Climate Change AI

NeurIPS 2025 workshop – Tackling Climate Change with Machine Learning

NEURAL INFORMATION PROCESSING SYSTEMS

Mass Conservation on Rails – Rethinking Physics-Informed Learning of Ice Flow Vector Fields

Kim Bente
School of Computer Science,
The University of Sydney

Roman Marchant
Human Technology Institute,
University of Technology Sydney

Fabio Ramos
NVIDIA
School of Computer Science,
The University of Sydney

THE UNIVERSITY OF SYDNEY
Australian Government
Australian Research Council
UTS UNIVERSITY OF TECHNOLOGY SYDNEY
NVIDIA
DARE
ARC Training Centre in
Data Analytics for Resources & Environments

Antarctica is about 1.5× the size of the US & the largest potential contributor to sea level rise

Ice thickness models

ice sheet models

→ sea level rise

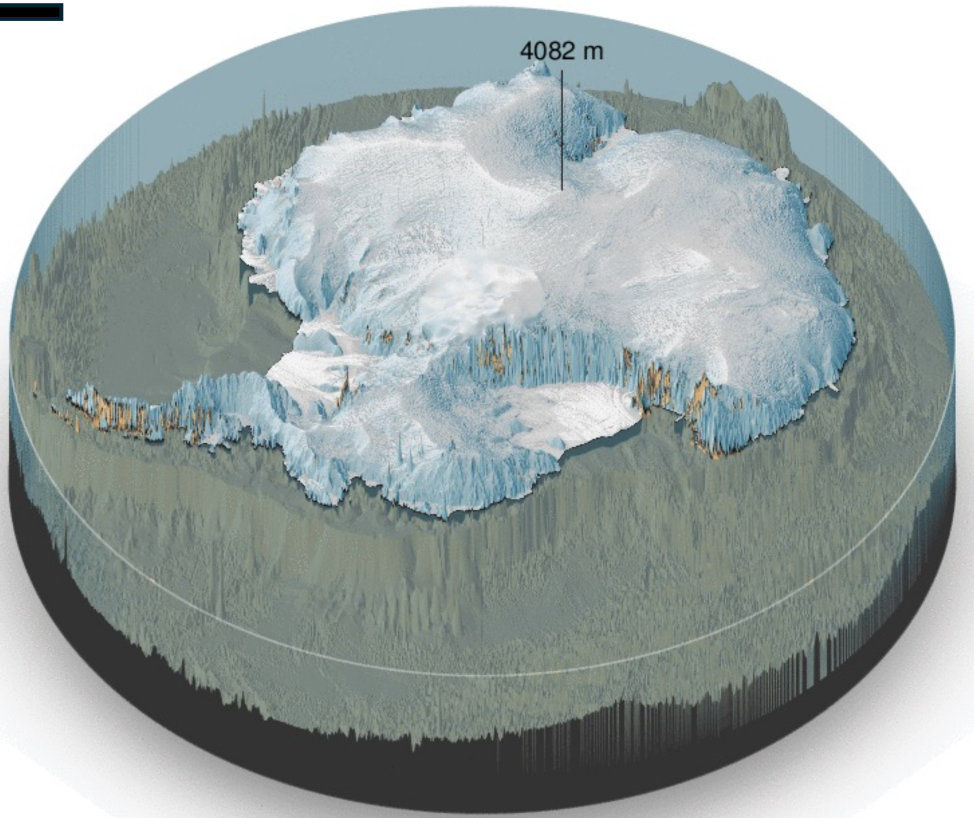

animation by Stef Lhermitte

data limitations

93% of 500 m grid cells rely on interpolation.
[Pritchard et al. 2025]

Interpolations require Mass Conservation (MC)

Flux: ice mass transport rate

2D flux $\boldsymbol{v} = $ ice thickness $h \times$ ice velocity $\boldsymbol{s}$

interpolation with flux divergences
blue = spurious sources
red = spurious sinks

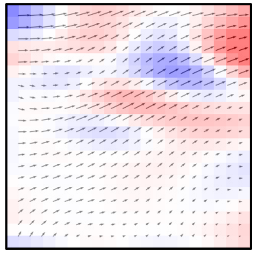

Ice Sheet Model

see Seroussi et al. 2011, Morlighem et al. 2011, and Nias et al. 2018.

divergence-free flux interpolation
inflow = outflow at any point

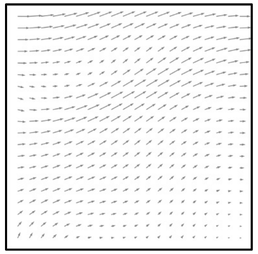

Ice Sheet Model

sea level rise projections that reliably inform climate action

→ Ice sheet models require ice thickness initialisations to respect **mass conservation**.

FEM ice flux Mass Conservation

- **Finite Element Method** (FEM) Mass Conservation introduced by Morlighem et al. 2011
- Applied to fast-flowing regions of Antarctica in BedMachine Antarctica (Morlighem et al. 2020)
  - **Finite Element Method** (FEM) drawbacks:
    - high computational cost
    - depends on fixed mesh/grid
    - integration of new measurements may be tricky


ARTICLES
https://doi.org/10.1038/s41561-019-0510-8
nature geoscience

Deep glacial troughs and stabilizing ridges unveiled beneath the margins of the Antarctic ice sheet

Mathieu Morlighem[1*], Eric Rignot[1,2], Tobias Binder[3], Donald Blankenship[4], Reinhard Drews[3,5], Graeme Eagles[3], Olaf Eisen[3,6], Fausto Ferraccioli[7], René Forsberg[8], Peter Fretwell[7], Vikram Goel[9], Jamin S. Greenbaum[4], Hilmar Gudmundsson[10], Jingxue Guo[11], Veit Helm[3], Coen Hofstede[3], Ian Howat[12], Angelika Humbert[3,6], Wilfried Jokat[3], Nanna B. Karlsson[3,13], Won Sang Lee[14], Kenichi Matsuoka[15], Romain Millan[1], Jeremie Mouginot[1,16], John Paden[17], Frank Pattyn[18],

Can we use Machine Learning for **mass conservation?**

Physics-Informed Machine Learning (ML)

1. Observational bias	2. Learning bias	3. Inductive bias
unconstrained models	soft-constrained models	hard-constrained models
learn from observations that embody physics	additional physics violation loss term promotes physics	model or kernel architectures enforce physics exactly
Implied in any ML model but observation are very limited	**Mass Conservation PINNs** Teisberg et al. 2021 and Steidl et al. 2025 - no physical guarantees - generalisation issues - convergence issues - spectral bias	**Our work** - exact physical guarantees - better generalisation - better convergence - less spectral bias

after Karniadakis et al. 2021

Divergence-free ice flux

sparse ice **height** point data h from airborne

dense **surface velocity** vectors $\boldsymbol{s}$ from satellites

from the Bedmap3 (Fremand et al. 2023)

ice height h

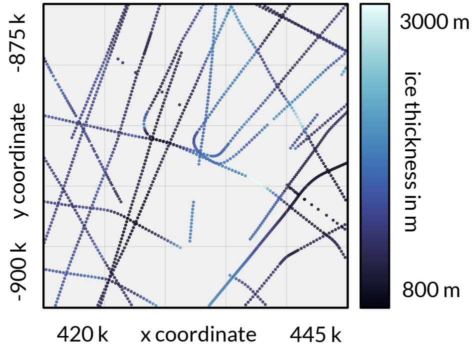

$h \cdot \boldsymbol{s}$

surface velocity $\boldsymbol{s}$

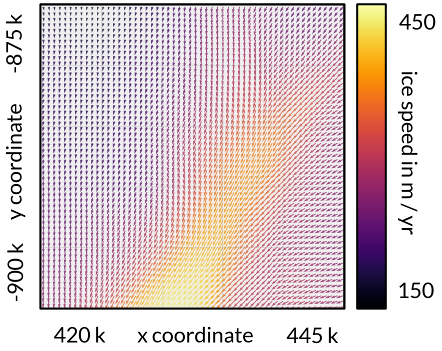

from MEaSUREs InSAR velocity map (Rignot et al. 2017)

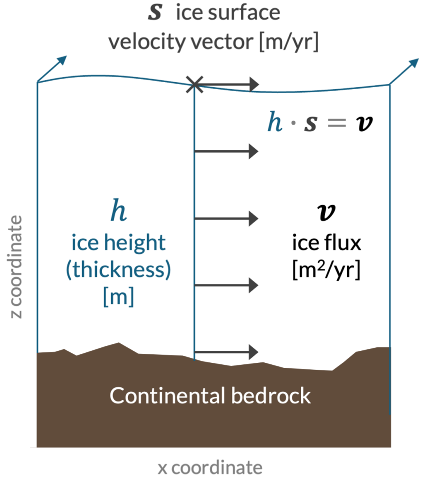

$$\boldsymbol{v} = h \cdot \boldsymbol{s}$$

Ice flux describes the transport of ice mass/volume per unit time for any point in 2D space.

ice flux $\boldsymbol{v}$

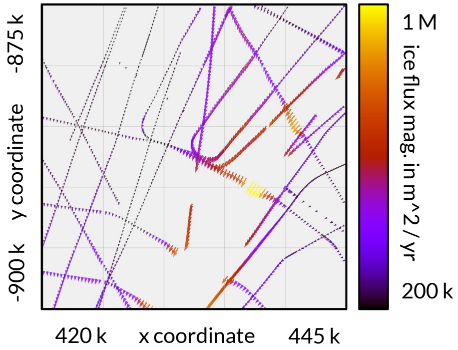

Remove negligible terms from the mass continuity equation (as in Teisberg et al. 2021)

$$\frac{\partial h}{\partial t} + \nabla \cdot \boldsymbol{v} = \text{smb} - \text{bmb}$$

$$\boxed{\nabla \cdot \boldsymbol{v} = 0}$$

→ Enforce MC as divergence-free constraint

A fluid dynamics trick

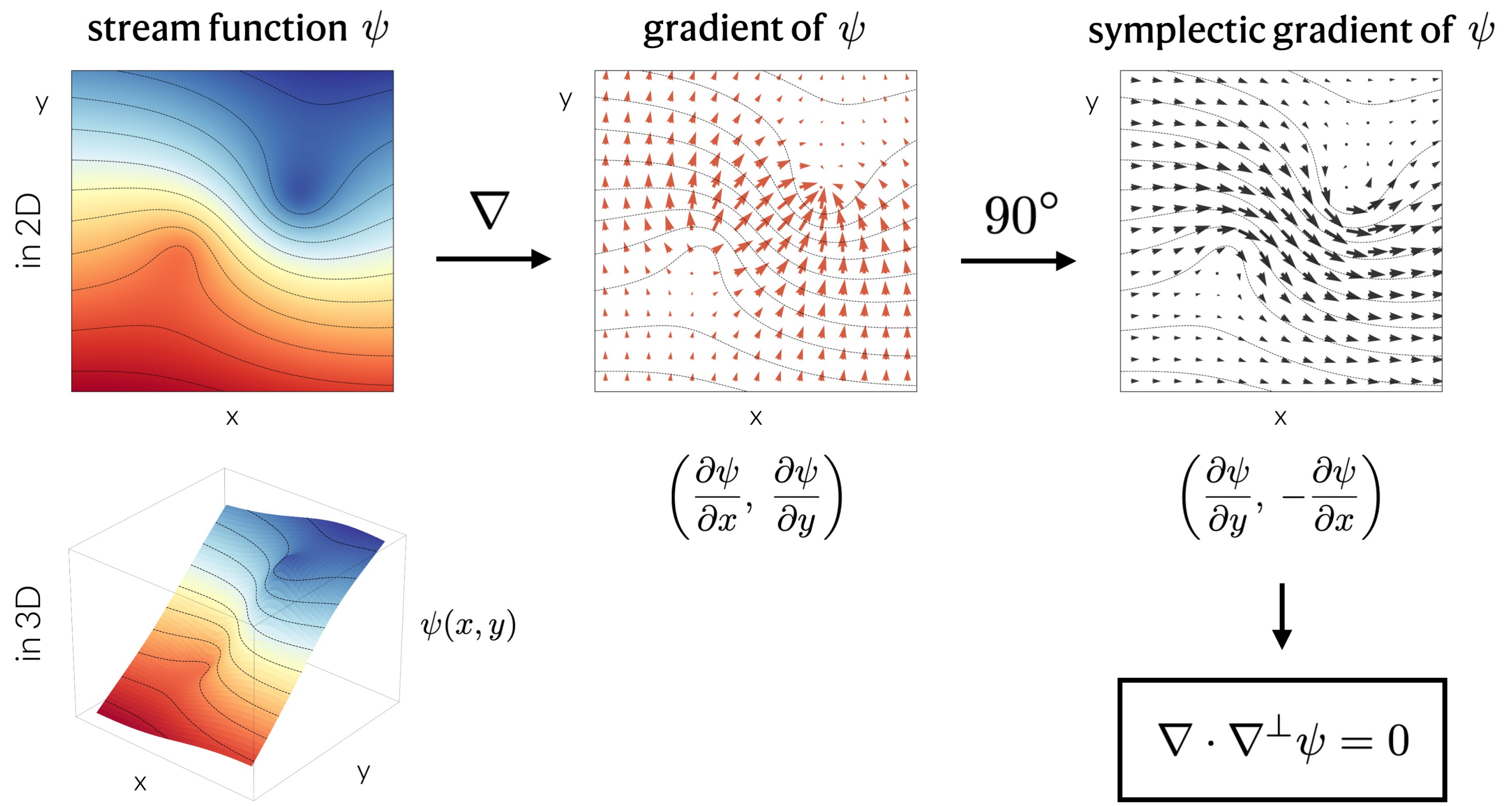

→ The symplectic gradient of the stream function is always **divergence-free**.

Exact Mass Conservation with divergence-free NNs

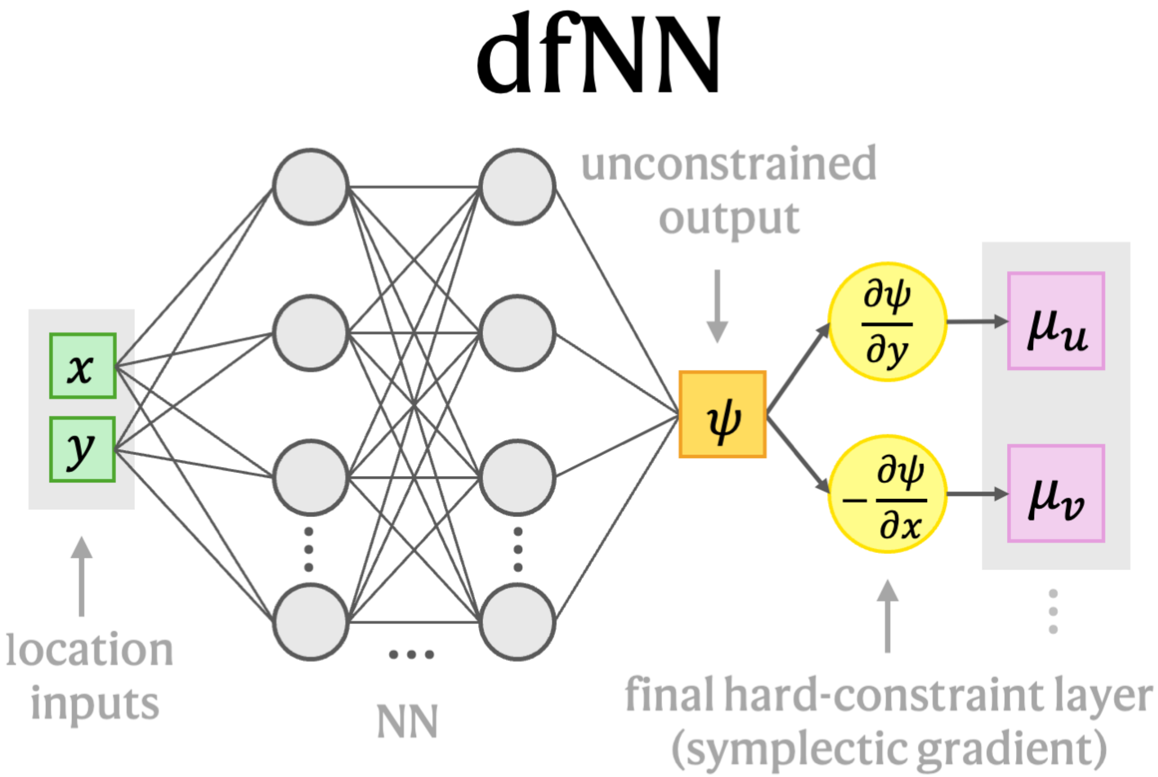

Based on

- **Model Inclusive Learning** [Kuroe et al. 1998]
- **Hamiltonian Neural Networks (HNNs)** [Greydanus et al. 2019]
- **Neural Conservation Laws (NCLs)** [Richter-Powell et al. 2022]

→ No glaciology applications yet.

→ No comprehensive comparison.

Mass Conservation on Rails

+ aux

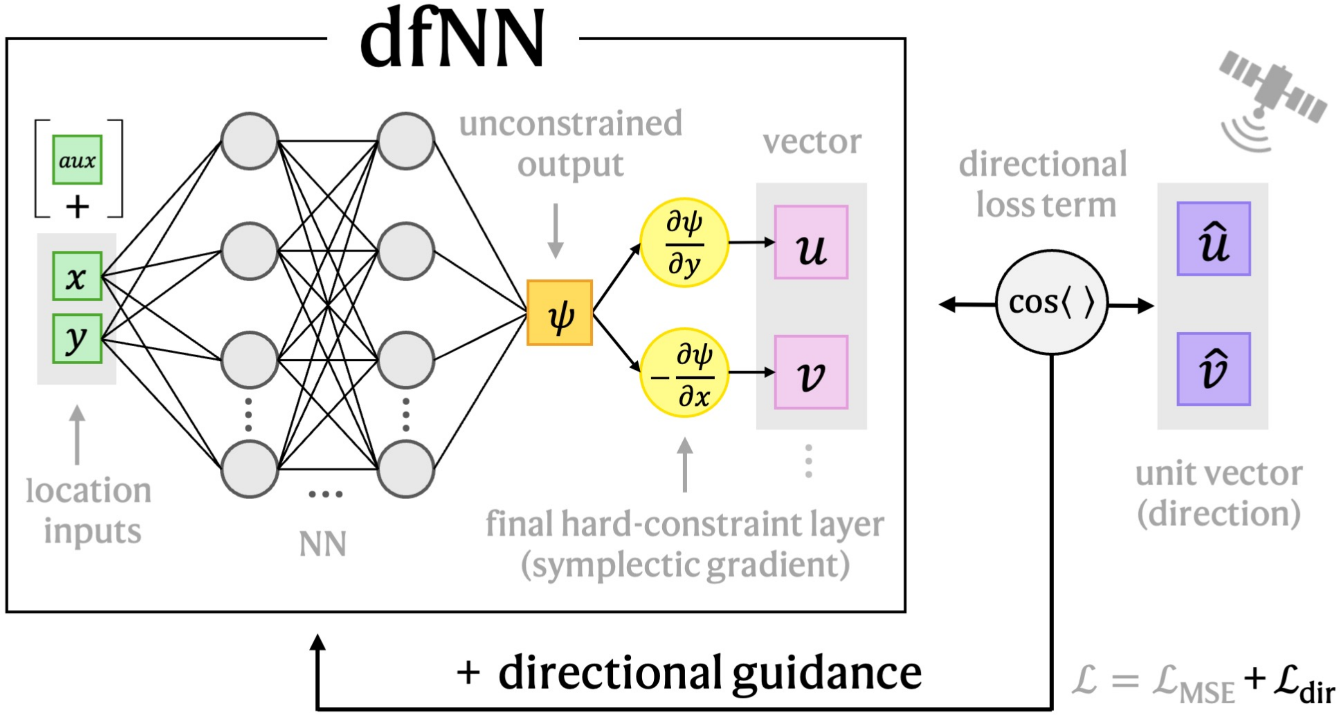

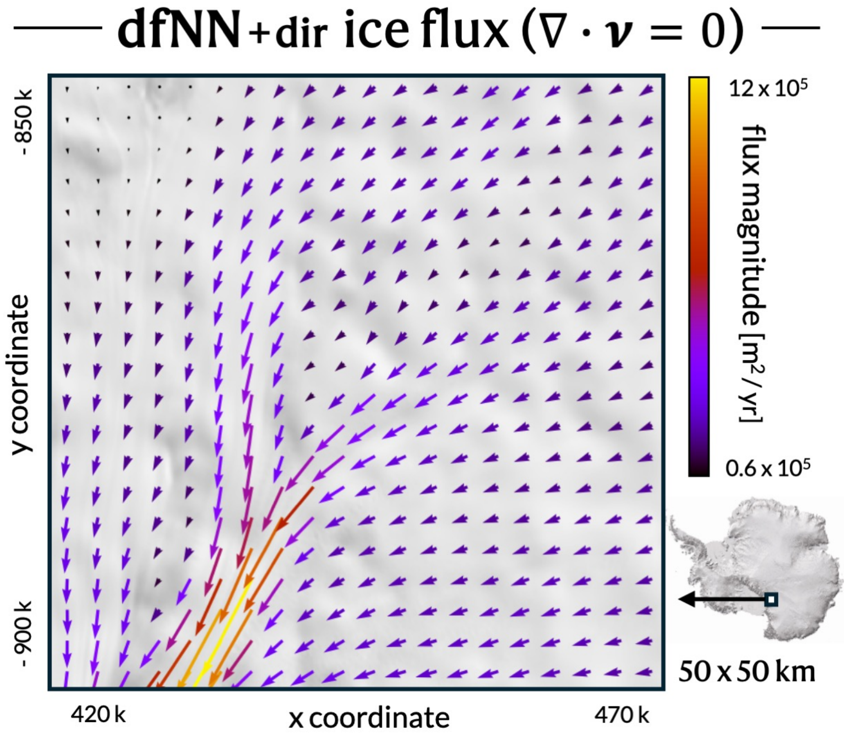

+ dir

Experiments

9 models:

- dfNN (hard-constrained) (proposed)
  + dir (directional guidance training step) (proposed)
  + aux (auxiliary input: surface elevation)

- PINN (soft-constrained)
  + dir (proposed)
  + aux

- NN (unconstrained)
  + dir (proposed)
  + aux

200 × 200 km region over Byrd Glacier, Antarctica, divided into train and test.

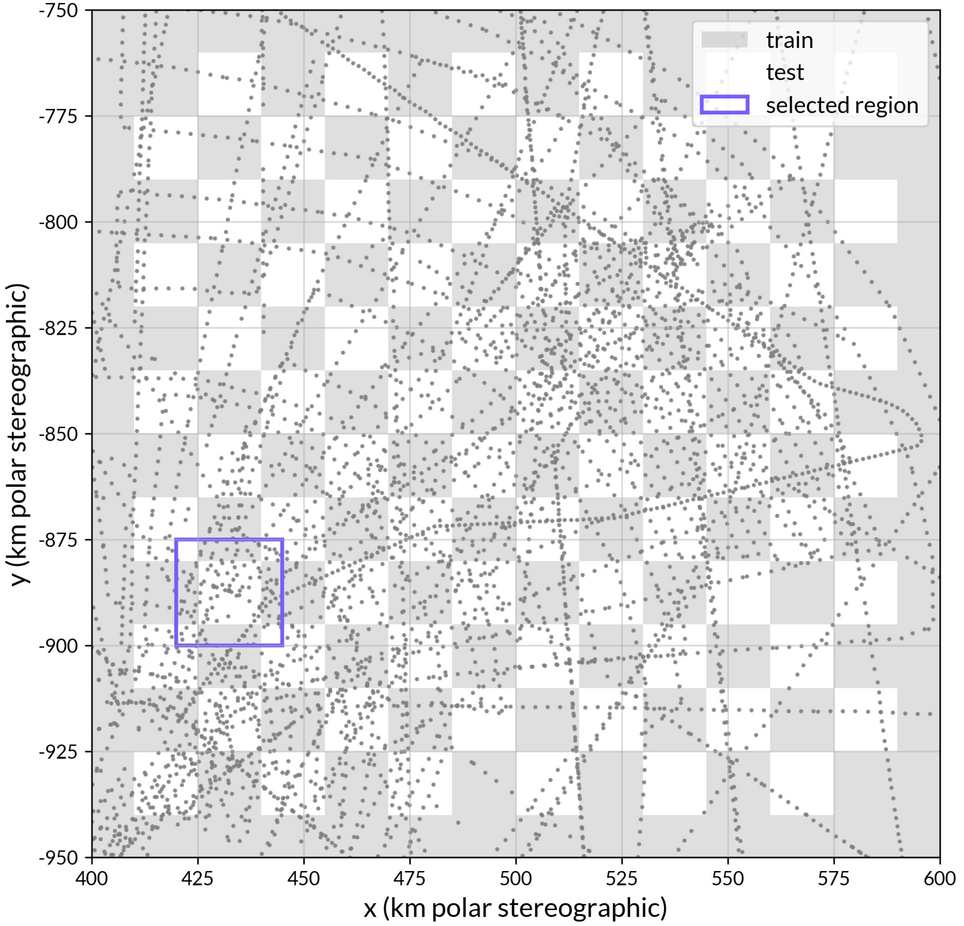

Results on test

MAD: Mean Absolute Divergence

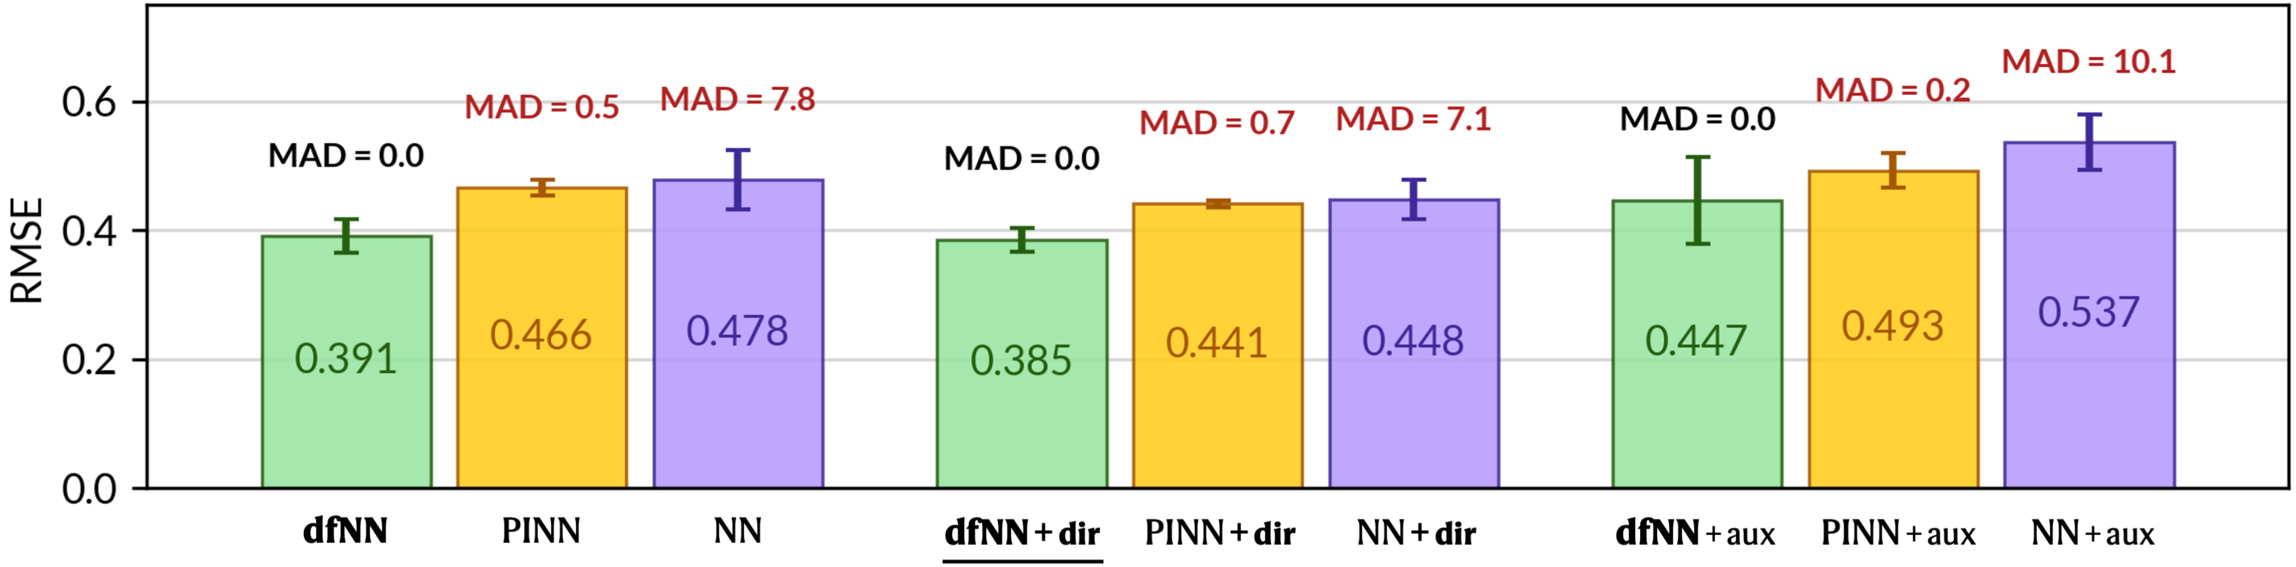

→ **Hard-constraining** yields higher accuracy reconstructions and better physical adherence.

→ Additional **directional guidance** training boosts all models.

→ Adding auxiliary surface predictors did not help.

Visual comparison of predictions

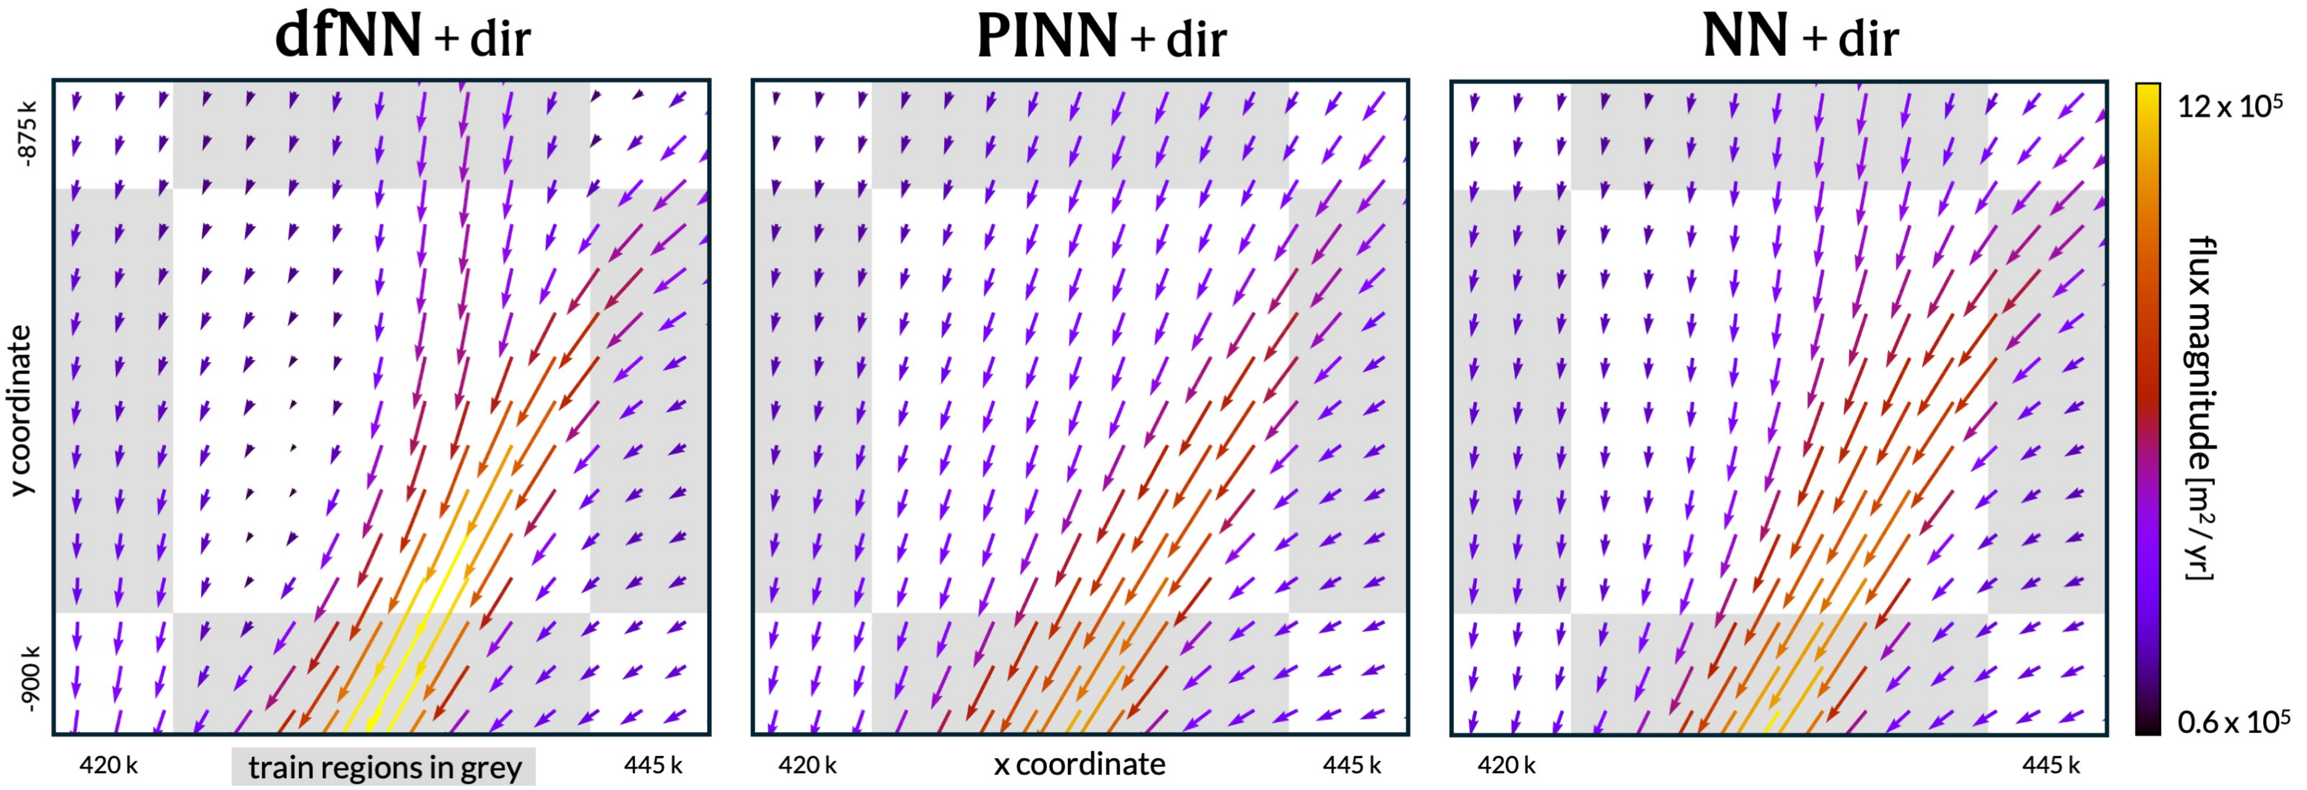

→ spectral bias

*showing small 25 x 25 km subregion

Summary – Mass Conservation on Rails

Our findings suggest that **hard-constrained PIML** models are preferrable, especially when data is **sparse & noisy.**

Future work may **apply dfNNs to other environmental flows** (e.g. ocean, groundwater, atmosphere) and extend dfNNs to higher dimensions.

Fully reproducible experiments on GitHub

- dfNN implementation in PyTorch
- Leverage autodiff in forward pass
- SiLU activation function improves convergence

dfNNs

- Physical guarantees ✓
- End-to-end model ✓
- Meshless ✓
- Fast ✓

Also find the workshop paper on arXiv

Ongoing work

Probabilistic model variant: dfNGPs

Integration of dfNNs with divergence-free Gaussian Processes (dfGPs) into dfNGPs (divergence-free Neural Gaussian Processes) for **Uncertainty Quantification**.

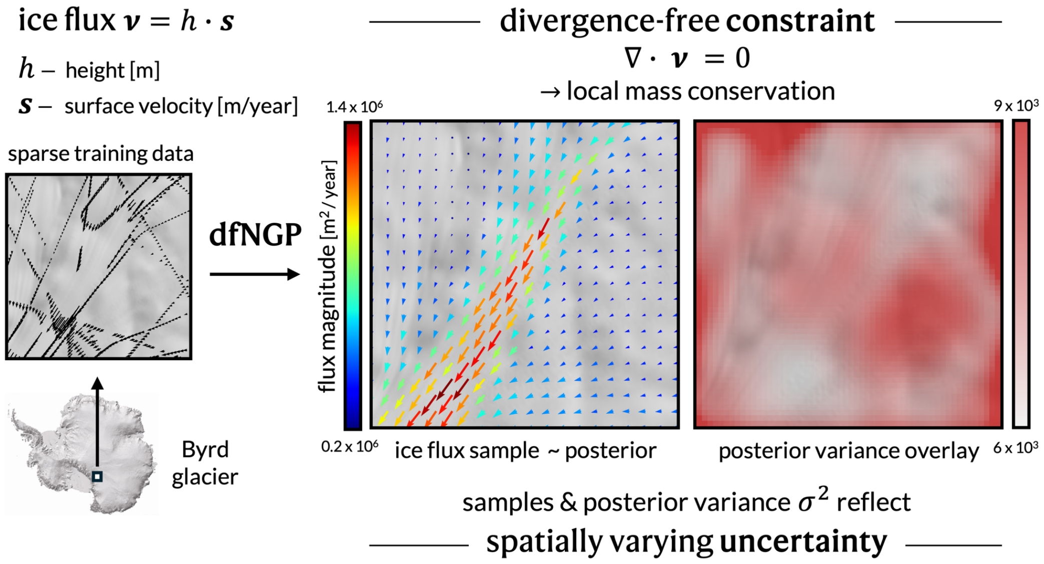

DHNNs for ice shelf basal melt estimation

We use a Dissipative Hamiltonian Neural Network (DHNN, Sosanya et al. 2022) to estimate basal melt over Ross ice shelf via **vector field decomposition**.

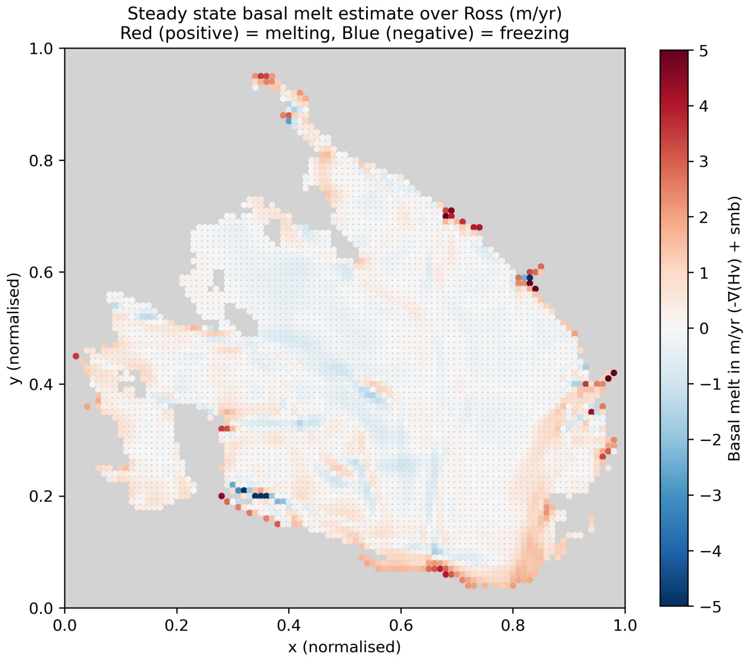

Join ML4Cryo (Machine Learning for the Cryosphere)

Julia Kaltenborn
Mila
McGill University

Andrew McDonald
University of Cambridge
British Antarctic Survey
Natural Environment Research Council

Kim Bente
The University of Sydney

ML4Cryo

A global research community, bridging the gap between machine learning and cryospheric science.

- slack community
- sessions at EGU and similar
- contributions to reports & blogs
- … review paper in the making

References I

Pritchard, Hamish D., Peter T. Fretwell, Alice C. Fremand, Julien A. Bodart, James D. Kirkham, Alan Aitken, Jonathan Bamber et al. "Bedmap3 updated ice bed, surface and thickness gridded datasets for Antarctica." Scientific data 12, no. 1 (2025): 414.

Seroussi, Hélène, Mathieu Morlighem, E. Rignot, Eric Larour, Denis Aubry, H. Ben Dhia, and Steen Savstrup Kristensen. "Ice flux divergence anomalies on 79north Glacier, Greenland." Geophysical Research Letters 38, no. 9 (2011).

Morlighem, Mathieu, E. Rignot, Hélene Seroussi, Eric Larour, H. Ben Dhia, and Denis Aubry. "A mass conservation approach for mapping glacier ice thickness." Geophysical Research Letters 38, no. 19 (2011).

Nias, I. J., S. L. Cornford, and A. J. Payne. "New mass-conserving bedrock topography for Pine Island Glacier impacts simulated decadal rates of mass loss." Geophysical Research Letters 45, no. 7 (2018): 3173-3181.

Morlighem, Mathieu, Eric Rignot, Tobias Binder, Donald Blankenship, Reinhard Drews, Graeme Eagles, Olaf Eisen et al. "Deep glacial troughs and stabilizing ridges unveiled beneath the margins of the Antarctic ice sheet." Nature geoscience 13, no. 2 (2020): 132-137.

Frémand, Alice C., Peter Fretwell, Julien A. Bodart, Hamish D. Pritchard, Alan Aitken, Jonathan L. Bamber, Robin Bell et al. "Antarctic Bedmap data: Findable, Accessible, Interoperable, and Reusable (FAIR) sharing of 60 years of ice bed, surface, and thickness data." Earth System Science Data 15, no. 7 (2023): 2695-2710.

Rignot, E., Mouginot, J. & Scheuchl, B. (2017). MEaSUREs InSAR-Based Antarctica Ice Velocity Map. (NSIDC-0484, Version 2). [Data Set]. Boulder, Colorado USA. NASA National Snow and Ice Data Center Distributed Active Archive Center.

References II

Karniadakis, George Em, Ioannis G. Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. "Physics-informed machine learning." Nature Reviews Physics 3, no. 6 (2021): 422-440.

Teisberg, Thomas O., Dustin M. Schroeder, and Emma J. MacKie. "A machine learning approach to mass-conserving ice thickness interpolation." In 2021 IEEE international geoscience and remote sensing symposium IGARSS, pp. 8664-8667. IEEE, 2021

Steidl, Viola, Jonathan Louis Bamber, and Xiao Xiang Zhu. "Physics-aware machine learning for glacier ice thickness estimation: a case study for Svalbard." The Cryosphere 19, no. 2 (2025): 645-661.

Y. Kuroe, M. Mitsui, H. Kawakami, and T. Mori. A learning method for vector field approximation by neural networks. In 1998 IEEE International Joint Conference on Neural Networks Proceedings. IEEE World Congress on Computational Intelligence (1998)

Greydanus, Samuel, Misko Dzamba, and Jason Yosinski. "Hamiltonian neural networks." Advances in neural information processing systems 32 (2019).

Richter-Powell, Jack, Yaron Lipman, and Ricky TQ Chen. "Neural conservation laws: A divergence-free perspective." Advances in Neural Information Processing Systems 35 (2022): 38075-38088.

Send me an email

kim.bente@sydney.edu.au

Github, links & more

Acknowledgements

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

Australian Government
Australian Research Council

DARE
ARC Training Centre in
Data Analytics for Resources & Environments

THE UNIVERSITY OF
SYDNEY