DISCRIMINATION AND CLASSIFICATION OF UXO USING MAGNETOMETRY: INVERSION AND ERROR ANALYSIS USING ROBUST STATISTICS

Stephen D. Billings, Leonard R. Pasion and Douglas W. Oldenburg
UBC-Geophysical Inversion Facility, Vancouver, BC

Abstract

Geophysical inversion using least-squares has been a very successful method for UXO discrimination of magnetics. However, the residuals do not follow a normal distribution, which means that the assumptions underlying the inverse problem are violated. Two consequences of this are (i) model bias and (ii) incorrect estimates of model parameter uncertainties. We have found that the residuals are better modeled by an Ekbloom distribution which can be designed to be more tolerant of statistical outliers than the Gaussian.

To reformulate the inverse problem we are left with the issue of determining the two parameters that define the Ekbloom distribution. We appeal to the concept of self-consistency to resolve these parameters; i.e. the statistical distribution used to determine the model parameters should be the same as the distribution of the residuals. To achieve this aim, we set the problem up as a two parameter inverse problem and use the maximum difference between the cumulative density functions of the Ekbloom distribution and the residuals as a misfit measure.

For magnetic data collected in Montana, the recovered dipole parameters using the Ekbloom distribution could be significantly different than those obtained by least-squares. In one case considered in detail, the 95.4% confidence regions for the different solutions didn't even overlap.

Introduction

A discrimination and classification strategy for interpretation of magnetic data over areas contaminated by UXO was successfully demonstrated in Billings et al. (2002a,b). The basis of the method is inversion for the dipole moment of an anomaly. To overcome dipole ambiguity and possible remanent magnetization the recovered dipole moment is compared to a library of ordnance items expected to occur in the area. The difference between the recovered moment and that from the closest ordnance item provides an estimate of the minimum remanent magnetization that the item must possess if it were a UXO. The lower the value of the remanence the more likely the source is an intact UXO. At sites in the Helana Valley, Montana it was found that prioritizing the dig-sheet on the basis of this remanent magnetization could return 100% of the detected UXO while leaving many identified anomalies in the ground. False-alarm rates (FARs) as low as two were found when using this method (that is two non-UXOs dug per UXO).

The discrimination and classification methodology depends on the magnitude of the recovered dipole moment and its angle relative to the Earth's field. The more accurate these parameters can be recovered by the inversion routine, the more reliable the method. In the methodology described in Billings et al. (2002a) a least-squares minimization procedure is used. Additionally a localized error analysis based on the assumption of normally distributed residuals is performed to show that the dipole parameters are tightly constrained by the inversion. However the observed distribution of residuals fails a Kolmogorov-Smirnov test for normality (e.g. Bhattacharyya and

Johnson, 1997) at the 0.05 significance level. This has two potential consequences. Firstly, outliers are more likely to be encountered than predicted by a Gaussian distribution and can significantly bias the values of the recovered dipole moment. Secondly, the assumptions underlying the error analysis are violated so that we cannot rely on the predicted confidence regions.

Alternatives to the normal-assumption in the form of robust-statistical norms have been investigated by a number of authors (see Farquharson and Oldenburg (1998) and references within). Robust-norms have been used on both the likelihood function (measuring the fit to the observed data) and on the prior (to bias the solution towards models with particular structure). The main criteria used to choose the norm for the likelihood has usually been the ability to down-weight the importance of outliers without introducing undue computational complexity into the inverse problem. There is often little additional justification in choosing the form of the norm (or likelihood function). In this paper we attempt to choose a norm that is *self-consistent*. That is, a norm that returns a statistical distribution of residuals that follows the same statistical distribution from which the norm was derived.

Once a self-consistent likelihood function has been found, we are able to conduct a sensible error analysis. Unlike Gaussian statistics, where there are well-developed techniques for local error analysis (through the use of tables of the χ^2 distribution), we are lead to fully non-linear analysis of the posterior probability density function. The NA-Bayes algorithm of Sambridge (1999b) is used for this purpose.

Dipole discrimination and classification methodology

A UXO discrimination and classification methodology for magnetics data was described in Billings et al. (2002a,b). For each identified anomaly in a magnetic survey, the method involves finding the dipole moment, position and dc-shift (7 parameters) that bests fits a $2 \text{ m} \times 2 \text{ m}$ window of data (the data window extent can vary depending on the size of the anomaly). Once the best fitting dipole moment has been found we then rank each anomaly and attempt to classify potential UXO into different ordnance classes.

In order to achieve this goal, we recognize that compact ferrous objects can have both a remanent magnetization, which is fixed, along with an induced component which varies with the orientation of the object relative to the Earth's field. Even without remanent magnetization, which serves to complicate matters further, it is not possible to unambiguously determine the dimensions of a compact body from the recovered dipole moment alone. This is because a given moment could be caused by an infinite number of objects with different orientations relative to the Earth's field. To reduce this non-uniqueness we build up a library of ordnance items expected to occur in the area and then find the best fitting object for each anomaly. We assume that the objects are radially symmetric, which means they can be characterized by two demagnetization factors, one perpendicular F_1 and the other parallel F_2 to the axis of symmetry. The objects in the library are usually modelled as spheroids (McFee, 1989) but can be any arbitrary axially-symmetric object for which we can calculate the demagnetization factors. When the symmetry axis makes an angle θ with the Earth's field $\mathbf{B}$ the magnitude of the induced moment $\mathbf{m}^{in}(\theta)$ is given by the expression

$$\|\mathbf{m}^{in}(\theta)\| = \frac{V\|\mathbf{B}\|}{\mu_o} \sqrt{F_1^2 \sin^2 \theta + F_2^2 \cos^2 \theta}, \quad (1)$$

where V is the volume of the object and μ_o is the permeability of free-space. The angle $\phi(\theta)$ between the Earth's field and the induced dipole moment is

$$\phi^{in}(\theta) = \arccos \left[\frac{F_1 \sin^2 \theta + F_2 \cos^2 \theta}{\sqrt{F_1^2 \sin^2 \theta + F_2^2 \cos^2 \theta}} \right]. \quad (2)$$

Figure 1: Polar plot of dipole moment (Am^2) and angle relative to the Earth's field showing the recovered moments from two items. Also shown are the induced magnetization curves for three ordnance items.

Ordnance Type	Item 1	Item 2
76 mm projectile	67.0 %	32.0 %
81 mm mortar	58.3 %	10.8 %
105 mm projectile	37.8 %	33.6 %

Table 1: Remanent magnetizations for items 1 and 2 in Figure 1.

Thus as the orientation changes both the magnitude and direction of the induced moment will change and will trace out a circular path as shown in Figure 1.

In order to classify each anomaly we calculate the percentage difference between the best-fitting dipole for each item in the library with the recovered dipole moment (e.g. Table 1). The closest fitting item becomes the predicted source of the anomaly and the size of the difference provides a measure of the minimum remanent magnetization the anomaly must possess if it were a UXO. For example in Figure 1, item 1 is most likely a 105 mm projectile while item 2 is more likely an 81 mm mortar (Table 1). Furthermore, the basis of the discrimination method is that item 2 with a lower remanent magnetization of 10.8 % is more likely a UXO than item 1 with 37.8 % remanence.

It is evident that accurate recovery of dipole parameters is crucial for the success of the method both in ranking anomalies according to UXO-likelihood and in reliable classification.

Formulation of the inverse problem

We now turn to the issue of determining the dipole parameters for each anomaly. From a set of N observations, $\mathbf{d}$, of the magnetic field, the inverse problem is to find the set of model parameters $\mathbf{m}$ (position, dipole moment and dc-shift) that *best-fits* the data. In geophysics this optimization

problem is usually achieved through maximization of the posterior probability density function,

$$P(\mathbf{m}|\mathbf{d}) = kp(\mathbf{m})L(\mathbf{d}|\mathbf{m}) \quad (3)$$

where $p(\mathbf{m})$ is the prior-PDF of the model parameters, $L(\mathbf{d}|\mathbf{m})$ is the likelihood-function that represents the fit to the data and k is a normalizing constant. This is the well known Bayes formula which is comprehensively described in the geophysical context by Tarantola (1987). The prior is especially important in regularizing under-determined problems but has less importance in an over-constrained problem such as the one under consideration. Apart from loose bounds on the magnitudes of the model parameters the prior contributes little information and can effectively be ignored. The posterior PDF is then determined entirely by the likelihood function

$$P(\mathbf{m}|\mathbf{d}) = kL(\mathbf{d}|\mathbf{m}) \quad (4)$$

If the residuals are assumed to follow a Gaussian distribution then the posterior PDF is

$$P(\mathbf{m}|\mathbf{d}) = k \exp \left[-\frac{1}{2}(\mathbf{d} - \mathbf{g}(\mathbf{m}))^T C_d^{-1}(\mathbf{d} - \mathbf{g}(\mathbf{m})) \right] \quad (5)$$

where $\mathbf{g}(\mathbf{m})$ is a forward-model that returns a set of predicted data and C_d is the covariance matrix of the data. The usual approach to the inverse problem is to maximize the posterior PDF, or equivalently in this case, minimize the log-likelihood,

$$-\log L(\mathbf{d}|\mathbf{m}) \propto \frac{1}{2} \sum_{n=1}^N \frac{[d_n - g_n(\mathbf{m})]^2}{\sigma_n^2} \quad (6)$$

where the last part of the equation holds if we make the assumption that the residuals are independently randomly distributed with variance σ_n^2 .

In Billings et al. (2002a) we implemented the above approach as a least-squares procedure and used it to find best-fitting dipoles for over a thousand magnetic anomalies in Montana. If the combination of the observation uncertainties and forward modelling errors do indeed follow a normal distribution, then the residuals should also follow a normal-distribution. This would justify our choice of a Gaussian-likelihood function. We applied the well-known Kolmogorov-Smirnov test of normality (e.g. Bhattacharyya and Johnson, 1997) to the residuals for each of about 360 anomalies at Limestone Hills and found that the majority fail the test at the 0.05 significance level.

Finding a self-consistent likelihood function

We can infer from the above example that the residuals are unlikely to be drawn from a normal-distribution. To rigorously formulate the inverse problem we need to find a likelihood function that is self-consistent. By this we mean a likelihood function, that when used to invert for model parameters, produces a distribution of residuals that follows the same statistical distribution as the likelihood function itself.

There are several types of statistical distributions in the literature that could potentially be used as a likelihood function. These include the numerous Pearson type distributions, the generalized Lambda distribution (Karian and Dudewicz, 2000) and the Ekblom distribution (Ekblom, 1973). This last distribution includes all the Lp-norms (including the L1-norm and the Gaussian) and was found to fit the distribution of the residuals following least-squares inversion of the Montana data. Furthermore, solution of the inversion problem with this distribution is simply achieved through an iteratively re-weighted least squares procedure (Farquharson and Oldenburg, 1998). We therefore use the Ekblom distribution in this paper.

The Ekblo PDF is parameterized by three different parameters p , ϵ and σ through the PDF,

$$P_E(r|p, \epsilon, \sigma) = k \exp\left(-\frac{(r^2 + \epsilon^2)^{p/2}}{p\sigma^p}\right) \quad (7)$$

where k is a constant that causes the integral of the PDF to be unity. The values of p and ϵ control the shape of the Ekblo PDF, while σ plays the part of the standard deviation and controls the width of the PDF. The parameter p most strongly influences the asymptotic behavior of the distribution (the tails of the PDF) with larger values making the PDF more tolerant of statistical outliers. The parameter ϵ controls the shape of the PDF near the origin. When $\epsilon = 0$ the Ekblo reduces to a Lp norm with the specified value of p .

The posterior PDF for the inverse problem, assuming independent errors σ_n , is given by the expression

$$P(\mathbf{m}|\mathbf{d}) = k \exp\left(-\frac{1}{p} \sum_{n=1}^N \frac{[(d_n - g_n(\mathbf{m}))^2 + \epsilon^2]^{p/2}}{\sigma_n^p}\right) \quad (8)$$

while the negative log-likelihood is

$$-\log L(\mathbf{d}|\mathbf{m}) \propto \frac{1}{p} \sum_{n=1}^N \frac{[(d_n - g_n(\mathbf{m}))^2 + \epsilon^2]^{p/2}}{\sigma_n^p} \quad (9)$$

This becomes the norm, or misfit function, that determines how well the specified dipole model fits the available data.

To find a self-consistent likelihood function, we need to solve for the dipole parameters that minimize the negative log-likelihood in Equation 9 for a given set of Ekblo parameters. We then calculate the maximum difference between the cumulative density functions (CDF) of the residuals $\{r_n, n = 1..N\}$ and the assumed Ekblo distribution,

$$\phi(p, \epsilon, \sigma) = \max_x \left| \frac{M(r_n < x)}{M(r_n < \infty)} - \int_{-\infty}^x P_E(r|p, \epsilon, \sigma) dr \right| \quad (10)$$

where $M(r_n < x)$ gives the number of residuals less than x ; for instance $M(r_n < \infty) = N$. We use $\phi(p, \epsilon, \sigma)$ as a measure of misfit and seek the value of p , ϵ and σ that minimizes this misfit. Thus we see that the search for a self-consistent likelihood function consists of a nested sequence of inversions; an outer layer for the determination of the optimum Ekblo parameters and an inner layer for the inversion for the best-fitting dipole parameters for each proposed Ekblo distribution. Changing the value of σ does not change the minimum of the negative log-likelihood so that another set of inner-loop inversions are not required. It has the effect of stretching or contracting the cumulative density function and the optimum value of σ for a given p and ϵ is easily found by a simple and fast one-dimensional optimization routine. The outer loop inverse problem for p and ϵ is solved by a Sequential Quadratic Programming method described in Schittkowski (1985).

Practical implementation of the method

To illustrate the practical considerations required to effectively implement the method, we applied the technique to 360 anomalies identified at Limestone Hills, Montana (Billings et al., 2002a). We first fit each of the anomalies by least-squares and retained the 316 anomalies with correlation-coefficients (of fitted to measured data) greater than 0.8. We are then faced with three choices:

1. Fit a separate Ekbloom distribution to each anomaly;
2. Fit a single Ekbloom distribution to all anomalies; or
3. Fit several Ekbloom distributions to groupings of anomalies.

We rejected the first choice because the smaller the number of residuals the lower the reliability of the cumulative distribution of the residuals and the easier it is to find a set of parameters that pass the Kolmogorov-Smirnov test of consistency. For the second choice there are two issues. The first is that the size of the errors for different anomalies will differ so that the raw residuals cannot be sensibly combined. To overcome this issue we can scale the residuals for each anomaly using the sample variance. This is calculated at the optimum set of model parameters, $\mathbf{m}^{opt}$ using

$$\sigma_s^2 = \frac{1}{N-7} \sum_{n=1}^N [d_n - g_n(\mathbf{m}^{opt})]^2, \quad (11)$$

where $N-7$ is the number of degrees of freedom (number of residuals minus the number of model parameters). The residual scaling is then achieved by

$$\hat{r}_n = \frac{[d_n - g_n(\mathbf{m})]}{\sigma_s}. \quad (12)$$

The second issue is that the statistical distribution for different anomalies may be different. This is best handled by the third listed choice: grouping the anomalies into classes with similar statistical properties.

To separate the anomalies into classes we propose to use the kurtosis of the residuals from the least-squares fit. For the residuals $\{r_n, n = 1..N\}$, the kurtosis is defined by

$$k = \frac{1}{(N-7)\sigma_s^4} \sum_{n=1}^N [r_n - \bar{r}]^4. \quad (13)$$

and provides a measure of the length of the *tails* of the distribution. For the Gaussian, $k = 3$, with smaller values of k indicating narrower tails and larger values much broader tails. For instance, the L1-norm is broader tailed and has $k = 24$.

The Limestone Hills data were grouped into 11 classes on the basis of the kurtosis, with at least 9 and at most 55 anomalies in a single class (Table 2). Estimation of the optimum Ekbloom parameters is a highly non-linear problem and several different starting models were often required before a good solution was found. The optimum values for each class are shown in Table 2 along with the significance level of the Kolmogorov-Smirnov test for the specified number of residuals at the 0.05 significance level. The p-values are all greater than 0.05 which confirms that the optimum Ekbloom parameters do indeed define a self-consistent norm for each class.

As one might expect, the values of p are strongly correlated with the kurtosis groupings; the larger the kurtosis (in general) the smaller the value of p . Interestingly, a non-zero value for epsilon is only required for five of the different groupings. We also note from that table that the sigma value decreases with increasing kurtosis (decreasing value of p), a consequence of the change in the definition of the spread of an Lp-norm with the value of p (Tarantola, 1987).

The cumulative density functions for four of the kurtosis groupings are shown in Figure 2. As one might expect, the tails of the Ekbloom distribution are narrower than the Gaussian for kurtosis values less than 3, while they become significantly wider for larger kurtosis values. Only those anomalies with kurtosis values around 3 can be sensibly modelled using the normal distribution and hence with least-squares.

Kurtosis	Number	Kol. Smir. p-value	p	ϵ	σ	Median change in dipole moment
<2.25	17	0.05	2.87	0	1.15	3.1%
2.25-2.75	52	0.30	2.12	0.017	1.04	1.5%
2.75-3.25	55	0.83	1.56	0.555	0.849	3.8%
3.25-3.75	44	0.21	1.46	0.4	0.788	5.5%
3.75-4.5	39	0.19	1.45	0	0.777	6.2%
4.5-5.5	40	0.05	0.85	0	0.525	16.0%
5.5-6.5	22	0.19	0.725	0	0.475	15.6%
6.5-7.5	11	0.43	0.55	0	0.437	15.5%
7.5-8.5	9	0.19	0.50	0.002	0.394	26.8%
8.5-9.5	12	0.16	0.576	0	0.4	15.6%
>9.5	15	0.47	0.587	0	0.3734	14.7%

Table 2: Ekbloom distribution fits to anomalies from the Limestone Hills area separated on the basis of the kurtosis of the residuals from a normal fit. The third-column lists the Kolmogorov-Smirnov level of significance between the likelihood function and the CDF of the residuals. The last column shows the median percentage change in the recovered dipole moment between the Ekbloom and least-squares solutions.

Effect of likelihood function on recovered parameters

The 316 anomalies at Limestone Hills were inverted for best fitting dipole parameters using both a least-squares procedure and the optimum Ekbloom parameters. For this second inversion, the Ekbloom parameters for the corresponding kurtosis class were used. While differences in all recovered parameters were noted, here we concentrate on the differences in the recovered dipole moment and angle relative to the Earth's field. It is these two parameters that control the classification and discrimination of an anomaly.

The last column in Table 2 shows the median percentage change in the recovered dipole moment for each kurtosis class, compared to the least-squares solution. For those anomalies with a kurtosis of around three, there are only small differences between the two solutions for each anomaly, as one would expect. For the classes with larger kurtosis values there are much larger differences between the two solutions, with a peak difference of over 25%. This result shows the importance of choosing the correct likelihood function, especially when the residuals deviate significantly from a normal-distribution.

There were 63 items validated as UXO in the data inverted. The recovered dipole moments for the least-squares and Ekbloom fits of a number of these UXO are displayed on a polar plot in Figure 3a. The polar plot has dipole magnitude increasing outwards in a radial direction (magnitudes are displayed in Am^2), with the angle relative to the Earth's field in the radial direction. Least-squares and Ekbloom fits for a given anomaly are joined by a black line. While the differences between the different solutions are quite small for many anomalies, there are a significant number where the discrepancy in recovered moment is quite large.

To determine the impact of the differences in solution on discrimination we first need to define a library of ordnance items. We note that previous surface sweeps of the Limestone Hills area revealed five different caliber projectiles had been used (76 mm AP/T, White Phosphorous (WP) and HE, 90 mm AP/T, WP and HE, 105 mm illumination, WP and HE, 4.2" illumination and HE, and 155 mm illumination, HE and WP). The size variability of the different types for a given

Figure 2: Cumulative density functions for four different kurtosis groupings.

caliber are quite small except for the 4.2" illumination and HE rounds. Therefore, a single set of dimensions were used for each ordnance caliber except for the 4.2" rounds.

Using this ordnance library we then calculated the minimum remanent magnetization required for each anomaly identified as a UXO. Figure 3b shows a histogram of the remanent magnetization for both the least-squares and Eklom fits. The Eklom has more anomalies with lower remanent magnetization or, in other words, better fits to the items in the library. This implies that the use of the Eklom norm has a positive impact on the accuracy of the recovered dipole parameters.

Error analysis

The normal Gaussian approach to error analysis involves the use of statistical tables and the model-covariance function (see Bard, 1974). Localized error analysis in this way is adequate as long as the objective function is quadratic about the optimum solution. Faced with a non-linear inverse problem and non-Gaussian statistics we turn instead to integration of the posterior PDF of Equation 4 for our error analysis.

Figure 3: Effect of likelihood function on recovered dipole moments.

The posterior PDF encapsulates all the statistical information we have on the model parameters. Two quantities in particular that we are interested in are the model covariance matrix $\mathbf{V}$ whose i -th and j -th component is

$$V_{ij} = \int_{\mathbb{R}^d} m_i m_j P(\mathbf{m}|\mathbf{d}) d\mathbf{m} - \bar{m}_i \bar{m}_j \quad (14)$$

and the marginal PDF. For each model parameter, the marginal is obtained by fixing the value of that parameter and then integrating the posterior PDF over all other model parameters. For example, the 2-D marginal PDF of the i -th and j -th parameters is,

$$M(m_i, m_j) = \int_{\mathbb{R}^{d-2}} P(\mathbf{m}|\mathbf{d}) \prod_{k \neq i, j}^d dm_k. \quad (15)$$

Thus we see that non-linear error analysis involves multi-dimensional integration. We could use Monte-Carlo integration for this purpose, but instead use a more efficient algorithm known as NA-Bayes (Sambridge, 1999b).

The NA-Bayes approach

NA-Bayes uses an ensemble of models where the PPD has already been calculated to approximate the multi-dimensional integrals required for non-linear error analysis. No additional forward modelling calculations are required because the PPD is approximated everywhere in parameter space using the Neighborhood approximation of the PPD from the input ensemble. In principle the input ensemble can follow any distribution and be generated by any search method (e.g. a genetic-algorithm, simulated-annealing or the neighborhood algorithm; Sambridge, 1999b). A standard Gibbs sampler is then used to provide estimates of the integrals of the PPD and their numerical error.

To generate the input ensemble we use the neighborhood algorithm (NA) of Sambridge (1999a). This algorithm uses geometrical constructs known as Voronoi cells to guide the search through parameter space. The best approximations to the integrals of the PPD were achieved by

Parameter	L2-soln	Ekblohm soln	Local L2 σ_m	Global L2 σ_m	Ekblohm σ_m
Easting (m)	23.108	23.118	0.027	0.036	0.030
Northing (m)	8.865	8.850	0.017	0.035	0.027
Depth (m)	0.502	0.588	0.017	0.036	0.026
Magnitude (Am^2)	0.321	0.407	0.023	0.037	0.033
Azimuth angle	-58.8°	-47.0°	3.9°	4.7°	3.3°
Dip angle	-37.8°	-57.6°	2.8°	3.5°	2.6°
Angle relative to B	49.3°	28.3°	2.9°	3.6°	2.4°
Remanence	20.5%	11.4%	-	5.3%	2.8%

Table 3: L2 and Ekblohm norm model parameters for anomaly 147, along with the predicted standard deviations on the model parameters.

forcing NA to sample widely in parameter space. If the search strategy were too aggressive (i.e. focussed on quickly locating the global minimum), insufficient information on the parameter space was gathered for the Neighborhood approximation to be representative of the PPD.

Application to the Limestone Hills data

To illustrate both the effect of changing the likelihood function and the application of the global error analysis we consider the recovered dipole parameters for target 147. This was an anomaly over a pre-emplaced 4.2" illumination round.

The sample standard deviation used for the error analysis of the least-squares solution for target 147 was $\sigma_s = 14.1$ and the kurtosis was 8.6. The parameters used for the Ekblohm norm solution for target 147 are $p = 0.576$ and $\epsilon = 0$. Once the optimum solution for this Ekblohm norm has been found an estimate of the standard deviation of $\sigma = 6.87$ nT was derived through the product of the sigma value for this kurtosis class, $\sigma_e = 0.4$, and the sample standard deviation at the Ekblohm solution, $\sigma_s = 17.2$ nT. We note that the residuals returned by the Ekblohm norm, pass a Kolmogorov-Smirnov test for consistency at the 0.05 significance level, while those for the least-squares norm fail the test with a p-value very close to zero.

The recovered model parameters and estimated standard deviations for target 147 are given in Table 3. The global least-squares standard deviations (obtained through equation 14) are between 20 and 120% larger than the local least-squares estimates, while the global Ekblohm estimates of the model standard deviations are around 10 to 30% smaller than the equivalent least-squares estimates. Of much more significance is the shift in model parameters, especially the dipole magnitude and angle relative to the Earth's field. The dipole moment changes by a total of 0.146 Am^2 which is equivalent to 40% of the magnitude of the least-squares solution. The shift in dipole moment causes the remanent magnetization to change from 20.5% ($\sigma_m = 5.3\%$) at the least-squares solution compared to 11.4% ($\sigma_m = 2.4\%$) at the Ekblohm solution. This change in remanent magnetization is quite small for such a large change in recovered dipole moment, and occurs because the two solutions lie on opposite sides of the induced magnetization curves (see Figure 3).

The 1-D marginal PDF's of the least-squares and Ekblohm solutions for dipole magnitude, angle relative to the Earth's field and azimuthal angle, and the remanent magnetization are given in Figure 4. We note that there is very little overlap in the marginal PDF's for any of the recovered model parameters. This is further evident in a plot of the 95.4% (two-sigma) confidence regions for dipole magnitude and angle (Figure 5) obtained by contouring the two-dimensional marginal PDF for these two parameters. The confidence regions for the different solutions don't even overlap.

Figure 4: Marginal PDF's for the L2 and Ekbloim norm solutions for target 147.

It is the remanent magnetization that determines the anomaly's position in a discrimination ranking list. In Figure 4d we plot the marginal PDF of remanence for the least-squares and Ekbloim solutions. The Ekbloim solution has a lower remanence but the difference is not great, given the very large change in recovered dipole parameters. The reason for this small difference is that the Ekbloim and least-squares solutions lie on opposite sides, but quite close to, the induced magnetization curves (see Figure 5).

Discussion

The parameters of the dipole model that *best fit* a magnetic anomaly are obtained by maximizing the likelihood function (or equivalently minimizing the negative log-likelihood). If the assumptions underlying the likelihood function hold, then the residuals should be distributed according to the same distribution as the likelihood function. For many anomalies at Limestone Hills, making an a-priori assumption of a Gaussian likelihood, and solving for the parameters by least-squares, results in residuals that are not normally distributed. In this paper, we reformulated the inverse problem to include an outer-loop to solve for the parameters of an Ekbloim PDF that

Figure 5: Two-sigma (95.4%) joint confidence region on dipole magnitude and angle relative to the Earth’s field for the least-squares and Ekblohm solutions to target 147.

result in a self-consistent norm; i.e. one where the residuals and likelihood function follow the same distribution according to a Kolmogorov-Smirnov test at the 0.05 significance level.

Anomalies at Limestone Hills were grouped into classes based on the kurtosis of the residuals of the least-squares fits. The kurtosis reflects the size of the tails of the statistical distribution underlying the residuals. Those anomalies with a kurtosis of around 3 (the value for a normal distribution) can be sensibly modelled using Gaussian statistics. Perhaps self-evidently, the larger the kurtosis, the greater the influence the change in norm has on the recovered dipole moment. When the kurtosis exceeds about 4.5 the inverse problem can no longer be sensibly formulated using Gaussian statistics. For instance, the effect of the change in norm on Target 147 was sufficiently marked that the 95.4% confidence regions on dipole moment and angle for the least-squares and Ekblohm solutions didn’t even overlap.

Those anomalies for which Gaussian statistics are adequate tend to have accurately located observations, properly levelled data on adjacent lines and fit the dipole model quite well. When the kurtosis exceeds about 4.5 one or more of the above conditions typically do not hold. When the anomaly is no longer purely dipolar, the forward model is unable to reproduce the observed data. In Billings et al. (2002a) we developed an inversion method that included the higher order magnetic moments (such as the quadrupole) in the forward model. Improving the forward model in this way would reduce the importance of the norm for certain of the anomalies (i.e. where there are no levelling or positional problems).

The error analysis conducted in this paper was performed through multidimensional integration of the posterior PDF, which is equivalent to the likelihood in the absence of a prior. The confidence regions on dipole magnitude and angle were considerably larger than suggested by the local error analysis performed in Billings et al. (2002a). In that paper, it was argued that uncertainty in the recovered dipole moment would not significantly influence classification decisions (i.e. assigning an ordnance type to each UXO candidate). Improper classification decisions were inferred to occur due to remanent magnetization and inaccuracies in the induced dipole model.

However, the present error analysis implies that uncertainty in the recovered dipole parameters is likely to be an important contributing factor.

By formulating the inverse problem using a self-consistent norm, bias in the recovered dipole parameters was reduced considerably. The remanent magnetizations required to fit the induced models were decreased, implying more accurate recovery of dipole parameters. Thus, by using a self-consistent likelihood function, our ability to discriminate and classify using magnetics were enhanced.

Acknowledgements

The authors would like to thank G-tek Australia and Clif Youmans of the Montana Army National Guard for the Limestone Hills magnetic dataset and Malcolm Sambridge for permission to use the NA-Bayes algorithm. Funding for the work presented in this paper was provided by the U.S. Army Engineer Research and Development Center and the Army Research Office.

References

- Bard, Y., 1974, Non-linear parameter estimation: Academic Press.
- Bhattacharyya, G. K., and Johnson, R. A., 1997, Statistical Concepts and Methods: John Wiley and Sons.
- Billings, S. D., Pasion, L. R., and Oldenburg, D. W., 2002a, Inversion of magnetics for UXO discrimination and identification: Proc. 2002 UXO Forum, Orlando September 3-6.
- Billings, S. D., Stanley, J. M., and Youmans, C., 2002b, Magnetic discrimination that will satisfy regulators: Proc. 2002 UXO Forum, Orlando September 3-6.
- Ekblom, H., 1973, Calculation of linear best L_p approximations: BIT, **13**, 293–300.
- Farquharson, C. G., and Oldenburg, D. W., 1998, Non-linear inversion using general measures of data misfit and model structure: Geophys. J. Int., **134**, 213–227.
- Karian, Z. A., and Dudewicz, E. J., 2000, Fitting statistical distributions: The generalized Lamda distribution and the generalized bootstrap methods: CRC Press.
- McFee, J. E., 1989, Electromagnetic remote sensing; Low Frequency Electromagnetics:, Technical Report 124, Defence Research Establishment Suffield, January 1989.
- Sambridge, M., 1999a, Geophysical inversion with a neighbourhood algorithm I. Searching a parameter space: Geophys. J. Int., **138**, 479–494.
- 1999b, Geophysical inversion with a neighbourhood algorithm II. Appraising the ensemble: Geophys. J. Int., **138**, 727–746.
- Schittkowski, K., 1985, NLQPL: A FORTRAN-subroutine solving constrained nonlinear programming problems: Annals Oper. Res., **5**, 485–500.
- Tarantola, A., 1987, Inverse problem theory: Methods for data fitting and model parameter estimation: Elsevier.