GT-IBTR-3D: Graph Transformer for 3D Ice-Bed Topography Reconstruction

Anonymous Author(s)
Affiliation
Address
email

Abstract

Accurately mapping subglacial bed topography is critical for understanding ice dynamics and climate impacts. Ice-penetrating radar provides direct returns from the bed, yet reliable bed picking remains difficult due to low signal-to-interference-and-noise ratio (SINR), strong spatial variability in subglacial relief, and acquisition artifacts. We propose GT-IBTR-3D, graph transformer for 3D ice-bed topography reconstruction. GT-IBTR-3D is a geometric deep learning approach that represents surface observations as graphs and reconstructs 3D ice-bed topography with a graph transformer. It combines the GraphSAGE inductive framework for localized structure with transformer layers for long-range dependencies, and it reformulates the prediction target from absolute bed elevation to surface-to-bed thickness. The bed is then recovered by subtracting the predicted thickness from the observed surface. This design isolates informative surface variation while avoiding outliers in full radargrams, stabilizes training by removing scene-level offsets, and unifies local and global reasoning. On Canadian Arctic Archipelago data, GT-IBTR-3D achieves substantially lower mean absolute error than prior probabilistic graphical methods and traditional deep neural networks.

1 Introduction

Accurate knowledge of Antarctic bed topography is essential for building ice sheet models, estimating ice volume, and projecting future sea level rise [2, 10]. Observations indicate rising Antarctic ice mass loss, driven largely by fast flow, retreat, and rapid thinning of major West Antarctic glaciers. [8, 11].

Airborne ice penetrating radar provides direct subglacial measurements by transmitting through thick ice and capturing two-dimensional cross-sectional images of ice sheets as radargrams. However, reliable ice bed topography reconstruction remains difficult because bed echoes are often weak due to a low signal to interference and noise ratio, subglacial topography varies sharply across regions, which requires robustness to spatial heterogeneity, and artifacts from system noise, surface clutter, and processing can obscure or mimic the true interface. Modern radar surveys also produce extensive datasets, which call for fast and scalable algorithms for automatic ice bed boundary detection.

Researchers have proposed many automatic probabilistic graphical models to identify ice-bedrock boundaries from radargrams[1, 6, 14, 7]. Crandall et al.[1] introduced a Markov random framework that automatically delineates the ice–air and ice–bedrock interfaces. Lee et al.[6] performed inference with Markov chain Monte Carlo and reported associated confidence intervals for both boundaries. Rahnemoonfar et al.[7] detected the two interfaces using distance-regularized level-set evolution. However, these models often suffer from limited generalization ability or low efficiency, making them unable to address a vast amount of radargrams. Under the umbrella of deep neural networks, Xu et al.[14] extended bed reconstruction to three dimensions by formulating a 3D topography problem within a Markov random field. Subsequently, Xu et al.[15] presented the first deep model for this task, combining 3D convolutional networks with recurrent layers to extract the ice–air and ice–bedrock boundaries directly from radargrams. Despite some success, these 3D convolutional networks process entire radargrams as input, so noise and outliers, together with the susceptibility of CNNs and RNNs to such artifacts, can substantially degrade prediction quality.

To address these challenges, we focus on accurate 3D ice bed topography reconstruction using ice air boundary coordinates and geometric deep learning. In this work, we proposed GT-IBTR-3D, namely **G**raph **T**ransformer for **3D I**ce-**B**ed **T**opography **R**econstruction. Rather than feeding the full radargram into the network, we use only the ice air boundary, which is more stable, has a higher signal to noise ratio, and contains fewer outliers. For each radargram, we represent this boundary as a graph and apply a graph transformer designed to learn from irregular geometry and be more robust to noise. Our graph transformer combines the GraphSAGE inductive framework [3] for localized feature learning and aggregation with transformer layers that capture long-range dependencies. We compare GT-IBTR-3D with previously reported probabilistic models and deep learning baselines. By reconstructing 3D ice bed topography from ice–air boundary elevations and utilizing a robust geometric deep learning framework, GT-IBTR-3D consistently achieves a lower mean absolute error than existing approaches.

2 Methods

In this work, we propose GT-IBTR-3D, a graph transformer designed to accurately reconstruct 3D bed elevation from ice–air boundary coordinates. As shown in Figure 1, the network input is a graph representation of the surface trace for each radargram together with its along-track neighbors. A GraphSAGE[3] encoder with hardswish activations and dropout learns local structure, while residual connections mitigate oversmoothing and ease optimization. After layer normalization, attention-based encoder layers capture long-range dependencies across the graph; residual connections are again applied to stabilize training. Finally, a two-layer linear head with GELU[4] activations and dropout produces the final prediction.

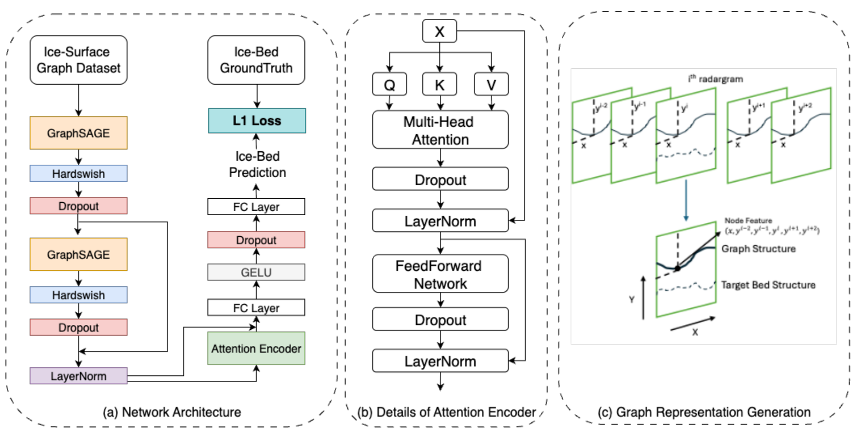

Figure 1: Schematic illustration of GT-IBTR-3D. (a) Our Network Architecture (b) Detail of attention-based encoder (c) Diagram of generating graph representations from ice-air boundary coordinates

2.1 Graph Representation Generation and Task Reformulation

Graph Representation Generation: To preserve temporal consistency and improve 3D reconstruction, each radargram's graph incorporates its temporal neighbors. As shown in Figure 1(c), the graph for the i^{th} radargram is augmented with frames $i-2$, $i-1$, i, $i+1$, and $i+2$ to encourage continuity over time. For each radargram, we build a graph with 64 nodes, one per horizontal pixel. Each node stores six features: the horizontal coordinate x and the surface elevations at that location from the five frames $\{y^{i-2}, y^{i-1}, y^i, y^{i+1}, y^{i+2}\}$.

Task Reformulation: Instead of predicting absolute bed elevation, GT-IBTR-3D treats 3D ice-bed topography reconstruction as thickness regression from the ice–air surface geometry. During training, thickness targets are computed from paired ice–air and ice–bed annotations; at inference, the model predicts thickness from surface data alone, and the bed is obtained by subtracting the prediction from the surface. This reformulation removes unknown scene-level offsets and other global biases, reduces the target's dynamic range, and stabilizes optimization.

2.2 GraphSAGE Inductive Framework

We adopt the GraphSAGE inductive framework to learn from the graph structure. GraphSAGE[3] produces node embeddings for unseen graphs by sampling and aggregating information from a node's neighborhood[16]. It does not rely on preassigned edge weights, making it well-suited to new datasets and practical settings where meaningful edge weights are difficult to define. For a node i with feature $\mathbf{X}_i$, the update rule of GraphSAGE with a mean aggregator is defined as:

$$\mathbf{X}'_i = \mathbf{W}_{self}\mathbf{X}_i + \mathbf{W}_{neigh}\text{mean}_{j\in\mathcal{N}(i)}\mathbf{X}_j \tag{1}$$

where $\mathbf{X}'_i$ is the learned embedding from GraphSAGE, $\mathcal{N}(i)$ is the neighbor set for node i, and $\mathbf{W}_{self}, \mathbf{W}_{neigh}$ are trainable matrices. With the mean aggregator, GraphSAGE can be viewed as a linear approximation to a localized spectral convolution [3].

2.3 Attention-Based Encoder

Graph neural networks excel at capturing local structure via message passing, but their receptive fields grow slowly, making it difficult to model long-range interactions. To complement this, we introduce an attention-based encoder that captures high-level global dependencies.

We adopt the standard multi-head self-attention mechanism[13]. Given input feature embeddings X, each head forms its queries, keys, and values via

$$Q_i = XW_i^Q, \quad K_i = XW_i^K, \quad V_i = XW_i^V, \tag{2}$$

where W_i^Q, W_i^K, W_i^V are learnable matrices. The attention score for head i is then calculated as

$$\text{Attention}(Q_i, K_i, V_i) = \text{softmax}\left(\frac{Q_i K_i^T}{\sqrt{d_k}}\right)V_i, \tag{3}$$

where d_k is the key dimension. Heads are computed independently and then concatenated together to obtain the output of the multi-head self-attention layer, defined as

$$\text{MultiHead}(X) = \text{Concat}(\text{Head}_1, \ldots, \text{Head}_n)W^O, \tag{4}$$

where $\text{Head}_i = \text{Attention}(Q_i, K_i, V_i)$ an d W^O is a learnable output projection applied after concatenation. As shown in Figure 1(b), our attention-based encoder layer also contains commonly used dropout layer, layer normalization, feedforward network and residual connections.

3 Experiment and Results

3.1 Dataset and Training Detail

In this study, we work with the Canadian Arctic Archipelago ice-sheet dataset [14, 15] collected using the Multichannel Coherent Radar Depth Sounder (MCoRDS) [9]. The dataset contains five tomographic sequences, each consisting of 3,332 consecutive radargram images. Each image has a width of 64 pixels and a depth of 824 pixels. Using these radargrams, scientists manually labeled the ice-air and ice-bedrock boundaries. We adopt 3 sequences (3×3332 radargrams) for training and 2 sequences (2×3332 radargrams) for testing. Throughout training, boundary coordinates are scaled to $[-1, 1]$ [12].

We trained our proposed GT-IBTR-3D on 8 NVIDIA A5000 GPUs. The L1 loss was used as the loss function. We used the Adam[5] optimizer with a learning rate of 0.00001. Dropouts with 0.1 probability and a weight decay of 0.0001 are used to ensure the robustness of the learning process. Based on the learning curve, we trained our model with 10 epochs.

Table 1: Ice-bed topography reconstruction results comparing our graph transformer with various baseline methods. Error is reported as the mean absolute column-wise difference from ground truth across radargrams (lower is better).

Method	**RMSE**
Crandall [1]	101.6
Lee [6]	35.6
Xu et al. (Prob. Model, w/o ice mask) [14]	30.7
Xu et al. (Prob. Model) [14]	11.9
Xu et al. (Neural Network) [15]	13.1
Ours (Graph Transformer)	**7.9**

3.2 Results Analysis

In order to highlight the performance of our proposed network, we compare it with various probabilistic models or traditional deep neural networks for bed topography reconstruction. As shown in Table 1, our proposed GT-IBTR-3D attains a substantially lower mean absolute error. We attribute this improvement to three factors. First, GT-IBTR-3D uses only the ice–air boundary coordinates as input and reconstructs 3D bed topography purely from surface variations. This design avoids the noise and outlier pixels that can contaminate predictions when the entire radargram image is fed to the network, leading to cleaner supervision and more stable learning. Second, by coupling graph convolutions for localized structure with transformers for long-range dependencies, GT-IBTR-3D learns fine-grained spatial detail while retaining broad contextual relationships. Third, we recast the target from absolute bed elevation to surface to bed thickness, which removes scene-level offsets and stabilizes learning, and we reconstruct the bed by subtracting the predicted thickness from the observed surface; taken together, these three design choices yield significant MAE improvement and consistent gains over all baselines. Figure 2 shows some qualitative visualizations of 3D bed elevation reconstructed by our proposed graph transformer.

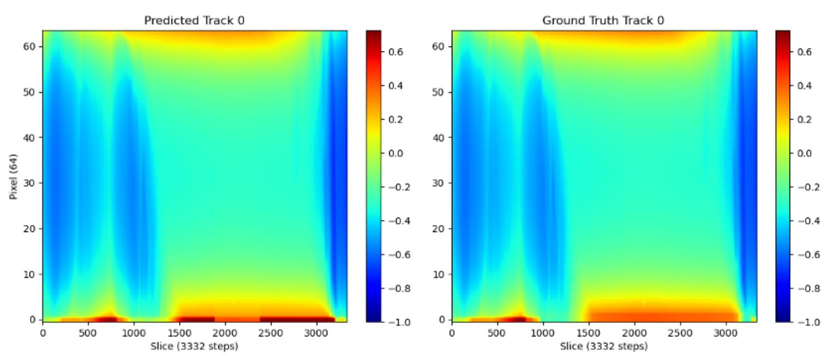

Figure 2: Qualitative results of GT-IBTR-3D predictions on one tomographic sequences; x is along-track distance, y is radargram width, and color encodes elevation (depth from the radar).

4 Conclusion

In this work, we introduce GT-IBTR-3D, a graph transformer for precise 3D ice-bed topography reconstruction. Rather than parsing full radargrams, our proposed geometric deep learning model focuses on the relationship between ice-air and ice-bed coordinates, which reduces sensitivity to noisy pixels and improves robustness. GT-IBTR-3D uses GraphSAGE to capture local spatiotemporal patterns and use an attention-based encoder to model long-range dependencies, enabling both fine-scale structure and broad context. On Canadian Arctic Archipelago data, it achieves consistently lower mean absolute column-wise error than probabilistic graphical models and conventional deep networks, highlighting the promise of graph-based methods for scalable, data-driven subglacial mapping.

References

[1] D. J. Crandall, G. C. Fox, and J. D. Paden. Layer-finding in radar echograms using probabilistic graphical models. In *Proceedings of the 21st International Conference on Pattern Recognition (ICPR2012)*, pages 1530–1533, 2012.

[2] P. Fretwell, H. D. Pritchard, D. G. Vaughan, J. L. Bamber, N. E. Barrand, R. Bell, C. Bianchi, R. G. Bingham, D. D. Blankenship, G. Casassa, G. Catania, D. Callens, H. Conway, A. J. Cook, H. F. J. Corr, D. Damaske, V. Damm, F. Ferraccioli, R. Forsberg, S. Fujita, Y. Gim, P. Gogineni, J. A. Griggs, R. C. A. Hindmarsh, P. Holmlund, J. W. Holt, R. W. Jacobel, A. Jenkins, W. Jokat, T. Jordan, E. C. King, J. Kohler, W. Krabill, M. Riger-Kusk, K. A. Langley, G. Leitchenkov, C. Leuschen, B. P. Luyendyk, K. Matsuoka, J. Mouginot, F. O. Nitsche, Y. Nogi, O. A. Nost, S. V. Popov, E. Rignot, D. M. Rippin, A. Rivera, J. Roberts, N. Ross, M. J. Siegert, A. M. Smith, D. Steinhage, M. Studinger, B. Sun, B. K. Tinto, B. C. Welch, D. Wilson, D. A. Young, C. Xiangbin, and A. Zirizzotti. Bedmap2: improved ice bed, surface and thickness datasets for antarctica. *The Cryosphere*, 7(1):375–393, 2013.

[3] W. L. Hamilton, R. Ying, and J. Leskovec. Inductive representation learning on large graphs, 2018.

[4] D. Hendrycks and K. Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016.

[5] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization, 2017.

[6] S. Lee, J. Mitchell, D. J. Crandall, and G. C. Fox. Estimating bedrock and surface layer boundaries and confidence intervals in ice sheet radar imagery using mcmc. In *2014 IEEE International Conference on Image Processing (ICIP)*, pages 111–115, 2014.

[7] M. Rahnemoonfar, G. C. Fox, M. Yari, and J. Paden. Automatic ice surface and bottom boundaries estimation in radar imagery based on level-set approach. *IEEE Transactions on Geoscience and Remote Sensing*, 55(9):5115–5122, 2017.

[8] E. Rignot, J. Mouginot, B. Scheuchl, M. van den Broeke, M. J. van Wessem, and M. Morlighem. Four decades of antarctic ice sheet mass balance from 1979–2017. *Proceedings of the National Academy of Sciences*, 116(4):1095–1103, 2019.

[9] F. Rodriguez-Morales, S. Gogineni, C. J. Leuschen, J. D. Paden, J. Li, C. C. Lewis, B. Panzer, D. Gomez-Garcia Alvestegui, A. Patel, K. Byers, R. Crowe, K. Player, R. D. Hale, E. J. Arnold, L. Smith, C. M. Gifford, D. Braaten, and C. Panton. Advanced multifrequency radar instrumentation for polar research. *IEEE Transactions on Geoscience and Remote Sensing*, 52(5):2824–2842, 2014.

[10] H. Seroussi, Y. Nakayama, E. Larour, D. Menemenlis, M. Morlighem, E. Rignot, and A. Khazendar. Continued retreat of thwaites glacier, west antarctica, controlled by bed topography and ocean circulation. *Geophysical Research Letters*, 44(12):6191–6199, 2017.

[11] A. Shepherd, E. Ivins, E. Rignot, B. Smith, M. van den Broeke, I. Velicogna, P. Whitehouse, K. Briggs, I. Joughin, G. Krinner, S. Nowicki, T. Payne, T. Scambos, N. Schlegel, G. A, C. Agosta, A. Ahlstrøm, G. Babonis, V. Barletta, A. Blazquez, J. Bonin, B. Csatho, R. Cullather, D. Felikson, X. Fettweis, R. Forsberg, H. Gallee, A. Gardner, L. Gilbert, A. Groh, B. Gunter, E. Hanna, C. Harig, V. Helm, A. Horvath, M. Horwath, S. Khan, K. K. Kjeldsen, H. Konrad, P. Langen, B. Lecavalier, B. Loomis, S. Luthcke, M. McMillan, D. Melini, S. Mernild, Y. Mohajerani, P. Moore, J. Mouginot, G. Moyano, A. Muir, T. Nagler, G. Nield, J. Nilsson, B. Noel, I. Otosaka, M. E. Pattle, W. R. Peltier, N. Pie, R. Rietbroek, H. Rott, L. Sandberg-Sørensen, I. Sasgen, H. Save, B. Scheuchl, E. Schrama, L. Schröder, K.-W. Seo, S. Simonsen, T. Slater, G. Spada, T. Sutterley, M. Talpe, L. Tarasov, W. J. van de Berg, W. van der Wal, M. van Wessem, B. D. Vishwakarma, D. Wiese, B. Wouters, and T. I. team. Mass balance of the antarctic ice sheet from 1992 to 2017. *Nature*, 558(7709):219–222, Jun 2018.

[12] A. Toshev and C. Szegedy. Deeppose: Human pose estimation via deep neural networks. *CoRR*, abs/1312.4659, 2013.

[13] A. Vaswani. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017.

[14] M. Xu, D. J. Crandall, G. C. Fox, and J. D. Paden. Automatic estimation of ice bottom surfaces from radar imagery. In *IEEE International Conference on Image Processing (ICIP)*, 2017.

[15] M. Xu, C. Fan, J. D. Paden, G. C. Fox, and D. J. Crandall. Multi-task spatiotemporal neural networks for structured surface reconstruction. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2018.

[16] J. Zhou, G. Cui, S. Hu, Z. Zhang, C. Yang, Z. Liu, L. Wang, C. Li, and M. Sun. Graph neural networks: A review of methods and applications. *AI Open*, 1:57–81, 2020.