Dear Colin,

Thank you for your insightful comments and review. We have implemented the changes in the manuscript that you requested and adjusted the text when we thought explanations could be made clearer in the light of your comments.

In particular, we have included 2 additional inversion runs in the synthetic case studies. These use uncertainty volumes where some interfaces suggested by $\boldsymbol{W_H}$ are in disagreement with geology and where uncertainty is exaggerated.

We think that we have answered your concern about the justification of the integration approach we follow.

In the following, we answer to your comments relating to the methodology and to your suggestions.

Best regards,

Jérémie Giraud and co-authors

Red: review and comments.
Blue: authors' answer.
Green: text modifications.

Comment, paragraph 1,p. C2

- It is exactly the case that if one weights the gradient (roughness) term in a minimum-structure term with the spatially varying weights shown in Figure A4b, then the gravity (or magnetic) inversion will construct a model for which the gradient is concentrated in the locations where $\boldsymbol{W_H}$ is small. The gravity (or magnetic) inversion is sufficiently non-unique that the data are quite happy for the gradients in the model to be put where $\boldsymbol{W_H}$ is small: the data will essentially never have a strong enough influence to overcome this effect of $\boldsymbol{W_H}$.

Answer

This is not exactly true. This is case-dependent in that you can well imagine a very complex geological $\boldsymbol{W_H}$ model with short wavelength variations of $\boldsymbol{W_H}$ where of course gravity (magnetic) inversion will not be able to overcome the effect of $\boldsymbol{W_H}$. In contrast, as long as variations in $\boldsymbol{W_H}$ are of similar or larger wavelength as what gravity (magnetic) data can resolve, then geophysical inversion may overcome the effect of $\boldsymbol{W_H}$. This is illustrated in Fig. 3 and 4 (added to manuscript), which we remind below:

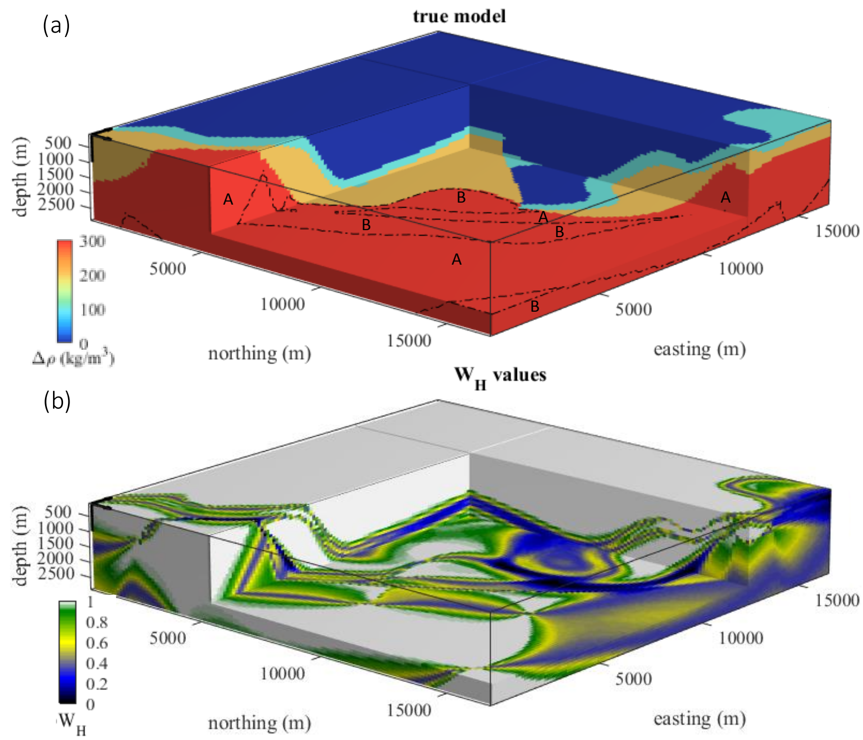

Figure 1. (a) true density contrast model with outline of the 'ghost' unit B (black dashed line), embedded in ultramafic unit A, and (b) local weights calculated from the PGM calculated using MCUE for model (a).

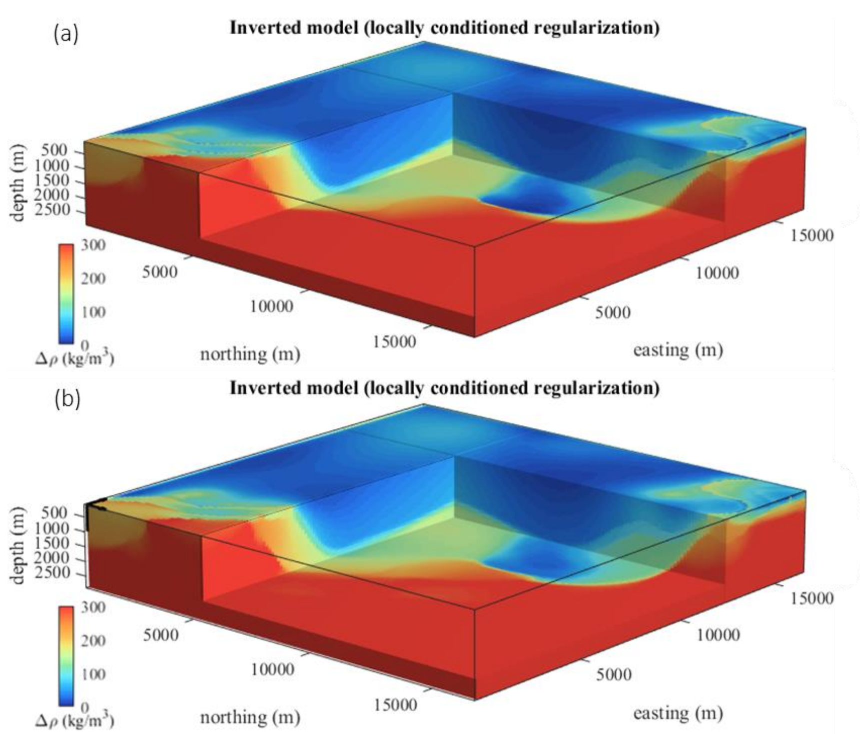

Figure 2. Comparison of inverted model using $\boldsymbol{W_H}$ *derived from a PGM considering the ghost unit (b) and without it (a); (a) is the inverse model obtained without bias in the PGM as per* ***Error! Reference source not found.*** *and is shown here for comparison with (b).*

Comment, last part paragraph 2, p. C2

- [...] you're specifying where you want this sharp interface.

Answer

This remark is true to some extent but needs to be amended.

The model covariance $\boldsymbol{W_m}$ is the same everywhere. This means that we do not encourage specific cells to vary independently from others. Rather, we reduce the influence between adjacent volumes of cells in the calculation of the gradient constraint by decreasing $\boldsymbol{W_H}$. This allows differences between neighbouring groups of cells to be stronger than otherwise by encouraging cells in between these volumes (e.g., around interfaces or poorly constrained lithologies) to have stronger gradients. These gradients make up the interfaces. This does not mean that we are specifying where interfaces should occur, but that we specify where they are more likely to occur or where the geological information suggests they could be. This distinction is important in that when geophysical data do not support the presence of contrasts strong enough to constitute an interface, geophysical inversion does not allow this interface, and it is not observed in the inverted model. We have tested this, as illustrated in Figure 3 and 4 where we basically inserted a fictitious interface by splitting a lithology that is subsequently treated as 2 separate lithologies in the MCUE process. We then let inversion decide whether it creates an interface or anomalous body in the inverted model around the uncertain area generated by the corresponding interfaces. The result is quite clear and supports this claim. It needs, however, to be put into the context of the data inverted for. In our case, for gravity data, vertical interfaces may be well discriminated but it is likely not to be the case for mostly horizontal interfaces where it is known to show poor sensitivity. For such horizontal structures, we expect that model update will basically just follow $\boldsymbol{W_H}$ and to accommodate geophysical data as you mentioned.

Comment, paragraphs 3, p. C2

- I don't have any problem with this process as such. However, I'm uneasy with the connection between the regions of low W_H and your quantification of geological uncertainty. Okay, the process you create, which uses geological uncertainty to locate the low values of W_H, works. But this is because these areas of geological uncertainty (happen to?) correspond to where the boundaries between the geological units are: it's not the geological uncertainty that's the true, fundamental piece of information, it's that this uncertainty in the geological modelling is indicative of a boundary between units, and it's this estimated location of the boundary between units that becomes the key information to provide to the gravity (or magnetic) inversion via the low values of W_H.

Answer

This is true, but to be more accurate, in general, low values of $\boldsymbol{W_H}$ indicate areas of the model that are poorly constrained from a geological point of view. This comprises interfaces between 2 lithologies but also triple points and areas insufficiently 'illuminated' by geological data. In general, this is reflective of how well geology can resolve specific parts of the model, which corresponds mostly to interfaces, but can comprise whole geological units.

Added to section 2.3:

"Fundamentally, geological uncertainty contained in $\boldsymbol{H}$ encapsulates information about possible locations of interfaces between units and areas where the geological data is insufficiently informative."

Comment, paragraphs 1, p. C3

- What if you were to consider a synthetic example in which there is essentially zero uncertainty in the location of the interfaces. And make a W_H that's pretty much 1 everywhere except zero for the cells straddling the interfaces. I'd expect the gravity inversion would give a nice density model that pretty much has sharp interfaces right where the geological model has it's interfaces. If you then broaden up the zones of low values in W_H, I'd expect the boundaries in the density model to pretty much stay in the same location but now start to be smeared out and more like an inversion result for constant W_H. If you have a broad region of constant low W_H, the constructed density will be smeared out and smoothly varying through here, it won'd be sharp at one end or the other.

Answer

I agree that there will be a transition between well-constrained, sharp model and the smeared model where $\boldsymbol{W_H}$ = 1 everywhere (e.g., no geophysical information in the constraints). We expect that there will not be a linear transition from sharp to smooth models when varying geological input uncertainty. We would expect results similar to Giraud et al. 2018, who use a model relatively similar to this one and vary uncertainty from uninformed geological models to very well constrained geology. In their case, the 'smearing' of the inverted model happens at high uncertainty (much higher than the levels tested). This work shows a non-linear relationship between the level of 'smearing' and geological uncertainty. The inversion approach they use is different (local petrophysical constraints) but the observations can be extrapolated, to a certain extent, to the present case.

For the presented case, we expect that the uncertainty level where transition between well-constrained inversion and non-constrained inversion happens corresponds to very high geological input measurement uncertainty in MCUE that are beyond what can realistically be expected. To assess this, we ran a simulation of a case where k=6, which corresponds to very uncertain data. The result is shown in Figure 5 (reproduced below).

Comment

- If you have a broad region of constant low W_H, the constructed density will be smeared out and smoothly varying through here, it won'd be sharp at one end or the other.

Answer

Figure 5 above answers partly to this remark. It holds true for the parts of the model where structures are mostly horizontal, for which gravity inversion shows little sensitivity, but not where the causative density model is shows steep structures.

Generally speaking, this observation is dependent on the sensitivity of data. For instance, replacing gravity by seismic, the opposite would be true: vertical parts would be strongly influenced by the constraints while horizontal structures would be well resolved simply because, as a rule of thumb, seismic resolves horizontal structures better than gravity.

The following has been added to the conclusion, renamed 'concluding remarks'.

"The quantitative integration technique we presented reduces uncertainty and ambiguity compared to qualitative interpretation technique or non-integrated workflows. However, despite its robustness to misplaced interface (e.g., bias) or to high geological uncertainty (e.g., sparse or very uncertain geological input measurements) as shown in the synthetic case, interpreters need to bear in mind the specificities of the geophysical data inverted for (resolution of specific geometries, depth of investigation) and the shortcomings of geological modelling workflows. As for all geological modelling, MCUE is oblivious to geological units or faults that are not sampled by field geological measurements, which can lead to biases in final models due to, for instance, inclusions not be accounted for."

Comment

- And if you have the same true synthetic model but try putting the low values of W_H in the incorrect locations (i.e., not where the interfaces in the true model are) then the constructed density model is going to have it's interfaces (sharp or diffuse depending on whether W_H is sharp or diffuse) pretty much where the lows in W_H are, not where the boundaries are in the true model. (Have you tried such a suite of examples?)

As hinted above, the horizontal part would be impacted and what you describe would occur. On the contrary, the vertical structures will be mostly preserved since sensitivity of gravity data is maximal for vertical anomalies. As stated earlier, this depends on the sensitivity of the method and geometry of the model. From the results we obtained, we think that, as a rule of thumb, it is better to have too many possible interfaces than too few.

We ran a series of tests, two of which are now in the manuscript (section presenting the synthetic case study); as you recommended, we move the synthetic case study directly in the text.

Comment, paragraphs 2, p. C3:

- So, yes, using spatially variable weights in the roughness term results in the interfaces in the density (or susceptibility) models occurring where you'd like them to occur. However, I don't agree with the thoughts that the gravity (magnetic) inversion is helping out, or overcoming, the geological uncertainty; rather, the uncertainty is mapping out parts of the subsurface on and close to the interfaces, but that it's simply this (fuzzy) location of the interfaces that you're using to tell the gravity (or magnetic) inversion where to put the boundaries in the density (or susceptibility) model.

Answer

As discussed above, we have run the inversion using a $\boldsymbol{W_H}$ that proposes interfaces in places where there are none in the true density contrast model.

From a more philosophical point of view, inversion is not 'correcting' geological uncertainty; rather, it accommodates and exploits it to produce an inverted model accounting for it while honouring geophysical data.

Comment

- The above is my main issue with this manuscript: the justification and motivation, not the mechanics of the workflow itself.Some further comments ... I think there has to be a typo in equation 3: You've got "max H - max H" on the denominator, i.e., a difference between identical things.

Answer

Absolutely. This has been corrected.

Comment

- In equation 1, what are you using the prior model for? This isn't a "starting" model, is it? (You say the starting and prior are the same thing on page 7, and use "starting model", without a mention of the prior model, in the caption to Figure 2.) The m_p in equation 1 is important, as the inversion is going to try to construct a model that is close as possible to m_p (which is the whole point of that second term on the right-hand side of equation 1). This is a linear inverse problem, so it won't matter at all if one uses a starting model and then solves for the model update (or just solves directly for the model from the observed data). So a "starting" model should have no influence in the inversion. I'd definitely not use the term "starting" at all, and be careful to always use "prior" when thinking about the m_p in equation 1.

Thank you for pointing this out. We followed your suggestion and have addressed this issue. The equation was modified accordingly and added a short description of the term. We have also changed the semantics a little bit to avoid similar confusions around 'starting', 'reference' and 'prior' models. We made it clearer and ensured consistency with the rest of the text.

Comment

- Do you now use a trade-off parameter for the "model term" in equation 1? Maybe you do in the code but it's simply been omitted from this equation when writing this manuscript?

Answer

We do use a trade-off parameter. The way we wrote equation implicitly suggests that the trade-off parameter for the model term is included in the $\boldsymbol{W_m}$ term multiplying the values initially assigned for each cell in the model covariance matrix. For clarity we rewrote equation (1) by separating the term of which the norm is calculated to stick to more conventional formulation of the cost function; the scalar weight α_m was introduced.

Comment

- Figure 2 and Figures 3 a & b are fine to show the whole model. But it's hard to make out the details in the parts of the model around features A, B & C. It's therefore hard to assess how much of a difference the "locally conditioned regularization" has made. I think it would be good and important to also show zoomed-in sections through the parts of the model around A, B & C.

Answer

We have tried this but adding a zoomed-section of model parts around A, B and C requires its own figure since there is not enough space left on the figure as it is. The alternative we propose is the one in the new version of the manuscript, where we added a top view of greentone belts A, B, and C to the right of the full volume view. This is helpful in providing the reader with additional information about how locally constrained regularization constrain the model and change in areas around the assumed greenstone belts.

Comment

- Finally, the structural comment: I really like the example in Appendix A. I think that should come in the body of the manuscript, between sections 2 & 3 (perhaps with a description of the geological modelling process, and the process of determining the "geological uncertainty").

Answer

We have moved it to a new section before the field case study application. In this section we also included additional examples. We have also extended this section to show some of the inversions from the suite of runs you mentioned and asked us about. This is now section 3, which has 4 subsections.

References

Giraud Jeremie, Pakyuz-Charrier Evren, Ogarko Vitaliy, Jessell Mark, Lindsay Mark, Martin Roland (2018) Impact of uncertain geology in constrained geophysical inversion. ASEG Extended Abstracts 2018, 1-6, https://doi.org/10.1071/ASEG2018abM1_2F