## S1 Understanding the SOM algorithm output

The input data to the SOM algorithm consists of a binary melt / no-melt distinction for each pixel across the shelf. However, the final values output by the SOM algorithm (i.e. Fig. 2) are not binary values, but rather a likelihood of melt being observed in the pixel. We explain this here using a simplified example.

The SOM algorithm is initiated with nine *nodes*, each corresponding to the shelf mask filled with randomlyassigned binary values. The first day of the first melt season ("Day 1") is compared against each of these nine nodes, and judged (for example) to be most similar to node 1, and second most similar to node 2. Day 1 is therefore assigned to node 1. In turn, node 1 slightly adjusts to better represent Day 1, its newest member, and node 2 also slightly adjusts to move closer to node 1. These adjustments can be thought of as a closer alignment of the nodes and the observed melt as vectors in n-dimensional space, where each dimension represents a single pixel within the shelf mask. The amount that nodes 1 and 2 adjust depends on the learning rate selected (Fig. S4).

After the adjustment, the values in nodes 1 and 2 are no longer only binary values, but rather contain values between 0 and 1. These values continue to change as the algorithm iterates through the input data (i.e. all melt season days across all 18 summers).

After a defined number of iterations through the full dataset (Fig. S4), the final values of node 1 are nearidentical ( $r^2 = 0.997$ ; RMSE = 0.022) to the proportion of times that melt was observed in the pixels on the days best matched by node 1. This equivalence holds across all of the nodes (Fig. S1b). We therefore discuss the pixel values as the likelihood of melt being observed in that pixel, and refer to the final nodes as melt "patterns" throughout the paper.



Figure S1: (a) Boxplot showing the difference between the SOM pixel values, and the proportion of days assigned to the pattern on which melt was observed. Values across all SOM patterns are plotted together. (b) Visual comparisons of patterns 2, 5 and 9 are shown; plots on the top row display the SOM values, and plots on the bottom row display the melting frequency within the pattern.



Figure S2: Visual comparison of SOM patterns and a selection of the melt observations they represent, shown for SOM patterns 2, 4, 7, and 8. Values in the upper left plot of each series correspond to the scale used in Fig. S1; in the remaining 21 plots, the dark colour indicates pixels in which melt was observed, and the beige colour pixels which remained dry. Dates of the observations are shown in the bottom right.



Figure S3: Boxplot of the Tanimoto Distance between each SOM pattern and the melt days it best represents.



Figure S4: SOM sensitivity testing, showing the result for (a) 500 iterations; (b) a different random initialisation (a starting seed of 2706); (c) a steeper learning rate, descending from 0.1 to 0.005; and (d) a requirement for 16 output patterns. The remaining parameters were held steady between sensitivity tests (100 iterations; learning rate of 0.05 to 0.01; 9 output patterns; a starting seed of 1609). Although the ordering of the patterns differs between these tests, the patterns remain largely consistent, and show that our results do not depend on the settings used in the algorithm. For the results in (a-c), the mean distance to the closest pattern is equal to the value (0.073) obtained from the parameters used in the study (i.e. those which produce Fig. 2). The mean value decreases in (d) to 0.067 when more output patterns are used (see also Fig. S5), but the additional patterns <sup>4</sup>are near-replicas and therefore provide no additional, physically-interpretable output.



Figure S5: Result of a SOM sensitivity test, when 100 output patterns were requested from the algorithm. The frequency of the pattern being observed is shown in the bottom-right of each plot. Despite an improved (i.e. lower) mean distance to the closest pattern (0.049) compared to smaller grids (0.073; Fig. S4a–c), many of these patterns are near-replicas of each other, differing only in a small number of pixels, and would therefore be clustered together for interpretation.



Figure S6: Equivalent to Fig. 3, but plotted against calendar dates. Comparing the two figure shows that although only patterns 1 and 2 happen in early-mid November, not all melt seasons initiate in this way.



Figure S7: Time series of the daily melt extent on the Shackleton Ice Shelf through the 2017/2018 melt season, plotted as a percentage of the full shelf area. No melt is observed for a week at the start of December 2017 after a clear spike in the melt extent at the end of November. Comparison between the two sensors shows that the extent is always greater in the AMSR observations, but that fluctuations in the melt extent tend to be observed by both sensors simultaneously. However, this is not always the case: in February, the AMSR-2 sensor observes melt missed by the SSMIS sensors, for which the melt season had already ended.



Figure S8: Relative occurrence of the nine melt patterns each summer. A value of 20 % of the y-axis indicates that 20 % of the days in a given melt season were best described by the given pattern. Years refer to the January date of the summer (2009 is 2008/09). No AMSR sensor operated in 2011/12. See also Fig. 5.



Figure S9: Examples of the differences in melt observed by the AMSR and SSMIS sensors over the Shackleton Ice Shelf. The central two columns display the binary melt observations for each sensor, and the columns to the left and right show the SOM pattern the observation was assigned to, as indicated by the value in the topleft of the plot; a value of 0 indicates that the melt season had not yet begun, or had already ended, according to the SSMIS observations.



Figure S10: The same as Fig. S9, for further dates.