Antarctic Bed Topography Super-Resolution via Transfer Learning

Kim Bente
School of Computer Science
The University of Sydney
Sydney, Australia
kim.bente@sydney.edu.au

Roman Marchant
Human Technology Institute
University of Technology Sydney
Sydney, Australia
roman.marchant@uts.edu.au

Fabio Ramos
NVIDIA
Seattle, USA
School of Computer Science
The University of Sydney
Sydney, Australia
fabio.ramos@sydney.edu.au

Abstract

High-fidelity topography models of the bedrock underneath the thick Antarctic ice sheet can improve scientists' understanding of ice flow and its contributions to sea level rise. However, the bed topography of Antarctica is one of the most challenging surfaces on earth to map, requiring aircrafts with ice penetrating radars to survey the vast and remote continent. We propose FROST, Fusion Regression for Optimal Subglacial Topography, a method that leverages readily available surface topography data from satellites as an auxiliary input modality for bed topography super-resolution. FROST uses a non-parametric Gaussian Process model to transfer local, non-stationary covariance patterns from the ice surface to the bedrock. In a controlled topography reconstruction experiment over complex East Antarctic terrain, our proposed method outperforms bicubic interpolation at all five tested magnification factors, reducing RMSE by 67% at $\times 2$, and 25% at $\times 6$ magnification. This work demonstrates the opportunity for data fusion methods to advance downstream climate modelling and steward climate change adaptation.

1 Introduction

'What happens in the polar regions, doesn't stay in the polar regions.' With average warming in the polar regions exceeding global mean warming (polar amplification) [1], [2], the Greenland and Antarctic ice sheets are expected to be major contributors to sea level rise [2], [3]. Greenland is estimated to hold 7.4 m SLE (sea level equivalent; increase in average global sea level if completely melted) [4], whereas the Antarctic ice sheet is estimated to hold 57.4 m SLE [5]. While there is considerable disagreement between projections from Antarctic ice sheet models [6], [7], sea level rise poses a threat to the large population living in coastal and low-lying regions, and thus reliable projections are critical to guide climate change adaptation strategies.

Meanwhile remote sensing has accelerated progress in environmental monitoring and climate modelling [8]. Satellite are instrumental in data collection for remote regions like Antarctica, and they enable monitoring at sufficient spatial and temporal scales. While this increase in data, particularly in near-surface data, opens up many pathways to impact for machine learning in climate [2], [9],

Tackling Climate Change with Machine Learning: workshop at NeurIPS 2023.

mapping Antarctica's subglacial terrain remains logistically and technologically challenging [5], [10]. Nonetheless, gridded maps of the Antarctic bed topography as well as the distinct local features it reveals (e.g. stabilising ridges or inland sloping [5]), are a fundamental input to many downstream Antarctic climate studies [10]. Due to data limitations, the fidelity of Digital Elevation Models (DEMs) of the subglacial landscapes remains a large source of uncertainty in ice volume estimation and ice flow models [11]. International project groups have been assimilating data from ongoing field campaigns of mostly airborne radar surveys to produce coherent topography grids of Antarctica like BedMachine [5], [12], Bedmap2 [13], or the recently announced Bedmap3 [10].

2 Multimodal bed topography super-resolution

Pushing beyond the limited capacity to collect field data during resource-intensive sensing expeditions, machine learning methods have been proposed [14], [15] to fill ice sheet modellers' needs for continent-scale, accurate, and detailed bed topography grids, given the data available. Super-resolution (SR) is an established computer vision task, which aims to enhance the resolution of a given input image. Viewing a topography grid as a 2D image indexed by x and y coordinates where elevation values are the pixel values, the task is to learn the mapping f from the low-resolution (lr) bed topography image b_{lr} to the high-resolution (hr) image b_{hr}, where $b_{hr} = f(b_{lr})$.

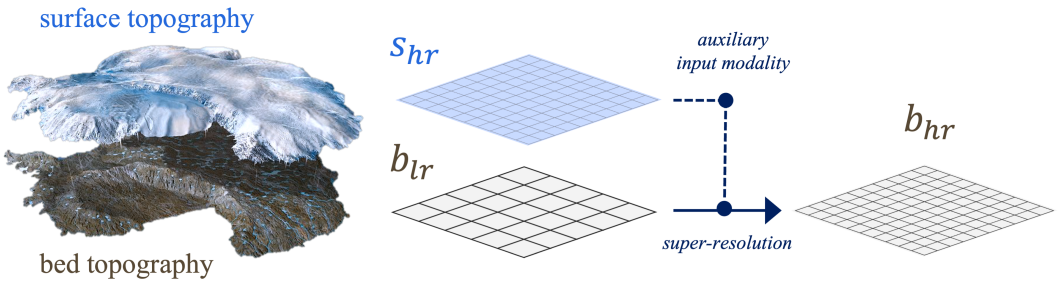

Figure 1: Schematic overview of multimodal bed topography super-resolution for Antarctica

Traditional single image SR, which takes the lr image as the sole input, is a challenging **inversion problem** with a poorly constrained solution space, particularly for data sparse domains. Probabilistic machine learning methods address this by estimating a probability distribution over reconstructions. The spread of the predictive distribution represents the level of confidence. Additional input modalities, if available and related to the target modality, can aid SR. Pan-sharpening is a **multimodal SR** technique commonly used in remote sensing, that merges multi-spectral images with spatial detail from hr panchromatic images [16]. Another multimodal method [17] adopts a probabilistic perspective and proposes a Gaussian Process (GP) framework to upscale[1] infrared images of farmland scenes by borrowing the covariance from multi-spectral hr drone images of the same scene.

Whereas previous machine learning methods only have access to limited training data at target resolution [14], [15] and thus produce biased estimates and artifacts, particularly in East Antarctica which holds most of the continents' ice, we adopt the multimodal Bayesian approach introduced by Reid, Ramos, and Sukkarieh [17]. Rather than using an alternative sensing modality that measures the same underlying landscape, we leverage surface elevation values s, available at very high resolutions from satellites, as the auxiliary hr input to assist bed topography SR, see Figure 1. While we presume a relationship between ice surface and bedrock, the challenge is to model the non-linear, non-stationary and non-deterministic relationship to achieve faithful hr bedrock representations.

3 FROST: Fusion Regression for Optimal Subglacial Topography

We propose FROST (acronym for *Fusion Regression for Optimal Subglacial Topography*), an adaptation of the multimodal Gaussian Process (GP) framework introduced by Reid, Ramos, and Sukkarieh (2013) [17]. FROST infers hr bed topography from the lr bed topography input as well as the

[1]We use the term *upscaling* to denote resolution increase, which is common computer vision terminology. This is contrary to terminology used in the climate community where *downscaling* describes resolution increase.

auxiliary hr surface topography input, with both inputs corresponding to the same area, see Figure 1. At the core of this kernel-based method is the design of the covariance function k that facilitates transfer learning. Rather than defining the covariance between pairs of points, we define covariance between pairs of grid cells i.e. pixels (P_{hr}, P'_{hr}) [17]. Furthermore, k is defined in the modality of the auxiliary input, the ice surface elevation values $s(P_{hr})$, to enable the GP to learn non-stationary covariance patterns at target resolution. The covariance matrix computed from the auxiliary variable acts as a high-resolution covariance substitute in the GP interpolation of the hr bed elevation values.

$$k\left(P_{hr}, P'_{hr}\right) = k_s\left(P_{hr}, P'_{hr}\right) \circ k_{aux}\left(s\left(P_{hr}\right), s\left(P'_{hr}\right)\right) \times \sigma_f^2 \tag{1}$$

The covariance function k is composed of a spatial covariance term k_s and the auxiliary covariance term k_{aux}, see Equation 1. The spatial covariance term k_s creates sparsity in k by controlling a threshold distance, defined by a hyperparameter, at which the covariance of pixel pairs reduces to zero. The sparsity yields computational efficiencies, especially in the inversion of large covariance matrices required for GP regression. k_s is implemented as a piece-wise polynomial kernel between spatial locations of grid cell mid-points, and thus acts as a **sparse, smooth**, and **local filter** for the second covariance term, k_{aux}. The auxiliary covariance term k_{aux} introduces **non-stationary behaviour**, critical for elevation modelling. A squared exponential kernel is used to compute the similarities between pairs of surface elevation values $(s(P_{hr}), s(P_{hr})')$, with another hyperparameter controlling the sensitivity. Hereby, distinct sub-input-grid features like discontinuities or roughness, present on the surface, can be captured in the covariance matrix. Finally the positive definite covariance matrix is computed as the element-wise product of k_s and k_{aux} and then scaled by σ_f^2. Refer to Figure 4 in the Appendix for an illustration of the covariance composition.

$$\mu^* = \widehat{b}\left(P_{hr}\right) = k\left(P_{hr}, P_{lr}\right)\left[k\left(P_{lr}, P_{lr}\right) + \sigma_n^2 I\right]^{-1}\left(b\left(P_{lr}\right) - \mu_b\right) + \mu_b \tag{2}$$

Within FROST's GP framework, the predictive mean function μ^*, see Equation 2, generates hr bed topography estimates, where hyperparameter σ_n^2 denotes the noise term and μ_b denotes the bed elevation mean constant. With k defined over hr pixel pairs, covariances $k\left(P_{hr}, P_{lr}\right)$ and $k\left(P_{lr}, P_{lr}\right)$ are derived by averaging those hr covariance values that the larger, lr pixels are spatially composed of. Covariance hyperparameters including σ_f are optimised for each magnification factor with stochastic gradient decent, by minimising the negative log marginal likelihood (NLML). Please refer directly to [17] and the GP literature [18] for further methodological details. Figure 2 demonstrates how

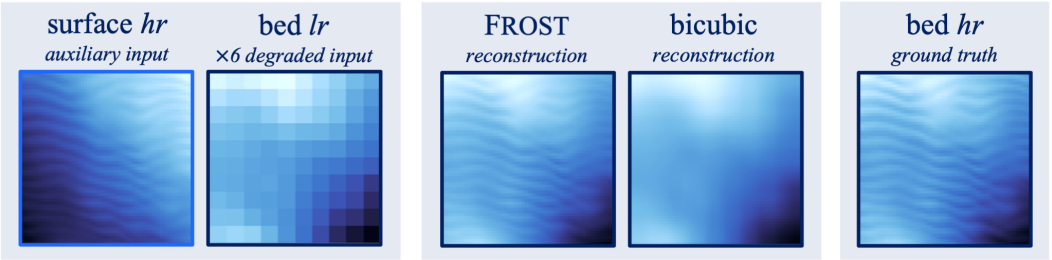

Figure 2: Inputs, $\times 6$ reconstructions and ground truth bed topography for exemplary scene in East Antarctica. Lighter hues indicate higher elevation.

FROST is able to transfer fine-scale, intricate patterns from the ice surface to the hr estimate of the bed topography. Although surface slope and bedrock slope are merely related at image level in this example, wave-like ridges correspond between surface and glacier base. The basal ridges are not visible in lr so that smooth bicubic interpolation fails to reconstruct these features. With well calibrated hyperparameters, our methods improves over standard single image interpolation.

4 Preliminary experiments

In this section we present controlled bed topography image reconstruction experiments at magnification factors $\times 2, \times 3, \times 4, \times 5$ and $\times 6$ for $N = 400$ images from the East Antarctic interior. Hereby we seek to evaluate FROST's super-resolution skill in comparison to baselines in a controlled setting with known ground truth, to determine its potential to produce faithful topography grids at higher resolutions than existing data products and to fill data voids. We use the MEaSUREs BedMachine

Antarctica, Version 3 [5], [12] data set, containing bed elevation and surface elevation data on a square 500 m grid. As a non-trivial test bed we select a region with complex topography and high ice thickness in the East Antarctic interior, around Dome A (see Figure 5 for the location). We then divide the domain into 400 images, each consisting of 60 x 60 pixels. The ground truth bed elevation images, b_{hr}, are degraded to five different coarse resolutions, producing input images b_{lr}. This results in reconstruction tasks of increasing difficulty, with magnification factors $m \in [2, 3, 4, 5, 6]$. For example, for $m = 6$, we are upscaling a low-resolution 10×10 pixel image to reconstruct a high-resolution 60×60 image, see Figure 2. The auxiliary input, surface elevation s_{hr}, is also given at 60×60 target resolution. The aim is to reconstruct b_{hr} given b_{lr} and s_{hr}. We use two qualitative metrics, Root Mean Square Error (RMSE) and Peak Signal-to-Noise Ratio (PSNR) to compare mean reconstruction errors of our proposed method to bilinear and bucubic interpolation baselines. Table 1 contains results of our experiments. Please also refer to Figure 3 for plots of the reported results and Appendix 5 for more information on metrics.

Table 1: Results from bed topography reconstruction experiments comparing FROST to baselines at 5 magnification factors. We report RMSE ↓ (lower is desired) and PSNR ↑ (higher is desired).

magnification	×2		×3		×4		×5		×6	
metric	RMSE	PSNR	RMSE	PSNR	RMSE	PSNR	RMSE	PSNR	RMSE	PSNR
Bilinear	4.94	47.6	9.64	41.8	15.8	37.5	21.7	34.7	28.4	32.3
Bicubic	3.18	50.8	5.99	45.7	10.2	41.2	14.8	38.0	20.1	35.4
FROST (ours)	**1.03**	**61.1**	**3.52**	**51.0**	**6.71**	**45.8**	**10.6**	**41.8**	**14.9**	**38.8**

5 Discussion

We present FROST, a probabilistic super-resolution model that addresses the mapping challenge of Antarctica's hidden bed terrain, by aiming to increase the fidelity of existing bed topography grids. In our experiments we upscale gridded lr bed elevation images up to a magnification factor of $\times 6$ (i.e. from 3000 m to 500 m resolution) using widely available surface elevation measurements as an auxiliary hr input. Despite high ice thickness separating surface and bedrock in the experiment's domain, our method is able to transfer local covariance patterns and thereby outperform interpolation baselines: We reduce bicubic RMSE by over 67% for $\times 2$ magnification and 25% for $\times 6$ magnification. These preliminary results demonstrate a proof of concept for our multimodal GP framework. Further experiments are needed to compare our proposed method against more sophisticated baselines on a continental scale. Future work should therefore explore an extension to non-stationary GP hyperparameters as well as the use of other probabilistic and meta-learning approaches like Neural Processes [19], [20]. These techniques may enable the model to adapt to different terrain types (e.g. mountains, basins, coastal areas) without local hyperparameter tuning. Additional satellite observed modalities, like ice surface velocity used in [14], could be incorporated and combined with physics-guided approaches for robust inversion.

Deep learning methods, including CNNs and more recently Transformer-based methods like HAT [21], [22] or SWIN [23], [24], are clearly dominating single image super-resolution benchmarks. Further work should assess their performance on bed topography super-resolution tasks and adapt these state-of-the-art methods to the challenges of the cryosphere domain including data limitations. Training on more accessible topographic data sets may benefit domain adaptation. In contrast to GAN-based DeepBedMap [14], which operates on a fixed set of input modalities to super-resolve Antarctic bed topography, kernel methods have the ability to quantify uncertainty, adapt more easily to new magnification factors and inputs, and provide an interpretable, flexible and robust covariance mechanism. Non-parametric models like FROST also do not rely on target-resolution gridded training data as deep learning models do. Nonetheless, fine-tuning approaches, utilizing embeddings from pre-trained computer vision models, and adaptations to state-of-the-art super-resolution architectures represent interesting research directions at the intersection of machine learning and climate. In conclusion, both deep learning and Gaussian Process models have vast opportunity to exploit diverse remote sensing data modalities for topography super-resolution and to thereby enhance ice sheet modelling, improve resulting sea level rise projections and inform climate adaptation strategies.

Acknowledgments and Disclosure of Funding

This research was supported by an Australian Government Research Training Program (RTP) Scholarship, as well as through the Australian Research Council's Industrial Transformation Training Centre in Data Analytics for Resources and Environments (DARE) (project ID IC190100031).

References

[1] D. M. Smith, J. A. Screen, C. Deser, *et al.*, "The polar amplification model intercomparison project (PAMIP) contribution to CMIP6: Investigating the causes and consequences of polar amplification," *Geoscientific Model Development*, vol. 12, no. 3, pp. 1139–1164, 2019.

[2] D. Rolnick, P. L. Donti, L. H. Kaack, *et al.*, "Tackling climate change with machine learning," *ACM Computing Surveys*, vol. 55, no. 2, 42:1–42:96, 2022.

[3] IPCC, "Summary for policymakers," in *IPCC Special Report on the Ocean and Cryosphere in a Changing Climate [H.-O. Pörtner, D.C. Roberts, V. Masson-Delmotte, P. Zhai, M. Tignor, E. Poloczanska, K. Mintenbeck, A. Alegría, M. Nicolai, A. Okem, J. Petzold, B. Rama, N.M. Weyer (eds.)]*. Cambridge, UK and New York, NY, USA: Cambridge University Press, 2022, pp. 3–35.

[4] M. Morlighem, C. N. Williams, E. Rignot, *et al.*, "BedMachine v3: Complete bed topography and ocean bathymetry mapping of greenland from multibeam echo sounding combined with mass conservation," *Geophysical Research Letters*, vol. 44, no. 21, pp. 11, 051–11, 061, 2017.

[5] M. Morlighem, E. Rignot, T. Binder, *et al.*, "Deep glacial troughs and stabilizing ridges unveiled beneath the margins of the antarctic ice sheet," *Nature Geoscience*, vol. 13, no. 2, pp. 132–137, 2020.

[6] H. Seroussi, S. Nowicki, A. J. Payne, *et al.*, "ISMIP6 antarctica: A multi-model ensemble of the antarctic ice sheet evolution over the 21st century," *The Cryosphere*, vol. 14, no. 9, pp. 3033–3070, 2020.

[7] H. Seroussi, S. Nowicki, E. Simon, *et al.*, "InitMIP-antarctica: An ice sheet model initialization experiment of ISMIP6," *The Cryosphere*, vol. 13, no. 5, pp. 1441–1471, 2019.

[8] M. Reichstein, G. Camps-Valls, B. Stevens, *et al.*, "Deep learning and process understanding for data-driven earth system science," *Nature*, vol. 566, no. 7743, pp. 195–204, 2019.

[9] C. Monteleoni, G. A. Schmidt, F. Alexander, *et al.*, "Climate informatics," in *Computational Intelligent Data Analysis for Sustainable Development*, Chapman and Hall/CRC, 2013.

[10] A. Fremand, H. Pritchard, P. Fretwell, and J. Bodart, "Bedmap3: New data and gridded products of antarctic ice thickness, surface and bed topography," Copernicus Meetings, EGU23-13665, 2023.

[11] D. Farinotti, M. Huss, J. J. Fürst, *et al.*, "A consensus estimate for the ice thickness distribution of all glaciers on earth," *Nature Geoscience*, vol. 12, no. 3, pp. 168–173, 2019.

[12] M. Morlighem, *MEaSUREs BedMachine antarctica, version 3*, Boulder, Colorado USA. NASA National Snow and Ice Data Center Distributed Active Archive Center., 2020.

[13] P. Fretwell, H. D. Pritchard, D. G. Vaughan, *et al.*, "Bedmap2: Improved ice bed, surface and thickness datasets for antarctica," *The Cryosphere*, vol. 7, no. 1, pp. 375–393, 2013.

[14] W. Leong and H. Horgan, "DeepBedMap: A deep neural network for resolving the bed topography of antarctica," *Cryosphere*, vol. 14, no. 11, pp. 3687–3705, 2020.

[15] Y. Cai, F. Wan, S. Lang, X. Cui, and Z. Yao, "Multi-branch deep neural network for bed topography of antarctica super-resolution: Reasonable integration of multiple remote sensing data," *Remote Sensing*, vol. 15, no. 5, p. 1359, 2023.

[16] G. Vivone, M. Dalla Mura, A. Garzelli, *et al.*, "A new benchmark based on recent advances in multispectral pansharpening: Revisiting pansharpening with classical and emerging pansharpening methods," *IEEE Geoscience and Remote Sensing Magazine*, vol. 9, no. 1, pp. 53–81, 2021.

[17] A. Reid, F. Ramos, and S. Sukkarieh, "Bayesian fusion for multi-modal aerial images," in *Robotics: Science and Systems IX*, Robotics: Science and Systems Foundation, 2013.

[18] C. E. Rasmussen, "Gaussian processes in machine learning," in O. Bousquet, U. von Luxburg, and G. Rätsch, Eds., Berlin, Heidelberg: Springer, 2004, pp. 63–71.

[19] S. Jha, D. Gong, X. Wang, R. E. Turner, and L. Yao, *The neural process family: Survey, applications and perspectives*, 2022. arXiv: 2209.00517.

[20] M. Garnelo, J. Schwarz, D. Rosenbaum, *et al.*, *Neural processes*, 2018. arXiv: 1807.01622.

[21] X. Chen, X. Wang, J. Zhou, Y. Qiao, and C. Dong, "Activating more pixels in image super-resolution transformer," in *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Vancouver, BC, Canada: IEEE, 2023, pp. 22 367–22 377.

[22] X. Chen, X. Wang, W. Zhang, *et al.*, *HAT: Hybrid attention transformer for image restoration*, 2023. arXiv: 2309.05239.

[23] M. V. Conde, U.-J. Choi, M. Burchi, and R. Timofte, "Swin2sr: SwinV2 transformer for compressed image super-resolution and restoration," in *Computer Vision – ECCV 2022 Workshops*, L. Karlinsky, T. Michaeli, and K. Nishino, Eds., Cham: Springer Nature Switzerland, 2023, pp. 669–687.

[24] Z. Liu, H. Hu, Y. Lin, *et al.*, "Swin transformer v2: Scaling up capacity and resolution," in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, ISSN: 2575-7075, 2022, pp. 11 999–12 009.

[25] I. M. Howat, C. Porter, B. E. Smith, M.-J. Noh, and P. Morin, "The reference elevation model of antarctica," *The Cryosphere*, vol. 13, no. 2, pp. 665–674, 2019.

Appendix

Data availability

In this paper we use the MEaSUREs BedMachine Antarctica, Version 3 [5], [12] data set. The data set can be downloaded from from the US National Snow and Ice Data Center, using the following link: `https://nsidc.org/data/nsidc-0756/versions/3` (last accessed 28/11/2023).

BedMachine's bed elevation data was inferred using mainly ice thickness measurements from airborne ice-penetrating radar surveys. BedMachine's surface elevation data originate from the Reference Elevation Model of Antarctica (REMA) [25], a satellite-derived digital elevation model (DEM) that originally is available at very high resolution < 10 m. Please refer to BedMachine's user guide for more details: `https://nsidc.org/sites/default/files/documents/user-guide/nsidc-0756-v003-userguide.pdf`.

Metrics and baselines

We use Root Mean Square Error (`RMSE`) and Peak Signal-to-Noise Ratio (`PSNR`) to evaluate the performance of our proposed method and to compare it against baselines. While other computer vision methods focus on perceptual quality, we focus on subjective reconstruction error, as bed topography images are intended as data input for use in downstream studies rather than as images for human perception.

The Root Mean Square Error (`RMSE`) between the ground truth bed topography b_{hr} and corresponding reconstruction $\hat{b}_{hr}$ is defined by Equation 3, where an image consists of n pixels.

$$\texttt{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(b_{hr} - \hat{b}_{hr})^2} \tag{3}$$

Peak Signal-to-Noise Ratio (`PSNR`) is defined by Equation 4, where R represents the maximum signal strength. For topography reconstruction we set R to the range of elevation values per image, as an upper bound for the error signal. Thereby `PSNR` normalises errors with respect to the spread of elevation values in each specific image.

$$\texttt{PSNR} = 20 \cdot \log_{10}\left(\frac{R}{\texttt{RMSE}}\right) \tag{4}$$

Figure 3 shows plots of the experimental results reported in Table 1.

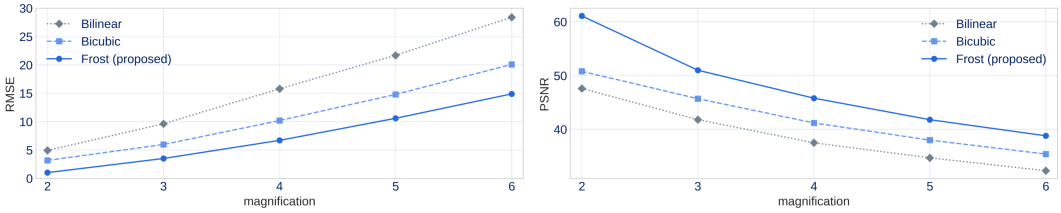

Figure 3: Plots of reconstruction RMSE (left, lower is desirable) and PSNR (right, higher is desirable) corresponding to Table 1

Implementation details and hyperparameters

As the degradation operator we apply an average pooling convolution (`torch.nn.AvgPool2d` in PyTorch) to simulate the hr to lr transformation. Scaling is only performed after degradation to avoid information

leakage from ground truth data to the lr inputs. We implement bilinear and bicubic interpolation baselines using PyTorch's `torch.nn.functional.gridsample` function. The baseline methods only take b_{lr} as input and they do not require the auxiliary surface modality.

The following hyperparameters, optimised for each magnification factor, were used to generate reconstruction with FROST:

Magnification factor	λ_s	λ_{aux}	σ_f^2
2	0.17	6.2	16.64
3	0.20	6.3	16.64
4	0.28	9.0	16.64
5	0.32	11.0	16.64
6	0.32	11.0	16.64

Figure 4 visualises components and composition of the point-to-plane covariance of an exemplary image.

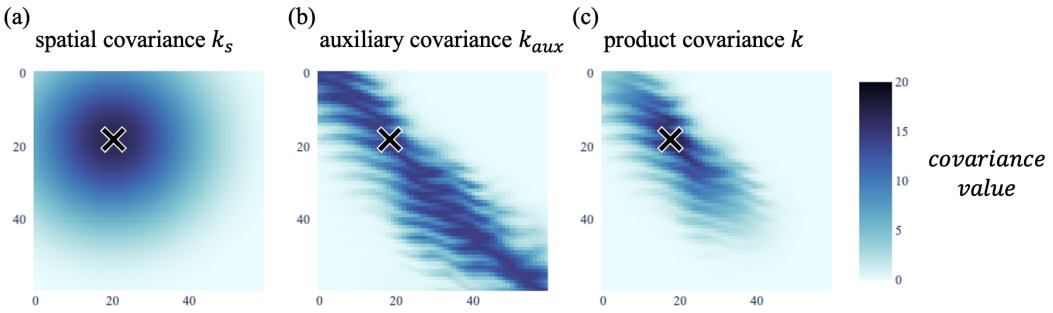

Figure 4: Exemplary composition of covariance matrix following Equation 1. We show the covariance matrices between a point (cross) and all other points in the image to represent the relationship in 2D. The selected image is the same as in Figure 2

Domain in East Antarctica

The domain selected for the experiment is a 600 × 600 km area around Dome A in the East Antarctic interior, located between the South Pole and the Amery Ice Shelf. The domain contains 20 × 20 (400) images with each image spanning 30 × 30 km. The domain lies between [898,750, 1,498,750] projected x coordinate (in meters), and [-250, 599,750] projected y coordinate (in meters), see 5. The projected coordinate system used in the MEaSUREs BedMachine Antarctica, Version 3 data set [5], [12] is WGS 84 / Antarctic Polar Stereographic (EPSG code: 3031). The median bed elevation of the domain is 747 meters, the median surface elevation is 3560 meters and the median ice thickness is 2679 meters. This domain does not contain any ice-free area, but only grounded ice (ice sheets), which makes it a suitable test bed of the proposed method.

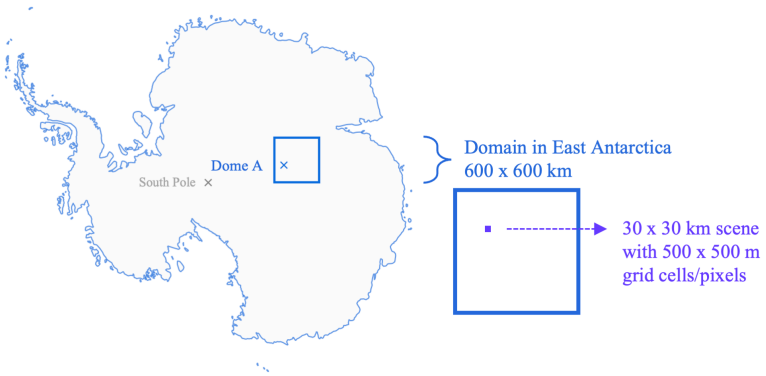

Figure 5: Location of experiment domain in East Antarctica