A Bayesian method for imaging voids using Gravimetry

By Gareth B.

Defence Science Technology Laboratory, Salisbury, Wiltshire, United Kingdom, SP4 0JQ

Abstract

We describe the implementation of a Bayesian model to assess novel gravimeters which employ recent advances in cold matter-wave interferometry to accurately measure micro-gravity. Potential applications of this technology include: improved identification of oil reservoirs; the detection of sink-holes to support civil engineering; detection of underground buildings; and geophysical mapping as a means to navigate in GPS denied environments. For the Defence and Security led challenge of imaging underground buildings at depths not amenable to ground penetrating radar, the utility of the posterior distribution for visualising the location of an underground void is shown by calculating *probability of excavation* heat maps. The impact of different sampling configurations and densities is investigated. It is shown that, when background gravitational noise is comparable to the signal strength of the anomaly, dense sampling of the local environment is necessary to accurately determine the location of the void.

1. Introduction

Since its inception following the Nobel prizing winning work of Steven Chu *et al.* (1997), the use of cold atom matter-wave interferometry to measure inertial forces has developed from lab-based measurements of fundamental physical constants to deployable sensors for high precision measurement of acceleration and rotation. The current capabilities of atom interferometers for measuring inertial forces are comparable to their state-of-the-art classical counterparts, see Louchet-Chauvet *et al.* (2011). However, in principal their precision can be improved by several orders of magnitude, see Dickerson *et al.* (2013).

With the potential for greatly improved sensitivity to tiny changes in acceleration, applications of gravimetry and gravity gradiometry include: improved identification of oil reservoirs worth billions to the energy sector; the detection of sewage services and sink-holes to support civil engineering (it is estimated that up to 4 million holes are cut into the UK road network each year in order to install or repair buried service pipes and cables Bongs *et al* (2012)); detection of underground buildings; and geophysical mapping as a means to navigate in GPS denied environments (see Stockton *et al* (2011)).

This paper describes the implementation of a Bayesian model to assess novel gravimeters which employ recent advances in cold atom matter-wave interferometry to accurately measure micro-gravity. The application is the resolution of underground structures close to the surface, as an array of measurements of acceleration can be used to infer what lies beneath. This is by virtue of the fact that one weighs less when above a void. More specifically the quality and number of measurements required to resolve simulated voids, of the order of 100m^3, at depths not amenable to standard surveying technology such as ground-penetrating radar, will be investigated. This paper will introduce and discuss the quality of spatially-resolved *probability of excavation* heat maps and use this as a metric for assessing sampling strategies.

This work is part of the field of geophysical inversion; many excellent texts describe the problems and challenges of determining the composition and morphology of the sub-surface (see for example Sen & Stoffa (2013) and Parker (1994)). For the problem under consideration it is clear that a method that accounts for all sources of uncertainty is critical. Without greatly constrained prior information, the number of possible physical models that could describe a finite set of inertial measurements is infinite. Therefore a Bayesian approach to this problem is the natural choice.

For the simple case of a gravitational anomaly caused by a single void, it is difficult to differentiate between the void's volume, its depth beneath the surface and the density contrast with respect to the soil. This is complicated greatly by the inhomogeneity of the soil; the location of interfaces between different rock and soil types; and numerous other environmental sources of noise including but not limited to: tidal variation, water table height variation, atmospheric pressure, seismic waves, terrain effects and other buildings, noise generated by the sensor itself and even vehicles and living creatures (for detail concerning measurement of the Earth's gravitational field see Hinze *et al* (2013)).

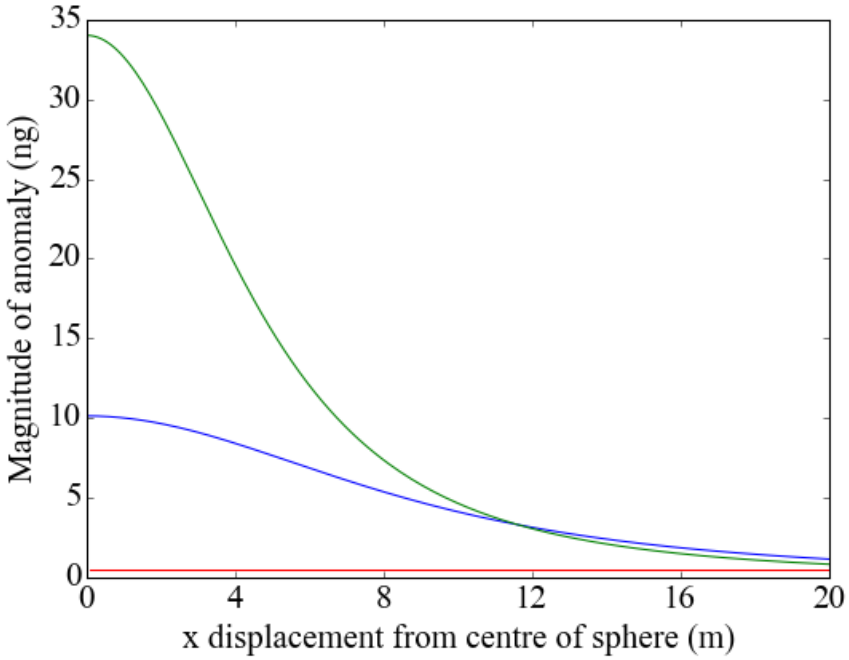

FIGURE 1. Magnitude of the anomalous gravitational acceleration of a spherical cavity as a function of horizontal displacement from the cavity's centre of mass (blue and green curves). The cavity is located at 10m (blue) and 5m (green) beneath the Earth, has a volume of 100m^3, and the local density of earth is assumed to be 1800kg/m^3 (see Equation 2.1). Measurements are made 1m above the Earth's surface. The noise floor of the atom interferometer gravimeter, 0.462ng, is given by the red curve for a 5min period of sampling at $8\text{ng/Hz}^{1/2}$ precision. It is clear that, in the absence of unaccounted sources of noise, the gravimeter would detect the anomaly at distances in excess of 20m from the cavity's centre.

2. Measuring small changes in gravity

To assess the capability of atom interferometers to locate underground voids, it is useful to take the demonstrated precision of cold atom gravimeters at $8\text{ng/Hz}^{1/2}$, see Muller *et al* (2008) (1ng is $10^{-9} \times 9.81\text{m/s}^2$), and compare this to typical signal strengths. The gravitational acceleration exerted by a 150kg mass at 1m is approximately 1ng. Therefore it is possible to detect a human being if they stand still and the sensor makes an integrated measurement over several seconds; an accuracy of $\pm$1ng could be achieved after 64s.

To assess the feasibility of detecting voids beneath the Earth, a simple comparison with the analytic expression for a spherical void can be made (see Hinze *et al* (2013) page 49). The z-component of acceleration is given by

$$g_z(x,z) = \frac{4\pi\, G\, \Delta\rho\, R^3}{3z^2\left[1+(x^2/z^2)\right]^{3/2}} \tag{2.1}$$

where x is the horizontal displacement from the centre of the sphere, and z is the depth of the centre of the sphere beneath the Earth. The radius of the sphere is given by R, the density contrast between the void and the surrounding soil is $\Delta\rho$ and G is the Gravitational constant.

The blue and green curves in Figure 1 show the expected magnitude of acceleration anomaly caused by a 100m^3 spherical cavity at $z = 10$m and $z = 5$m depth respectively. The red curve shows the noise floor of the atom interferometer for a 5min integrated signal (a magnitude of 0.462ng for a sensor with a precision of $8\text{ng/Hz}^{1/2}$). The local density of earth is assumed to be 1800kg/m^3 (see Hinze *et al* (2013) for representative values). From this simple assessment it can be concluded that atom interferometry based technology is capable of detecting the presence of such voids. However, this result is without due consideration of sources of environmental and geophysical noise encountered in the real world. The spatial resolution of more complex cuboid voids, and the consideration of background noise sources, will be described in terms of a Bayesian model in the following section.

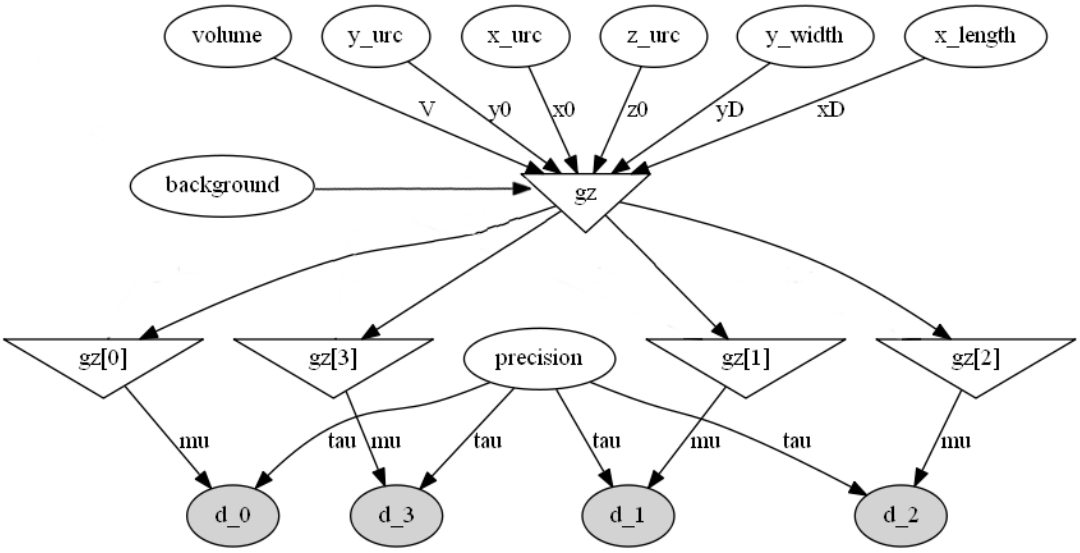

FIGURE 2. Bayesian gravimeter model with background noise. Ellipses define stochastic variables: white are inferred, grey are observed. Triangles define deterministic functions. The measurement model is normally distributed and centred on the forward model; the statistical precision (inverse of the variance; not to be confused with measurement precision) is inferred. The function gz computes the anomalous z-component of acceleration due to a cuboid plus a background offset correction (analytic expression for a cuboid can be found in Chappell *et al* (2012)). The background offset can account for a shift in the pre-calculated local value of g, the acceleration due to Earth's gravity. As an example, gz[0] is the value obtained at the measurement location defined by the observed data d_0. The parameters of the void are defined by the upper-right co-ordinates, x_urc, y_urc, z_urc, the volume, width and length (the use of upper-right co-ordinates, rather than the centre of mass, prevents the proposal of voids that lie partially above the surface). The reader may notice that the cuboid is constrained to be aligned the Cartesian co-ordinate system.

3. A Bayesian approach to gravity anomalies

To determine the spatial extent of a void using geophysical measurements such as gravity, there exists extensive literature on the use of methods in linear algebra (for example Parker (1994)). Whilst these are powerful and can quickly provide a solution, they often suffer from great sensitivity to small changes in the measurement data. The geophysical inversion problem is generally known to be ill-posed so methods are applied to regularise the equations to provide stability (see for example Tickhonov regularisation Sen & Stoffa (2013) page 44). In contrast, a Bayesian formulation of the problem provides a natural framework for accounting for all sources of error; it is not limited to linear dependence on model parameters.

The Bayesian solution to the inversion problem is given by the posterior probability density function (pdf): the space of all possible solutions describable by the model, m, (see for exampleSen & Stoffa (2013)). The posterior pdf, $p(\theta \,|\, x, m)$, over model parameters, θ, given the data x, is related to the prior, $p(\theta \,|\, m)$, through:

$$p(\theta \,|\, x, m) = \frac{p(\theta \,|\, m)\, p(x \,|\, \theta, m)}{p(x \,|\, m)}, \tag{3.1}$$

where $p(x \,|\, \theta, m)$ is the likelihood function, and the denominator can be regarded as a normalising constant for the purpose of computing marginal distributions and probability of excavation maps.

3.1. *Bayesian model*

The model used throughout this paper is given in Figure 2. It describes the dependencies between the measurement data, measured stochastic variables, and the parameters, unobserved stochastic variables. The unobserved variables define a single cuboid void contained by a homogeneous medium of soil with an assumed density of 1800kg/m^3. Whilst not accounting for access, this would be typical of a simple underground building and is a good starting point before moving to more complex geometries and correspondingly complex models. The measured data consists of a regular two dimensional grid of gravitational measurements at a set of locations above a featureless surface (only 4 measurement locations, d_0, d_1, d_2 and d_4 are shown in the figure).

This simple model does not directly consider terrain effects and rock soil interfaces that would give rise to strong spatial variations in the background gravity field not describable by the cuboid void. Smaller scale spatial variations, due to inhomogeneities in the soil for example, can be accounted for by the precision variable (inverse of the variance; not to be confused with measurement precision). As variability in the soil density

Sampler spacing	Normally distributed background noise / ng		
	0	20±2	200±20
5m	*coarse sampling density zero noise background*	*coarse sampling density low noise background*	*coarse sampling density high noise background*
2m	*medium sampling density zero noise background*	*medium sampling density low noise background*	*medium sampling density high noise background*
1m	*fine sampling density zero noise background*	*fine sampling density low noise background*	*fine sampling density high noise background*

TABLE 1. MCMC simulation matrix identifiers.

increases, the effective sensor noise increases and the inferred precision variable should decrease. Comparison of the inferred precision with calculated values of the precision can be used as a method of determining if all sources of gravitation have been accounted for. For example, for confidence in the model's representation of the environment, stochastic errors due to sensor noise models and soil noise models should agree with the inferred precision.

The background offset variable shown in Figure 2 accommodates systematic shifts on all measured data as caused by, for example, water table and tidal variations. Such effects typically occur on time-scales much greater than the sensor integration time and their spatial length scales are typically much greater than the domain over which inference is being made. This may not be true for large scale geophysical inversion problems, but is a reasonable approximation for inferring the location of underground voids as described in this paper.

With regard to Equation 3.1, the white ellipses encode prior knowledge about unobserved variables and the network describes the likelihood function (i.e. how the unobserved stochastic variables are related to the observed stochastic variables). Following work described in JafarGandomi & Binley (2013) the priors are uniform to reduce bias (parameters are however constrained by minimum and maximum values) and following McCalman *et al* (2014) the likelihood function is given by a multivariate normal distribution centred on the forward model. The forward model gives the z-component of acceleration at a specified location as

$$g_z(x,y,z) = \xi + G\,\Delta\rho\left[\acute{x}\ln\left(\acute{y}+r\right) + \acute{y}\ln\left(\acute{x}+r\right) - \acute{z}\arctan\frac{\acute{x}\,\acute{y}}{\acute{z}\,r}\right]\Bigg|_{\acute{x}=\hat{x}-L-x}^{\hat{x}-x}\Bigg|_{\acute{y}=\hat{y}-W-y}^{\hat{y}-y}\Bigg|_{\acute{z}=\hat{z}-H-z}^{\hat{z}-z} \tag{3.2}$$
$$r = \sqrt{\acute{x}^2+\acute{y}^2+\acute{z}^2}$$

where ξ is the inferred background variable, $\hat{x}$ is the upper-right x co-ordinate and similarly for y and z. The length, width and height of the cuboid are specified by the variables L, W and H. The remaining parameters are the same as in Equation 2.1. For a derivation of the analytic expression for a cuboid see Chappell *et al* (2012).

The measurement covariance matrix is diagonal and specified completely by the inferred scalar statistical precision variable (inverse of the variance; not to be confused with measurement precision). Limiting the covariance matrix to no correlation between sensor measurements is a reasonable approximation in the limit of no background noise. This will be discussed further in later sections.

3.2. *Inference engine*

The computation of an approximation to the posterior pdf is implemented using standard Metropolis-Hastings Markov chain Monte Carlo (MCMC) as described in Brooks *et al* (2011). During the burn-in period of the sampler, Adaptive Metropolis (see Haario *et al* (2001)) is used to account for strong correlations between the volume of the void and its depth beneath the Earth's surface.

The model was developed using PyMC by Patil *et al* (2010). For each MCMC simulation presented, 80,000 MCMC steps were taken. A burn-in period of 20,000 simulation steps with adaptive sampling preceded standard Metropolis-Hastings. Following the burn-in, a thinning of 1 in 50 leaves 1,200 samples.

4. Verification of the model

In order to test the ability of the inference engine to locate underground voids, a number of MCMC experiments using synthetic data were performed. In addition to the intrinsic sensor noise of ±0.462ng (representative of a 5min integration period for a gravimeter with a precision of 8ng/$\text{Hz}^{1/2}$), a simple background comprising a

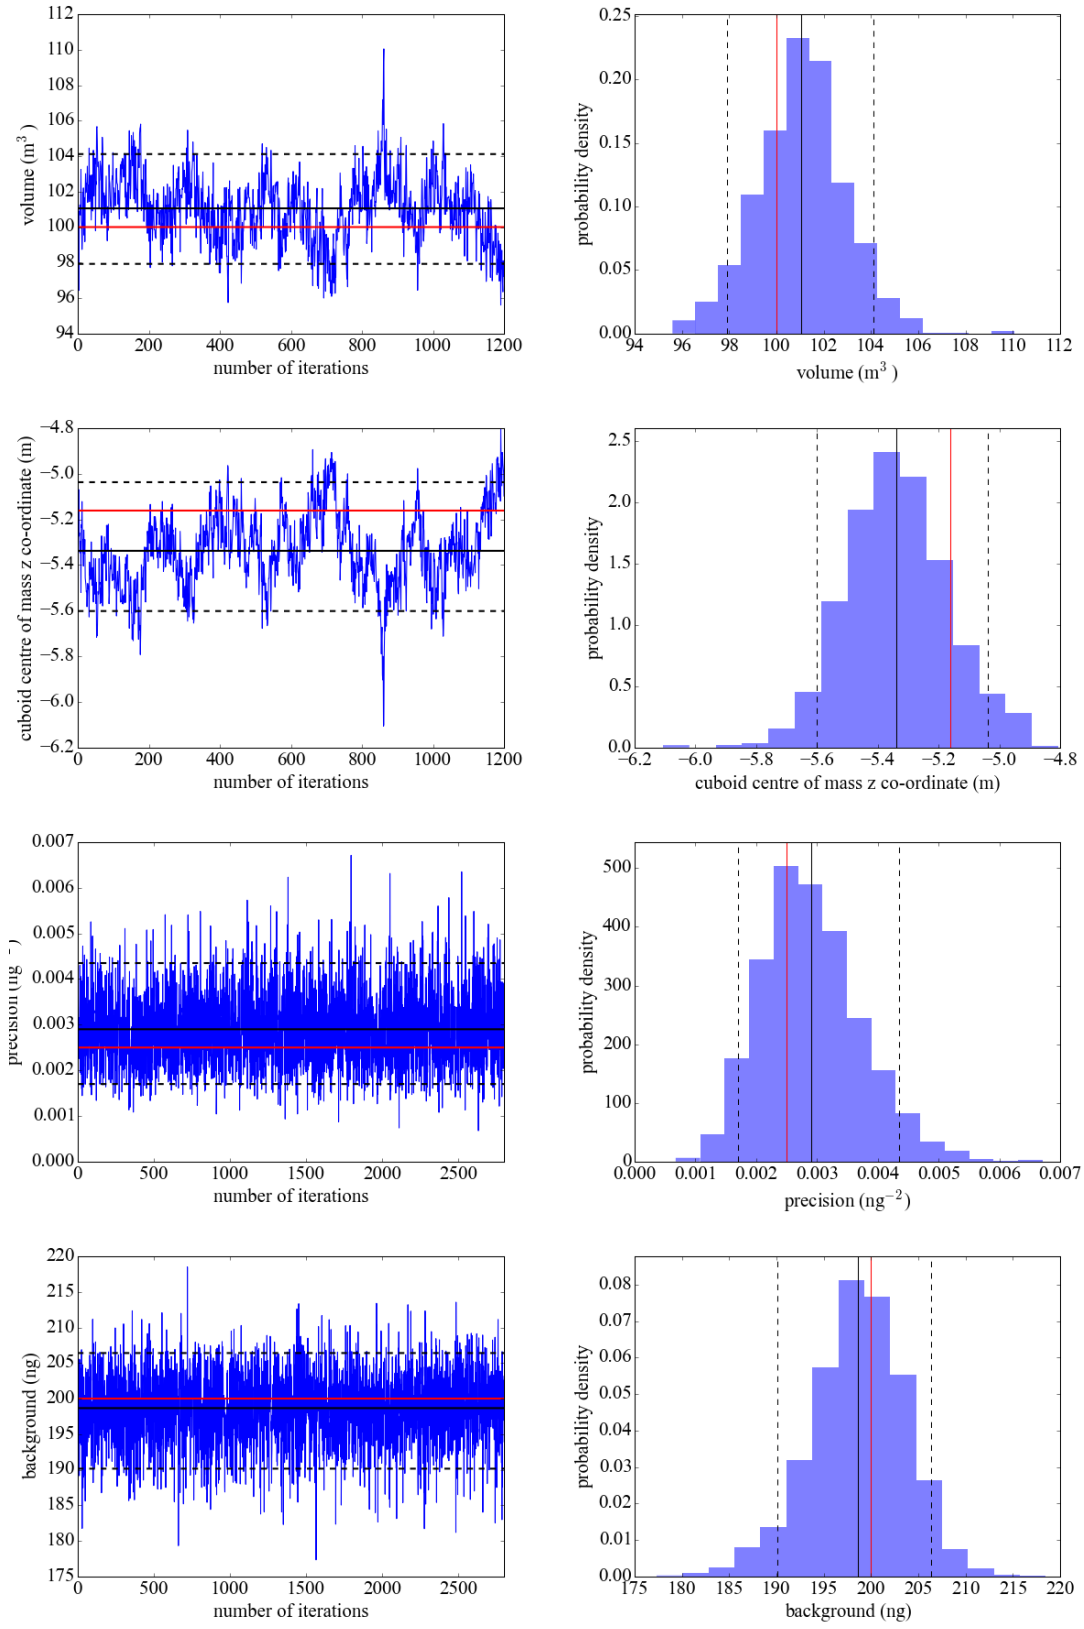

FIGURE 3. Selected simulation MCMC output trace and corresponding histograms (approximate marginal distributions) for the cuboid volume and the cuboid centre of mass z co-ordinate, the precision and background variable. With reference to Table 1, the volume and the cuboid centre of mass z co-ordinate are from simulation *fine sampling zero noise background*; the precision and background are from *coarse sampling high noise background*. The synthetic data was generated using the likelihood function forward model, so that the fit to the model should be excellent. For each plot, the mean value is shown by the black line and the dotted lines are the 5% and 95% population boundaries; the true value is shown in red. A strong correlation between the cuboid's depth and volume can be easily seen in the traces.

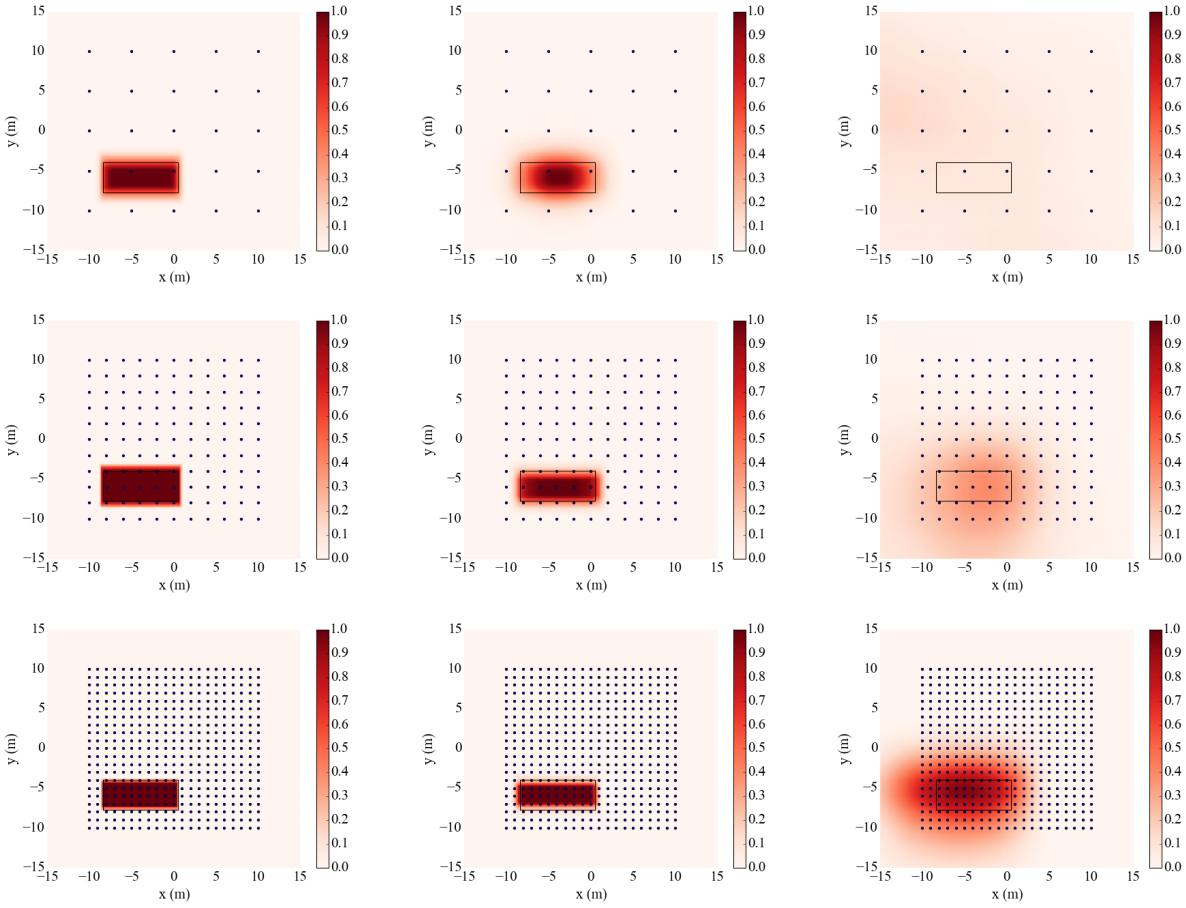

FIGURE 4. Probability of excavation maps in the x-y plane as a function of sampler density and uncorrelated background noise. Sampler locations are on a regular grid identified as points within the domain. Sampling densities ranges from coarse (top) to fine (bottom) and measurements were simulated ranging from zero (left) to high (bottom); see Table 1. The projection of the ground-truth cuboid in the plane is identified by the rectangle. It can be seen that increased sampling density improves agreement with the true location of the anomaly. The simulated background noise clearly has a detrimental impact on the ability of the inference engine to estimate the location of the true anomaly. Figure 5 shows the corresponding heat map in the y-z plane.

constant offset and a normally distributed noise source were modelled. Without developing detailed models of the environment, this is an attempt to model both ideal and more realistic environments where uncertainties due to temporal variations, such as water table and tidal motions (which produce a shift in the field), and noise sources such as vibration and soil inhomogeneity (which add to the intrinsic sensor noise) exist.

4.1. *Inference engine traces and histograms*

The test scenario consisted of an anomaly due to a 100m^3 cuboid void, assuming the soil density to be 1800kg/m^3, with centre of mass 5.16m beneath the Earth. The response of a gravimeter making 5min integrated measurements on a regular 2D grid, 1m from the surface of a flat Earth was calculated. Three different background noise sources and three different sampling densities were investigated; unique identifiers for each simulation are specified in Table 1.

Figure 3 shows a selection of traces and histograms (approximate marginal distributions) from simulation experiments *fine sampling density zero noise background* and *coarse sampling density high noise background* (with reference to Table 1). The first two trace and histogram pairs, of cuboid volume and depth of centroid, are taken from simulation *fine sampling density zero noise background*. As there is no background noise source, the fit to the true values is excellent; strong correlation between the two parameters can be clearly seen. The last two traces and histogram pairs are from *coarse sampling density high noise background* and show excellent fits to the sensor precision and the background offset. For a simulated sensor noise of ± 0.462ng, with a background simulation noise of ± 20ng, the precision parameter's true value, given as the inverse of the variance, is 0.00250ng^{-2}.

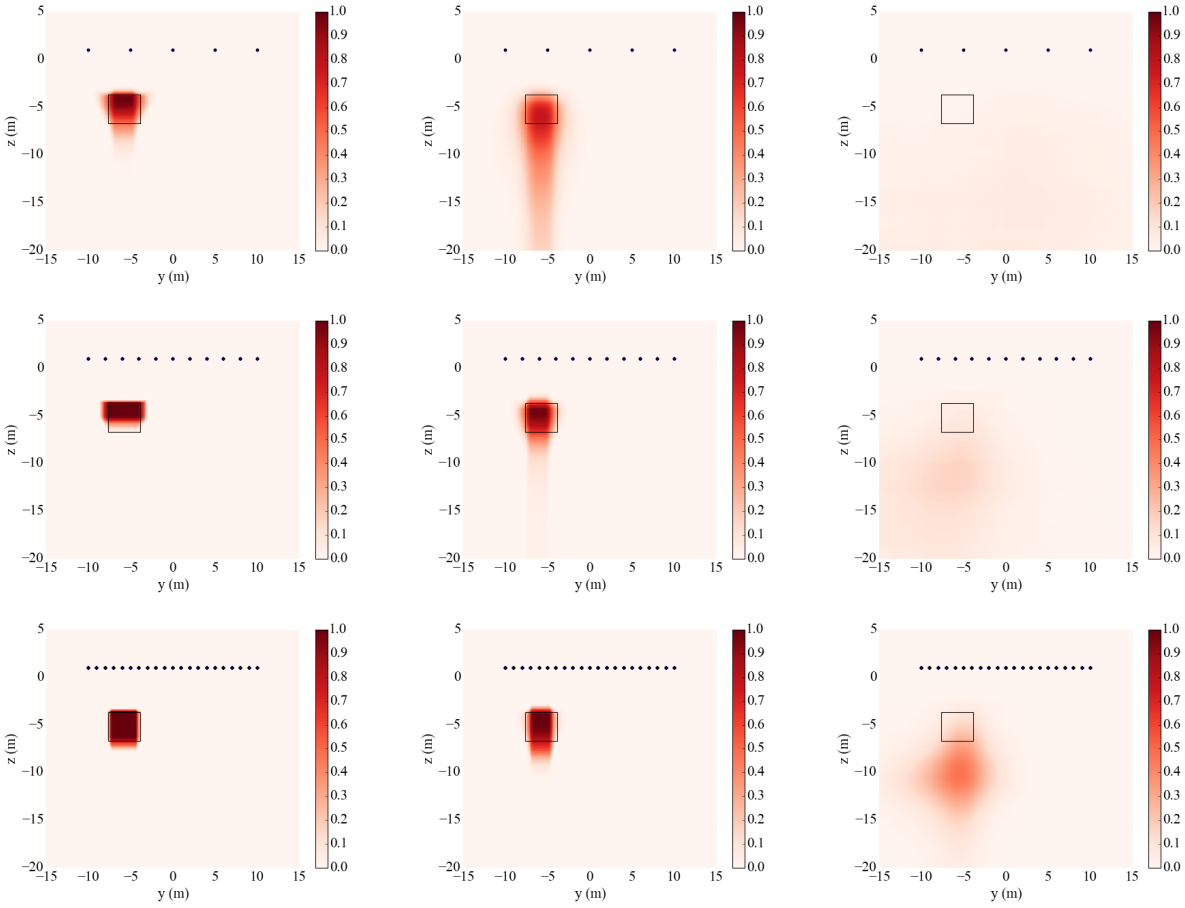

FIGURE 5. Probability of excavation maps in the y-z plane as a function of sampler density and uncorrelated background noise. Sampler locations are on a regular grid identified as points within the domain. Sampling densities ranges from coarse (top) to fine (bottom) and measurements were simulated ranging from zero (left) to high (bottom); see Table 1.; see Table 1. The projection of the ground-truth cuboid in the plane is identified by the rectangle. It can be seen that increased sampling density improves agreement with the true location of the anomaly. The simulated background noise clearly has a detrimental impact on the ability of the inference engine to estimate the location of the true anomaly. Figure 4 shows the corresponding heat map in the x-y plane.

Whilst trace and histogram information is useful to understand the performance and convergence of the MCMC sampler, it is not a particularly useful output for an operator of a gravitational imaging system. This is exacerbated when the model dimension becomes greater and, as a consequence, review of a few marginal distributions could easily lead to erroneous conclusions. A more useful visualisation, *probability of excavation* heat maps, is introduced in the next section.

4.2. *Probability of excavation heat maps*

A considerable weakness of inspecting marginal distributions is the inability to easily see correlations between the inferred parameters. For this application, probability of excavation heat maps can be used as a method that accounts for correlation and is invariant to changes in the dimensionality of the inference. Using Markov chain samples that approximate the pdf, each probability of excavation heat map pixel, a_{ij}, is given by

$$a_{ij} = \sum_{\mathrm{k=1}}^{\mathrm{N}} \frac{\mathrm{I_{xy}}\left(m_k\right)}{\mathrm{N}}, \tag{4.1}$$

where m_k is a sample from the simulated Markov chain of length N (this could of course be generalised to voxels in 3 dimensions). The indicator function $\mathrm{I_{xy}}$ returns 1 if the pixel is partially or completely contained by the cuboid as projected into the xy plane (or otherwise) of the 2D heat map. Thus, each sample from the chain describes a cuboid; if every cuboid in the chain overlays a given particle pixel, when projected into the specified plane, then that pixel has a probability of 1 of being part of the void.

Figure 4 shows the resulting x-y plane probability of excavation heat maps for the experiment matrix and

Figure 5 respectively shows the y-z projection. These images were created by summing the 1200 chain members for each simulation and averaging, as given by Equation 4.1. Sampling density (top to bottom) and background noise (left to right) are varied as described in subsection 4.1 and Table 1. Moving from top to bottom, improved agreement with the ground-truth can be seen as the sampling density is increased. For the case of zero background offset and zero normally distributed background noise, it can be seen that the inference is not adversely affected and a coarse sampling of the environment would likely provide sufficient information for a user to make a decision.

As background noise is introduced into the model, the fit to the true cuboid is degraded and coarse sampling is shown to be ineffective for the purposes of gravity imaging. Taking simulation *coarse sampling density high noise background* as an example (top right image in Figure 4 and Figure 5), despite the poor fit to the cuboid void, the fit to the precision and background nuisance parameters is good (see Figure 3). This implies that the model inference is working well with the simulated forward model data and that the noise is too great to infer the location of the void, so the cuboid parameter values are dominated by the prior.

The Bayesian model has been tested with scenarios involving multiple cuboid voids. As expected, inference is poor in these scenarios; however it is useful to note that the inferred statistical precision could be used as a diagnostic. When the inferred statistical precision is much lower than the predicted value (including all sources of sensor, soil and environmental noise models), this indicates failure of the model to account for all sources of gravitation. Operationally, large deviations of the statistical precision marginal distribution from the model predicted value would lead the user to question the generated probability of excavation heat maps.

From this limited scenario it can be concluded that dense sampling, at 1m spatial resolution, is necessary when the normally distributed background noise is of similar magnitude to the strength of the anomaly (see the green curve in Figure 1 where the peak signal is 34.03ng). The inferred background variable, whilst increasing the dimensionality of the problem, accommodates unaccounted for shifts in the mean gravitational field and or systematic error in the sensor measurement. Furthermore, through simulation, it has been shown that the background shift and anomaly can be inferred despite the shift being much greater than the peak signal of the anomaly.

5. Conclusions

A simple model for the detection of underground structures at depths not amenable to ground penetrating radar, using measurements of acceleration, has been presented. Through a single simulated scenario, the quality of the computed probability of excavation map as a function of sampler density and environmental noise has been examined. It has been shown that sampler density is important for visualising a 100m^3 cuboid void, with centre of mass 5.16m beneath the Earth, when environmental noise is comparable to the peak signal. In low noise environments, coarse sampling at 5m spacing gave comparable performance to a finer 1m grid. Through these simulations, the probability of excavation heat maps have been shown as an appropriate way of presenting multi-dimensional information encapsulated within the posterior pdf. It is believed that this form of visualisation could be usefully employed in a gravitational imaging system.

The utility of inferring the background off-set and the sensor precision have been discussed. Inference of the background off-set accommodates unknown systematic shifts due to, for example changes in the water table or variations due to the tides. The variable can also accommodate errors in classical gravimeters that do not affect cold atom matter-wave interferometers, in particular errors such as sensor temperature dependence and linear creep on sensor springs.

Inference on the sensor precision naturally accommodates all sources of normally distributed noise without necessarily knowing the origin or magnitude of the source. Thus, provided the assumption of spatially uncorrelated normally distributed noise is valid, then an imaging system should be able to resolve structure if the signal strength is comparable to the standard deviation of the inferred noise. This assumption will be violated, for example gravitational noise due to variations in the soil density will have an associated length scale. One approach to account for this would be to infer this length-scale. Future work will assess the impact of this approximation by simulating spatially correlated noise.

It may therefore be possible to calibrate a sensor by first making measurements of the local environment in an area where it is known that no local anomaly exists. Then, assuming approximate homogeneity of nearby regions, an exploration for anomalies could be usefully undertaken. Finally, comparing the inferred precision with a pre-calculated model precision may prove a useful diagnostic for detecting model error and therefore assessing the reliability of the posterior pdf and resulting heat maps.

With the ultimate aim of developing a gravity imaging system, a key problem to be addressed is accounting for all sources of noise and better understanding the trade-off between number and quality of measurements.

Work will focus on realistic models of the environment including geophysical effects and all other sources of gravitation. Optimised sampling strategies will be investigated: for example using the current probability of excavation to direct sampling to more informative regions of space.

Acknowledgements

The author would like to thank Kevin Ridley, Philip Milsom, Anthony Rodgers, Daniel Boddice, Michael Holynski, Kai Bongs, Stephen Till, Jonathan Pritchard, Geoffrey Garnett, Brian Barber, Richard Hollins and Peter Robins for useful discussions and their continuing support. This work was supported by the UK Ministry of Defence.

REFERENCES

Peters, A., Chung, K. Y., Young, B., Hensley, J. & Chu, S. 1997 Precision atom interferometry *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **335**, 2223–2233.

Louchet-Chauvet, A., Merlet, S., Bodart, Q., Landragin, A., Pereira Dos Santos, F., Baumann, H.,D'Agostino, G. & Origlia, C. 2011 Comparison of 3 Absolute Gravimeters Based on Different Methods for the e-MASS Project *IEEE Transactions on Instrumentation and Measurement* **60**, 2527–2532

Dickerson, S., Hogan, J., Sugarbaker, A., Johnson, D. M. S. & Kasevich, M. 2013 Multiaxis inertial sensing with long-time point source atom interferometry *Physical Review Letters* **111**, 1–5

Bongs, K. & Krüger, P. 2012 Unveiling the underground *Physics World* **25**, 22–26

Stockton, J. K., Takase, K. & Kasevich, M. A. 2011 Absolute Geodetic Rotation Measurement Using Atom Interferometry *Physical Review Letters* **107**, 13–30

Sen, M. K. & Stoffa, P. L. 2013 Global Optimization Methods in Geophysical Inversion *Cambridge University Press*

Parker, R. L. 1994 Geophysical Inverse Theory *Princeton University Press*

Hinze, W. J., von Frese, R. R. B. & Saad, A. H. 2013 Gravity and Magnetic Exploration: Principles, Practices, and Applications *Cambridge University Press*

Müller, H., Chiow, S., Herrmann, S., Chu, S. & Chung, K. 2008 Atom-Interferometry Tests of the Isotropy of Post-Newtonian Gravity *Physical Review Letters* **100**

Chappell, J. M., Chappell, M. J., Iqbal, A. & Abbott, D. 2012 The gravitational field of a cube *arXiv*

JafarGandomi, A. & Binley, A. 2013 A Bayesian trans-dimensional approach for the fusion of multiple geophysical datasets *Journal of Applied Geophysics* **96** 38–54

McCalman, L., O'Callaghan, S. T., Reid, A., Shen, D., Carter, S., Krieger, L., Beardsmore, G. R., Bonilla, E. V. & Ramos, F. T. 2014 Distributed bayesian geophysical inversions *Thirty-Ninth Stanford Geothermal Workshop*

Brooks, S., Gelman, A., Jones, G. L. & Meng, X. 2011 Handbook of Markov Chain Monte Carlo *Chapman & Hall/CRC Handbooks of Modern Statistical Methods*

Haario, H., Saksman, E., & Tamminen, J. 2001 An Adaptive Metropolis Algorithm *Bernoulli* **7** 223–242

Patil, A., Huard, D. & Fonnesbeck, C. J. 2010 PyMC: Bayesian Stochastic Modelling in Python *Journal of statistical software* **4** 1–81