Deep learning to predict Antarctic ice shelf melt rates
Sebastian H. R. Rosier and Christopher Y. S. Bull

1. Aims and Introduction

- The current main driver of change in Antarctica is ocean driven melting of the floating portions of the ice sheet (ice shelves).
- Ice sheet models model this process either (1) with coupled ocean models that are very computationally expensive or (2) with simple parameterisations that generally fail to capture the correct spatial distribution of melt.
- Here, we propose a deep learning methodology (MELTNET) as a middle ground between these two approaches.
- Due to the dearth of observations and as a first step, we train and measure the success of this approach on synthetically generated input data.
- An ocean model (NEMO) calculates melt rate fields for thousands of input geometries and these are considered our observations for the purposes of training and validation.

2. Synthetic input generation

- The required inputs for both NEMO and MELTNET are ice shelf geometry and ocean conditions (temperature and salinity).
- We generate entirely random synthetic geometries with a high degree of variation, with ice shelf thickness profiles based on analytical solutions.
- Temperature and salinity forcing originate from World Ocean Atlas (WOA) data in the deep ocean around Antarctica.
- To generated an unlimited number of possible ocean conditions but retain physically plausible profiles, these finite observations are used to train a GAN type network that outputs synthetic forcing.

3. Deep learning methodology

- Synthetic inputs are sent to the NEMO ocean model which calculates a melt rate field. This melt rate field is converted to a discrete number of melt rate labels.
- Separately, a denoising AutoEncoder network is trained to map from an image of labelled melt rates to a continuous melt rate field.
- A segmentation network, based on the SegNet architecture, takes synthetic inputs in the form of a 4-channel image learns to reproduce the labelled melt rates from NEMO.
- Finally, the labelled melt rates output by the segmentation net are input to the autoencoder network to output the melt rate prediction.

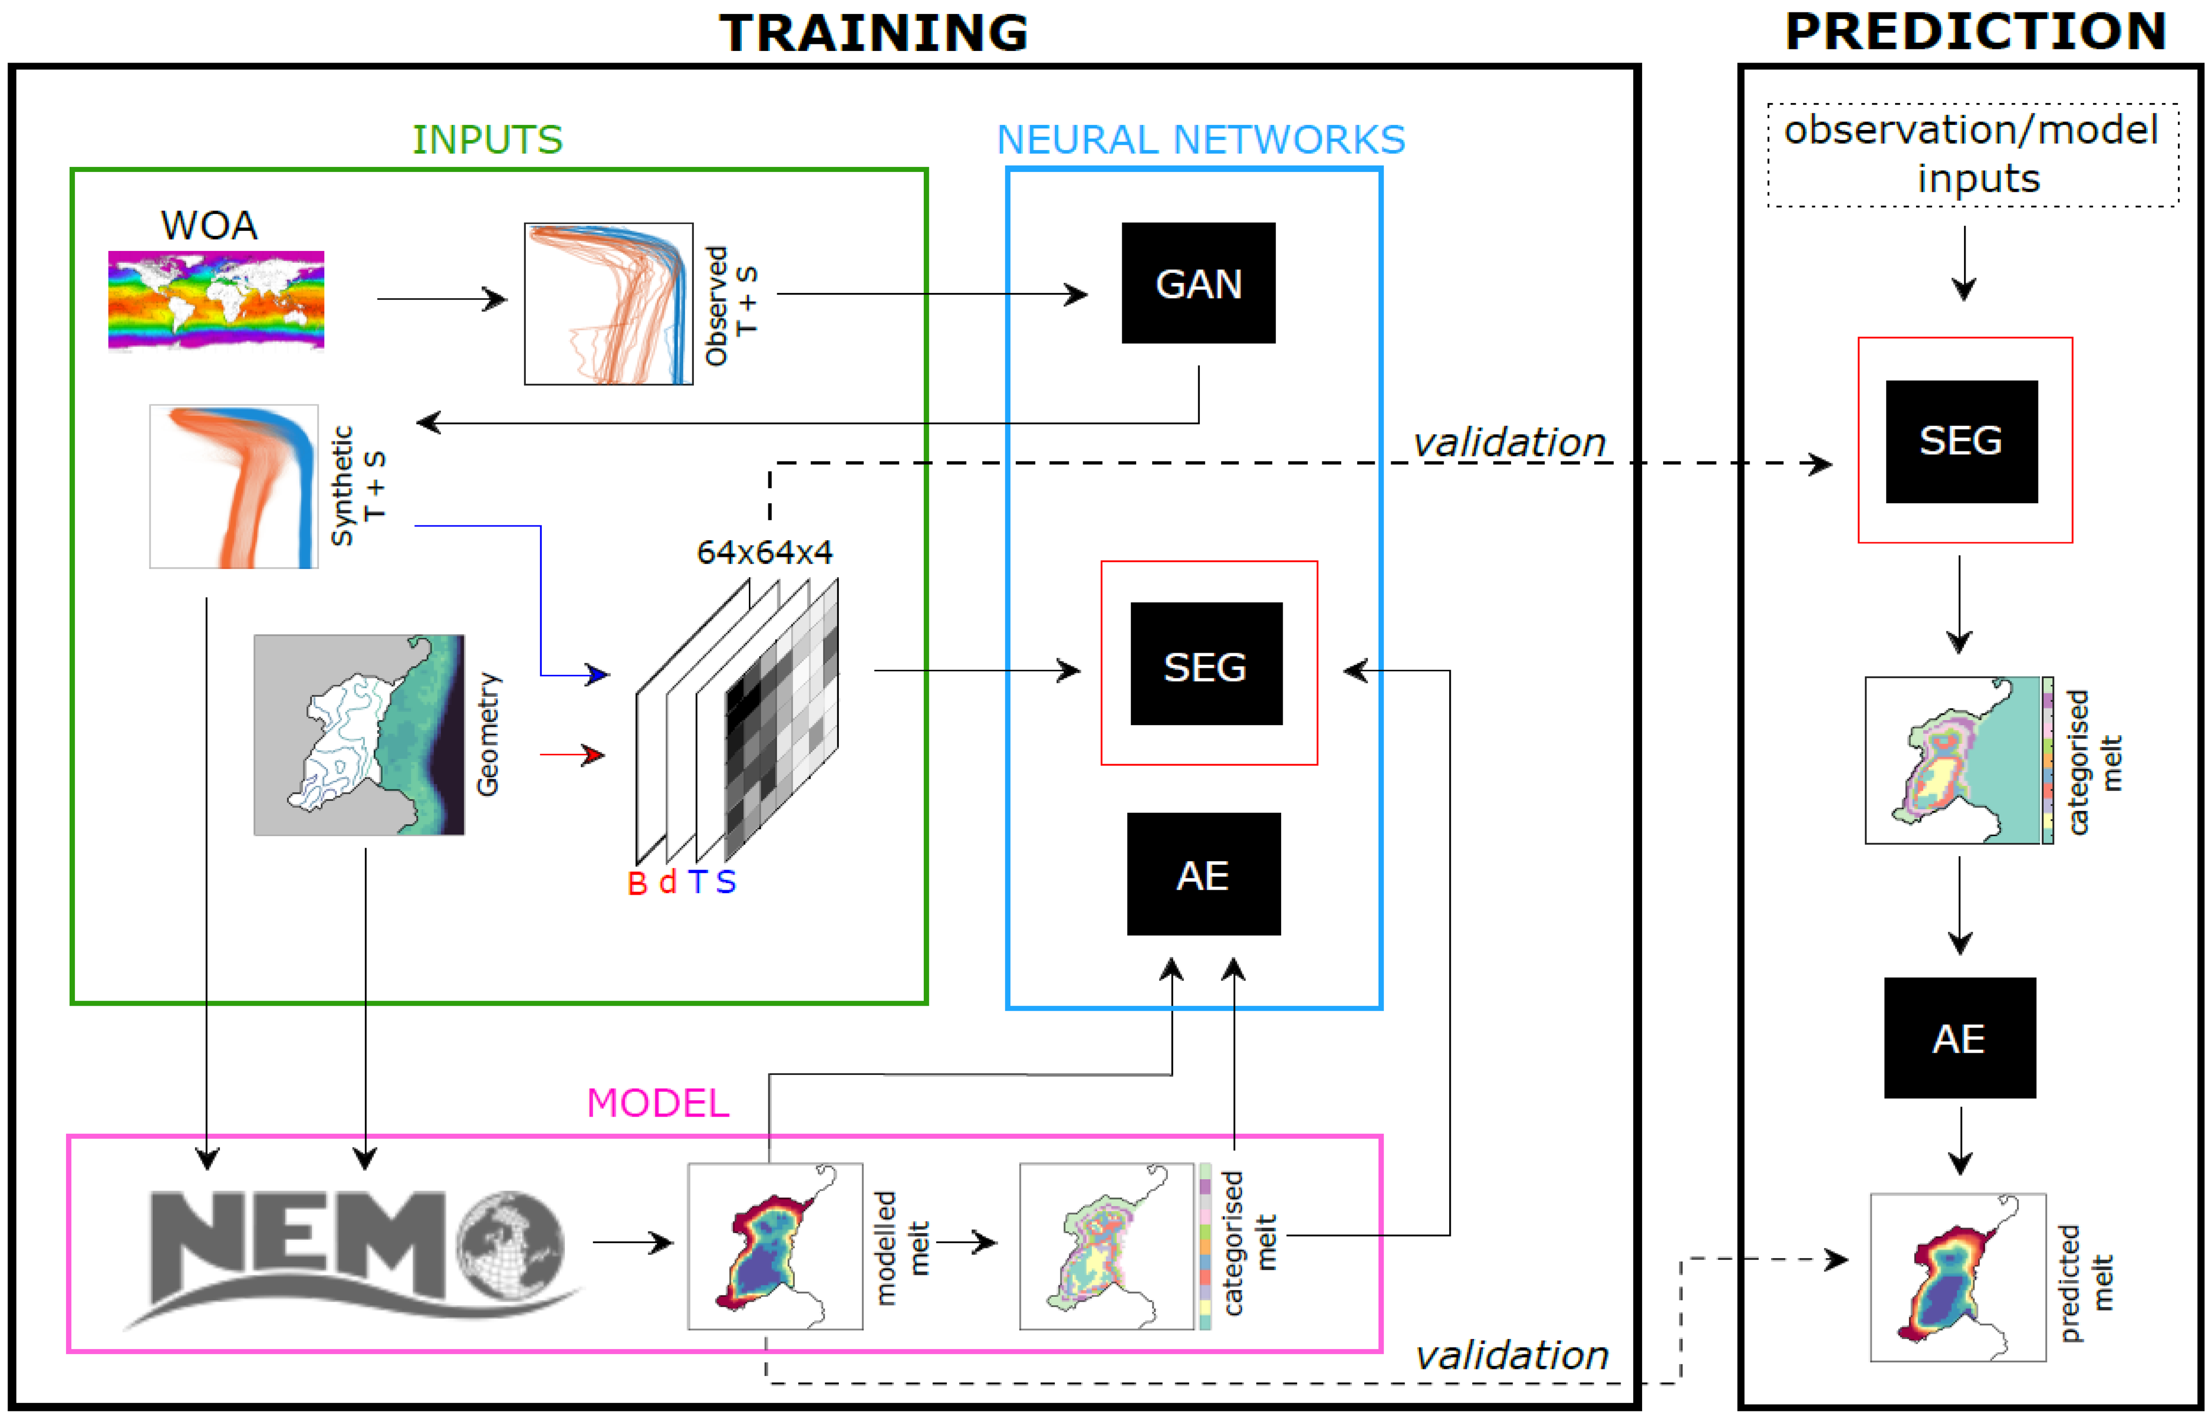

Figure 1. Workflow for the training and application of MELTNET. Inputs (green box) are given to the NEMO ocean model which predicts melt rates (magenta box) and together the inputs and melt rates are used to train the neural networks (blue box). Once trained, MELTNET takes input geometry, temperature and salinity and outputs a continuous melt rate field.

4. Results

- We compare MELTNET results to the ocean model output and two commonly used melt rate parameterisations (PICO and PLUME, Figure 2).
- Two tunable parameters in the PICO and PLUME models were tuned to minimise the misfit to melt rates of the training set.
- MELTNET outperformed PICO and PLUME in terms of NRMSE for 95% of the validation set and in terms of correlation coefficient for 99% of the validation set (Figure 3).

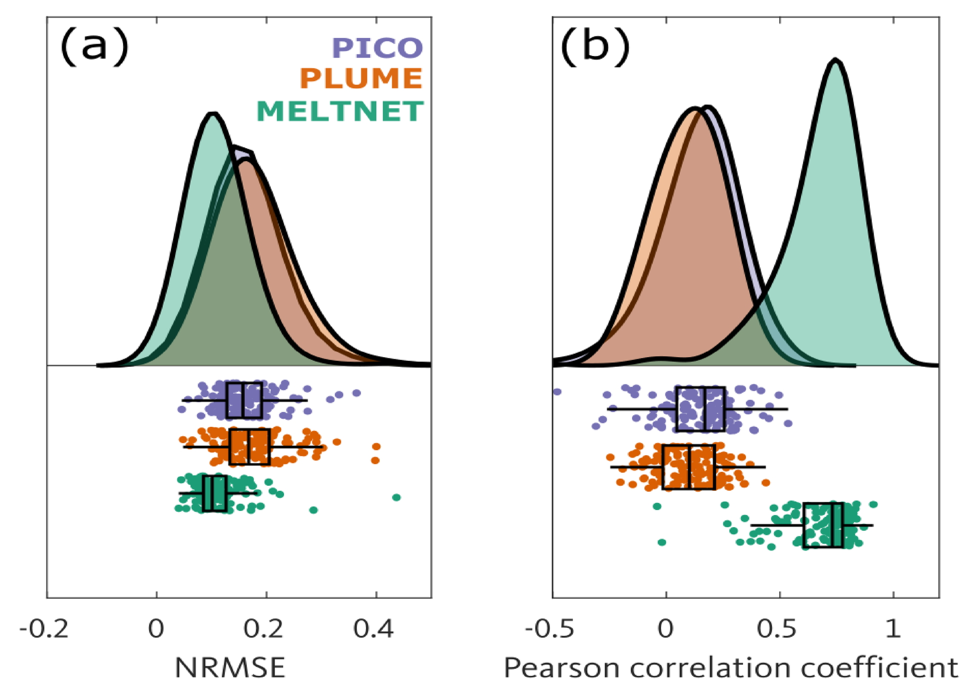

Figure 3. NRMSE (panel a) and correlation coefficient (panel b) distributions for all members of the validation set for MELTNET, PLUME and PICO compared to NEMO.

Figure 2. A sample of input geometries (first row) and the resulting melt rate predictions by the NEMO ocean model (second row), MELTNET (third row), PICO parameterisation (fourth row) and PLUME parameterisation (fifth row). Both the PICO and PLUME parameterisations have been optimised to minise the misfit to the NEMO model. The number in the bottom left corner of each image tile represents the total melt rate for each synthetic ice shelf. The validation set consists of 132 geometries and for each geometry we score MELTNET based on the misfit and correlation compared to NEMO. The nine geometries that form this figure are selected by evenly sampling the from the distribution of MELTNET scores.