

S1 Frontal area change

S1.1 Model prediction postprocessing scheme

We develop an original postprocessing routine to prepare the model's predictions for monthly averaging and subsequent area change computation. The postprocessing scheme includes general cleaning of apparent model mispredictions, initial outlier prediction detection, restoration of ice pixels lost to adjacent glaciers, and reclassification of nodata predicted pixels. To assist the postprocessing workflow, vector and raster layers are generated for each RGI box. First, a postprocessing RGI box mask is defined to distinguish glacier territory between spatially overlapping RGI boxes and to sample predicted pixels only corresponding to ice and ocean zones, since we assume a constant land extent throughout the study period. As 87 of the 147 RGI boxes overlap an adjacent box, we manually modify these boxes to ensure each box covers its own glacier outlet area. Most boxes are adjusted by using the constrained fjord morphology to define exclusive outlet regions; however, boxes sharing a common, near-ocean ice boundary must be separated by the predefined RGI v6 basins. From the postprocessing mask, we create a raster version for each RGI box, as well as a combined zone label raster from the ice, ocean, and land geometries (i.e. the training labels) within the box (hereafter referred to as the zone label).

To begin postprocessing, all predictions are converted from PNG format to georeferenced data (GeoTIFF) with the RGI box's reference transform. The reclassification and cleaning of predicted pixels is applied on connected pixel clusters, rather than individual pixels, to optimize processing efficiency. Inland ocean clusters that have no contact with a zone label ocean pixel are reclassified as ice, while ice clusters that do not touch a predicted nodata or ice zone label pixel are reclassified as ocean. After the initial ice and ocean pixel cleanup, we attempt to detect obvious prediction outliers by identifying connected ice clusters that carry the model's patch size footprint along a predicted ice-ocean boundary. Because the model predicts on square, 256x256 pixel patches, an outlier prediction may imprint fragments of the patch footprint as a morphologically implausible straight row or column of ice pixels (Fig. S1a, left panel). The zone label is used to define an adaptive threshold for detecting such artifacts, based on the glacier's size and known frontal morphology, both of which are inherent information in the ground truth zone label. In the zone label, we first count the number of pixels in the longest contiguous ice-ocean pixel boundary (i.e. straight row or column of pixels). If the pixel count is less than 20 (i.e. 200 m) for the given glacier, its predictions are defined as having a patch footprint artifact when there is a straight ice-ocean boundary greater than 64 pixels (i.e. 25% of the model's patch size); if the zone label pixel count is 20 or more, patch size artifacts are flagged when the prediction's straight pixel boundary is greater than 128 pixel (i.e. 50% of the model's patch size). The 20 pixel threshold was a data-driven decision that generally separates small and large glaciers, as well as glaciers that have complex vs simple frontal geometry in the zone label.

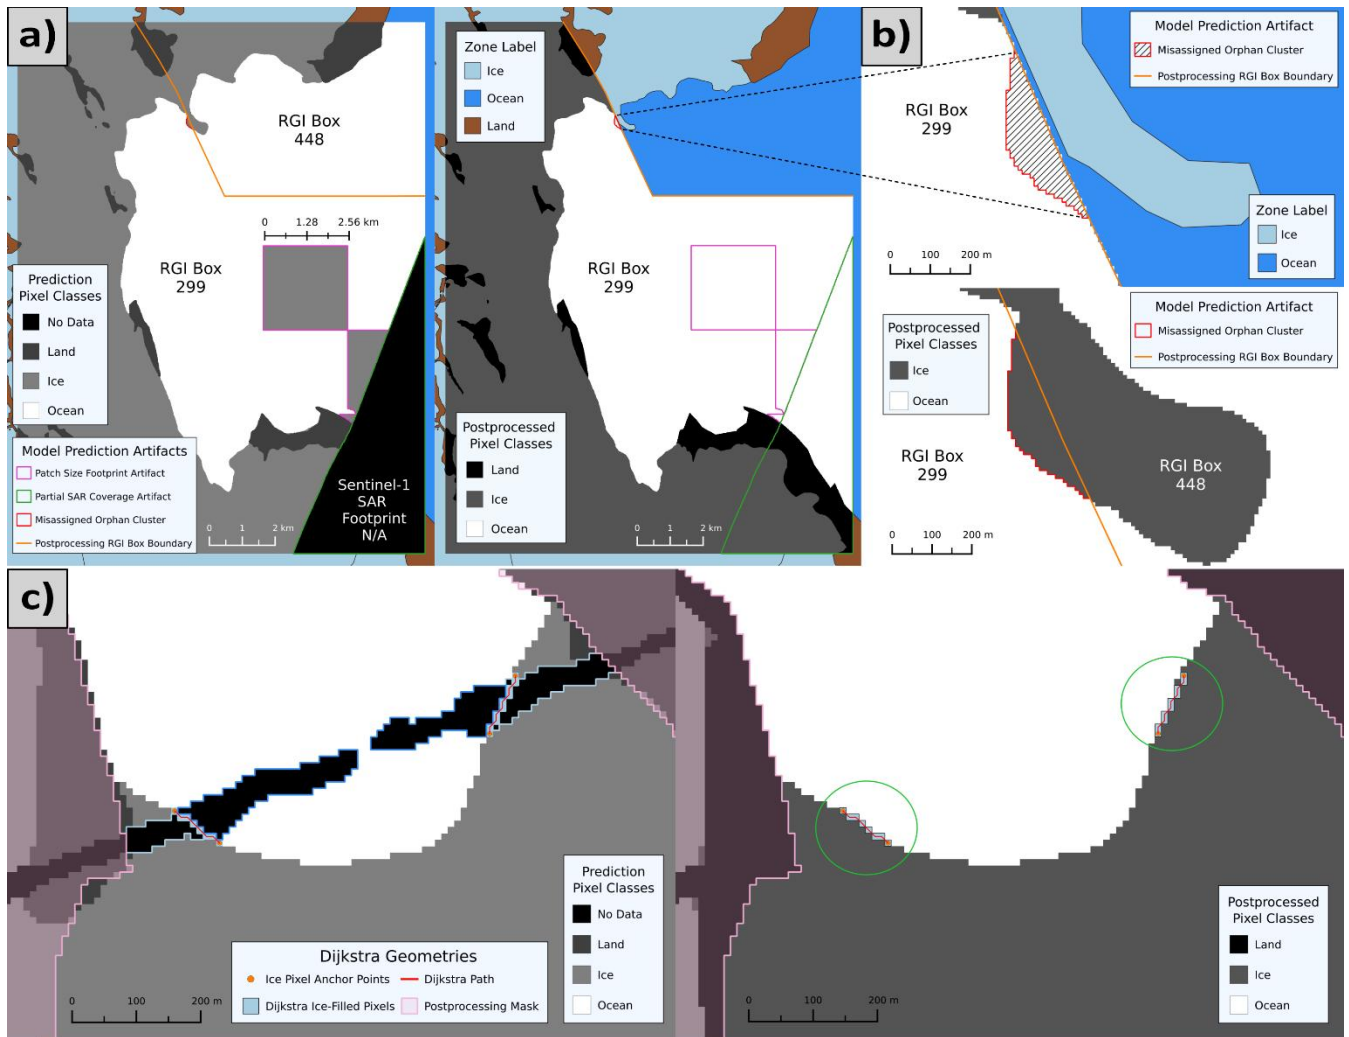
To handle detected artifacts, ice clusters with this morphology are converted to ocean if they do not touch a zone label ice pixel, and if touching labeled ice, the prediction is excluded from further postprocessing; we detect 3,315 artifact predictions meeting the exclusion criteria.

The next postprocessing step is reassigning ice pixels that moved into an adjacent glacier's postprocessing mask, back to the correct neighboring glacier (referred to as the parent glacier). Since the postprocessing bounds of some glaciers share a common lateral ice-ocean boundary, terminus changes can result in predicted ice pixels migrating into the adjacent glacier's territory (Fig. S1b). To identify such misassigned clusters (referred to as orphan clusters), the number of all ice clusters are counted and labeled, then the RGI box's postprocessing mask is applied, and the ice clusters are recounted and relabeled, then compared with the initial ice cluster count and labels. New ice clusters created after applying the mask are polygonized and stored in a table for possible reconstruction to the parent glacier's prediction at the final postprocessing step (i.e. after all cluster cleaning; see last paragraph in this section); for clarity, these new orphan ice clusters are always reclassified to ocean in the prediction of the falsely attributed RGI box, to ensure we do not double count them if they are later reconstructed in the parent glacier's prediction. When the postprocessing mask is applied (during the orphan cluster identification), predicted land pixels are replaced with a temporary value and we force all land pixels in the zone label to land in the prediction, as land is assumed to be constant. Like the initial ocean and ice cluster cleanup, nodata and the temporary land clusters are reclassified. Nodata clusters that only border ice or ocean are reassigned to ice or ocean, respectively. The mispredicted temporary land clusters are reclassified as ice, ocean, or nodata, respectively, if the cluster only touches that specific label type. If the land cluster borders both ice and ocean, it is classified to the majority bordering label, and if the cluster is surrounded by all three other label types, it is forced to nodata.

50

The remaining nodata predicted clusters, which are in contact with the predicted ice front, are reassigned with a complex set of operations. To simplify, the general approach is to connect pixel paths through a nodata cluster using Dijkstra-based pathing (Dijkstra, 1959), then infilling the pixels along the path with ice values (Fig. S1c); this process may eventually result in a split of the nodata cluster into two distinct clusters, which are then classified using the majority surrounding pixel class. The Dijkstra path is connected to ice pixel anchor points on opposing sides of the target nodata cluster, which are defined at a triple junction of ice, ocean, and nodata predicted pixels (Fig. S1c). The path structure is built with a cost function where each pixel type is assigned a price, and the chosen route minimizes the cost while preferring the shortest straight path of pixels through the nodata cluster. An infilled ice pixel path therefore represents a plausible terminus morphology for undefined segments along the ice front. The targeted application of this process is to reconstruct predicted nodata areas that split or discontinue the ice-ocean pixel boundary, which primarily occurs for SAR images that were mosaicked during the preprocessing and had a small (<100 m) nodata gap along the interface where two images were merged, or for nodata areas resulting from radar shadow caused by the ice cliff height at the terminus. After the Dijkstra pathing, a final cluster cleanup is performed on all pixel types to address any unresolved clusters.

60



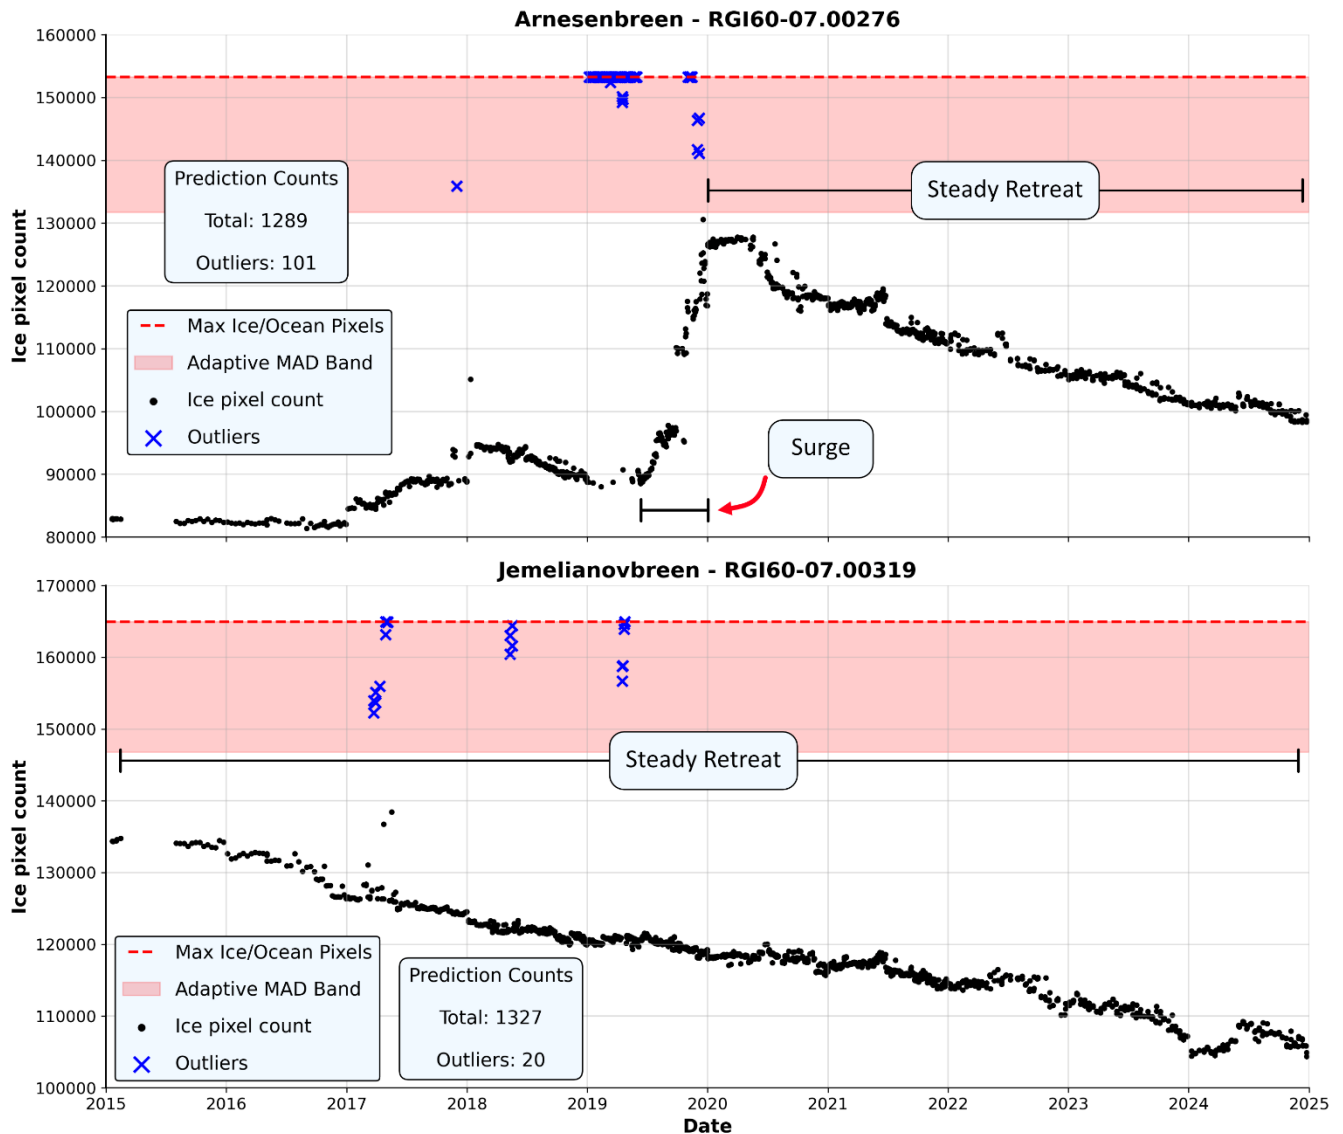
65 **Figure S1.** Shows several examples of a model segmentation prediction and the postprocessed output. **a)** Collection of model prediction artifacts (patch size footprint, partial coverage SAR imagery, and orphan ice clusters) and their handling from the prediction (left) to the postprocessed image (right). RGI box 299 and 448 share a common boundary, and as such, are split with the postprocessing box geometry and postprocessed separately. **b)** Zoomed window of the orphan ice cluster along the RGI 299/448 boundary. Top panel shows the missing ice cluster in RGI Box 299 that was naturally removed during postprocessing but saved and sutured onto the RGI box 448 postprocessed prediction in the bottom panel. **c)** Dijkstra's path reconstruction of a prediction's nodata areas (left) along the ice front with reasonable morphology in the postprocessed image (right panel; green circles).

70

With the cleaned predictions, for the last postprocessing step, the stored polygonized orphan ice clusters are burned back on to their parent glacier's corresponding prediction in the N/A area outside the postprocessing bounds (Fig. S1b, bottom panel). We remove any orphan clusters that originate from predictions with model patch footprint artifacts and ensure an orphan cluster has at least one pixel with a physical edge connection to an ice cluster in the parent prediction. Finally, an additional statistical outlier detection step is performed, before monthly averaging of the predicted ice front positions. For each RGI basin, an ice area time series is quantified by counting pixels belonging to connected ice clusters that intersect fluxgate locations (see Sect. 3.2.1 in main paper for fluxgates). To identify implausible segmentations with respect to intra- and

75

interannual variability, we apply an adaptive, glacier-specific outlier detection scheme based on a median absolute deviation (MAD) envelope, anchored to a physical area constraint (Fig. S2). The maximum number of valid pixels in an RGI box's postprocessing raster defines a physical upper bound for plausible ice extent. A lower bound tolerance margin is derived from the relative variability of the time series, quantified using the MAD, such that glaciers with higher variability yield a narrower admissible range, while more stable time series permit a wider range. Predictions whose ice pixel counts fall within the MAD-constrained envelope are flagged as outliers and validated manually by inspecting the pixel count time series and MAD envelope (Fig. S2). Flagged predictions often reflect poor model segmentations arising from homogenous SAR backscatter between ocean and ice. In total, 1,636 predictions meet the outlier criteria, leaving 198,343 predictions for monthly averaging across all RGI basins.



90 **Figure S2.** Shows two examples of the MAD adaptive band envelope to prefilter model prediction outliers, before the monthly averaging of predictions (see Sect. 3.1.5 in main paper). Top panel shows a more variable ice pixel count time series at Arnesenbreen glacier due to a surge event in 2019, compared to a more stable time series at Jemelianovbreen glacier in the bottom panel. Ice pixel count variability affects the relative size of the adaptive MAD band per glacier.

S1.2 Frontal displacement rate uncertainty table

We include additional information for the glaciers where the assessment of the calving front predictions was carried out (see Sect. 3.1.6 in main paper). Geographic distribution is varied with three glaciers in AF and SS, and one glacier each in BE, NE, 95 NW, and VF (Table S1). Terminus lengths range from 1-50 km during the evaluation period in 2016, with eight glaciers between 3-15 km length. From the updated surge-type fields in RGI version 7 (RGI 7.0 Consortium, 2023), four glaciers have observed surge behavior and another four have no surge evidence, while two are tagged as probable surge-type. Strongbreen glacier began actively surging in mid-2016 (Benn et al., 2019; Koch et al., 2023), for which we have a ground truth and model prediction pair in July, August, September, and December.

100

Table S1. Mean spatial front uncertainty for all 10 evaluated glaciers. Includes geographic identifiers, terminus length, and surge-type classification from RGI version 7. The count of available ground truth and model-predicted polygon pairs (max of 10) is also provided. The mean regional frontal uncertainty is 38.1 meters and applied to all glaciers in Svalbard.

Name	Region	RGI-Id (v6)	Mean Front Uncertainty (m)	Number of Pairs	Terminus Length (km)	Surge-Type
Fonndalsbreen	AF	RGI60-07.00552	10.0	10	1.4	No Evidence
Storøyjøkulen	AF	RGI60-07.00071	19.5	10	11.8	No Evidence
Lillichøøkbreen	NW	RGI60-07.00661	20.5	10	14.1	No Evidence
Strongbreen	SS	RGI60-07.00296	22.4	8	8.1	Observed
Nuddbreen	SS	RGI60-07.00422	30.2	9	3.0	Probable
Sonklarbreen	NE	RGI60-07.00892	32.9	6	12.4	Observed
Bråsvellbreen	AF	RGI60-07.00025	38.1	10	49.1	Observed
Gimlebreen	VF	RGI60-07.01534	56.5	10	8.4	No Evidence
Doktorbreen/Nathorsbreen	SS	RGI60-07.01474	58.5	10	4.4	Observed
Deltabreen	BE	RGI60-07.01546	69.9	6	9.8	Probable

S2 Ice discharge

105 S2.1 Monthly velocities and flow angles

To obtain a continuous monthly time series, gaps in the monthly velocity magnitude, flow direction, and the associated uncertainties are interpolated. ITS_LIVE datacubes only contain Sentinel-1 derived velocities through December 2022, meaning observations over polar night periods between January 2023 and December 2024 are currently unavailable. We record a data coverage flag for each flux point and month, indicating whether at least one velocity observation contributes to the
110 monthly estimate or if the value is interpolated. Velocity magnitudes (v) are linearly interpolated in two stages: (1) spatially across a fluxgate's flux points for each month, and (2) temporally along each flux point's time series to fill any remaining gaps. Flow directions (θ_m) undergo circular interpolation, where missing monthly values are filled by linearly interpolating the sine and cosine components of the monthly circular mean angles to avoid artifacts from angular discontinuity. Finally, the velocity magnitude and flow direction uncertainties (σ_v and σ_{θ_m} , respectively) are linearly interpolated along the flux point
115 time series. These interpolations provide an initial complete monthly time series, which is subsequently masked using several quality control filters.

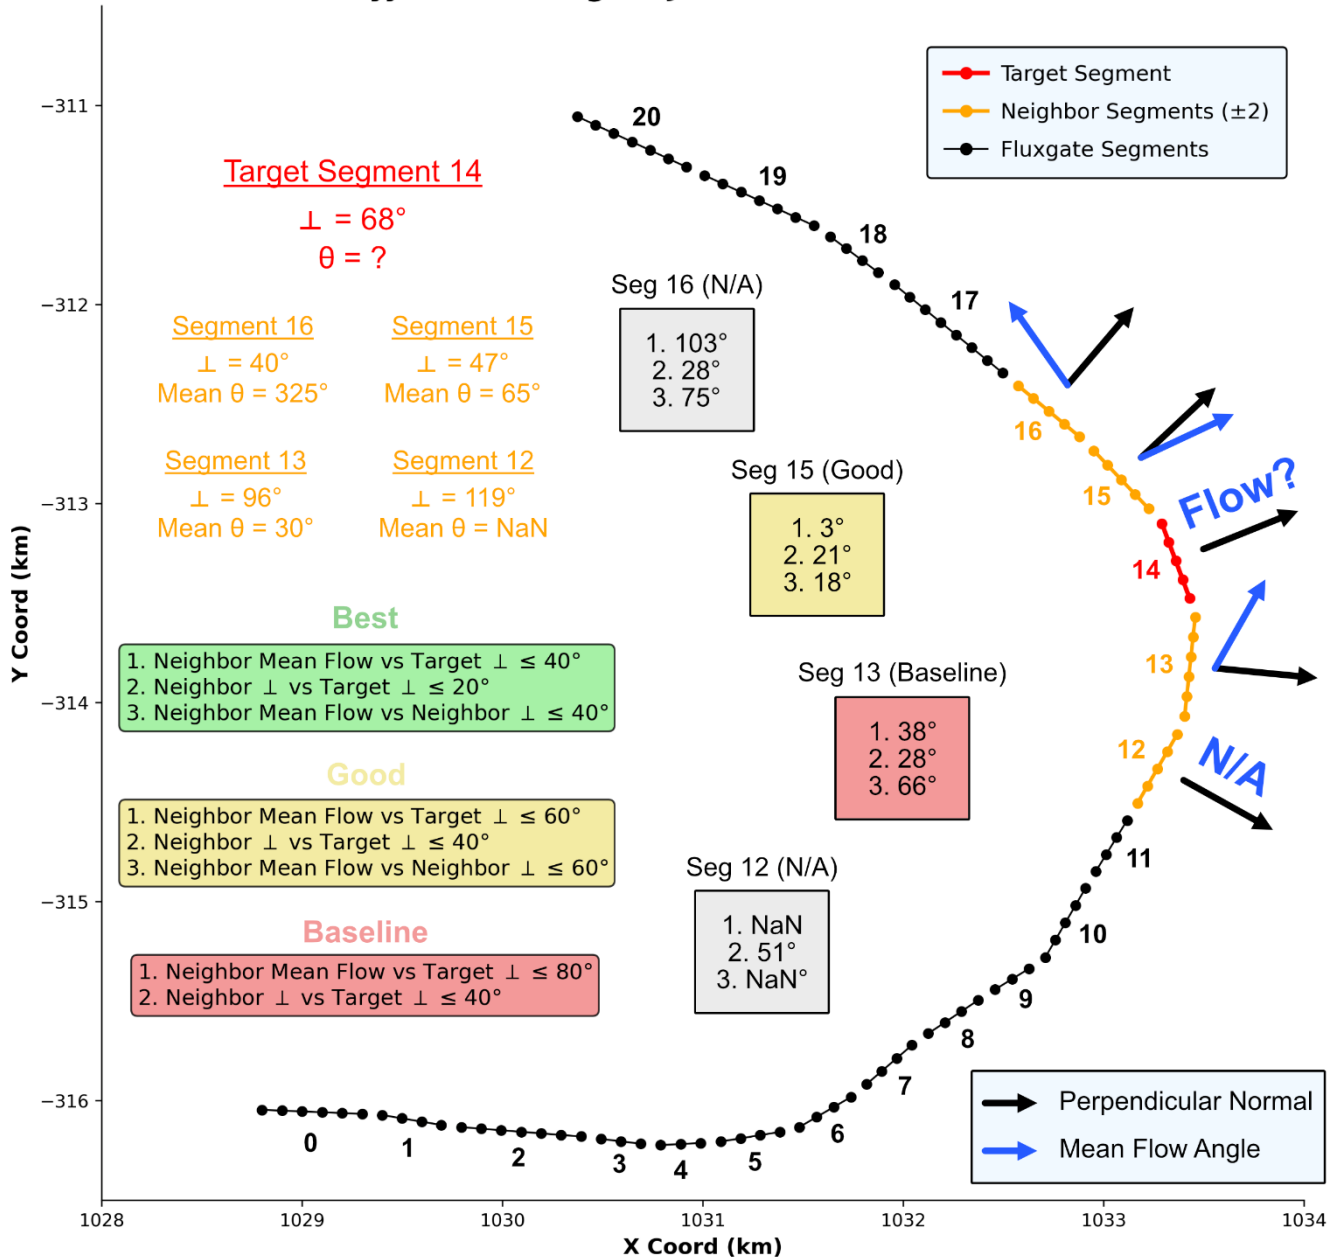
To remove physically inconsistent and poorly constrained measurements, the velocity, flow direction, and uncertainty time series are masked according to four criteria. First, the angular difference between the measured monthly flow direction and the
120 perpendicular normal flow angle is computed for each flux point and month, and values with an angular deviation $\geq 80^\circ$ are masked, as they are inconsistent with the expected flow orientation towards the ocean. Second, measurements with a perpendicular-normal velocity component below 5 m yr^{-1} are masked to remove negligible or stagnant flow. Third, months with insufficient spatial coverage – where fewer than 10% of flux points contain valid measurements after the previous masks – are removed across all points and variables, ensuring that sparsely sampled periods do not bias downstream discharge
125 calculations. Finally, flux points located within the regional land zone label are masked, since they do not represent valid glacier flow. The masked, non-land values are reconstructed using the remaining well-constrained originally sampled ITS_LIVE measurements.

S2.2 Filtered reconstruction

Masked velocity magnitudes are reconstructed using the same two-stage interpolation strategy applied during the initial
130 monthly infilling, where spatial interpolation fills values along the flux points per month, then any persisting gaps are temporally interpolated along a flux point's time dimension. The velocity magnitude uncertainty is also infilled using spatial interpolation along the flux points for each month. If originally sampled ITS_LIVE uncertainties exist, we use a conservative 75th percentile of the sampled distribution as the spatial fill value; however, if only initial interpolated values exist, the 75th percentile of the interpolated value distribution is used to fill. After the spatial interpolation, any missing uncertainties are
135 infilled temporally using linear interpolation for each flux point across the monthly time series.

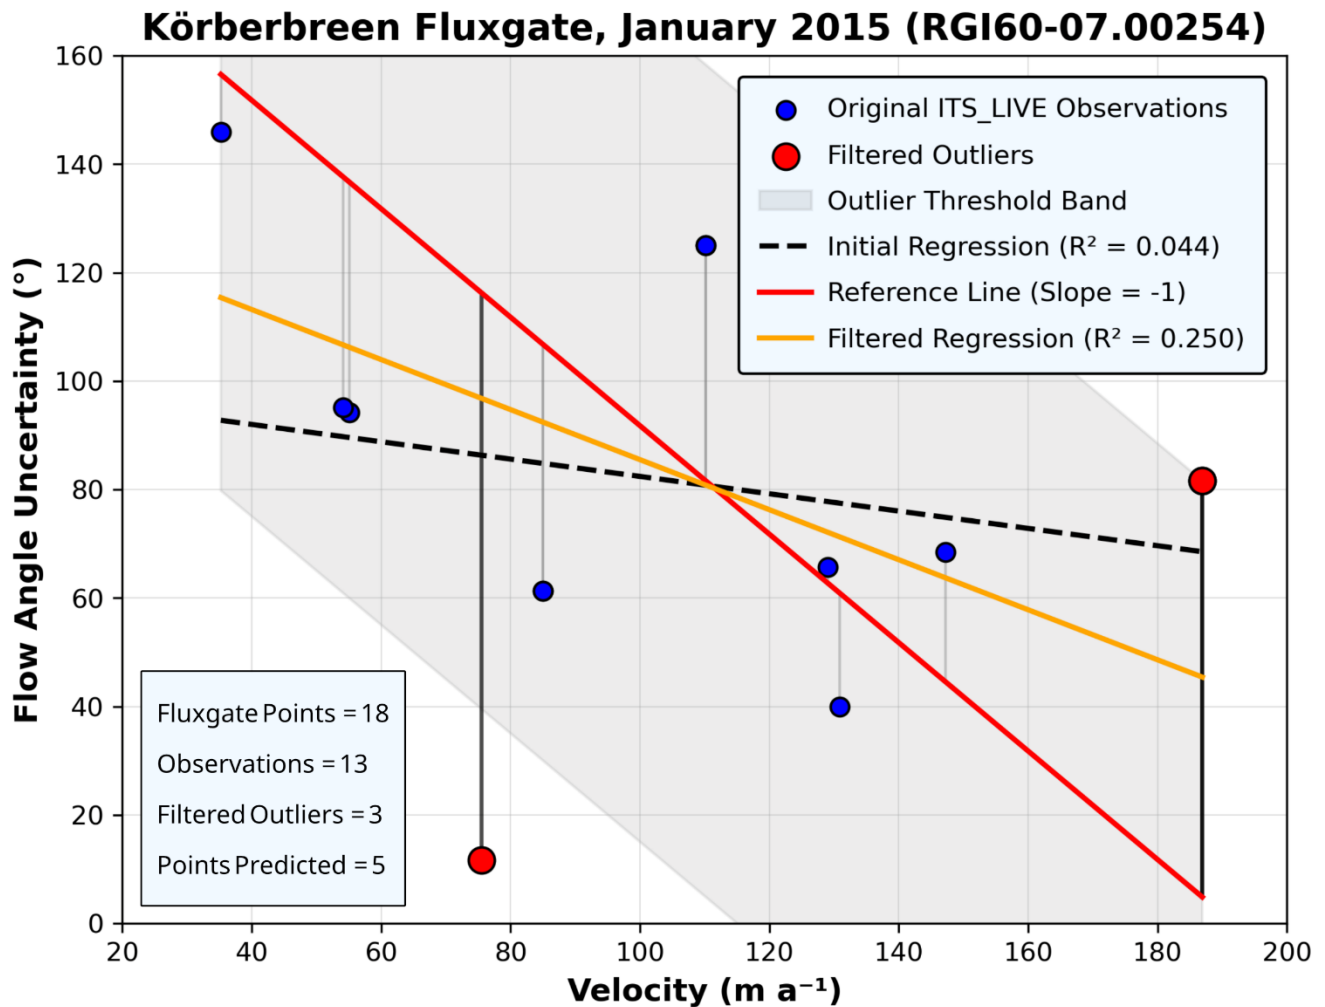
The reconstruction of masked flow directions is a more complex multi-pass infilling process, designed to preserve and extend the physically well-constrained ITS_LIVE monthly flow measurements. All reconstructed flow angles for all infill pass strategies are required to remain within $\pm 80^\circ$ of the perpendicular normal flow direction; infilled values violating this constraint are discarded at every infill step. For each fluxgate, flux points are grouped into spatially contiguous, straight segments (i.e. continuous slope) along the gate's geometry (Fig. S3). The first three infill passes attempt to reconstruct the flow directions spatially. In the first infill pass, we reconstruct a segment's missing flow directions locally using the available flow angle observations within the segment. If at least two valid observations exist, a distance weighted circular interpolation is applied, and if only a single valid angle exists, that angle is propagated across the segment. For the second infill pass, incomplete local segments are reconstructed borrowing flow angle information from neighboring segments; candidate neighbors are defined as being within ± 2 segments of the targeted infilling segment (Fig. S3). These adjacent segments are evaluated based on their mean flow direction, mean perpendicular-normal direction, and the target segment's perpendicular normal direction. The quality of neighboring segments is classified into three groups – best, good, and baseline – using increasing tolerance thresholds of the evaluated flow direction and perpendicular metrics (Fig. S3). The best and good classified segments are preferred, but baseline neighbors are used if higher ranking segments are not available. Missing flow angles on the target segment are infilled with the circular mean flow angle of the originally sampled ITS_LIVE measurements in the chosen neighbor segment, or as a fallback, the circular mean of the neighbor's available flow directions, which may have been infilled locally in the first infill pass. In the third spatial infill pass, like the first pass, another distance weighted circular interpolation is performed locally on a segment's valid flow angles, leveraging any new directional information kept from the neighbor fills in the second pass. For segments without a single valid flow angle after the first two passes, flow direction is inferred from the circular mean flow angle of the immediately adjacent segment(s) along the fluxgate. In the fourth infill pass, very few flow angles are missing, nevertheless we fill these remaining gaps temporally along a flux point's monthly time series using circular temporal interpolation. If any flow measurements are still unresolved because of the enforced $\pm 80^\circ$ to the local perpendicular-normal requirement, the fifth and final fallback infill pass assigns the local perpendicular flow direction to the missing value, ensuring a complete monthly time series. Finally, the angular difference ($\Delta\theta$) between the updated monthly flow direction and the perpendicular normal direction is recomputed for all flux points and months.

Storøyjøkulen Fluxgate June 2019 (RGI60-07.00071)



165 **Figure S3.** Schematic of flow angle reconstruction at the Storøyjøkulen glacier fluxgate in June 2019. Continuous slope segments are split and annotated (0-20) along the fluxgate geometry. The best, good, and baseline quality classification thresholds for neighboring segments are provided in their respective text boxes. This example is focused on infilling missing flow angles for flux points at target segment 14, by evaluating the flow angle information of neighbor segments ± 2 from the target segment (i.e. 12/13 and 15/16). Perpendicular normals (\perp) and mean flow angles (θ) are given for all relevant segments. Each of the evaluated neighbor segments are classified based on the threshold calculations (numbered 1-3 in segment boxes); segments 12 and 16 are unclassified as they do not meet the baseline thresholds, while segment 15 (good) and 13 (baseline) may attempt to infill the segment 14 missing flow angles. Flow vectors for segments are plotted in black (\perp) and blue (θ), with 0° defined as north on the plot.

170 Flow direction uncertainty is reconstructed first using a velocity-dependent regression, followed by conservative spatial and
temporal infilling. Velocity magnitude is used as a proxy predictor because directional uncertainty in ITS_LIVE arises from
image-pair correlation quality; small, noisy displacements at low velocities produce larger angular uncertainty, whereas
stronger, well-defined correlation peaks at higher velocities yield more stable flow direction estimates (Gardner et al., 2025).
This implies a physically motivated inverse relationship between velocity magnitude and angular uncertainty, which is used
175 as a reference in the reconstruction procedure. For each fluxgate and month, an initial linear regression between velocity
magnitude and flow angle uncertainty is fit using only originally sampled ITS_LIVE observations (Fig. S4). To reduce
sensitivity to outliers, observations are filtered using percentile-based residual thresholds defined relative to a median-centered
reference line with a fixed slope of -1 , representing the expected inverse scaling between velocity and angular uncertainty
(Fig. S4). The filtering threshold is adapted based on the sign and strength of the initial regression slope, allowing more
180 aggressive filtering when the observed relationship deviates from the expected physical behavior. Importantly, this physical
relationship is only used to identify outliers and is not enforced in the regression itself; if the relationship remains positive or
weakly negative after outlier filtering, the regression is kept. After outlier removal, a second regression is fit to the retained
observations and missing flow direction uncertainties are infilled using the prediction (Fig. S4). Regression-based infilling is
applied only when at least three valid ITS_LIVE observations are available; months with insufficient data or zero-slope
185 regressions after outlier filtering are excluded. All reconstructed uncertainties are constrained to physically meaningful bounds
between 0° and 180° . Remaining gaps after regression-based infilling are filled using a conservative 75th percentile of the
original ITS_LIVE sampled distribution. If original samples are unavailable, the distribution of all interpolated values is taken
as the monthly spatial fill value. Finally, the flow direction uncertainty time series is completed by linear interpolation along
the monthly dimension at each flux point.



190 **Figure S4.** Example of flow angle uncertainty reconstruction at Körberbreen glacier in January 2015. The initial regression is produced from
 all fluxgate points with an original (i.e. non-interpolated) ITS_LIVE velocity acquisition. The outlier threshold band is defined from the
 reference line (slope -1), which assumes a physical relationship of increasing velocity magnitude and decreasing flow angle uncertainty.
 Residuals are computed for each observation relative to the reference line to distinguish outliers. The outlier filtered regression predicts the
 flow angle uncertainty for missing flux point measurements in the target month. In this example, the fluxgate has 18 flux points and 13 valid
 195 observations, with 3 outliers exceeding the percentile-based residual threshold. Therefore, 5 flux points receive predicted flow angle
 uncertainties from the filtered regression based on the 10 non-filtered observations.

S3 Climatic mass balance

A valid near-ocean MAR pixel within the ice polygon may have a CMB value of 0, which is physically unlikely and indicative
 of a pixel whose area is dominated by ocean. The CMB for these pixels is estimated using the average mean monthly CMB of
 200 adjacent pixels that share a common edge with the 0-value pixel; neighboring pixels that only connect by corner contact are
 not included in this average. For basins containing multiple polygons or domains, the CMB is summed. Storøyjøkulen is the

only glacier missing reference CMB data, as the entire outlet region is covered by two MAR pixels that estimate a CMB of 0 due to ocean presence and the small size of the island glacier. The Storøyjøkulen climate data is filled using the data of the two nearest glaciers – Austfonna Worsleybreen (20km) and Austfonna Basin 8 (25 km; AF Basins defined in Dowdeswell et al., 2008; Zheng, 2022) – as the climatic conditions are presumably similar; these neighbor glaciers filled all monthly estimates except for December 2015, which is linearly interpolated.

References

- Benn, D. I., Jones, R. L., Luckman, A., Fürst, J. J., Hewitt, I., and Sommer, C.: Mass and enthalpy budget evolution during the surge of a polythermal glacier: a test of theory, *J. Glaciol.*, 65, 717–731, <https://doi.org/10.1017/jog.2019.63>, 2019.
- 210 Dijkstra, E. W.: A note on two problems in connexion with graphs, *Numer. Math.*, 1, 269–271, <https://doi.org/10.1007/BF01386390>, 1959.
- Dowdeswell, J. A., Benham, T. J., Strozzi, T., and Hagen, J. O.: Iceberg calving flux and mass balance of the Austfonna ice cap on Nordaustlandet, Svalbard, *J. Geophys. Res.*, 113, 2007JF000905, <https://doi.org/10.1029/2007JF000905>, 2008.
- 215 Gardner, A. S., Greene, C. A., Kennedy, J. H., Fahnestock, M. A., Liukis, M., López, L. A., Lei, Y., Scambos, T. A., and Dehecq, A.: ITS_LIVE global glacier velocity data in near-real time, *The Cryosphere*, 19, 3517–3533, <https://doi.org/10.5194/tc-19-3517-2025>, 2025.
- Koch, M., Seehaus, T., Friedl, P., and Braun, M.: Automated Detection of Glacier Surges from Sentinel-1 Surface Velocity Time Series—An Example from Svalbard, *Remote Sensing*, 15, 1545, <https://doi.org/10.3390/rs15061545>, 2023.
- 220 RGI Consortium: Randolph Glacier Inventory - A Dataset of Global Glacier Outlines, Version 7, <https://doi.org/10.5067/F6JMOVY5NAVZ>, 2023.
- Zheng, W.: Glacier geometry and flow speed determine how Arctic marine-terminating glaciers respond to lubricated beds, *The Cryosphere*, 16, 1431–1445, <https://doi.org/10.5194/tc-16-1431-2022>, 2022.