

Analysis of isothermal remanent magnetization acquisition curves using the expectation–maximization algorithm

D. Heslop, M. J. Dekkers, P. P. Kruiver and I. H. M. van Oorschot

Palaeomagnetic Laboratory 'Fort Hoofddijk', Faculty of Earth Sciences, Utrecht University, Budapestlaan 17, 3584 CD, Utrecht, the Netherlands.
E-mail: heslop@geo.uu.nl

Accepted 2001 July 15. Received 2001 July 15; in original form 2001 March 8

SUMMARY

The acquisition of isothermal remanent magnetization (IRM) curves through the application of stepwise-increasing uniaxial fields to a rock-magnetic sample provides an important non-destructive tool for the investigation of coercivity spectra (Dunlop & Özdemir 1997). We show that, through the use of an automated procedure based on the expectation–maximization algorithm (Dempster *et al.* 1977), both saturated and non-saturated IRM acquisition curves can be effectively modelled into their individual coercivity contributions.

Key words: coercivity, IRM acquisition, remanent magnetization, rock magnetism.

INTRODUCTION

Robertson & France (1994) showed that the individual magnetic mineral phases contributing to a bulk IRM curve each have a cumulative log-normal acquisition path that can be described using three parameters, namely

- (1) $B_{1/2}$: the applied field at which the mineral phase acquires half of its saturation IRM (SIRM), providing a measure of the mean coercivity of that population,
- (2) M_{ri} : the magnitude of the phase distribution, providing an indication of the component SIRM and therefore its contribution to the bulk IRM curve, and
- (3) DP: the dispersion parameter, expressing the coercivity distribution of a mineral phase and corresponding to one standard deviation of the log-normal function.

Therefore, at any given field, B , the IRM intensity of an individual magnetic mineral component can be approximated using the function

$$\text{IRM}(B) = \frac{M_{ri}}{DP(2\pi)^{1/2}} \int_{-\infty}^{\infty} \exp\left[-\frac{(\log B - \log B_{1/2})^2}{2DP^2}\right] d \log B. \quad (1)$$

Fitting is normally performed on the first derivative of the IRM acquisition curve with respect to the \log_{10} field (McIntosh *et al.* 1996; France *et al.* 1999). Presenting the acquisition data in such a manner allows the investigator to readily fit a number of normal probability density functions (p.d.f.s) to the curve, with each function corresponding to a single mineral phase. Recently, the process of fitting IRM acquisition models in a consistent, unbiased manner has been addressed by Stockhausen (1998)

and Kruiver *et al.* (2001). Both of the resulting techniques, however, still require the investigator to work interactively towards the final model fit using the minimization of pre-defined statistics (e.g. the magnitude of the residuals between the measured and modelled curves). Here we present an automated fitting method based on the expectation–maximization (EM) algorithm (Dempster *et al.* 1977; Jones & McLachlan 1990), which only requires the user to define whether the sample has reached saturation during the acquisition procedure, and the number of individual mineral phases to be modelled in the final solution.

METHOD

The work of Robertson & France (1994) demonstrated that, providing no magnetic interactions occur within a sample, the first derivatives of IRM acquisition curves could be represented by the combination of a number of separate log-normal probability density functions. An IRM curve should therefore be considered as a finite mixture distribution; that is, an acquisition curve represents a bulk distribution that is composed of a finite number of log-normal populations. Under this assumption, a mixture of g separate IRM populations can be represented at a field of intensity B by the frequency function

$$f(B) = \sum_{i=1}^g M_{ri(i)} k(B; B_{1/2(i)}; DP_{(i)}), \quad (2)$$

where k corresponds to a log-normal p.d.f. on the field axis, with mean coercivity $B_{1/2(i)}$, standard deviation $DP_{(i)}$ and non-negative mixing proportion $M_{ri(i)}$. Measured IRM gradient values are therefore a discrete realization of the continuous bulk frequency function and as such should be considered as an incomplete data set. The bulk frequency function, $f(B, \theta)$, can

be used to describe the IRM gradient curve, where the parameters in the matrix θ represent the form of the mixture distribution (i.e. the $B_{1/2}$, DP and M_{Ti} of each of the IRM components). However, for the bulk frequency function of a measured sample the parameters of θ are all unknowns.

The likelihood, L , provides an indication of how likely a set of observed data points would have been, had they been selected at random from the bulk frequency function described by θ . The process of maximum likelihood estimation (MLE) involves the determination of values for the θ matrix that will maximize the probability that the sample data came from the calculated bulk frequency function. In order to perform this maximization procedure and determine the θ parameters associated with IRM bulk distributions we used the value of log-likelihood and the expectation–maximization (EM) algorithm of Dempster *et al.* (1977).

THE EM ALGORITHM

Using a two-step procedure (expectation and maximization), the EM algorithm iteratively determines the MLE of the θ matrix that describes the bulk distribution of a given incomplete data set. Before EM iteration can begin, it is necessary to provide an initial estimation of the θ parameters. Expectation (E-step) is performed first, and involves the determination of the complete data log-likelihood constrained by the observed (incomplete) data and the previously made estimation of θ . MLE (the M-step) is then performed for the estimated complete data log-likelihood (obtained during the E-step). From this maximization procedure, new estimates of θ are determined. The E- and M-steps are repeated, with the new θ -values produced during each M-step being utilized in the complete data likelihood determination in the subsequent E-step. By the stepwise improvement of the θ vector the log-likelihood of the observed data is increased until a predefined convergence criterion is reached. We have not included a formal derivation of the EM algorithm, because the essential information on the algorithm and its application to finite mixture models can be found in Dempster *et al.* (1977); Jones & McLachlan (1990); McLachlan & Krishnan (1997) and McLachlan & Peel (2000).

APPLICATION OF THE EM ALGORITHM TO IRM MIXTURE MODELS

To test the applicability of the EM algorithm to the problem of IRM ‘unmixing’ we wrote a Fortran90 program utilizing the MLE routine of Jones & McLachlan (1990). It also implements the method of Kruiver *et al.* (2001) to determine the optimum number of IRM components to include in a bulk distribution. For a typical IRM acquisition curve consisting of 30–40 data points, the IRMUNMIX program, which can be downloaded from <http://www.geo.uu.nl/~forth/Software/soft.html>, required less than 5 s to reach convergence.

In order to model the individual log-normal populations that contribute to an IRM bulk distribution it is necessary to provide the algorithm with the following information.

- (1) The number of components believed to be contributing to the bulk curve.
- (2) An initial estimate for the values in the θ matrix.
- (3) Whether or not the sample reached saturation during the IRM acquisition procedure.

In trial runs of the EM algorithm on modelled data we found that the characteristics of the resolved solutions were extremely consistent, irrespective of the starting parameters assigned to θ (the required number of iterations to reach convergence, however, increased when the algorithm was provided with initial θ estimates differing greatly from their real values). Because of the ability of the algorithm to handle poor estimates of θ , we suggest that fitting models should be initialized with systematic values assigned to θ , representing equal component contributions (i.e. $M_{Ti(i)} = 1/g$), equal dispersions, and mean coercivities evenly distributed across the \log_{10} field axis.

Determining the number of components contributing to an IRM bulk distribution is not trivial, because the goodness of fit of a finite mixture model will always improve as the number of components in the mixture is increased. To assess the number of individual components that should be included in a model we adopted the technique of Kruiver *et al.* (2001), which is based on a comparison of the residuals (calculated between the measured and modelled curves) for fits involving different numbers of components. The technique compares the variances and means of the residual arrays for two competing models. If the inclusion of an additional component does not significantly reduce the variance and mean of the residual array (assessed using an F -test and Student’s t -test, respectively) then the more complex (higher-component) model is unlikely on a statistical basis.

Saturation of a sample during IRM acquisition is an important consideration in the modelling of the measured curves because a non-saturated sample represents a truncated distribution. In the application of the EM algorithm to saturated and non-saturated samples we follow the method of Jones & McLachlan (1990). In the case of non-saturated samples the maximum likelihood estimates can be obtained for the truncated distribution using the EM algorithm without modification to the IRM data set. In the case of saturated curves (data are not truncated), Jones & McLachlan (1990) demonstrated that the EM algorithm requires the addition of an extra class (applied field interval in the case of IRM analysis), corresponding to (B_r, ∞) , where B_r is the maximum field applied to the sample. A further class, $(-\infty, B_0)$, was also included in the procedure in order to conform to the method of Jones & McLachlan (1990). In practice, the inclusion of the above class is useful because the acquisition curve should not be truncated in the low-coercivity portion of the curve, as the IRM components cannot pass into negative field values. In the following, we present the analysis of saturated and non-saturated IRM acquisition curves derived from both modelled data and natural samples.

MODELLED DATA (SATURATED)

We constructed an overlapping three-component IRM acquisition curve from the individual contributions described in Table 1. The input data was processed a total of four times, with the number of fitted components, g , increasing each time from one in the first run to four in the final run (Fig. 1). In each case the characteristics of the starting distributions given in θ were selected systematically under the criteria of equal contributions (i.e. $M_{Ti(i)} = 1/g$), equal dispersions, and mean coercivities evenly distributed across the \log_{10} field axis (Table 1). When the input value of g was left unspecified, our implementation of the Kruiver *et al.* (2001) technique to

Table 1. Component parameters used in the construction of the modelled one-, two-, three- and four-component IRM acquisition curves (A, B, C and D, respectively). The fitting procedure was initialized with systematically chosen starting components and in each case took less than 1 s to reach convergence. Range represents $B_{1/2} \pm DP$.

Comp.	Constructed Components				Starting Components				Fitted Components			
	M_{ri}	$B_{1/2}$ log mT (mT)	DP log mT	range mT	M_{ri}	$B_{1/2}$ log mT (mT)	DP log mT	range mT	M_{ri}	$B_{1/2}$ log mT (mT)	DP log mT	range mT
A												
1	0.25	0.90 (8)	0.25	4–14	1	1.75 (32)	0.10	45–56	1	1.82 (66)	0.73	12–355
2	0.30	1.7 (50)	0.4	20–125	N/P	N/P	N/P	N/P	N/P	N/P	N/P	N/P
3	0.45	2.40 (251)	0.27	134–467	N/P	N/P	N/P	N/P	N/P	N/P	N/P	N/P
B												
1	0.25	0.90 (8)	0.25	4–14	0.5	1.2 (16)	0.10	13–20	0.39	1.08 (12)	0.36	5–28
2	0.30	1.7 (50)	0.4	20–125	0.5	2.4 (251)	0.10	200–316	0.61	2.29 (195)	0.33	91–417
3	0.45	2.40 (251)	0.27	134–467	N/P	N/P	N/P	N/P	N/P	N/P	N/P	N/P
C												
1	0.25	0.90 (8)	0.25	4–14	0.33	1.50 (32)	0.10	25–39	0.25	0.91 (8)	0.25	4–14
2	0.30	1.7 (50)	0.4	20–125	0.33	2.00 (100)	0.10	79–125	0.34	1.77 (58)	0.43	21–157
3	0.45	2.40 (251)	0.27	134–467	0.33	2.50 (316)	0.10	252–398	0.41	2.41 (257)	0.26	141–474
D												
1	0.25	0.90 (8)	0.25	4–14	0.25	0.7 (5)	0.10	4–6	0.21	0.87 (7)	0.24	4–13
2	0.30	1.7 (50)	0.4	20–125	0.25	1.4 (25)	0.10	20–32	0.21	1.42 (26)	0.38	11–63
3	0.45	2.40 (251)	0.27	134–467	0.25	2.1 (126)	0.10	100–158	0.28	2.14 (138)	0.35	62–309
4	N/P	N/P	N/P	N/P	0.25	2.8 (631)	0.10	501–794	0.30	2.45 (282)	0.25	158–501

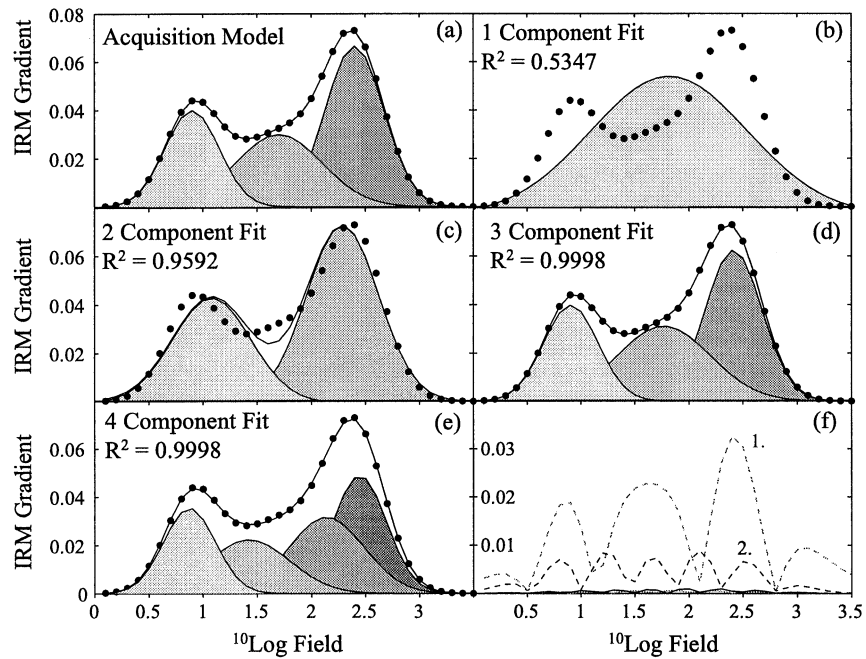


Figure 1. (a) Modelled three-component IRM acquisition curve. The closed symbols show the bulk IRM distribution, and the solid line represents the combined signals of the log-normal components (indicated by different degrees of shading). (b) (c) (d) and (e) Final fits produced by the EM algorithm (shaded components and combined signal) compared with the input distribution (solid symbols) for the one-, two-, three- and four-component fits, respectively. (f) Absolute residuals calculated between the input and EM-fitted distributions for the one- to four-component models.

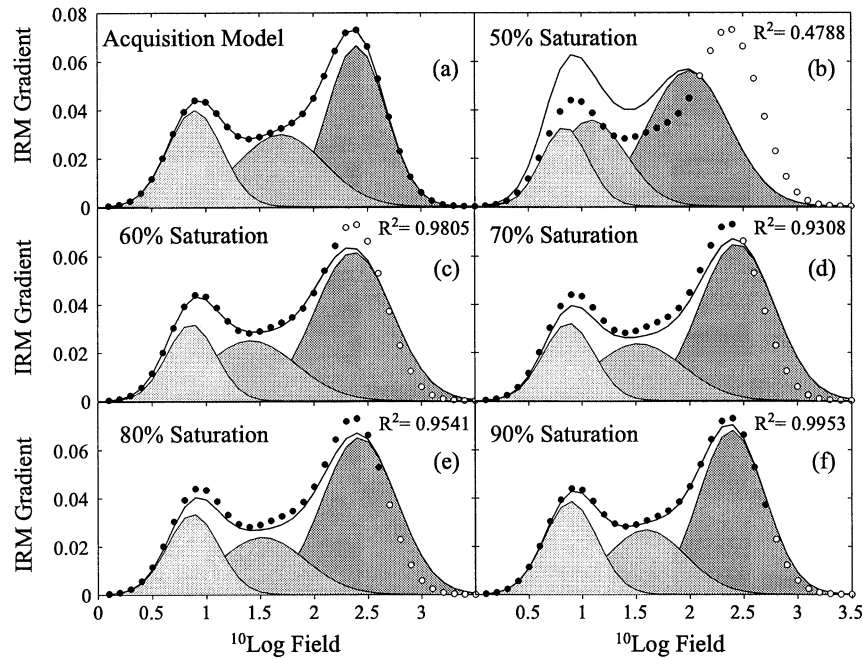


Figure 2. (a) Modelled three-component distribution used in the investigation of the effects of non-saturation on the EM fitting procedure. (b) (c) (d) and (e) Fitted solutions at different levels of saturation. The closed symbols represent the data points of the bulk distribution that were provided to the algorithm, and the open symbols correspond to points that were removed from the input data in order to simulate non-saturation. R^2 values correspond to the linear correlation coefficient calculated between the full IRM distribution [solid symbols in (a)] and the modelled distribution for the five non-saturated solutions.

determine model complexity successfully determined that g should be set equal to three.

In the case of the three-component fit, a squared correlation coefficient of $R^2=0.9998$ was determined for the linear correlation of the final fit to the input model. The one- and two-component absolute residual arrays are approximately an order of magnitude greater than those of the three- and four-component solutions. The bulk curves for the three- and four-component models are effectively identical, demonstrating that, although the fit of a finite mixture model will always improve as the number of components is increased, it is generally best to favour simplicity over complexity, for example in this case a three-component model.

MODELLED DATA (NON-SATURATED)

To investigate the effects of non-saturation on the EM-fitted models we used the same IRM acquisition curve as above but

only supplied the algorithm with data corresponding to 50, 60, 70, 80 and 90 per cent of total saturation remanence. A three-component model was fitted for each truncated data set, and the correlation coefficient between the final fit over the full field range and the input model was calculated. Fig. 2 shows the results of the fitting procedures. A good approximation of the three-component model was produced when the algorithm was supplied with data up to the 60 per cent saturation level. As more data were provided (70, 80 and 90 per cent) the fitted model gradually converged towards the solution for the full data set shown in Fig. 1(d).

Comparison of the fitted models with the input curve, shown in Table 2, demonstrates that, although the R^2 value of the 60 per cent saturation fit is higher than that of the 70 per cent model, the parameters of the individual components in the 70 per cent curve are closer to those of the input distribution. In the models where few data points are provided for the high-coercivity tail of the curve there is a tendency for the algorithm to overestimate DP for the final component. The provision

Table 2. Comparison of the EM-derived component parameters for the non-saturated model curves with the known population characteristics of the input distribution.

Model	Comp1 $B_{1/2}$ Log mT	Comp2 $B_{1/2}$ Log mT	Comp3 $B_{1/2}$ Log mT	Comp1 DP Log mT	Comp2 DP Log mT	Comp3 DP Log mT
Input Curve	0.90	1.70	2.40	0.25	0.40	0.27
50 per cent Saturation	0.84	1.09	1.99	0.22	0.32	0.38
60 per cent Saturation	0.86	1.41	2.37	0.23	0.42	0.36
70 per cent Saturation	0.87	1.52	2.44	0.24	0.41	0.35
80 per cent Saturation	0.88	1.52	2.42	0.24	0.40	0.35
90 per cent Saturation	0.89	1.59	2.39	0.25	0.37	0.30
100 per cent Saturation	0.91	1.77	2.41	0.25	0.43	0.26

of data beyond the maximum gradient point of the last component improves the fit; however, even at 90 per cent saturation, DP is overestimated in the final model. We therefore recommend that, in the case of curves where the maximum gradient point of the highest-coercivity phase has not been passed during acquisition, an interactive fitting procedure, for example Kruiver *et al.* (2001), should be utilized in preference to the EM algorithm.

Although we have tested the fitting procedure on both saturated and non-saturated modelled data, both of these investigations assumed a noise-free environment. The quality of any fit will degrade as the relative magnitude of any contaminating noise increases. Therefore, when dealing with real samples, more caution is required in the interpretation of fits obtained from noisy data sets.

MEASURED DATA (SATURATED)

A small suite of 20 samples from the ODP609 core (Ruddiman *et al.* 1989) was selected to provide examples of saturated IRM acquisition curves. The samples were taken from a section of the core spanning the climate transition from marine isotope stage 55 (warm) to stage 54 (cold) at ~ 1.6 Ma. In all cases the samples reached saturation before application of the maximum field of 1 T (Figs 3a and b). A number of the obtained IRM gradient curves were extremely noisy at low field values, and it was therefore necessary to smooth the data sets with a cubic spline before they were fitted. Each curve was modelled with two components, a decision made after the majority of the preliminary models returned two-component fits based on the criteria of Kruiver *et al.* (2001). Fig. 3(c) shows a fit for one of

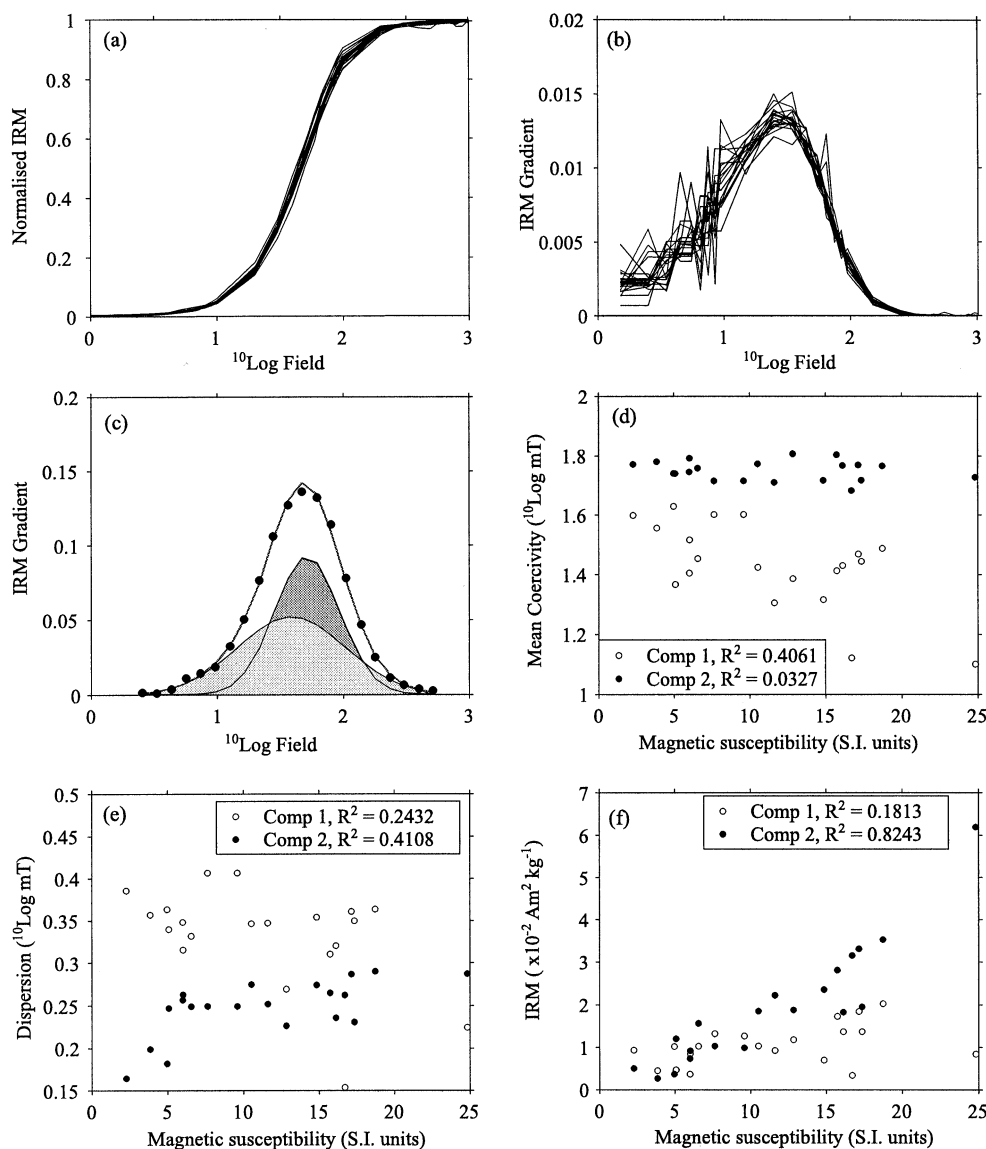


Figure 3. (a) IRM acquisition plots for the 20 samples selected from Marine Oxygen Isotope Stage (MIS) 55 to the transitional zone of MIS54 in ODP core 609. (b) IRM gradient plots for the investigated samples (the gradient curves are often fairly ‘noisy’ because the measurement error is visualized better in this way of representation). (c) Two-component EM solution for a bulk IRM distribution: component 1 (2) light (dark) shading. (d) (e), and (f) Comparison of the distribution parameters for the fitted components with the proxy climate index provided by the magnetic susceptibility record.

the samples after smoothing, revealing that the coercivity spectra for the two populations are quite similar (see Table 3 for the mean characteristics of the two populations).

This analysis enables us to compare the characteristics of the IRM components with the climatic signal as provided by magnetic susceptibility (the fidelity of magnetic susceptibility as an accurate recorder of palaeoclimatic change was determined by its comparison with the lower-resolution carbonate content record of Ruddiman *et al.* 1989). The linear regression coefficients reveal that climate had little influence over the mean coercivity or dispersion of the two magnetic-mineral components (Figs 3d and e). The absolute magnetization of component 1 shows no relationship with magnetic susceptibility (χ). The magnetization of component 2, however, demonstrates a linear trend with respect to the χ signal (Fig. 3f). This pattern indicates that, in this portion of the ODP609 core, climate was controlling the absolute concentration of component 2 (higher concentrations in warmer conditions), but not its magnetic properties, whilst component 1 remained unaffected by the ambient climatic conditions. The mean $B_{1/2}$ values of components 1 and 2 are 27 and 56 mT, respectively, indicating that both mineral populations consist of single-domain magnetite.

Table 3. Mean component parameters for the two magnetic-mineral populations fitted to the ODP609 sample collection.

Component	Mean SIRM $10^{-2} \text{ Am}^2 \text{ kg}^{-1}$	Mean $B_{1/2}$ log mT (mT)	Mean DP log mT	Mean Range mT
1	1.11	1.43 (27)	0.34	12–59
2	1.96	1.75 (56)	0.25	32–100

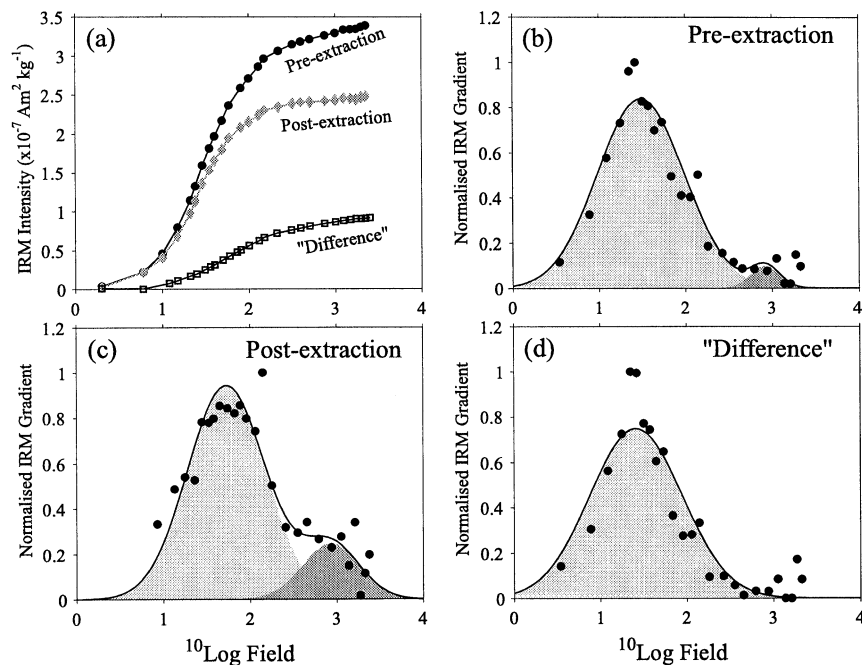


Figure 4. (a) IRM acquisition of Czech palaeosol sample B-180 ‘pre-’ and ‘post-’ AAO-Fe(II) treatment, and the ‘difference’ between the two curves. (b) (c), and (d) EM fits for each data set, where the two mineral components (magnetite/haematite) are represented by different degrees of shading (light/dark).

The apparent sensitivity of component 2 and insensitivity of component 1 to climatic change raises a number of interesting questions concerning the processes controlling the magnetic mineral assemblage in the ODP609 core. Although a more detailed rock-magnetic investigation is required to understand the observed IRM variations fully, we tentatively suggest that component 2 may correspond to a biogenic contribution. Kruiver & Passier (2001) have suggested that magnetosomes in Mediterranean sediments are characterized by magnetites with increased coercivities, typically $B_{1/2}$ is ~ 65 mT, and reduced dispersions, $DP < 0.3$. These characteristics indicate that the IRM component 2 in the ODP609 core may correspond to a magnetosome population, the magnitude of which is at least partially controlled by climatic conditions.

MEASURED DATA (NON-SATURATED)

In order to investigate the effectiveness of the fitting algorithm on IRM data obtained from non-saturated natural samples we determined the acquisition curve for a Czech palaeosol sample (B-180), which was subsequently subjected to an acid ammonium oxalate/ferrous iron [AAO-Fe(II)] extraction procedure (van Oorschot *et al.* 2001). The palaeosol was red to brownish red (Munsell: 7.5 YR6/6) in colour and was formed in a subtropical (Mediterranean) climate. In such an environment one would expect the formation of both haematite and goethite in a palaeosol, with a preference towards haematite production in sufficiently warm and dry conditions (Cornell & Schwertmann 1996). The extraction was expected to remove fine-grained iron oxides, and after treatment a second IRM acquisition curve was obtained from the sample. Fig. 4 shows that the fitted IRM curve of the ‘pre-extraction’ sample is dominated (~ 96 per cent contribution) by a magnetite component with a $B_{1/2}$ of 30.7 mT. In addition, there is a high-coercivity haematite phase (goethite was rejected on the basis of

Table 4. Fitted IRM component parameters for the palaeosol sample B-180, 'pre-' and 'post-' AAO-Fe(II) treatment, and for the 'difference' between the two acquisition curves representing the minerals dissolved during the procedure. N/P = not present.

Comp.	'Pre-extraction'				'Post-extraction'				'Difference'			
	M_{ri}	$B_{1/2}$ log mT (mT)	DP log mT	range mT	M_{ri}	$B_{1/2}$ log mT (mT)	DP log mT	range mT	M_{ri}	$B_{1/2}$ log mT (mT)	DP log mT	range mT
1	0.963	1.487 (31)	0.502	9–97	0.839	1.726 (53)	0.447	19–148	1.000	1.412 (26)	0.535	7–88
2	0.037	2.920 (832)	0.168	565–1124	0.161	2.931 (853)	0.331	398–1828	N/P	N/P	N/P	N/P

thermal IRM decay experiments) contributing less than 4 per cent to the bulk curve (Table 4). The bulk IRM curve was thus close to saturation and the EM algorithm returns reliable output.

After extraction, the relative contribution of the magnetite phase has decreased (~84 per cent) and its mean coercivity has increased (53.2 mT). The low-coercivity component is still sufficiently dominant for the EM algorithm to perform smoothly. Indeed, the fit of the 'difference' curve—calculated by subtracting the 'post-extraction' curve from the 'pre-extraction' one to provide a representation of the IRM acquisition of the minerals dissolved during the AAO-Fe(II) treatment—shows that only the magnetite phase of the sample was affected by the treatment. The coercivity of the removed magnetite appears to be lower than that of the original population. Attempts to fit the 'difference' curve with a two-component model produced an unstable second population that was inconsistent with the mineral phases defined in the 'pre-extraction' curve, suggesting that no haematite was present in the removed mineral component.

CONCLUSIONS

(1) The use of the fitting procedure based on the EM algorithm provides an effective method for determining the contributions and characteristics of individual magnetic mineral populations in both saturated and non-saturated bulk IRM curves. If the maximum of the highest-coercivity component in the gradient curve has not been passed, however, interactive fitting (Kruiver *et al.* 2001) is preferred.

(2) It has been shown that the procedure is particularly powerful when separating a number of mineral populations with overlapping coercivity spectra.

(3) Implementation of the Kruiver (2001) criteria in the modelling software provides a robust indication of the number of mineral components that should be fitted to a bulk IRM curve.

ACKNOWLEDGMENTS

The authors are grateful to Christoph Geiss and Derek France for their constructive reviews. This work was conducted under the programme of the Vening Meinesz Research School of Geodynamics.

REFERENCES

- Cornell, R.M. & Schwertmann, U., 1996. *The Iron Oxides*, VCH Publishers, Weinham.
- Dempster, A.P., Laird, N.M. & Rubin, D.B., 1977. Maximum likelihood from incomplete data via the EM algorithm, *J. R. statist. Soc. B.*, **39**, 1–38.
- Dunlop, D.J. & Özdemir, Ö., 1997. *Rock Magnetism: Fundamentals and Frontiers*, Cambridge University Press, Cambridge.
- France, D., Hu, Y., Snowball, I., Rolph, T., Oldfield, F. & Walden, J., 1999. Additional rock magnetic measurements, in *Environmental Magnetism: a Practical Guide*, pp. 197–211, eds Walden, J., Oldfield, F. & Smith, P., Quaternary Research Association, London.
- Jones, P.N. & McLachlan, G.J., 1990. Algorithm AS 254: Maximum Likelihood Estimation from grouped and truncated data with finite normal mixture models, *Appl. Statist.*, **39**, 273–282.
- Kruiver, P.P. & Passier, H.F., 2001. Different magnetic phases related to variations in redox conditions (Eastern Mediterranean Sea) indicated by IRM analyses: a cautionary note on the S-ratio, *Geochem., Geophys., Geosys.*, in press.
- Kruiver, P.P., Dekkers, M.J. & Heslop, D., 2001. Quantification of magnetic coercivity components by the analysis of acquisition curves of isothermal remanent magnetisation, *Earth. planet. Sci. Lett.*, **189**, 269–276.
- McIntosh, G., Rolph, T.C., Shaw, J. & Dagley, P., 1996. A detailed record of normal-reversed-polarity transition obtained from a thick loess sequence at Jiuzhoutai, near Lanzhou, China, *Geophys. J. Int.*, **127**, 651–664.
- McLachlan, G.J. & Krishnan, T., 1997. *The EM Algorithm and Extensions*, John Wiley & Sons, New York.
- McLachlan, G. & Peel, D., 2000. *Finite Mixture Models*, Wiley-Interscience, New York.
- Robertson, D.J. & France, D.E., 1994. Discrimination of remanence-carrying minerals in mixtures, using isothermal remanent magnetisation acquisition curves, *Phys. Earth planet. Inter.*, **84**, 223–234.
- Ruddiman, W.F., Raymo, M.E., Martinson, D.G., Clement, B.M. & Backman, J., 1989. Pleistocene Evolution: Northern hemisphere ice sheets and North Atlantic Ocean, *Paleoceanography*, **4**, 353–412.
- Stockhausen, H., 1998. Some new aspects for the modelling of isothermal remanent magnetisation acquisition curves by cumulative log Gaussian functions, *Geophys. Res. Lett.*, **25**, 2217–2220.
- van Oorschot, I.H.M., Dekkers, M.J. & Havlicek, P., 2001. Selective dissolution of magnetic iron oxides in the acid-ammonium-oxalate/ferrous-iron extraction method—II. Natural loess-palaeosol samples, *Geophys. J. Int.*, submitted.