Supplement of The Cryosphere, 18, 1733–1751, 2024
https://doi.org/10.5194/tc-18-1733-2024-supplement

The Cryosphere Open Access EGU

Supplement of

Alpine topography of the Gamburtsev Subglacial Mountains, Antarctica, mapped from ice sheet surface morphology

Edmund J. Lea et al.

Correspondence to: Edmund J. Lea (edmund.j.lea@durham.ac.uk)

S1 Automated Mapping Procedure

The full automated mapping procedure is summarised in Fig. S1 and detailed in the following sections.

S1.1 Pre-processing

All input data were first clipped to the extent of the study area in a GIS, with subsequent operations performed in Matlab (version 2021a), to allow for additional flexibility in some operations (such as the adaptive method of binary thresholding). The entire script is available in an online repository at https://doi.org/10.5281/zenodo.10550538. The input data and their spatial referencing information are read into Matlab using the "readgeoraster" function, allowing the data to be operated on using image processing techniques. The following pre-processing steps were required:

1. (REMA only) Data gaps were filled with zero values, as a necessary condition for constructing an adaptive threshold.
2. The data were smoothed using a 5-by-5 pixel median filter, in order to remove noise at the pixel level while preserving the sharpness of edges or contrasts within the data (Fig. S2b).
3. (REMA only) Negative curvature values were converted to zero; this was so that only positive values (representing convex slopes) were considered in the thresholding process. Since the convex portion of a step-change in surface elevation occurs up-flow of the concave portion (Fig. 4), this allows for consistent identification of the maximal slope as the downflow edge of a positive region.

S1.2 Adaptive thresholding

Before applying edge detection, a binary threshold was applied to each of the (pre-processed) input datasets. This produced a simplified representation of the data (mask) in which pixels were deemed either to have "high" (1) or "low" (0) curvature/brightness value, facilitating the identification of edges, particularly in regions of more gradual contrast. Rather than a global threshold value, an adaptive threshold matrix was constructed, in which a unique threshold was determined for each pixel based on the average intensity within a local neighbourhood (Fig. S2c). This method accounted for any spatial variability in contrast within the data, allowing for small but locally significant changes to be identified while reducing the impact of noise in high-variability regions.

The adaptive thresholding algorithm used has several parameters that can be tuned, including the neighbourhood size, the statistic applied, and a sensitivity factor. In order to identify the most suitable values to use, each parameter was systematically varied (Table S1), such that several different maps were produced from each of the two input datasets (Fig. S3). Differences introduced by the variation of factors include particularly the identification of small and/or connecting features. When selecting which results to investigate further, extracts from the maps produced (after further processing and edge detection were applied) were inspected and compared, with those in which these smaller features were more often identified judged to be more useful. There was to an extent a trade-off made here between detail and the presence of potentially erroneous/artefact features; settings

conducive to the identification of smaller features (low sensitivity, small neighbourhood size) were also more likely to misinterpret artefacts or noise in the data as genuine features. Based on this qualitative assessment, those maps selected as giving the optimal balance between representation of smaller features and oversensitivity to noise used the tuning parameters identified in bold in Table S1 (corresponding to extracts a and f in Fig. S3).

Table S1 – Tuning parameters for adaptive thresholding on the two datasets used.

Input Dataset	Parameter	Range / options	Increment	Optimal value
REMA mean curvature	Neighbourhood statistic	Mean, median, Gaussian weighted mean	N/A	**Mean**
	Sensitivity factor	0-1	0.1	**0.5**
	Neighbourhood size (N-by-N pixels / km)	25-275 / 5-35	50 / 10	**75 / 15**
RAMP image mosaic	Neighbourhood statistic	Mean, median, Gaussian weighted mean	N/A	**Mean**
	Sensitivity factor	0-1	0.1	**0.6**
	Neighbourhood size (N-by-N pixels / km)	25-275 / 5-35	50 / 10	**25 / 5**

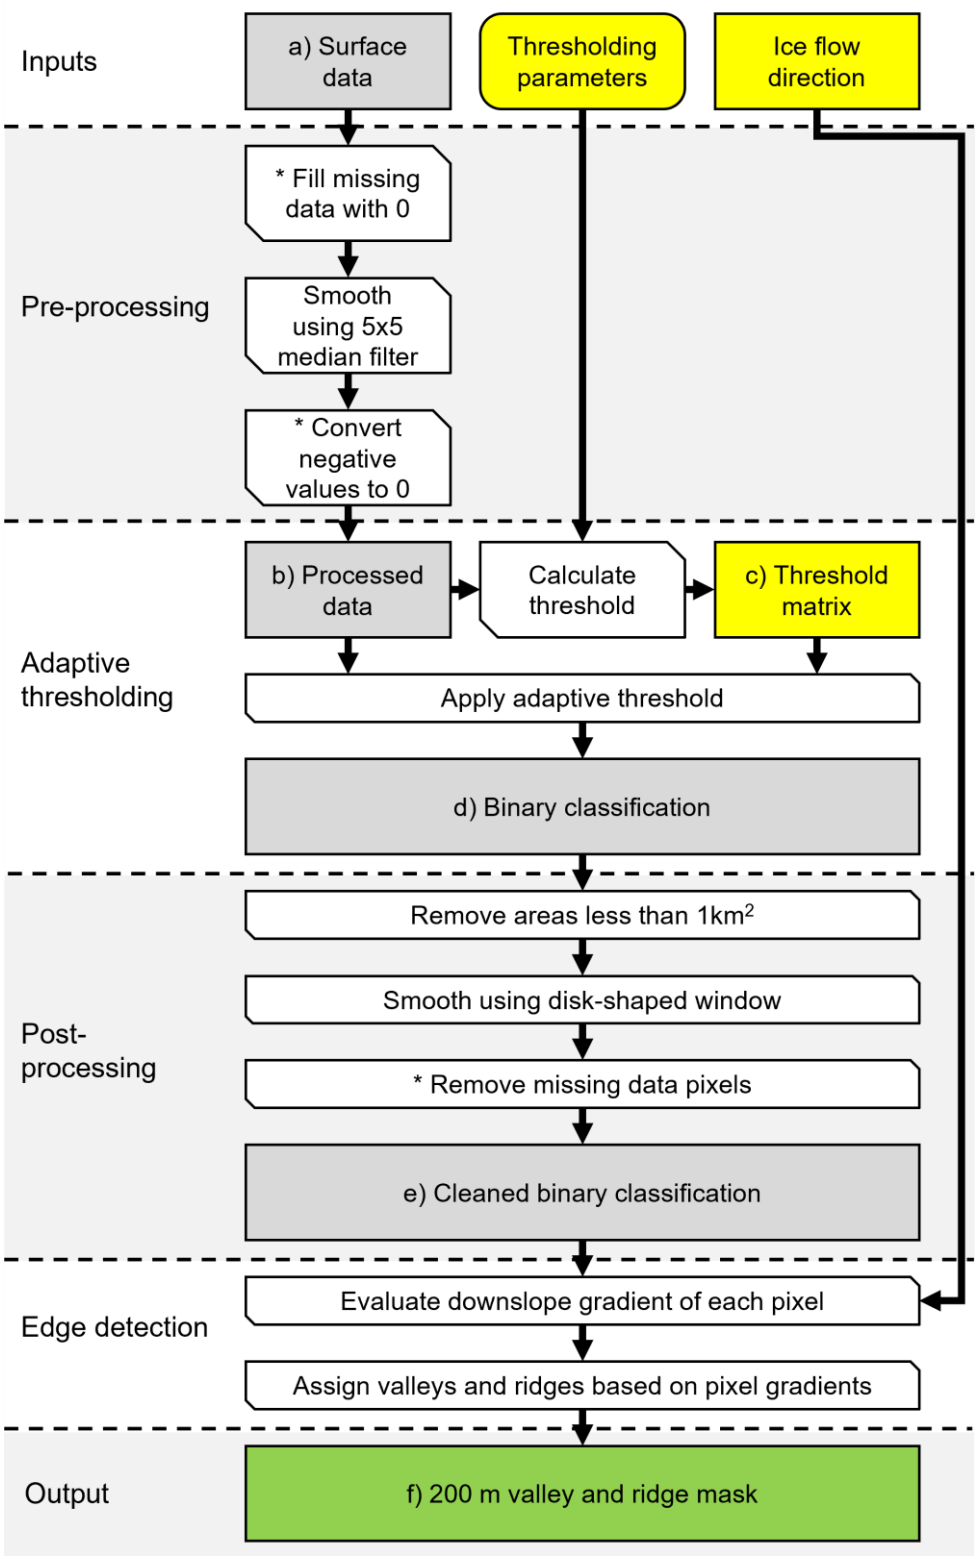
Inputs
a) Surface data
Thresholding parameters
Ice flow direction
Pre-processing
* Fill missing data with 0
Smooth using 5x5 median filter
* Convert negative values to 0
Adaptive thresholding
b) Processed data
Calculate threshold
c) Threshold matrix
Apply adaptive threshold
d) Binary classification
Post-processing
Remove areas less than 1km2
Smooth using disk-shaped window
* Remove missing data pixels
e) Cleaned binary classification
Edge detection
Evaluate downslope gradient of each pixel
Assign valleys and ridges based on pixel gradients
Output
f) 200 m valley and ridge mask

Figure S1 – Summary of automated mapping procedure. Grey boxes = mapping inputs; yellow = auxiliary inputs; green = mapping outputs. Square-ended boxes = raster datasets; round-ended boxes = vector datasets; cut-cornered boxes = operating steps. Steps marked with * were necessary only when using REMA curvature as the input surface dataset (due to data gaps).

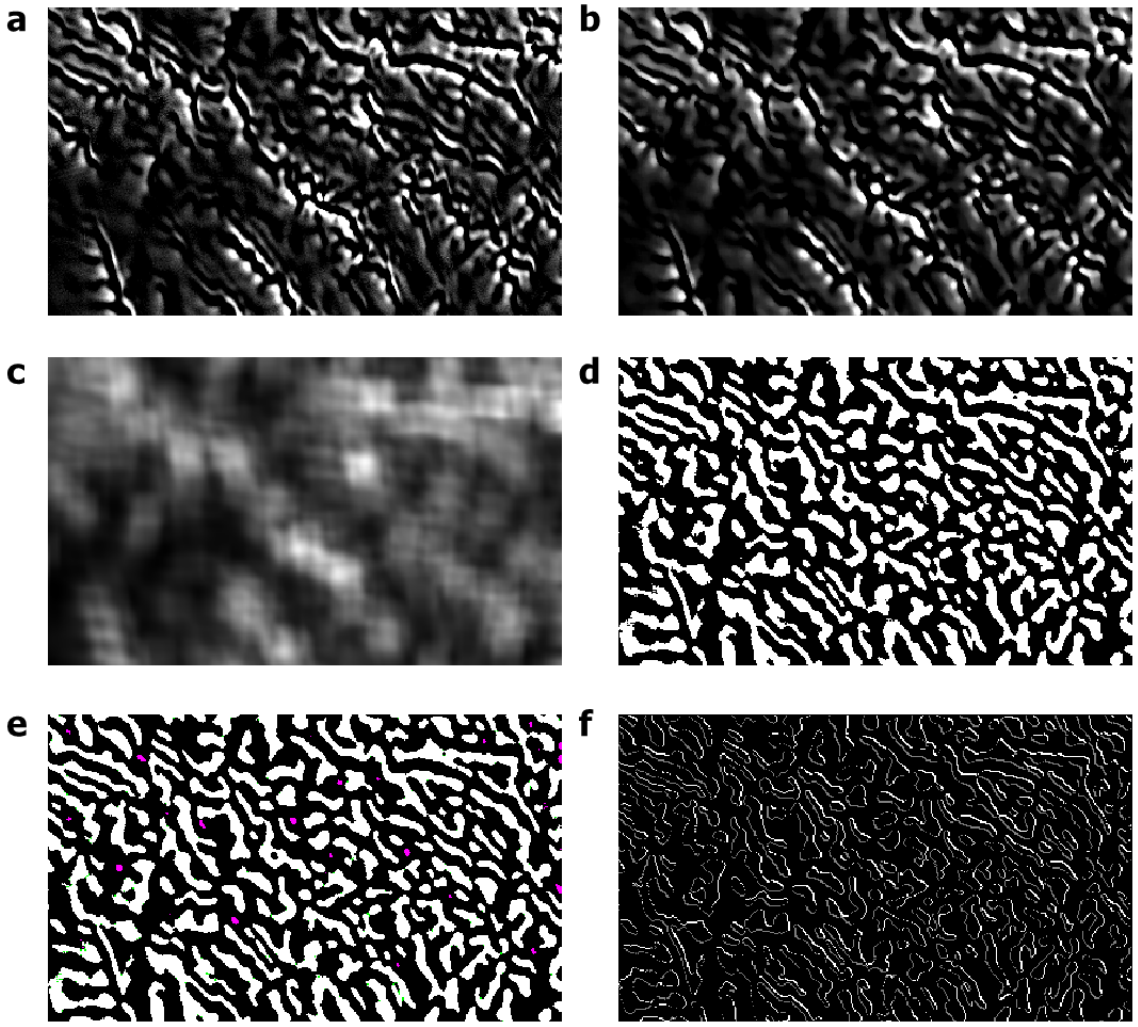

Figure S2 – Visualisations of different stages of the automated mapping procedure, using a 243-by-405 pixel (~ 4000 km2) extract from REMA mean curvature as the input data (a). Subsequent panels show: b) the input after application of a 5-by-5 pixel median filter; c) the adaptive threshold matrix (brighter pixels indicating a higher threshold); d) the initial binary classification; e) the cleaned binary classification (magenta = pixels removed due to filtering of small regions, green = pixels added due to smoothing); and f) the final valley and ridge mask (black = background (0), grey = valley (1), white = ridge (2)). Note that values in images a-c have been arbitrarily adjusted and stretched over a black-to-white scale for ease of visualisation. Panel lettering corresponds to marked stages in flowchart in Fig. S1.

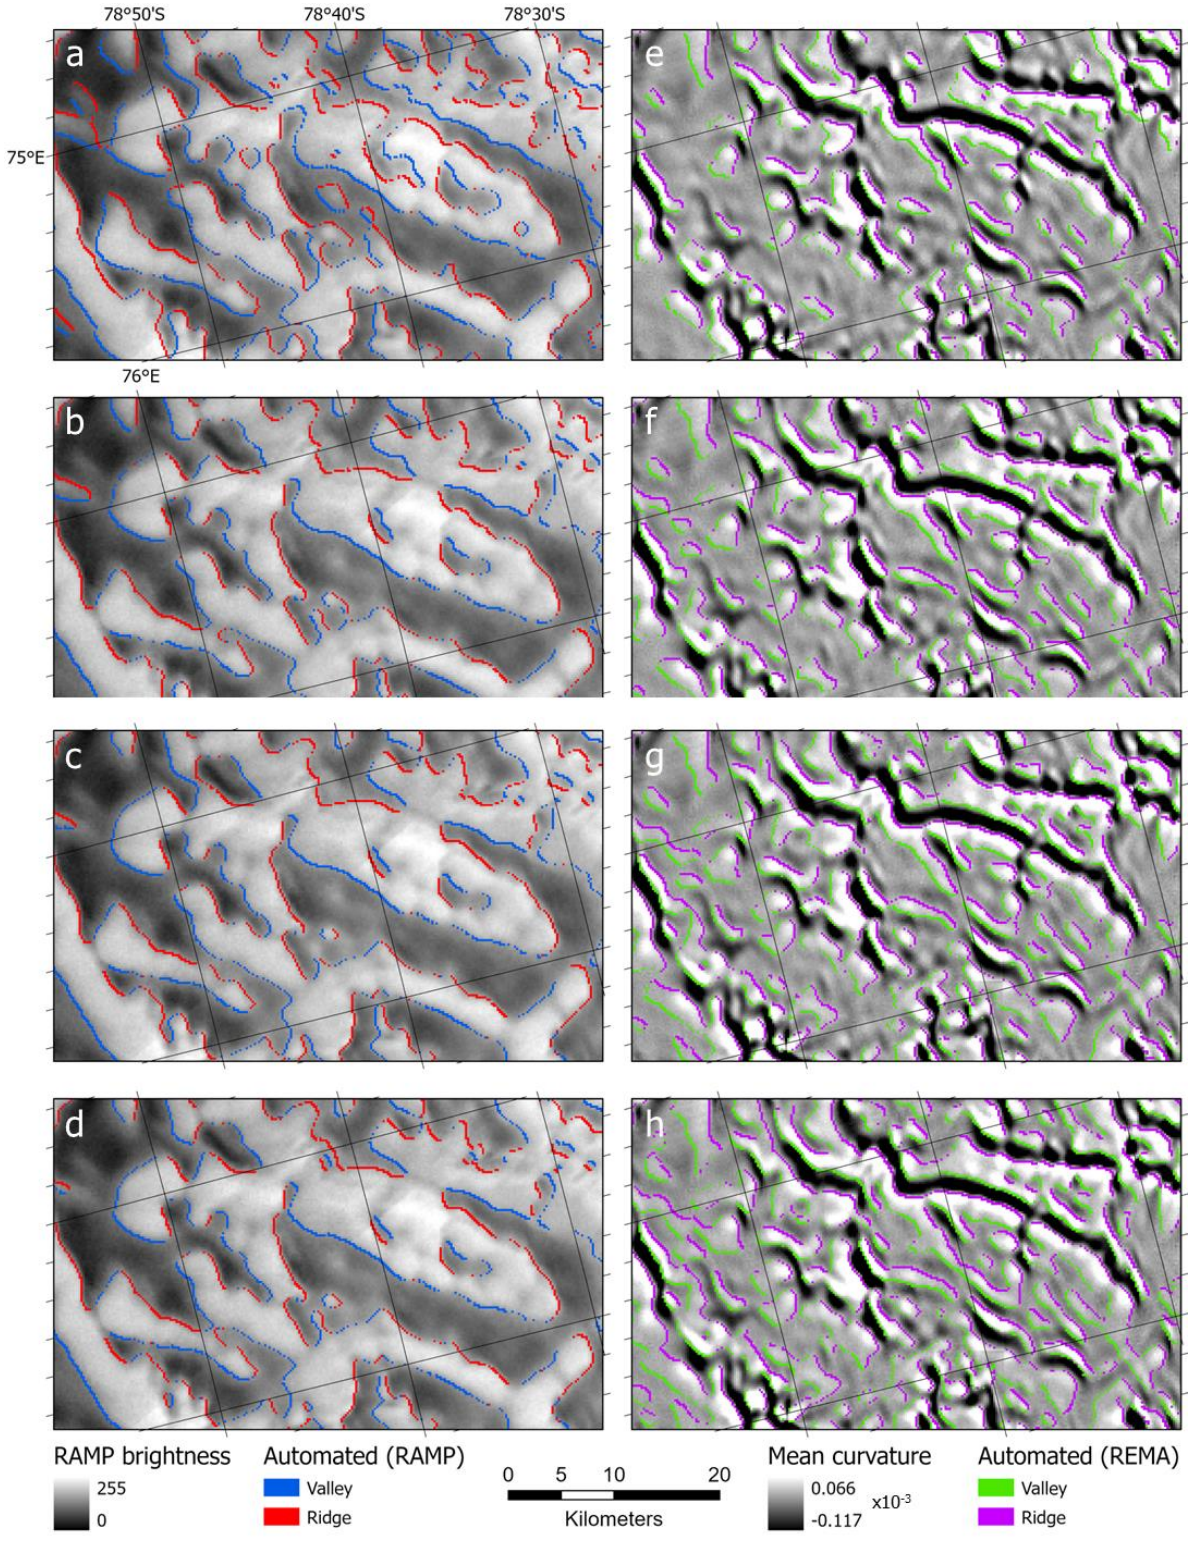
78°50'S
78°40'S
78°30'S
75°E
76°E
a
b
c
d
e
f
g
h
RAMP brightness
255
0
Automated (RAMP)
Valley
Ridge
0 5 10 20
Kilometers
Mean curvature
0.066
-0.117
x10-3
Automated (REMA)
Valley
Ridge

50

Figure S3 – Extracts from automated ridge and valley maps created using different input data, thresholding sensitivity and neighbourhood sizes, overlaid over the input data used (a-d: RAMP image mosaic; e-h: REMA mean curvature). For a-d: sensitivity = 0.5; neighbourhood size = a) 5; b) 15; c) 25; d) 35. For e-h: sensitivity = e) 0.1; f) 0.5 g) 0.7 h) 1; neighbourhood size = 15. Mean was used as the averaging statistic for all examples.

S1.3 Post-processing

The following clean-up operations were applied to improve the mask quality before moving on to edge detection (Fig. S2e):

1. All positive regions of the mask with areas of less than 1 km2 (25 pixels) were removed, as likely artefacts due to noise. This cut-off was selected because it had already been identified as the minimum appropriate window size for calculating curvature, due to noise dominating variability on smaller scales.
2. The mask was smoothed using a disk-shaped moving window of radius 1 pixel, to fill in small gaps likely caused by noise.
3. (REMA only) Every pixel for which original data were missing was removed from the mask.

S1.4 Edge detection

A custom-built edge-detection algorithm was developed in which the gradient in the cleaned mask was evaluated for each pixel according to the mean flow direction – derived from MEaSUREs - within the corresponding 200-by-200-m area. For the purposes of this comparison, flow direction was rounded to the nearest 8-directional value (i.e. 45° 'sectors'), and each pixel in the mask was compared with its neighbour immediately downflow. If there was no change in value (i.e. both pixels had value 1 or both pixels had value 0), the pixel was marked as neither a ridge nor a valley (0). If there was a positive change downflow (from 0 to 1, representing increasing curvature/brightness), the pixel was marked as a valley (1). If there was a negative change downflow (from 1 to 0, representing decreasing curvature/brightness), the pixel was marked as a ridge (2). Every pixel for which original data were missing was then converted back to a missing data marker (-1). This procedure is summarised in Table S2. The resulting edge mask (Fig. S2f) was saved as a GeoTIFF file using the same spatial referencing information as the original data.

Table S2 – List of values and their meanings in the final feature masks.

Change downflow	Value assigned	Meaning
N/A	-1	Original data missing
None	0	Neither a ridge nor a valley
Positive (0 to 1)	1	Valley
Negative (1 to 0)	2	Ridge

Only the immediate downflow neighbour was considered when assigning pixel values, and not proximity to other pixels of the same type. Such a check would be beneficial in a more advanced version of the procedure, but is not straightforward to implement in two dimensions, and was not attempted for the purposes of this paper.

S2 Manual Mapping Procedure

Additional clarifications on the manual mapping procedure are provided here, expanding upon section 2.4 in the main paper.

S2.1 Ridges

Ridge features proved to be easiest to observe and trace in the RAMP image and in the profile (along-slope) version of REMA surface curvature (Fig. S4, 1). The two datasets are broadly consistent in the shape and position of these features, though in places there is a small (generally less than 1 km) spatial offset between inferred ridge locations. In such cases, the location according to RAMP was preferred, due to the greater sharpness of the transitions representing the features in this image than in the curvature dataset making them easier to digitise consistently. In some areas, particularly near the ice divide, however, contrast in RAMP is greatly diminished; in these cases all positions were inferred from REMA. Additionally, while individual ridgelines were more clearly represented in RAMP, the connections between them were often much clearer in REMA profile curvature. Where possible, therefore, those ridges in an area able to be mapped from RAMP were digitised first, then those only clearly visible in REMA.

S2.2 Valleys

Valley lines were easiest to trace from RAMP, and from the plan (across-slope) version of REMA curvature (Fig. S4, 2), though many were also inferred from the positions of the surrounding ridges (Fig. S4, 3). Where valleys appeared narrowest (~ 1-3 km), the transition from dark to bright in RAMP often coincided with a distinct maximum in plan curvature. Where valleys were slightly wider (~ 4-6 km), this transition would sometimes coincide with the centre of a broader plateau in plan curvature, bounded up-flow by a maximum and down-flow by a minimum. The widest valleys, which constitute the main trunk of the inferred palaeo-drainage system, were the least well defined, being represented for the most part by smoothness in both RAMP and REMA plan curvature. Each valley was digitised from its upper end towards its terminus (usually where it joins another valley), such that the direction of each line would reflect the most likely direction of ice-free palaeo-drainage. As much as was practicable, each valley line was extended as far as possible rather than segmenting what were likely to be continuous drainage routes, in order to preserve valley connectivity and the usefulness of valley length as an informative measure of the landscape.

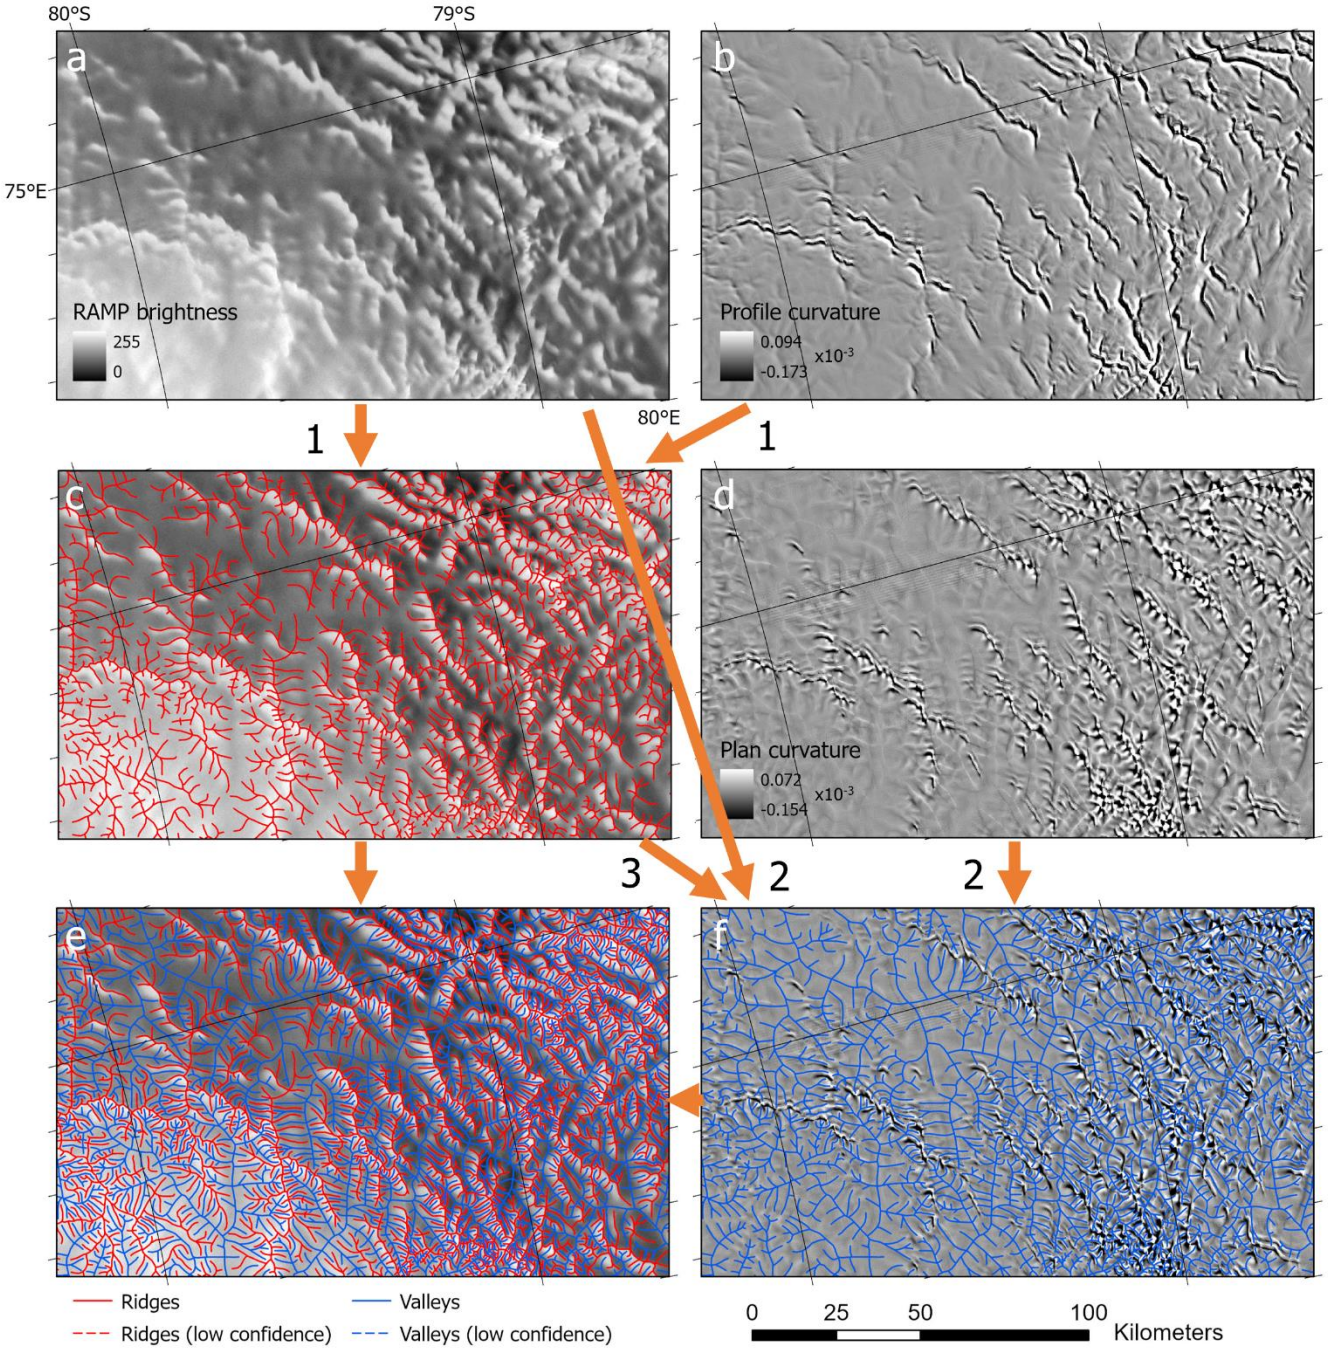

Figure S4 – Extracts from manual mapping demonstrating the datasets and procedure used to digitise ridges and valleys. Datasets: RAMP image mosaic (a, c, e); REMA profile curvature (b); REMA plan curvature (d, f). Orange arrows and numbers indicate steps referred to in text.