced by the n grains in question.

One problem with the above is that $\bar{\bar{\xi}}$, the average dose rate for the n grains in our sample, may differ somewhat from the average environmental dose rate that is actually measured, which is an average for a much wider population of grains. In other words, the dose rates for the n sampled grains need to be representative of those for the population for which the average dose rate is actually measured (i.e., the population of all grains in the sample location and its near environment), which might or might not be the case in practice.

Equations (1) and (2) above are analogous to Equations (1) and (2) in Guérin et al. (2013) but are expressed here as a statistical model rather than as relationships between 'true' values. The above argument (with the caveats mentioned) supports their suggestion that, when all grains have the same true age, then this age may be estimated using an average or central age, regardless of how the individual dose rates vary. However, that does not mean that this is the *best* method or even that it is necessarily a good one.

4. Modelling the dose rate distribution

The model equation (1) could be extended to specify a distribution for the single-grain dose rates ξ_i. In practice there is usually very little independent information about what form this distribution might take, other than what can be seen in the observed equivalent doses. Let us consider two simple possibilities.

1. A normal distribution. For example, in the absence of other information, we might assume that ξ_i was drawn from a normal distribution with mean μ_ξ and standard deviation σ_ξ, independently for each grain. Then equation (1) can usefully be expressed as

$$y_i = t\mu_\xi + t(\xi_i - \mu_\xi) + e_i. \qquad (3)$$

We might regard the measured environmental dose rate as an estimate of μ_ξ, i.e., the mean dose rate for the population from which the grains were sampled. Then the quantity $t(\xi_i - \mu_\xi)$ would be another component of error, from

a normal distribution with mean 0 and variance $t^2\sigma_\xi^2$. If e_i was also from a normal distribution, and independent of ξ_i, then the overall error would be from a normal distribution with mean 0 and variance $t^2\sigma_\xi^2 + \sigma_i^2$. In that case, the optimal OSL age estimate would be that obtained from the 'unlogged' version of the central age model.

In other words, making this quite natural assumption about the unobserved single-grain dose rates leads directly to the central age model for optimal estimation of the burial dose and hence the burial age. The dispersion parameter in this central age model will include variation between single grain dose rates, as the above analysis shows.

2. A two component mixture. Another possibility might be to assume that the dose rates come from a two-component mixture distribution where ξ_i takes the value $\mu_{\xi 1}$ with probability p or $\mu_{\xi 2}$ with probability $1 - p$, say. Then equation (1) could be written as

$$y_i = t\mu_{\xi 1}u_i + t\mu_{\xi 2}(1 - u_i) + e_i \qquad (4)$$

where u_i is a bernoulli random variable that takes the value 1 with probability p and 0 with probability $(1 - p)$. So the equivalent doses also have a two component mixture distribution with component means $\mu_1 = t\mu_{\xi 1}$ and $\mu_2 = t\mu_{\xi 2}$. The common age t is given by several expressions, including

$$t = \frac{p\mu_1 + (1 - p)\mu_2}{p\mu_{\xi 1} + (1 - p)\mu_{\xi 2}} = \frac{\mu}{\mu_\xi},$$

where μ and μ_ξ are the mean equivalent dose and mean dose rate, respectively, for the population from which the grains were drawn. So, as usual, the age t can be estimated by dividing an estimate of μ by an estimate of μ_ξ.

One could imagine estimating μ either by fitting a two-component mixture to estimate p, μ_1 and μ_2, and hence $\mu = p\mu_1 + (1 - p)\mu_2$, or by simply using an average equivalent dose, ignoring the two-component mixture structure. In the latter case one might consider using either an unweighted average or a weighted average with weights proportional to the reciprocals of the error variances. What are the relative merits of these methods?

In theory, if the data really do come from a two-component mixture distribution, with a well specified error distribution, then it must be optimal to use that model for estimation by maximum likelihood, say. Nevertheless it is useful to consider this in more detail.

Suppose that, as before, e_i has a normal distribution with mean 0 and standard deviation σ_i for grain i. Consider first the hypothetical case that σ_i is the same for all grains ($\sigma_i = \tau$, say) whether or not they are from the same component. Then weighting by $1/\sigma_i^2$ is the same as weighting by $1/\tau^2$, i.e., using a straight (unweighted) average. In that case, it can be shown that this gives exactly the same estimate of μ as that obtained by fitting the two-component mixture.

Now suppose that σ_i is the same for all grains within the same component, but different for grains in different components. That is, $\sigma_i = \tau_1$ if grain i is from component 1 and $\sigma_i = \tau_2$ if grain i is from component 2. Then it can

be shown that using an unweighted average of the equivalent doses will still produce exactly the same estimate of μ as that obtained by fitting the two-component mixture. But weighting by $1/\sigma_i^2$ (i.e., by $1/\tau_1^2$ and $1/\tau_2^2$) will produce a different estimate, and it is easy to construct cases where such weighting produces a grossly biased estimate. This is because the unweighted mean implicitly combines the observations from the different components in the same proportions as those estimated by fitting the two-component mixture, whereas the weighted mean does not, unless τ_1^2 and τ_2^2 happen to be in the ratio of $1-p$ to p (which is unlikely).

The usual situation in practice is that the error standard deviations σ_i differ across grains, both within and between components. In that case, the estimate of μ obtained by fitting a two component mixture does differ from the unweighted mean equivalent dose — but typically not by very much if the differences between σ_is in the same component are small compared to those in different components. On the other hand, weighting by $1/\sigma_i^2$ will generally produce a rather different estimate.

Of course a more detailed numerical analysis would be needed to quantify these differences, but a general message is that if the y_is come from a two component mixture, then it could be misleading to combine them by weighting them by the reciprocals of their error variances ignoring the mixture structure, which is what the central age model does.

Finite mixture models are often used to estimate the parameters of specific sub-populations. The use here of a two component mixture as a form of 'error' distribution differs in concept, though it is not unknown in statistical applications where it has sometimes been used to deal with samples containing small numbers of outliers or 'unusual' values.

In the 'adjusted dose rate' method of Jacobs et al. (2008), a finite mixture model was fitted to the equivalent doses but the age was estimated from just one of the component means (μ_1, say) which was divided by an 'adjusted dose rate' (i.e., an estimate of $\mu_{\xi 1}$). That method, as I understand it, was used as an attempt to deal with equivalent doses that were thought to come from well-bleached grains that were buried at the same time but looked as if they were from a finite mixture distribution with one component containing a large majority of the grains. The rationale behind that method merits discussion, but there is no reason in principle why it should not give a valid estimate. Reasons for focusing on just one of the mixture components might include the possible unreliability of data in the other components, though this also raises the question of how to best estimate a dose rate specific to that component.

5. Models with multiplicative errors

Sometimes the dominant source of error variation in observed equivalent doses is multiplicative rather than additive. Signs of this are a strong positive skewness in a histogram of the y_is and a strong positive association in a scatter plot of s_i against y_i. Then an equation analogous to (1) may be expressed as

$$y_i = t\xi_i e^{u_i}, \qquad (5)$$

where u_i is a random error drawn from a distribution with mean 0 and standard deviation τ_i, say. This could be rewritten as an additive model for the log doses:

$$\log y_i = \log t + \log \xi_i + u_i. \qquad (6)$$

A parallel analysis to that in the previous section can be made for the model equation (6). For example, assuming that the log dose rates are a random sample from a normal distribution would lead to the usual 'logged' version of the central age model as being the optimal method.

It should be emphasised that equations (1) and (5) represent the same physical relationship. The only difference between them is in their 'error' distributions, which lead to different methods of estimating the burial dose μ. It would be possible to have either an additive error or a multiplicative error also for estimating the environmental dose rate μ_ξ, and there is no reason in principle why you should not have an additive error for the estimated dose rate and multiplicative errors for the equivalent doses. It would be wrong to say, for example, that because the dose rate is an arithmetic average of physical quantities then one should use an arithmetic average of equivalent doses to estimate the burial dose μ. In short, the appropriate method of estimation depends not only on the definition of the parameters but also on the *error* distributions.

6. Summary remarks

I have tried to highlight some statistical issues relating to how OSL ages are estimated when dose rates vary between grains, as is often the case for the beta dose rate contribution. Most of the statistical concepts I have used here are explained briefly in Galbraith & Roberts (2012, Appendix A). A key aspect is to identify the relevant parameters: not only the equivalent dose that corresponds to the burial dose (or the dose of interest), but also the relevant dose rate that corresponds to that equivalent dose. A second key aspect is that appropriate methods of estimating these parameters depend of the errors associated with individual measurements.

I have used models with additive errors for illustration. In that case, the assumption that the unobserved single grain dose rates follow a normal distribution leads to the 'unlogged' central age model for the equivalent doses. But the central age model is not necessarily appropriate for other dose rate disributions. Models with multiplicative errors may often be more appropriate in practice — and are in some ways simpler because relative errors do not depend on the scale of measurement — and the same issues arise there. In particular, the assumption that the unobserved single grain dose rates follow a log-normal distribution leads to the 'logged' central age model for the equivalent doses. Such models could be generalised to include situations where the OSL ages vary between grains. Their usefulness depends partly on how, or how much, any heterogeneity in dose rates

affects the estimation of the relevant equivalent dose, and indeed whether there is other relevant information about the dose rates that could be obtained, particularly about the nature or form of the dose rate distribution in any given application.

However, I would take issue with the idea that "modelling of the dose rate distribution is both unnecessary and undesirable" (Guérin et al., 2013, page 315). The simple analysis above has produced some useful insight, not only in suggesting appropriate methods of estimation in different situations but also in identifying a component of error variance that might be reduced if it were possible to measure single-grain dose rates. The assertions that such modelling "cannot improve accuracy, and must introduce additional uncertainties" (Guérin et al., 2013, page 315) are also incorrect, as explained above.

Statistical models are important not only for providing a basis for assessing the merits of any given method of estimation, but also for assessing sources of variation. The relevance of a statistical model to a given situation depends on the scientific context and it is not always easy to judge what models are most appropriate. Sometimes it can help to try more than one method. If different methods result in similar estimates, then some reassurance may be achieved; if not, then it may be illuminating to find out why.

Acknowledgments

References

Galbraith, R.F. *On the mis-use of mathematics: a comment on "How confident are we in the chronology of the transition between Howieson's Poort and Still Bay?" by Guérin et al. (2013)*. Journal of Human Evolution, 80: 184–186, 2015. doi: http://dx.doi.org/10.1016/j.jhevol.2014.10.006.

Galbraith, R.F. and Roberts, R.G. *Statistical aspects of equivalent dose and error calculation and display in OSL dating: An overview and some recommendations*. Quaternary Geochronology, 11: 1–27, 2012.

Guérin, G., Murray, A.S., Jain, M., Thomsen, K.J., and Mercier, N. *How confident are we in the chronology of the transition between Howieson's Poort and Still Bay?* Journal of Human Evolution, 64: 314–317, 2013.

Jacobs, Z., Wintle, A.G., Roberts, R.G., and Duller, G.A.T. *Equivalent dose distributions from single grains of quartz at Sibudu, South Africa: context, causes and consequences for optical dating of archaeological deposits*. J. Archaeol. Sci., 35: 1808–1820, 2008.

Reviewer

Regina DeWitt