Interactive comment on "DeepBedMap: Using a deep neural network to better resolve the bed topography of Antarctica" *by* Wei Ji Leong et al.

Wei Ji Leong et al.

weiji.leong@vuw.ac.nz

General comments

This paper introduces a new method, based on Machine Learning, namely a Generative Adversarial Network (GAN), to add short-scale roughness to the bed of Bedmap2. The paper is well written, easy to follow and well illustrated, I really enjoyed reading it. I recommend publication after minor revisions. My main problem while reading the manuscript was that I felt like the authors were overselling their approach and the performance of the GAN.

> What the GAN is doing is to essentially try to reintroduce basal rough-
> ness in the smooth bed of Bedmap2 based on surface features. While
> the method is different, the goal of this study is very similar to the paper
> of Graham et al. 2017 (www.earth-syst-sci-data.net/9/267/2017/) or Goff et
> al. 2017 (https://doi.org/10.3189/2014JoG13J200), papers that are barely
> mentioned in the text.

We thank the reviewer for their considered review and comments. Thank you for high-
lighting the work of Goff et al. (2014) and Graham et al. (2017). In regard to the
publication by Graham et al. (2017), we have actually compared their Synthetic HRES
product at some earlier conferences (see Leong and Horgan (2019a) and Leong and
Horgan (2019b)), but decided to focus on the newer BedMachine Antarctica product
for this manuscript. For completeness, we have now reproduced a 3D image of this
Synthetic HRES product here (see Fig 1), using the same Pine Island Glacier extent in
Fig. 3 of the manuscript.

We acknowledge that the goal of this paper is similar to the two aforementioned pa-
pers, and fall in the broad category of using spatial statistics to derive a higher spatial
resolution bed. Specifically, the conditional simulation method applied by Goff et al.
(2014) is able to resolve both fine-scale roughness and channelized morphology over
the complex topography of Thwaites Glacier, and make use of the fact that roughness
statistics are different between highland and lowland areas. Graham et al. (2017) uses
a two-step approach to generate their synthetic HRES grid, with the high frequency
roughness component coming from the ICECAP and Bedmap1 compilation radar point
data, and the low frequency component coming from BEDMAP2. In DeepBedMap,
we attempt to capture bed topography directly from gridded pixels, while incorporating
extra knowledge from satellite remote sensing datasets to fill in larger gaps between
flightlines, much like in BedMachine Antarctica (Morlighem et al., 2019). Neither one
method is perfect, and we see all of them as complementary.

Specific comments

Being able to capture both long wavelength and short wavelength bed features is the goal. We do however rely on the BEDMAP2 surface as a reference for this super resolution task, which limits the generated topography to within a tolerance of the surface. If we don't use BEDMAP2, then the modelled bed elevation could diverge significantly from the actual bed elevation. Ideally we would be able to run the model independent of BEDMAP2, however, this would no longer be a super resolution model.

Note that the provided DeepBedMap DEM model is only one 'possible' version, generated from one model training run we deemed best according to our training metric, and we may have biased our model towards resolving short wavelength features, compared to BedMachine Antarctica which recovers large scale features like ridges and valleys well. That is not to say we cannot combine super resolution with inversion techniques, and as mentioned in text, the DeepBedMap model architecture should be applicable to any reference bed, be it BedMachine Antarctica or the upcoming BEDMAP3.

This is a valuable exercise and using machine learning to do this is definitely a good idea and worth publishing, but I don't think we are there yet. The training dataset is extremely small and probably not representative of all
the different types of terrains under the Antarctic ice sheet (as mentioned by the authors).

There is certainly more work to do on both the modelling and data collection side (see our reply to Reviewer 1). It should be mentioned though that bed interpolation exercises such as ours and BedMachine Antarctica help tell us where the data gaps are. As more datasets are gathered from targeted acquisitions, marine swath bathymetry, etc, these method will become even more powerful.

We see a lot of artifacts in the solution and many of these artifacts are discussed in the text: dunes and missing mountains around Byrd (4h), Terraces (4i), Speckle (4a), etc. In the maps of figure 4, I could not find a bed that seemed realistic.

In Figure 4 of the manuscript, we have chosen to highlight different locations, some of which are unrealistic as acknowledged in the text. The example we provide in our reply to Reviewer 1 (see Fig. 1 there) provides an example of a realistic bed as does Fig 5e over the non-mountainous areas of Rutford Ice Stream.

If we able to quantify precisely what is wrong with the generated bed topography, this can be incorporated into the Discriminator component of the Generative Adversarial Network. Currently we use a basic Discriminator designed for standard computer vision tasks. That is not to say that we cannot incorporate glaciology specific criteria such as ice flow direction into the Discriminator model design, which would push the Generator model to produce more realistic results. Alternatively, we can adjust the loss function weights to dampen the effects of the REMA ice surface elevation input, as our model may have overfitted to the REMA surface DEM.

> Even along the flight line of OIB (figure 6) the roughness of DeepBedMap seems exacerbated and not necessarily representative of the actual roughness measured by the radar. And again, the authors make it clear, I just find the title/abstract and parts of the paper a bit misleading in the sense that I don't think this approach achieves the objectives of this work, and that's ok! I would not say that the GAN "better resolves" the bed topography for example.

We may have been overly enthusiastic in some of our language and will do our best to temper this in revision. In regard to roughness, our neural network model was trained by minimizing the error between the generated bed elevation and the bed elevation of the groundtruth training data, rather than the roughness parameter which is a derived statistic. Incorporating roughness into the loss function would be a useful exercise. "Better" is indeed a subjective term that is dependent on the current baseline, and we will consider using another title for the formal publication.

> Another problem is that it is not straightforward to constrain the model with radar data, and this is not mentioned in the text. The roughness of the bed that is captured (and known) by the radar data along flightlines cannot be preserved. This is an important limitation.

We agree that the pixel-based DeepBedMap model is unable to constrain itself easily to point-based radar data. The along track resolution of radar bed picks are much smaller than the 250 m pixels, and it it not easy to preserve roughness from radar unless smaller pixels are used. This may change once we start using swath radar data for training instead of interpolating our own grid from radar point data collected along flightlines.

> I also did not understand the paragraph line 204-205: why would we use the inferred bed under ice shelves when clearly surface features do not reflect the shape of the bathymetry? It is not because the authors "can" do it that they should do it.

The intention was to provide a means for others to more easily interpolate their own bathymetry grid with the DeepBedMap grid. There is a choice of different grounding lines, and rather than enforce one, we would prefer to let others cut and blend it with their own bathymetry dataset, smoothed out over any selected distance. We now intend to provide a mask file with the final product, allowing the user to apply this ice shelf mask directly, or use one of their own. We will also clarify this intention better in text so as not to suggest that we have managed to super resolve the under ice shelf bathymetry.

> That said, much of these issues can be addressed in future work. I still think this is a good piece of work and look forward to seeing the modified version.

There is always potential to improve this work further, and one that we have faced over the year developing this methodology, with better techniques and new data coming in all the time. Hopefully this paper can serve as a good starting point, and we are excited to see what others will come up with in the future.

Technical corrections

> Other than that, the paper is easy to follow and really well written, I only found one typo: line 297: care has been taking → taken

Once again, thank you very much for your constructive feedback.

References

Fretwell, P., Pritchard, H. D., Vaughan, D. G., Bamber, J. L., Barrand, N. E., Bell, R., Bianchi, C., Bingham, R. G., Blankenship, D. D., Casassa, G., Catania, G., Callens, D., Conway, H., Cook, A. J., Corr, H. F. J., Damaske, D., Damm, V., Ferraccioli, F., Forsberg, R., Fujita, S., Gim, Y., Gogineni, P., Griggs, J. A., Hindmarsh, R. C. A., Holmlund, P., Holt, J. W., Jacobel, R. W., Jenkins, A., Jokat, W., Jordan, T., King, E. C., Kohler, J., Krabill, W., Riger-Kusk, M., Langley, K. A., Leitchenkov, G., Leuschen, C., Luyendyk, B. P., Matsuoka, K., Mouginot, J., Nitsche, F. O., Nogi, Y., Nost, O. A., Popov, S. V., Rignot, E., Rippin, D. M., Rivera, A., Roberts, J., Ross, N., Siegert, M. J., Smith, A. M., Steinhage, D., Studinger, M., Sun, B., Tinto, B. K., Welch, B. C., Wilson, D., Young, D. A., Xiangbin, C., and Zirizzotti, A.: Bedmap2: Improved Ice Bed, Surface and Thickness Datasets for Antarctica, The Cryosphere, 7, 375–393, https://doi.org/10.5194/tc-7-375-2013, 2013.

Goff, J. A., Powell, E. M., Young, D. A., and Blankenship, D. D.: Conditional Simulation of Thwaites Glacier (Antarctica) Bed Topography for Flow Models: Incorporating Inhomogeneous Statistics and Channelized Morphology, Journal of Glaciology, 60, 635–646, https://doi.org/10.3189/2014JoG13J200, 2014.

Graham, F. S., Roberts, J. L., Galton-Fenzi, B. K., Young, D., Blankenship, D., and Siegert, M. J.: A High-Resolution Synthetic Bed Elevation Grid of the Antarctic Continent, Earth System Science Data, 9, 267–279, https://doi.org/10.5194/essd-9-267-2017, 2017.

Leong, W. J. and Horgan, H. J.: DeepBedMap: Using a Deep Neural Network to Better Resolve the Bed Topography of Antarctica, in: EGU General Assembly, Vienna, Austria, 2019a.

Leong, W. J. and Horgan, H. J.: DeepBedMap: A Super-Resolution Deep Neural Network for Resolving the Bed Topography of Antarctica, in: IGS Five Decades of Radioglaciology Symposium, Stanford, California, United States, 2019b.

Morlighem, M.: MEaSUREs BedMachine Antarctica, Version 1, https://doi.org/10.5067/C2GFER6PTOS4, 2019.

Morlighem, M., Rignot, E., Binder, T., Blankenship, D., Drews, R., Eagles, G., Eisen, O., Ferraccioli, F., Forsberg, R., Fretwell, P., Goel, V., Greenbaum, J. S., Gudmundsson, H., Guo, J., Helm, V., Hofstede, C., Howat, I., Humbert, A., Jokat, W., Karlsson, N. B., Lee, W. S., Matsuoka, K., Millan, R., Mouginot, J., Paden, J., Pattyn, F., Roberts, J., Rosier, S., Ruppel, A., Seroussi, H., Smith, E. C., Steinhage, D., Sun, B., van den Broeke, M. R., van Ommen, T. D., van Wessem, M., and Young, D. A.: Deep Glacial Troughs and Stabilizing Ridges

Unveiled beneath the Margins of the Antarctic Ice Sheet, Nature Geoscience, https://doi.org/10.1038/s41561-019-0510-8, 2019.
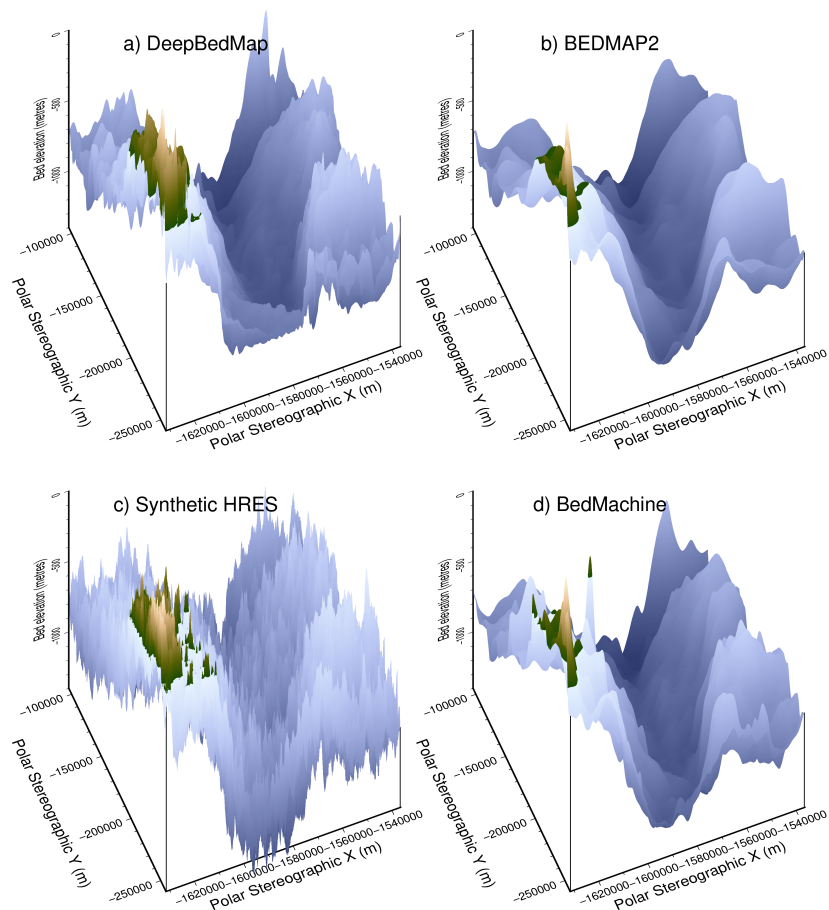
Fig. 1. Comparison of interpolated bed elevation grid products over Pine Island Glacier. a)
DeepBedMap (ours) at 250 m resolution. b) BEDMAP2. c) Synthetic HRES product. d) Bed-
Machine Antarctica.