

**ICE, CLOUD, and Land Elevation Satellite
(ICESat-2) Project**

**Algorithm Theoretical Basis Document
(ATBD)**

for

**Land-ice DEM (ATL14) and Land-ice height
change (ATL15)**

Prepared By: Benjamin Smith University of Washington



**Goddard Space Flight Center
Greenbelt, Maryland**

**National Aeronautics and
Space Administration**

Abstract

This document describes the algorithms that generate the ATL14 (land-ice DEM) and ATL15 (land-ice height change) products. These products derive a reference Digital Elevation Model (DEM) and quarterly elevation-difference surfaces that specify temporal surface-height variations relative to this DEM, for regions around the Arctic and for the Antarctic ice sheet. The products also provide reduced-resolution versions of the height-change datasets and of their temporal derivatives, and error estimates for each, which are intended to allow error propagation for mass-change estimates. The products (ATL14 and ATL15) are generated together using a least-squares-based algorithm that fits changing surface height estimates to surface-elevation data from the ATL11 data product. We describe the algorithm, describe how the parameters controlling the smoothness of the solution were selected, and present tests based on simple datasets that demonstrate how the solution parameters affect the precision and resolution of the data products.

CM Foreword

This document is an Ice, Cloud, and Land Elevation (ICESat-2) Project Science Office controlled document. Changes to this document require prior approval of the Science Development Team ATBD Lead or designee. Proposed changes shall be submitted in the ICESat-II Management Information System (MIS) via a Signature Controlled Request (SCoRe), along with supportive material justifying the proposed change.

In this document, a requirement is identified by “shall,” a good practice by “should,” permission by “may” or “can,” expectation by “will,” and descriptive material by “is.”

Questions or comments concerning this document should be addressed to:

ICESat-2 Project Science Office
Mail Stop 615
Goddard Space Flight Center
Greenbelt, Maryland 20771

Preface

This document is the Algorithm Theoretical Basis Document for the TBD processing to be implemented at the ICESat-2 Science Investigator-led Processing System (SIPS). The SIPS supports the ATLAS (Advance Topographic Laser Altimeter System) instrument on the ICESat-2 Spacecraft and encompasses the ATLAS Science Algorithm Software (ASAS) and the Scheduling and Data Management System (SDMS). The science algorithm software will produce Level 0 through Level 4 standard data products as well as the associated product quality assessments and metadata information.

The ICESat-2 Science Development Team, in support of the ICESat-2 Project Science Office (PSO), assumes responsibility for this document and updates it, as required, as algorithms are refined or to meet the needs of the ICESat-2 SIPS. Reviews of this document are performed when appropriate and as needed updates to this document are made. Changes to this document will be made by complete revision.

Changes to this document require prior approval of the Change Authority listed on the signature page. Proposed changes shall be submitted to the ICESat-2 PSO, along with supportive material justifying the proposed change.

Questions or comments concerning this document should be addressed to:

Tom Neumann, ICESat-2 Project Scientist
Mail Stop 615
Goddard Space Flight Center
Greenbelt, Maryland 20771

Review/Approval Page

Prepared by:

Ben Smith

*Physicist / Principal Investigator
University of Washington
Applied Physics Lab Polar Science
Center
Box 351310
Seattle, WA, 98105*

Reviewed by:

Leigh Stearns

*University of Kansas
Department of Geology
1414 Naismith Drive, #354A
Lawrence, KS 66045 USA*

Noel Gourmelen

*School of GeoSciences
University of Edinburgh
Drummond Street
Edinburgh EH8 9XP (UK)*

Approved by:

Tom Neumann

*ICESat-2 Project Scientist
Cryospheric Sciences Lab, Code 615
NASA Goddard Space Flight Center
Greenbelt, MD*

***** Signatures are available on-line at: [https:// icesatiimis.gsfc.nasa.gov](https://icesatiimis.gsfc.nasa.gov) *****

Change History Log

Revision Level	Description of Change	SCoRe No.	Date Approved
1.0	Initial Release		
2.0	Updates for rel002		

List of TBDs/TBRs

Item No.	Location	Summary	Ind./Org.	Due Date
1	Section 3.4.6	Update 80x80 figures to 60x60		

Table of Contents

Abstract	1-i
CM Foreword.....	ii
Preface	iv
Review/Approval Page	vi
Change History Log	viii
List of TBDs/TBRs.....	x
Table of Contents	xii
List of Figures.....	14
List of Tables	15
1.0 INTRODUCTION	17
2.0 OVERVIEW AND BACKGROUND INFORMATION.....	18
2.1 Limitations of the ATL14/15 product	18
2.1.1 Sampling limitations.....	18
2.1.2 Product resolution limitations.....	19
2.1.3 Grounding lines and ice fronts	19
3.0 Algorithm theory	20
3.1 Overview	20
3.2 Input data	20
3.2.1 Geographic data selection	21
3.2.2 Parameter-based selection.....	22
3.2.3 Data density reduction near the pole hole	22
3.2.4 Tide and atmospheric models	23
3.3 Time reference and coordinate systems.....	23
3.3.1 Time reference.....	23
3.3.2 Spatial reference systems	24
3.4 Mathematical formulation.....	24
3.4.1 Interpolating functions.	24
3.4.2 Least-squares solution	25
3.4.3 Regularization.....	26
3.4.4 Including the effects of correlated errors	27

3.4.5	Solution domains	28
3.4.6	Selecting weighting values	28
3.4.7	Iterative data selection.....	34
3.4.8	Error estimates	35
3.4.9	Tile-matching constraints.....	36
3.4.10	Tile mosaicking	37
3.4.11	Data averaging	37
3.4.12	Temporal derivatives	39
4.0	Data products	40
4.1	ATL14 product	41
4.1.1	DEM parameters.....	41
4.1.2	Attributes and metadata	41
4.2	ATL15 product	43
4.3	Parameters common among groups	43
4.3.2	ATL15 delta_h group	43
4.3.3	ATL15 dhdt_lag1 group	44
4.3.4	ATL15 dhdt_lag4 group	45
4.3.5	ATL15 dhdt_lag(K) groups	46
4.3.6	Metadata and attributes	46
5.0	Solution procedure	48
5.1	Computational details of the solution.....	48
6.0	ALGORITHM IMPLEMENTATION	51
6.1.1	Top-level module : ATL11_to_ATL15.py	51
6.1.2	Reading ATL11 data.....	54
6.1.3	Decimating data: <i>decimate_data</i> function	54
6.1.4	Applying tides and atmosphere correction: <i>apply_tide_model</i> function....	54
6.2	Smooth-surface fitting: <i>smooth_xytb_fit_aug</i> function	55
6.3	Mosaicking the tiles: the <i>regen_mosaics</i> script	58
6.4	Final product generation: ATL14_write_nc and ATL15_write_nc	59
7.0	TEST DATA AND RESULTS.....	61
	GLOSSARY/ACRONYMS.....	67

List of Figures

<u>Figure</u>	<u>Page</u>
Figure 1. Selection of DEM smoothness constraint values.	30
Figure 2. Map of DEM smoothness weighting required to give unit scaled data-residual magnitude.	31
Figure 3. Attenuation ratio as a function of period.	33
Figure 4. Attenuation as a function of wavelength for short-wavelength signals.	34
Figure 5. Tiling scheme.	36
Figure 6. Distribution of averaging domains within a tile.	38
Figure 7. Model time response to simulated data.	61
Figure 8. Recovered long-wavelength amplitude as a function of period.	62
Figure 9. Analytic and computed amplitude response to sinusoidal signals.	63

List of Tables

<u>Table</u>	<u>Page</u>
Table 1. ATL14 parameters	41
Table 2. ATL14 attributes	41
Table 3. ATL14/15 tile_stats group parameters	42
Table 4. ATL15 delta_h group parameters	44
Table 5. ATL15 dhdt_lag1 group parameters	45
Table 6. ATL15 dhdt_lag4 group parameters	45
Table 7. ATL15 dhdt_lag8 group parameters	46
Table 8. ATL15 metadata and attributes	47
Table 9. Inputs to smooth_xytb_fit.py	52
Table 10. Correspondence between ATL11 parameters and smooth_xytb variables	55

1.0 INTRODUCTION

This document covers the production of the ICESat-2 gridded height products over land ice, ATL14 (Antarctic and Greenland ice-sheet surface height) and ATL15 (Antarctic and Greenland ice-sheet surface height change). It describes an algorithm that converts a set of collocated height measurements provided by ATL11 into a DEM (Digital Elevation Model) and a gridded set of height-change maps. This algorithm produces the ATL14 and ATL15 products simultaneously. Therefore, the steps involved in their production are not described separately for the two products; however, the product descriptions are given independently, based on variables generated by the algorithm.

2.0 OVERVIEW AND BACKGROUND INFORMATION

ATL14 and 15 bring the time-varying height estimates provided in ATL11 into a gridded format. ATL14 provides a high resolution (100 m) digital elevation models (DEM) which is a “snapshot”, i.e. spatially continuous views of surface height for the ice sheet for a specific reference time (Jan 1, 2020), which can be used to initialize ice-sheet models, as boundary conditions for atmospheric models, or to help with the reduction of other satellite data such as optical imagery or synthetic aperture radar (SAR) data. ATL15 provides coarser resolution (1 km) height-change maps at 3 month intervals to enable visualization of height-change patterns, and calculations of integrated regional volume change. In this document we describe the two products together because they are generated as part of the same algorithm, that generates a self-consistent set of DEMs and height-change maps from the ATL11 data.

2.1 Limitations of the ATL14/15 product

ATL14 and ATL15 are both limited as to the features they can resolve by (1) the spatial resolution of ICESat-2 tracks, (2) the temporal sampling of the tracks, and (3) the resolution of the grids chosen for the products.

2.1.1 Sampling limitations

The spatial resolution of ICESat-2 is primarily limited in the across-track direction, with 3-km spacing between beam pairs, and RGT (repeat ground track)-to RGT spacing as large as 15 km in southern Greenland. This implies that some small glaciers that run parallel to RGTs will not be sampled at all by ICESat-2, and that sampling of larger glaciers will resolve large-scale changes, but often not the smaller details. The effects of this coarse sampling is somewhat compounded by the temporal sampling of ICESat-2 measurements: Each RGT is measured at most four times per year, and clouds often obscure as many as half of these measurements. This means that measurements should not be expected to resolve ice-sheet changes at sub-seasonal time scales, and in some cases, when the timing of measurements relative to short-term surface-height changes is different between adjacent tracks, the spatiotemporal sampling may produce apparent large rates of height change.

Because of these limitations, we expect the products to resolve year-to-year change rates at scales of a few km or better, but to provide a noisier and less well resolved picture of season-to-season changes. We present analysis in section 3.4.6.2 of the maximum spatiotemporal resolution of the product allowed by the smoothness constraints that we impose on the solution, and show an example of the temporal resolution allowed by the spatiotemporal sampling in section 7.1.1.1 Each of these confirm that the products should be treated with some caution at short (few-km or less) spatial scales, especially at short (semiannual or less) temporal scales. Users may also expect to see striping in the height-change-rate products that reflect the location and timing of the RGT measurements, especially when episodic snowfall or melt events produce apparent high rates of change. These effects should be much less noticeable in longer-term (annual or longer) height-change rate averages.

2.1.2 Product resolution limitations

The resolution of the DEM (100 m) and the height-change grids (1 km) each also impose limits on the resolution of the products, although this is only the dominant limitation of the product only in areas where the data sampling is dense (i.e. directly along the pair tracks).

2.1.3 Grounding lines and ice fronts

Grounding lines and ice fronts represent two kinds of features that may not be resolved well by ATL15. We expect ice fronts to advance and retreat over the course of the ICESat-2 mission, which, if not taken into account, could produce large apparent height changes in the surface height. For release 002, we use constant masks for Greenland and the Arctic, and use an annually updated mask for Antarctica that covers changes through early 2021. Users should treat large height changes close to ice fronts with caution, and should consult the *ice_area* variable provided by the data product to see whether their areas of interest might have been affected by ice-front changes.

Grounding lines can introduce problems in ATL15 because ice flexes smoothly in response to the tidal motion of water beneath the ice on the floating side of the grounding line, while the tide model that we use provides a sharp discontinuity, with the full tide amplitude on the floating side of the grounding line and zero amplitude on the grounded side. In release 002, we have implemented a refinement to the tidal model for Antarctica that smoothly modifies the tidal amplitude across each grounding line, based on tidal flexure measured in lower-level products (ATL11). We expect that some artifacts will remain in the data after this correction.

3.0 ALGORITHM THEORY

3.1 Overview

The ATL14/15 algorithm fits height and height-change maps to ATL11 repeat-track-corrected height estimates. Broadly, we proceed in four steps:

1. Select high-quality ATL11 data
2. Generate model fitting matrices for the selected ATL11 data, and regularization matrices
3. Use least-squares fitting techniques to derive the simplest model that fits the data, and to reject statistically outlying measurements
4. Calculate model errors.

The combined ATL14/15 model is the DEM, calculated for a reference date (decimal year 2020.0, equivalent to 12 AM GMT, Jan 1, 2019 in release 001, see 3.2) and a set of quarter-annual height-change maps for dates before and after the reference date that give the accumulated surface-height change between the ice surface and the DEM. The height for any point in space and time within the domain of the model can be found as the sum of the height for the reference date, as interpolated from the DEM, and the height difference interpolated in space and time from the height-difference maps. To derive time-varying height model, we formulate a set of linear equations that express this interpolation for the data points, and a set of equations that express the spatial and temporal variability of any candidate model. We then minimize sum of squares of the error-weighted misfit between the model heights and the measured heights and the square-integrated variability of the model. Techniques of this type have been used in glaciology to map changes in glaciers and Greenland based on ICESat altimetry, combined ICESat and Operation Icebridge altimetry, and Cryosat-2 altimetry.

We will provide granules for Antarctica, for Greenland, and for seven major Arctic regions that are covered by ICESat-2's repeat-track measurements (see list in 3.2.1)

3.2 Input data

The ATL14/15 product is derived from the ATLAS ATL11 Land Ice H(t) product. ATL11 data are available from two groups of datasets within each file:

Along-track measurements: measurements from different cycles on the same reference pair track (RPT). They are spatially dense (every 30 m) along each RPT, and measurements from different cycles are separated by 91 days. Consistent repeat-track measurements began in April, 2019, near the start of ICESat-2's Cycle 3. Along-track data include information about the heights from each cycle for which ICESat-2 made a successful surface-height measurement, about the surface-shape correction used to compare the different cycles, and about the errors in the data implied by the uncertainty in the reference surface.

Crossover measurements: measurements from different reference ground tracks (RGTs) and RPTs where the ground tracks intersect. Each crossover difference uses a surface-shape correction derived from multiple cycles of along-track data acquired along a particular RPT (the reference-track measurement), and a measurement on a different RPT (the crossing-track

measurement). They include a range of time differences, and may include both within-cycle differences and cycle-to-cycle differences.

For Cycles 1 and 2, we use crossover measurements only; for cycles 3 onwards we use along-track measurements only. ATL11 files can contain several crossover data points for each beam pair-to-beam pair crossover, because a crossover is identified whenever a measurement from a crossing track is within 65 m of a reference point, which allows many reference points to contain data from the same crossing track when the tracks cross obliquely. To ensure that crossover data points are statistically independent from one another, we select the crossover data point from each set of crossing beam pairs for each cycle that has the smallest error estimate.

From the input ATL11 data, we select a subset that we expect to give the most reliable height estimates. This is done through a combination of geographic data selection, which restricts the data under consideration to ice sheets and ice shelves only, parameter-based selection, which removes data flagged in ATL11 as potentially problematic, and iterative data selection, which uses the solution itself to identify outlying data points. The first two of these are described in this section, the third is described in 3.4.7.

3.2.1 Geographic data selection

A potential source of error in ATL11 height-change estimates is the assumption of smooth spatial and temporal variations in surface height in places where this is not the case; in particular, we expect to see large discontinuities in the surface height at the seaward edge of ice shelves, and between glaciers and the surrounding rock. For ice shelves, an advancing ice front could appear as a sudden vertical step change in surface height, while for ice-rock contacts, inferring a smooth height-change pattern could lead to biases (toward zero change) for estimates of change in rapidly thinning outlet glaciers. To avoid this, we reject reference points that are not on ice sheets, and reject measurements from ATL11 that are not on ice shelves or grounded ice.

For this masking we use three sets of data:

1. In Greenland, we use the BEDMACHINE-v3 ice mask, accepting all data points identified as ‘grounded ice’ or ‘floating-ice’ in the 125-m mask.
2. In the Arctic outside of the extent of the Bedmachine mask, we use the Randolph Glacier Inventory (RGI, version 6.0, [Consortium, 2017]) mask to identify points marked as ‘glacier.’
3. In Antarctica, we use a combination of the BEDMACHINE-v2 ice-thickness and ice-mask fields [Morlighem, 2020], and an annually updated mask of ice-front locations [Greene et al., 2022]. We reject points identified as having ice thickness less than 10 m and as being land in the BEDMACHINE mask, and reject points for which either the previous or the subsequent time-varying mask field indicates open water. The 10-m threshold is selected to remove thin areas of ice land ice from consideration, because we cannot be confident that BEDMACHINE can resolve ice thickness well enough to ensure that these areas are truly ice covered.

We expect the ice fronts to change over the course of the ICESat-2 mission, and will apply further updates to the masks and polygons as we become aware of these changes.

We will provide solutions for contiguous ice areas with an area greater than 100 km², so before use in ATL14 and 15, the Randolph Glacier Inventory polygons are consolidated along adjacent borders, and the consolidated polygons are accepted only if their area is greater than 100 km².

The Arctic RGI contains several regions, of which we use:

- 03: ArcticCanadaNorth
- 04: ArcticCanadaSouth
- 06: Iceland
- 07: Svalbard
- 09: RussianArctic

The other regions (e.g. 02:WesternCanada) are not well covered by ICESat-2 repeat data. The areas covered by these masks, plus Antarctica and Greenland, define the granules for ATL14 and ATL15 (one set of granules will be provided for each product and each region)

In Greenland, we use the Bedmachine mask in place of the RGI mask.

3.2.1.1 Antarctic mask updates for release 002

Because the Antarctic mask updates that run through early 2021 missed a significant calving event on the Ronne Ice Shelf, we created a custom ice mask for 2022 using ice front locations provided by the DLR Icelines Download Service [Baumhoer *et al.*, 2021]. These ice front locations estimated using Sentinel-1 radar imagery were converted to raster masks by semi-automatically updating a prior mask. The mask was updated by advecting the front location forward in time to estimate the invalid non-ice points, and backwards in time to estimate the valid ice points. The preliminary masks were then visually inspected and manually edited for complex regions, such as zones with heavy crevassing.

3.2.2 Parameter-based selection

Our analysis uses all ATL11 data points for which the reference-surface fit quality (*/ptx/ref_surf/fit_quality*) has values of 0 or 2, which eliminates the points that are identified as having large errors in at least one component of the surface-shape correction polynomial. For along-track data, the fit-quality parameter is provided for each reference point, while for crossover data it is provided for the reference-track measurement.

3.2.3 Data density reduction near the pole hole

ATL14/15 calculations are carried out on 81x81-km subdomains. The computational efficiency of our algorithm suffers as the number of input data points increases beyond 1 million. This is only a problem towards the northern and southern limits of ICESat-2 coverage, where tracks converge. To avoid this problem, after data have been loaded and filtered using the parameter selection and the rock mask, the computational domain (of width W_{domain}) is divided into squares of dimension W_{sq} . In each square, if the density of data points in the square (ρ_{sq} , equal to $\frac{N_{\text{sq}}}{W_{\text{sq}}^2}$) is larger than the density that would be obtained if the full domain contained 10^6 points (ρ_{1M} , equal

to $\frac{10^6}{W_{domain}^2}$), the number of data in the square is reduced by removing data from selected reference points until $\rho_{sq} < \rho_{1M}$: for example, if $\rho_{sq} = 1.1 \rho_{1M}$, the data from every 10th ATL11 reference point is removed, and if $\rho_{sq} = 3 \rho_{1M}$, two out of every three reference points are removed.

3.2.4 Tide and atmospheric models

Dynamic Atmospheric Corrections (DACs) are derived from the models available from ATL03, and carried through the ATL06 and ATL11 products. Tide corrections are derived from the 2008 Circum-Antarctic Tidal Simulation (CATS2008) tide model in Antarctica, and from the GR1km-v2 model in Greenland. These corrections are applied only in regions where the ice-shelf mask indicates floating ice, and in places where the ice-shelf mask indicates floating ice, but the tide model does not provide tide-constituent values (places that the model developers believed were grounded, but the tide mask indicates floating) we interpolate tide constituents using a nearest-neighbor algorithm with a distance cutoff of 200 km.

To account for ice flexure on in the vicinity of Antarctic grounding zones, we apply an empirical scaling factor to the tide model and DAC heights. The scaling factors are calculated by creating a bounded-variable least-squares model of the surface height variability using the combined DAC and CATS amplitudes as one of the designated variables. This bounded-least squares solution solves for average height, average height change and the scaling factor from the uncorrected ATL11 heights for every reference point within 20km of the grounding zone. The scaling factors are bounded in the fit to be between 0 (for fully grounded) and 1 (for fully floating). Assuming that the phases of the tidal constituents are continuous from the hydrostatic boundary through the flexure zone and to the grounding line, this scaling of the tidal amplitudes will compensate for most of the non-hydrostatic effects. These pointwise scaling factors were then interpolated to a regular 200m x 200m grid using Gaussian radial basis functions, and, for points that are more than 20 km from the grounding zone grid values were set to 0 or 1, for grounded or ungrounded points respectively. The estimated tidal height for each point input to the ATL14/15 algorithm is then the product of the value interpolated from the empirical scaling grid and the (potentially) extrapolated CATS height value.

3.3 Time reference and coordinate systems

3.3.1 Time reference

ATL14 and 15 use a time system that is based on the time stamps used in lower-level products, but expressed in more convenient units. Lower-level products use *delta_time* variables, which express time in seconds after the ATLAS reference epoch (00:00:00 1 Jan 2018). For ATL14 and 15 algorithm calculations, time is measured in decimal years, calculated as seconds after the ATLAS reference epoch, divided by 24x3600x365.25, plus 2018.0. For the final data products, the internal time values are expressed in days after the ATLAS reference epoch. This choice of units simplifies the use of some software tools (e.g. Python implementations of the xarray package) that efficiently convert dates between formats.

The date of the ATL14 DEM is chosen to be on a round-numbered date that it is not at the beginning or the end of the available data, so that edge effects in the solution do not produce large errors in the DEM. In principle, any date that corresponds to a height-difference surface could be chosen; we chose 2020.0, equivalent to 00:00, January 1, 2020 for releases 001 and 002 of the data, but might update later versions to use a later date that is closer to the center of the measurement time series.

3.3.2 Spatial reference systems

We provide ATL14 and 15 in polar-stereographic coordinates based on the WGS-84 ellipsoid, with WGS-84 as the vertical datum. In the Northern Hemisphere, the coordinate system is the NSIDC Greenland-centered sea-ice grid, a polar-stereographic projection with a standard latitude of 70°N and a central meridian of 45°W (EPSG:3413). In the Southern Hemisphere, the coordinate system is the Antarctic polar-stereographic projection, with a standard latitude of 71°S and a central meridian of 0° (EPSG:3031). Because our products are defined on grids that are square in these projected coordinates, the area of each grid cell is different from the square of the grid spacing. For ATL14 and ATL15, we provide an *ice_area* variable that gives the area of each grid cell that is covered by ice; this is derived from the input ice mask, as interpolated to the ATL14 (100-m) resolution, and the cell areas for the coarser-resolution products (ATL15, both at its native 1-km resolution and for the 10, 20, and 40-km averages) are determined by calculating adding the areas of the cells within each coarse-resolution average. When the ice extent changes, whether because the input mask changes (for Antarctica) or because the estimated ice-surface height falls below the geoid (for all regions), the cell area indicates the minimum ice extent for the time period included in each cell.

3.4 Mathematical formulation

Our solution has two main components. The first is the interpolating functions, which express the mathematical relationship between the heights in our grids and the measured (ATL11) heights, the combination of which, and a set of estimated height biases, constitutes a model. The second is a regularized least-squares solution that finds the simplest model that fits the height data. We address these in Sections 3.4.1 -3.4.4, address the selection of domain sizes in 3.4.5, address how parameters are chosen to control the smoothness of the solution in 3.4.6, address error estimates and data selection in 3.4.7, and describe how mosaicked solutions are constructions in 3.4.8-3.4.12.

3.4.1 Interpolating functions.

We express the interpolation between model parameters (i.e. the DEM heights and the height-change surfaces) and measurements (ATL11 heights) as the weighted sum of a set of nodal basis functions. For example, in two dimensions:

$$f(x, y) = \sum_{k=1}^M \sum_{j=1}^N B_{jk}(x, y) W_{jk} \quad 1$$

For a linear interpolation, the interpolation is equal to the weighting values at the nodes, so W_{jk} are the nodal values at points (x_j, y_k) .

$$B_j(x) = \begin{cases} 0 & x < x_{j-1} \text{ or } x > x_{j+1} \\ \frac{x - x_{j-1}}{x_j - x_{j-1}} & x_{j-1} < x < x_j \\ \frac{x_{j+1} - x}{x_{j+1} - x_j} & x_j < x < x_{j+1} \end{cases} \quad 2$$

This basis function is nonzero only on the interval (x_{j-1}, x_{j+1}) , and is continuous everywhere. The sum of two adjacent basis functions is unity on the interval between their nodes, so the sum of a set of basis functions $B(x_1) \dots B(x_N)$ is unity on the interval $[x_1, x_N]$. We construct two-dimensional bilinear basis functions from the product of one-dimensional basis functions in x and y , such that:

$$B_{jk}(x, y) = B_j(x)B_k(y) \quad 3$$

These functions have similar properties to the one-dimensional functions: Four adjacent functions $B_{jk}, B_{j+1k}, B_{jk+1}, B_{j+1k+1}$ add to unity in their overlapping interval, and each function is continuous everywhere in (x, y) . Likewise, we express three-dimensional basis functions, with the third dimension representing time, as the product of one-dimensional basis functions in x , y , and t .

Given a set of measurement points at (x_p, y_p) , we can precalculate the values B_{jk} , so for a grid of nodes with M rows and N columns, we can rewrite (1) as:

$$f(x, y) = \mathbf{I}_2(x, y)\mathbf{W}(x, y) = \mathbf{I}_2(x, y)\mathbf{W} \quad 4$$

We follow the same procedure to express a three-dimensional interpolation in x , y , and time, with a greater number of subscripts in the equation corresponding to (1), and with the \mathbf{I}_3 operator in place of the \mathbf{I}_2 operator in (4).

3.4.2 Least-squares solution

The interpolations described above let us express the heights, \mathbf{h}_m , for a candidate model at a set of data points, \mathbf{x} , \mathbf{y} , and \mathbf{t} . Here, we write these as vectors, because we expect them to have a large number of data points to help us determine our model. The model is then:

$$\mathbf{h}_m = \mathbf{I}_2(\mathbf{x}, \mathbf{y})\mathbf{z}_0 + \mathbf{I}_3(\mathbf{x}, \mathbf{y}, \mathbf{t})\boldsymbol{\delta}\mathbf{z} \quad 5$$

where, \mathbf{z}_0 and $\boldsymbol{\delta}\mathbf{z}$ represent DEM heights and deviations from the DEM, respectively. If the measured heights are \mathbf{h}_d , then the misfit, r , between the data and the model is:

$$\mathbf{r} = \mathbf{h}_d - (\mathbf{I}_2(\mathbf{x}, \mathbf{y})\mathbf{z}_0 + \mathbf{I}_3(\mathbf{x}, \mathbf{y}, t)\delta\mathbf{z}) \quad 6$$

We simplify this expression by substituting \mathbf{G}_i for the horizontal catenation of \mathbf{I}_2 and \mathbf{I}_3 and \mathbf{m} for the vertical catenation of \mathbf{z}_0 and $\delta\mathbf{z}$, to give:

$$\mathbf{r} = \mathbf{h}_d - \mathbf{G}_i\mathbf{m} \quad 7$$

To estimate values for \mathbf{z}_0 and $\delta\mathbf{z}$ we can attempt solve (7) for the \mathbf{m} that minimizes the chi-squared misfit between the data and the model, $\mathbf{r}^T\mathbf{C}_d^{-1}\mathbf{r}$, where \mathbf{C}_d is the covariance matrix for the height measurements. However, unless the data are very dense, with one or more data for each grid cell, this is an ill-posed problem, because not every cell in the DEM and the height-change grids contains measurements. This means that more than one choice of \mathbf{m} can minimize the misfit. Further, a model that exactly minimizes the chi-squared misfit might be very rough because of noise in the data.

3.4.3 Regularization

To obtain a solution in cases where the data do not specify values for every grid cell, and to avoid fitting noisy data values, we add regularization constraints to the problem. This implies minimizing the sum of the chi-squared misfit and a measure of the complexity (e.g. the squared roughness) of the model:

$$R = (\mathbf{G}_i\mathbf{m} - \mathbf{d})^T\mathbf{C}_d^{-1}(\mathbf{G}_i\mathbf{m} - \mathbf{d}) + (\mathbf{G}_c\mathbf{m})^T\mathbf{C}_r^{-1}(\mathbf{G}_c\mathbf{m}) \quad 8$$

Here \mathbf{C}_r is a set of weights for the rows of \mathbf{G}_c . We choose \mathbf{G}_c so that minimizing R gives the smoothest model that fits the data reasonably well: Specifically, we choose \mathbf{G}_c based on the “total energy” of the model in x and y for the reference DEM, and the “total energy” of the temporal derivative of height (here “total energy” is a term used in thin-plate spline theory, and describes the bending energy of an elastic surface [Wahba, 1990], the scare quotes indicating that no actual energy is involved). This gives:

$$\begin{aligned} & (\mathbf{Fm})^T\mathbf{C}_r(\mathbf{Fm}) \\ & \approx \iint \frac{1}{\sigma_{xx}^2} \left[\left(\frac{\partial^2 z_0}{\partial x^2} \right)^2 + 2 \left(\frac{\partial^2 z_0}{\partial x \partial y} \right)^2 + \left(\frac{\partial^2 z_0}{\partial y^2} \right)^2 + \frac{1}{L^2} \left[\left(\frac{\partial z_0}{\partial x} \right)^2 + \left(\frac{\partial z_0}{\partial y} \right)^2 \right] \right] dA \\ & \quad + \iiint \left[\frac{1}{\sigma_{xxt}^2} \left[\left(\frac{\partial^2 \delta z_t}{\partial x^2} \right)^2 + 2 \left(\frac{\partial^2 \delta z_t}{\partial x \partial y} \right)^2 + \left(\frac{\partial^2 \delta z_t}{\partial y^2} \right)^2 \right] \right. \\ & \quad \left. + \frac{1}{L^2} \left[\left(\frac{\partial \delta z_t}{\partial x} \right)^2 + \left(\frac{\partial \delta z_t}{\partial y} \right)^2 \right] + \frac{1}{\sigma_{tt}^2} \left(\frac{\partial^2 \delta z}{\partial^2 t} \right) \right] dAdt \quad 9 \end{aligned}$$

Here δz_t is the first temporal derivative of δz . In our solutions, we minimize the constraint residual ($\mathbf{G}_c\mathbf{m}$ in (8)), which combines the sum of the squares of the second derivatives of the reference DEM in space, the sum of the squares of the second derivatives of the rate of height

change in space, and the square of the second derivative of height with respect to time. The additional terms related to the gradient of height change and the DEM yield solutions that tend to be flat in space over distance L , preventing gradients in the data from being extrapolated indefinitely into regions where no data are present.

The values σ_{xx} , σ_{xxt} , and σ_{tt} express the relative importance of the smoothness of the DEM and the height-change rate, and of the second temporal derivative of the height change. These are tunable parameters that allow the properties of the solution to be tailored to different regions of the ice sheets where different kinds of signals are expected. The z_0 tuning parameters have units of $[\text{derivative}] [A]^{1/2}$, while the δz tuning parameters have units of $[\text{derivative}] [A T]^{1/2}$ where $[\text{derivative}]$ is the derivative that the tuning parameter scales, and A and T are units of area and time. These units are comparable to units of spectral density (which also involve half powers of area and time), and indicate the expected magnitudes of square-integrated derivatives. As an example, $\frac{\partial^2 \delta z_t}{\partial x^2}$ has units of $m^{-1} yr^{-1}$ (because $\frac{\partial^2}{\partial x^2}$ has units of m^{-1} and z_t has units of $m yr^{-1}$), so σ_{xxt} has units of $yr^{-1/2}$. By similar logic, σ_{xx} is unitless.

3.4.4 Including the effects of correlated errors

We expect a dominant error source in ATL11 height estimates to be the effect of geolocation errors in the data over sloping surfaces on the estimated surface heights. We expect these geolocation errors to be consistent over time scales of several seconds or more, equivalent to spatial scales of tens of km or more. This means that the data collected for any given pair track from a single RGT during any given cycle over a small region (e.g. an 61 x 61-km ATL14/15 subdomain, see 3.4.5) will have a bias component with a magnitude equal to the inner product of the mean surface slope for the ice surface measured during that cycle and the geolocation error. For any RGT, pair, and cycle, the expected magnitude of the vertical error due to geolocation errors is thus the magnitude of the mean surface slope for that cycle times the expected geolocation error magnitude.

In principle, the effect of biases due to geolocation errors could be included in the inversion by modifying the estimated data covariance matrix in (8). However, this would entail forming a matrix with a large number of nonzero entries, which is potentially costly in memory and computations. Instead, we include the data biases as parameters in the solution. In combination with the smoothness constraints, this lets the solution choose bias values that result in a smooth model, that are statistically consistent the expected magnitudes of the biases. To describe the effects of these biases, we append additional parameters to the interpolation matrix to form \mathbf{G}_{ib} , so that any row of \mathbf{G}_{ib} gives the interpolated model value plus the bias value for each point, and we modify (8) as:

$$R = (\mathbf{G}_{ib} \mathbf{m}_b - \mathbf{d})^T \mathbf{C}_d^{-1} (\mathbf{G}_{ib} \mathbf{m}_b - \mathbf{d}) + (\mathbf{G}_c \mathbf{m})^T \mathbf{C}_r^{-1} (\mathbf{G}_c \mathbf{m}) + \mathbf{b}^T \mathbf{E}_b^{-1} \mathbf{b}^T \quad 10$$

Here \mathbf{b} gives the bias magnitudes, \mathbf{m}_b gives the combined DEM, interpolation surfaces, and biases, and \mathbf{E}_b is a diagonal matrix whose entries give the expected squared magnitudes of the biases.

We approximate each term in (10) with a finite-difference operator multiplied by \mathbf{m} , and because these finite-difference and interpolation operators generally involve six or fewer coefficients,

each operator is a sparse, albeit large, matrix. This allows this potentially intractable minimization problem to be solved using standard mathematical libraries. We describe how the values of the matrices are calculated, and how the equations are solved, in section 5.0 of this document.

3.4.5 Solution domains

The matrix equations required to minimize (10) are implemented as finite-difference equations on two sets of grids. The finer grid specifies the DEM elevation (for decimal year 2020.0) at 100-meter resolution. This resolution is modestly coarser than the 60-meter averaging scale implied in the fitting procedure for ATL11, so the ATL14 DEM can resolve most of the ice-sheet structure resolved by ATL11. The coarser grids specify height differences relative to the DEM at 1-km resolution, with one grid defined for each quarter-year interval between the start of ICESat-2 cycle 1 in October 2018 and the time at which the ATL15 product was generated. The 1-km resolution of these grids is fine enough to resolve moderately-sized Greenland outlet glaciers, while still allowing ATL14 and 15 to be generated on moderately large grids so that the data biases are well constrained.

The computational costs involved in minimizing (10) on these grids increase rapidly with the size of the domain over which the computation is carried out, and is not practical (as of 2021) to carry out for the whole ice sheet simultaneously. Instead, we divide the ice sheet into overlapping subdomains, or *tiles*, and carry out the computation for each independently. Choosing a smaller tile width makes the computations faster, but choosing a larger width allows us to better take into account the correlation of the data biases over larger spatial scales. We have chosen 61x61 km as a size that balances these two considerations. It allows a minimization of (10), including six iterations to identify outlier data points, in about 20 minutes, while taking into account errors correlated on the scale of small Greenland drainage basins. The domain width is chosen to contain an odd number of height-change grid points so that the 1-km grid points are on integer kilometer values, and the grid is centered on an integer km multiple (i.e. the grid extends from -30 km to 30 km around its center in the x and y directions).

Within each tile, the solution is best constrained far from the edges, because nodes near the edge have fewer adjacent data points available to inform their heights, and because edge nodes are constrained by fewer regularization terms. This problem is minor for nodes surrounded by large numbers of data points, but in the gaps between repeat tracks or in gaps where clouds have prevented data acquisition, particularly at the start and end of the mission, solutions may be substantially different between adjacent tiles. To minimize the influence of edge effects on the overall solution, we perform our calculations on tiles centered on a 40-km grid of points, so that each tile overlaps its neighbors by ~20 km. We then combine the solutions for these tiles using a weighting scheme that assigns highest weights to the portion of the solutions near the center of each tiles, and ignores the solutions at the edges of each tile.

3.4.6 Selecting weighting values

The ATL14/15 model contains three weighting vectors whose values must be selected before running:

- i. σ_{xx} , the DEM smoothness constraint,

- ii. σ_{xxt} , the height-change-smoothness constraint, and
- iii. σ_{tt} , the temporal-variability constraint.

Each of these vectors provides one weight value for each node of the solution. However, the computations for the first data release will only use these weights to specify one two possible values for each constraint and each node: a tight constraint for nodes within the ice-sheet domain, and a loose constraint for nodes outside the domain.

3.4.6.1 DEM smoothness constraints

We selected DEM smoothness constraint values (σ_{xx}) based on a one-dimensional analysis of heights estimated in ATL11 data, in which we fitted a one-dimensional model to data from single epochs and single RPTs of ATL11 data, and adjusted the fitting weights to find the weight for which the model significantly deviated from the data. We considered a model to match the data if the RDE (Robust Dispersion Estimator, equal to half the difference between the 16th and 84th percentiles of a distribution) of its residuals divided by their estimated errors was less than unity. We performed this analysis using release 003 ATL11 data from Greenland, selecting 40-km transects of data and performing fits to the corrected heights (h_{corr}) as a function of the along-track coordinate (x_{atc}) of the ATL11 points. Figure 1 shows an example of this procedure for a transect in central Greenland over undulating ice-sheet terrain. For a loose smoothness constraint ($\sigma_{xx} = 0.5$), the recovered z_0 values match the ATL11 data to within a few cm, and RMS scaled residuals on the order of 0.7 m; for a very tight smoothness constraint ($\sigma_{xx} = 5 \times 10^{-6}$) the recovered z_0 is smoother than the data, leading to half-meter-scale residuals, and RMS scaled residuals on the order of 30 m. A plot of the misfit as a function of the σ_{xx} value shows that for loose constraints, the scaled misfit reaches a constant, asymptotic value likely defined by the grid resolution. For tighter constraints, the scaled misfit increases, reaching a value of unity for $\sigma_{xx} \approx 2 \times 10^{-4}$. For the tightest constraint values, some of the data were removed from the fit by the three-sigma-editing procedure (see 3.4.7), but in this case, no data were removed for the constraint value that produced a scaled misfit of unity.

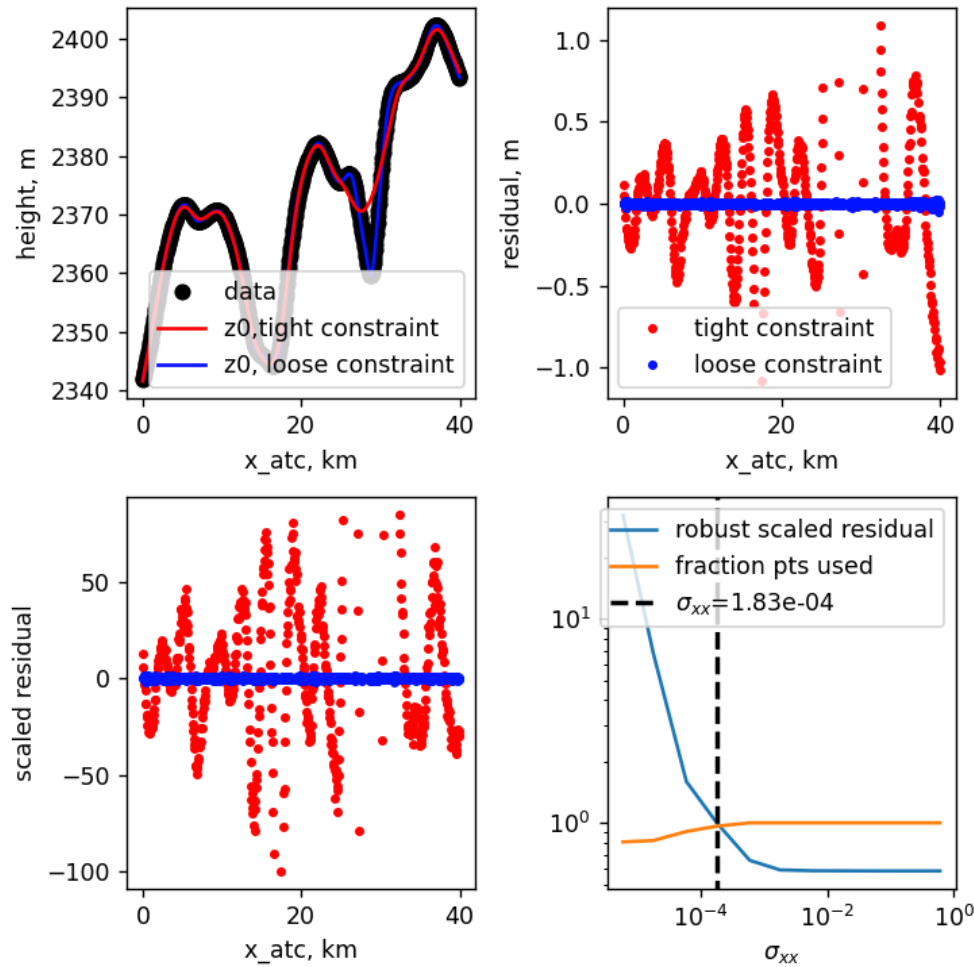


Figure 1. Selection of DEM smoothness constraint values.

Upper left: Height profile from a transect in the Greenland interior, with two z_0 fits for loose and tight constraints. Upper right: residuals to fits for the two constraint values. Lower left: Scaled residuals for the two constraint parameters. Lower right: RMS scaled residuals, and fraction of the data used, as a function of the constraint magnitude.

To select a σ_{xx} value that will allow models to match the data for a wide variety of ice-sheet terrains, we performed this procedure for all available 10-km ATL11 transects in Antarctica and Greenland, finding for each the constraint value that gave a scaled RMS misfit of unity, indicating that the solution was consistent with our estimate of the data uncertainties. A larger RMS misfit would indicate that the model constraints were introducing apparent error into the misfit because the solution could not match the data, a smaller misfit would indicate that the solution was potentially matching errors in the data. Figure 2 shows a map of the recovered values across Greenland. There is considerable scatter driven by the roughness of the ice sheet surface, with uniformly small required constraint values in the interior of the ice sheet, giving way to much larger values near the coasts. To help reduce the effects of small-scale scatter, we have generated grids of σ_{xx} values based on the statistics of 80x80-km subdomains of the ice sheets, where the σ_{xx} value is equal to the 95th percentile of all the values from all 10-km transects within that subdomain, limited to minimum and maximum values of 10^{-4} and 10^{-2} . This procedure produces constraint values equal to the lower-limit value (10^{-4}) for most of the

the area of the ice sheets, with larger values in a few rough coastal areas. In Antarctica and Greenland, the σ_{xx} value we use for any 60x60-km subdomain is then interpolated out of these grids. For the Arctic ice caps, we assigned a value equal to the upper limit of 10^{-2} . The value for each subdomain may be found in the *tile_stats/sigma_xx* field of each ATL14 and ATL15 file.

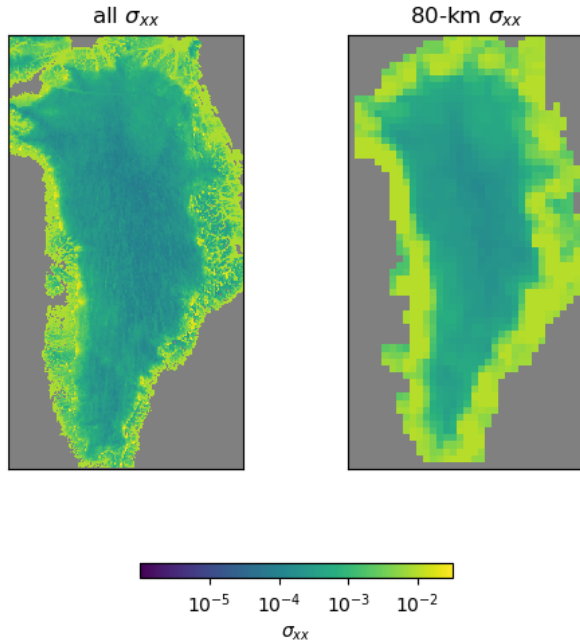


Figure 2. Map of DEM smoothness weighting required to give unit scaled data-residual magnitude.

Left: block-median value of σ_{xx} required to give unit scaled residuals for 20-km grid cells. Right: values downsampled to 80 km, based on the 95th percentile of values for 80x80-km subdomain.

3.4.6.2 Height-change smoothness constraints

In selecting the weighting for the z_0 parameter, we relied on fitting surfaces to data from single ATLAS cycles to separate the effects of variations in the σ_{xx} from changes in the other parameters. The effects of variations in the parameters related to height change (σ_{xxt} and σ_{tt}) are not as simple to separate, because they both affect the recovered height change fields. Instead, we choose these parameters based on an analytic version of (10), specifying harmonic functions for the data, and solving for the resulting model solution.

If we suppose that we have a solution domain that is a square of size L on a side, for a duration T , that we have measurements distributed over the domain with density ρ , and that each measurement has an estimated error of σ_d . If we specify that the data are drawn from a function:

$$d = A_d \sin(2\pi x/\lambda_x) \sin(2\pi y/\lambda_y) \sin(2\pi t/\tau) \quad 11$$

The equivalent residual to (10) is then:

$$R = L^2 T \left(A_m^2 \left(\frac{2\pi^4}{\tau^4 \sigma_{tt}^2} + \frac{8\pi^6}{\lambda^4 \tau^2 \sigma_{xxt}^2} \right) + \frac{\rho (A_d - A_m)^2}{8\sigma_d^2} \right) \quad 12$$

Here $\lambda = (\lambda_x^{-2} + \lambda_y^{-2})^{-2}$.

Setting the derivative of R with respect to A_m to zero and solving for A_m gives:

$$A_m = \frac{A_d}{1 + \frac{16\pi^4 \sigma_d^2}{\rho} \left(\frac{1}{\sigma_{tt}^2 \tau^4} + \frac{4\pi^2}{\sigma_{xxt}^2 \lambda^4 \tau^2} \right)} \quad 13$$

This equation shows that best-fitting model amplitude is equal to a fraction of the data amplitude that depends on the data error estimate, the data density, the wavelength and period of the data signal, and on the constraint values, with smaller amplitudes for stronger constraints, larger errors, shorter periods, and shorter wavelengths. The same results obtain for a two-dimensional analysis with respect to x and time, although the interpretation of data density is different, being instead the linear density of points.

This formulation represents the data as uniformly distributed and unbiased, while ICESat-2 data are sampled a four times per year along repeat tracks, and have correlated biases that are distinct from cycle to cycle and from track to track. As we try to match the theoretical response to the inversion responses, we can take the sampling into account by interpreting the error and density parameters differently for different components of the model.

3.4.6.2.1 Height-change smoothness constraints at long wavelength

For very long-wavelength signals, the last term in the denominator of Eqn. 13, is approximately equal to zero, so the balance of terms that determines A_m involve the signal period, the data density, and the data errors. Further, averaged over a full subdomain, the most significant errors should be the correlated errors, so we can approximate σ_d by the correlated error (0.03 m for flat surfaces, or a larger value if surface slopes are larger), and can approximate $\rho \approx N_{RGT}/L_{domain}^2/0.25 \text{ yr}$, where N_{RGT} is the number of RGTs in the domain, and L_{domain} is the width of the subdomain. We choose a value for σ_{tt} by choosing a period below which we want signals in the data to be attenuated by at least a factor of 2. Choosing this period to be 0.25 years (so that signals not well represented in the data sampling are attenuated), the correlated error to be the flat-surface value 0.03 m, and N_{RGT}/L_{domain}^2 to be consistent with the 10-km RGT-to-RGT spacing typical for central Greenland gives $\sigma_{tt} \sim 189000 \text{ m}^2 \text{ yr}^{-3/2}$, which we round to $200000 \text{ m}^2 \text{ yr}^{-3/2}$. Figure 3 shows the expected attenuation as a function of period for this constraint value, for different systematic error magnitudes.

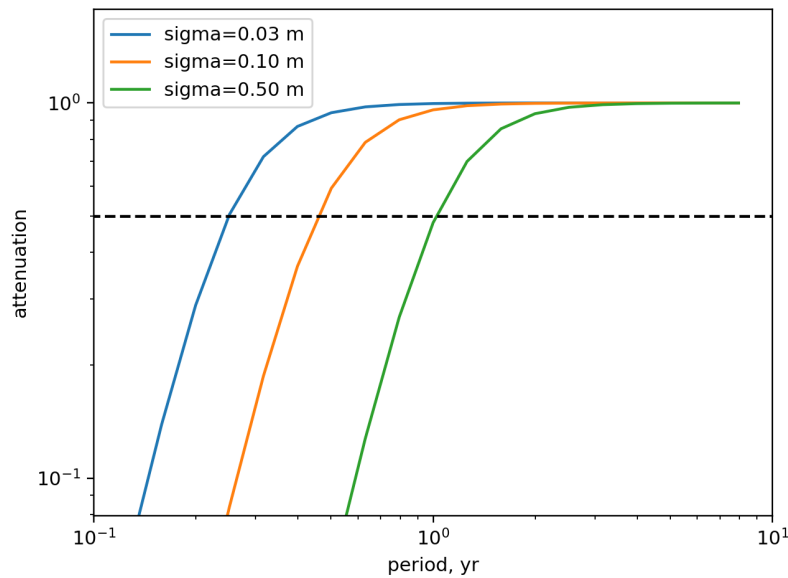


Figure 3. Attenuation ratio as a function of period.

The attenuation ratio for chosen constraint value ($\sigma_{it}=200000 \text{ m}^2\text{yr}^{-3}$) is plotted as a function of the signal period for three different systematic-error magnitudes. The dashed line indicates 50% attenuation.

As the error magnitudes increase, the solution is increasingly attenuated at short periods, with 50% attenuation at annual periods for the largest error values. These largest error magnitudes are unrealistic for most of Greenland and Antarctica, where even in coastal areas, most tiles have weighted average error less than 0.1 m; with this error magnitude, the solution should still admit around 50% of signals with a semi-annual period. A numerical experiment to test whether our assumptions about the significance of the correlated and uncorrelated errors in the solution is described in section 7.1.1.1.

3.4.6.2.2 Height-change smoothness constraints at short wavelength

The constraint on the second spatial derivatives of the height-change rate controls much of the behavior of the model at short wavelengths. For these scales, the most important role of the constraints is to attenuate signals directly along the RPTs that cannot be captured by the track-to-track spacing. Directly along the tracks the effective data density is equal to the inverse of the point spacing times the four measurements per year. We determine a value for σ_{xxt} based on a cutoff wavelength of half the distance between beam pairs, 1.5 km, and a cutoff period of 0.25 yr, which gives $\sigma_{xxt} = 4.83 \times 10^{-5} \text{ yr}^{-1/2}$, which we round to $5 \times 10^{-5} \text{ yr}^{-1/2}$.




Figure 4. Attenuation as a function of wavelength for short-wavelength signals.

The attenuation ratio is plotted as a function of wavelength for signals with five different periods. The dashed line indicates 50% attenuation.

Figure 4 shows the attenuation as a function of wavelength for errors of 0.02 m, for three different signal periods. The cutoff wavelength (attenuation < 0.5) decreases strongly as the signal period increases. For annual signals, the cutoff wavelength is 700 m, and $>95\%$ of the amplitude is recovered for wavelengths > 1.5 km.

3.4.6.3 Weighting based on ice masking

Minimizing (10) produces solutions whose height-change values are smooth over large scales (at most equal to the 61-km solution tile width). In the interior of the ice sheet, these constraints make sense, because we expect that adjacent locations will share common dynamic and climatic forcing. In coastal regions, however, adjacent glaciers separated by bare rock are expected to have different dynamic forcing and to experience different rates of height change. To allow spatially disjoint glaciers to have less strongly coupled height-change rates, we relax (by a factor of 10) the smoothness constraints in (10) for nodes that are identified in the Greenland mask as rock or water, or are identified in the Antarctic ice-thickness map as having 10-m or less ice thickness.

3.4.7 Iterative data selection

To account for outlier values within the data set whose errors are not well accounted for by a Gaussian distribution with standard deviation σ_d , we iteratively remove data points from the solution if their scaled residuals (equal to the residuals divided by their uncertainties) are large compared to other nearby scaled residuals. This selection is carried out during the fitting process for each tile, and is repeated to ensure that as outlying values are removed, the solution improves.

Our least-squares fitting produces surfaces that are smoother than the input data. This results in misfit between the data points and the best fitting model. We assume that for most data points, this misfit will reflect the errors estimated in ATL11, but that a minority of data points will be inconsistent with the solution because a failure in the lower-level product fitting processes (e.g. an ATL06 segment fitted to noise photons), to an extent not captured by the error estimates. We attempt to identify and remove these inconsistent points based on their residuals to the best-fitting model. A simple scheme for removing these data would be to calculate the scaled residual, equal to $r_i/\sigma_{est,i}$, and eliminate those data for which the scaled residual was larger than a fixed threshold. This scheme, however, neglects the misfit that results from smooth model surfaces that cannot capture local variations in the surface height. To help capture these local variations, we calculate a variable, σ_{extra} , that satisfies the formula:

$$RDE \left(\frac{r}{(\sigma_{est}^2 + \sigma_{extra}^2)^{1/2}} \right) = 1$$

Because we expect the roughness of the surface to vary on a glacier-by-glacier scale, we calculate σ_{extra} for overlapping 20x20-km subregions within each tile, on 10-km centers, and, for each data point, assign the value equal to the average of all subregions that overlap that point. To avoid allowing points with excessively large errors to remain in the solution, we limit the maximum value of σ_{extra} to 2m. Having established a local estimate for σ_{extra} , we then edit the data points by rejecting all points for which:

$$\frac{r}{(\sigma_{est}^2 + \sigma_{extra}^2)^{1/2}} > 3$$

The list of points that remain after the outlier removal are stored in the *three_sigma_edit* variable. We repeat this procedure for six iterations, or until it converges such that the difference between model heights in subsequent steps is less than 1 cm. The *three_sigma_edit* variable is stored with the solution for each tile in the as part of internal ATL14/15 calculations.

3.4.8 Error estimates

We assign each data point an error estimate based on the uncorrelated error estimate from ATL11 (*/ptx/h_corr_sigma*). These errors are assumed to be uncorrelated between cycles and between ATL11 reference points. We assign correlated errors based on the surface slope and geolocation-error estimates via the bias-parameter constraints, as described in 3.4.4. Because the error estimates from ATL11 and the correlated-error estimates do not necessarily account for all the processes that can produce scatter in the data (for example, due to large spatial or temporal gradients in surface-change processes) we account for any remaining error by adding the σ_{extra} values in quadrature to the per-point error estimates (see 3.4.7). The uncertainty in a model obtained by the minimization of a least-squares problem (e.g. 10) is straightforward but, in our case, numerically difficult; we describe our calculation in 5.1. The uncertainty calculation is carried out on the preliminary set of tiles (see 3.4.9) and does not include the effects of tile matching, because the error estimates for the matched solution appear to be artificially small.

3.4.9 Tile-matching constraints

We expect the iterative editing scheme to produce different results depending on the tracks selected in any solution domain, and on the range of terrain in the domain. For example, a tile containing large surface-slope variations will have larger bias estimates than an adjacent tile with a smoother surface, and both tiles will have solutions that are poorly constrained near the tile edges. This can lead to mismatched results in the overlapping area between adjacent tiles. To reduce this effect, we perform a preliminary (the *prelim* stage) fit for all of the tiles, iterating the solution up to six times to eliminate outlying data points, and calculate the model error estimate for this solution. We then rerun the solution for each tile (the *matched* stage) including a set of constraint equations that minimize the difference between the height-difference solutions at the edges of the tile and the solutions in the interior of the tiles it overlaps.




Figure 5. Edge matching. A central tile (black, center at (x_{tile}, y_{tile})) and the eight overlapping neighboring tiles (centers indicated by colored stars), showing the nodes from each overlapping tile that are used to constrain the height changes in the solution for the central tile. The nodes constrained by each neighboring tile are colored to match that tile’s outline and central point.

Figure 5 shows a tile that overlaps eight adjacent tiles. For each of these tiles, we construct a set of constraint equations that minimize the difference between the surface-change values in the central tile within 1 km of the tile border, and those in the adjacent tile that are more than 10 km away from that tile’s border. These constraints are weighted by the formal errors in the adjacent

tiles. We perform a single iteration fitting algorithm including these constraints, to produce a single, matched, solution. After this final iteration, we update the ice mask variables to exclude all cells for which the inferred surface elevation (DEM+height change) is below the EGM2008 geoid, and recompute the *ice_area* fields to match these changes.

3.4.10 Tile mosaicking

To form a final ATL14/15 product, we mosaic the results of all the domains together using a set of data weights as follows: The outer 5 km of each tile is assigned a weight of 0. Between 5 km and 15 km from the edge, the weight varies from 0 to 1 on a raised cosine curve, and in the central 31x31-km area of the domain the weight is 1. These distances are chosen so that the weights from overlapping regions of adjacent tiles sum to unity. The value for each pixel in the output mosaic is then found:

$$v_{mos} = \frac{\sum_{tiles} w_i v_i}{\sum_{tiles} w_i} \quad 14$$

Here w_i are the weights for the pixel in the tiles that overlap the pixel, and v_i are the values being mosaicked.

3.4.11 Data averaging

The ATL15 output products include both full-resolution (1-km) height-difference and height-rate estimates, and reduced-resolution averages of the height differences and rates, at resolutions of 10, 20, and 40 km. These reduced-resolution estimates are primarily provided because the error estimates for these larger areas can take into account interplay between per-track-correlated errors in the ICESat-2 data. These averages are calculated for each tile individually and mosaicked together to form the whole-ice-sheet reduced-resolution maps. Figure 6 shows how the averaging areas are distributed within an 60-km tile. The 10-km and 20-km averages have corners that are aligned with the corner of the tile, while the 40-km average centers are shifted by 10 km, so that the center of the 40-km average matches the center of the tile.

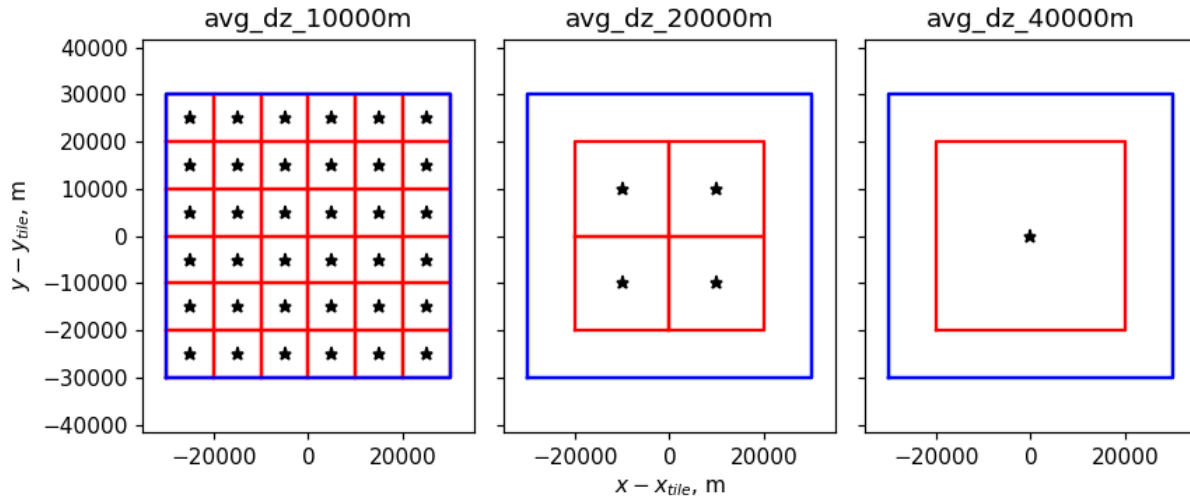


Figure 6. Distribution of averaging domains within a tile.

The blue outline shows the extent of the 60-km tile, the red outlines show the averaged areas at each scale, and the black stars show the centers of the averaged areas.

These averaged values represent the area-weighted average of the full-resolution quantities. The grids used here are in polar-stereographic projections, for which the scaling between grid area and true area varies with distance from the true-scale latitude (70 N for the northern hemisphere, 71 S for the southern), and the averaging cells may overlap some areas that are within the ice mask, and some that are not. To calculate these averages, we first sample the ice-rock mask to the z_0 grid at 100 m, and calculate the cell area of each 100-m grid cell. We then average the mask values for each 1-km grid cell:

$$M_{1km} = \frac{\sum M_{100m} A_{100m} W_{100m}}{\sum A_{100m} W_{100m}} \quad 15$$

Here M_{100m} is the ice mask, A_{100m} is the area of each 100-m grid cell, and W_{100m} is a weighting function for nodes in the 600 m on either side of the averaging-cell center. W is equal to 1 for the nodes the interior of the averaging cell, equal to 0.5 for the nodes on the edges of the cell, and is equal to $\frac{1}{4}$ for the nodes on the corners of the cell. This weighting implies that the weights of adjacent averaging cells add to unity.

Averaging 1 km height-difference and height-rates follows the same process:

$$Z_{coarse} = \frac{\sum Z_{1km} A_{1km} W_{1km}}{\sum A_{1km} W_{1km}} \quad 16$$

Here Z_{coarse} is the quantity being averaged, A_{1km} is the area of the 1-km grid cell, and W_{1km} is a weighting function analogous to that in (15).

3.4.12 Temporal derivatives

The ATL15 product includes estimates of the temporal derivatives of the height-difference surfaces, primarily because the error estimates for these derivatives can be calculated using the covariance structure of the ATL15 solution, and it is not practical to provide this covariance information in the product because the covariance matrix is a dense matrix with one row and column for each model point.

Temporal derivatives are calculated for quarterly height changes, for annual height changes, and for biennial (two-year) and longer (3+ year) height changes; each is equal to the difference between the height-difference surfaces separated by the specified time interval, divided by the length of the time interval. The time value specified for each height-change rate is equal to the temporal midpoint of the two epochs contained in the difference.

The temporal derivative groups include an associated *ice_area* variable, to indicate the area of each cell that was ice covered for the averaging period. Where the ice front advanced or retreated, the cell area indicates the minimum ice extent between the first and last grid cells included in the average.

4.0 DATA PRODUCTS

Both the ATL14 and the ATL15 products are provided in geographic subsets based on the standard divisions of the Randolph Glacier Inventory [Pfeffer *et al.*, 2014].

ATL14 and 15 product names follow a format:

ATLxx_rr_CCCC_[res]_RRR_VV.nc

Here:

- xx is the product (14 or 15).
- rr is the region for which the product was generated:
 - AA Antarctica
 - CN Arctic Canada North
 - CS Arctic Canada South
 - GL Greenland and peripheral ice caps
 - IS Iceland
 - SV Svalbard
 - RA Russian Arctic
- CCCC gives the first and last cycles of repeat-track data included in the file (i.e. 0309 include cycles 3 through 9, inclusive).
- [res] expresses the product's resolution.
 - For ATL14, the resolution will always be 100m
 - For ATL15, the resolution will be one of: 01km, 10km, 20km, or 40km, depending on the resolution at which the product has been averaged.
- RRR gives the release of the data. This number will be incremented each time there is a change in the software used to generate the product, or when there is a new major release in the lower-level products.
- VV gives the version of the product, which will be incremented if a new version of a granule is generated using the same software.

4.1 ATL14 product

4.1.1 DEM parameters

ATL14 DEM parameters are described in Table 1.

Table 1. ATL14 parameters

parameter	units	dimensions	ATBD section	description
x	meters	Nx	3.2	x coordinate of the DEM cell centers, in projected coordinates
y	meters	Ny	3.2	y coordinate of the DEM cell centers, in projected coordinates
h	meters	Nx, Ny	3.2	DEM surface height, referenced to WGS84
h_sigma	meters	Nx, Ny	4.1	Uncertainty in the DEM surface height
ice_area	meters ²	Nx, Ny	3.3.2	Area of each grid cell covered by ice, accounting for the area distortion in the polar-stereographic projections
data_count	counts	Nx, Ny	5.2.4.4	Weighted number of data contributing to each node in the DEM.
misfit_rms	meters	Nx, Ny	5.2.4.4	Root-mean_square of the residuals associated with each DEM node
misfit_scaled_rms	counts	Nx, Ny	5.2.4.4	Root-mean_square of the error-scaled residuals associated with each DEM node

4.1.2 Attributes and metadata

ATL14 granules include the following attributes:

Table 2. ATL14 attributes

Attribute	ATBD section	Description
sigma_xx	3.4.3	Constraint value for the second spatial derivative of the DEM
L_gap	3.4.3	Assumed gap scale between data points
Time	3.3.1	Time (in days since the ATLAS epoch) for which the DEM was generated

4.1.2.1 Tile_stats group

The *tile_stats* group contains variables pertaining to the fitting process for the 81x81-km tiles on which the solution was carried out. These parameters are included in both the ATL14 and the ATL15 product

Table 3. ATL14/15 tile_stats group parameters

Parameter	ATBD section	Description
x	3.4.5	Tile-center coordinates, in projected coordinates
y	3.4.5	Tile-center coordinates, in projected coordinates
N_data	3.4.7	Number of data used in the fit
RMS_data	3.4.3	root mean of squared, scaled data misfits
RMS_bias	3.4.4	root mean of squared, scaled bias values
N_bias	3.4.4	number of bias values solved for
RMS_d2z0dx2	3.4.3	root mean square of the constraint equation residuals for the second spatial derivative of z0
RMS_d2zdt2	3.4.3	root mean square of the constraint equation residuals for the second temporal derivative of dz
RMS_d2zdx2dt	3.4.3	root mean square of the constraint equation residuals for the second temporal derivative of dz/dt
sigma_xx0	3.4.6	Weighting values for the constraint equations on the second spatial derivatives of the DEM
sigma_xxt	3.4.6	Weighting values for the constraint equations on the second spatial derivatives of the height-change rate
sigma_tt	3.4.6	Weighting values for the constraint equations on the second temporal derivatives of the surface height

4.2 ATL15 product

The ATL15 height-change product includes both height-difference surfaces (the *delta_h* surfaces, giving the height difference relative to the DEM as a function of time), and height-change rates (the time derivative of the *delta_h* surfaces, *dhd_t*). These products are provided at the full resolution (1 km) and at reduced resolution (10 km, 20 km, and 40 km), each resolution provided in its own granule for each region. The time values for the height-difference surfaces are simply the time stamps for each surface, while the time values for the height-change rates are equal to the midpoints of the times for the surfaces included in the difference. Time values are expressed in days since the ATLAS epoch, (Midnight at the start of Jan 1 2018) which relate to the *delta_time* values included in other products by the formula $days = \text{delta_time} / 24 / 3600$. Location parameters (*x,y*) provide the coordinates for the DEM grid. These values describe the position of the grid-cell centers.

Each ATL15 file includes groups for the *delta-h* fields (height differences relative to the DEM), and for dh/dt at different temporal resolutions: quarterly (*dhd_t_lag1*), annual (*dhd_t_lag4*), biennial (*dhd_t_lag8*), and triennial (*dhd_t_lag12*). As the mission duration increases, more groups will be added to span the full length of the mission.

4.3 Parameters common among groups

4.3.1.1 Time variables

Each group within each granule includes a time value. These are the time stamps for height difference values and height-change rates. They relate to the *delta_time* variables in lower-level products as $time = \text{delta_time} / 24 / 3600$, so that they give time in days since the ATLAS epoch (00:00:00, Jan 1 2018). The time values in the *delta_h* group gives the time stamps for the height differences with respect to the reference epoch. The time values in the height-change groups (*dhd_t*) give the centers of the differences between the first and second value in the differences (e.g. in the *dhd_t_lag4* group, the time value for the difference between 2019 and 2020 is 2019.5).

4.3.1.2 Position variables

Position variables (*x,y*) describe the grid-cell center coordinates in polar-stereographic coordinates (see 3.3.2).

4.3.2 ATL15 *delta_h* group

The ATL15 *delta_h* group includes height differences between the model surface at any time and the DEM surface. It is provided in different granules at resolutions of 1 km, 10 km, 20 km, and 40 km.

Table 4. ATL15 delta_h group parameters

parameter	units	dimensions	ATBD section	description
x	meters	Nx	3.3.2	x coordinate of the height-change grid cell centers, in projected coordinates
y	meters	Ny	3.3.2	y coordinate of the height-change grid cell centers, in projected coordinates
time	days	Nt	3.3.1	time for the height-difference estimates since the ATLAS epoch (Midnight, Jan 1, 2018).
delta_h	meters	Nt,Nx, Ny	3.2	quarterly height difference with respect to the DEM
delta_h_sigma	meters	Nx, Ny	Error! Reference source not found.	uncertainty in the height differences
ice_area	meters^2	Nt, Nx, Ny	3.4.10	area of each grid cell covered by ice, accounting for the area distortion in the polar-stereographic projections
data_count	counts	Nx, Ny	6.2.1.4	weighted number of data contributing to each node in the DEM.
misfit_rms	meters	Nx, Ny	6.2.1.4	root-mean_square of the residuals associated with each DEM node
misfit_scaled_rms	counts	Nx, Ny	6.2.1.4	root-mean_square of the error-scaled residuals associated with each DEM node

4.3.3 ATL15 dhdt_lag1 group

The *dhdt_lag1* group gives height difference rates between subsequent quarterly height-difference surfaces. It is provided primarily because the error estimates can be calculated rigorously during the ATL14/15 calculation, using model covariance information that is not provided in a data product.

Table 5. ATL15 dhdt_lag1 group parameters

parameter	units	dimensions	ATBD section	description
x	meters	Nx	3.3.2	x coordinate of the height-change grid cell centers, in projected coordinates
y	meters	Ny	3.3.2	y coordinate of the height-change grid cell centers, in projected coordinates
time	days	Nt-1	3.3.1	time for the temporal center of the quarterly dh/dt estimates, in days since the ATLAS epoch (Midnight, Jan 1, 2018).
dhdt	meters	Nt-1, Nx, Ny	3.2	quarterly rate of change (dh/dt)
dhdt_sigma	meters	Nt-1, Nx, Ny	Error! Reference source not found.	uncertainty in the quarterly dh/dt estimates
ice_area	meters ²	Nt-1, Nx, Ny	3.4.12	Area of each grid cell covered by ice, corrected for polar-stereographic distortion and for changing ice fronts

4.3.4 ATL15 dhdt_lag4 group

The *dhdt_lag4* group provides annual height-change-rate estimates. Its variables are identical to those in the *dhdt_lag1* except that there are fewer temporal values provided, and that the rates refer to annual, not quarterly differences.

Table 6. ATL15 dhdt_lag4 group parameters

parameter	units	dimensions	ATBD section	description
x	meters	Nx	3.3.2	x coordinate of the height-change grid cell centers, in projected coordinates
y	meters	Ny	3.3.2	y coordinate of the height-change grid cell centers, in projected coordinates
time	days	Nt-3	3.3.1	time for the temporal center of the annual dh/dt estimates, in days since the ATLAS epoch (Midnight, Jan 1, 2018).
dhdt	meters	Nt-3, Nx, Ny	3.2	annual rate of change (dh/dt)
dhdt_sigma	meters	Nt-3, Nx, Ny	Error! Reference	uncertainty in the annual dh/dt estimates

			source not found.	
ice_area	meters^2	Nt-3,Nx,Ny	3.4.12	Area of each grid cell covered by ice, corrected for polar-stereographic distortion and for changing ice fronts

4.3.5 ATL15 dhdt_lag(K) groups

The *dhdt_lag(K)* groups provides biennial-and-longer height-change-rate estimates. Their variables are identical to those in the *dhdt_lag1* except that there are fewer temporal values provided, and that the rates refer to biennial-and-longer lags, not quarterly differences.

Table 7. ATL15 dhdt_lag8 group parameters

parameter	Units	dimensions	ATBD section	description
x	Meters	Nx	3.3.2	x coordinate of the height-change grid cell centers, in projected coordinates
y	Meters	Ny	3.3.2	y coordinate of the height-change grid cell centers, in projected coordinates
time	Days	Nt-K+1	3.3.1	time for the temporal center of the biennial dh/dt estimates, in days since the ATLAS epoch (Midnight, Jan 1, 2018).
dhdt	Meters	Nt-K+1,Nx, Ny	3.2	multiannual rate of change (dh/dt)
dhdt_sigma	Meters	Nt-K+1,Nx, Ny	1.1.1	uncertainty in the multiannual dh/dt estimates
ice_area	meters^2	Nt-K+1,Nx,Ny	3.4.12	Area of each grid cell covered by ice, corrected for polar-stereographic distortion and for changing ice fronts

4.3.6 Metadata and attributes

Attributes for each ATL15 file include:

Table 8. ATL15 metadata and attributes

Attribute	ATBD section	Description
L_gap	3.4.3	Assumed maximum gap between measurements
Reference_epoch_time	3.3.1	Time (days since the midnight at the start of Jan 1 2018) corresponding to the DEM
Reference_epoch_index	3.3.1	Zero-based count indicating which delta_h epoch corresponds to the DEM
Tide_model	3.2.4	Tide model used in processing

4.3.6.1 Tile_stats group

The *tile_stats* group contains variables pertaining to the fitting process for the 81x81-km tiles on which the solution was carried out. These parameters are included in both the ATL14 and the ATL15 product, and are described in 4.1.2.1. See Table 3.

5.0 SOLUTION PROCEDURE

Once we have chosen constraint values and have selected data, we can build the fitting matrices and attempt a solution for (10). The steps in this section describe the numerical linear algebra needed for the minimization of the combined residual (3.4.4), the three-sigma editing process (3.4.7), error estimation (**Error! Reference source not found.**), and data averaging (3.4.11).

We construct a *design matrix*, \mathbf{G} in two parts: the fitting matrix, \mathbf{G}_d , used in calculating the model values at the data, and the constraint matrix, \mathbf{G}_c , used in evaluating the regularization constraints. Thus, we seek the model that minimizes

$$\begin{aligned} & \boldsymbol{\epsilon}^T \mathbf{C}^{-1} \boldsymbol{\epsilon} \\ & \text{for} \\ \boldsymbol{\epsilon} = & \begin{bmatrix} \mathbf{G}_d \\ \mathbf{G}_c \end{bmatrix} \mathbf{m} - \begin{bmatrix} \mathbf{h}_d \\ \mathbf{0} \end{bmatrix}. \end{aligned} \tag{17}$$

Here \mathbf{C} is a matrix whose diagonal contains the squared data errors and the weighting values for the constraint equations. Thus, the solution procedure is:

- S1: Calculate \mathbf{G}_d for all data points and Calculate \mathbf{G}_c , and calculate \mathbf{C} , giving each data a weight of σ_i^{-2} .
- S2: minimize $\boldsymbol{\epsilon}^T \mathbf{C}^{-1} \boldsymbol{\epsilon}$ in (17) to find \mathbf{m} .
- S3: Calculate $\hat{\sigma}$ based on the RDE of r/σ for data included in the fit.
- S4: Evaluate r for all data based on \mathbf{G}_d and the current \mathbf{m} .
- S5: Choose the points for which $|r/\sigma| < 3 \hat{\sigma}$
- S6: minimize $\boldsymbol{\epsilon}^T \mathbf{C}^{-1} \boldsymbol{\epsilon}$ in (12) find \mathbf{m} based on the subset of \mathbf{G}_d for points included in the solution.
- S7: Check for convergence: If \mathbf{m} differs from \mathbf{m} at the previous iteration by less than 1 m, or if more than 6 iterations have taken place, terminate the iterations, otherwise return to step S3.

Once iterations are terminated, we estimate the errors in \mathbf{m} based on an error propagation of the solution to (12) using the most recently derived set of data weights.

5.1 Computational details of the solution

We note that our problem (17) is of the form:

$$\min ((\mathbf{d}_0 - \mathbf{G}_0 \mathbf{m})^T \mathbf{C}^{-1} (\mathbf{d}_0 - \mathbf{G}_0 \mathbf{m})) \tag{18}$$

Because our \mathbf{C} matrix is diagonal, with all positive entries on the diagonal. Because this is true, we can find the matrix square root of \mathbf{C}^{-1} such that:

$$\mathbf{C}^{-1} = \mathbf{T}^{-1} \mathbf{T}^{-1T} \tag{19}$$

This lets us rewrite the problem as:

$$\min ((\mathbf{T}^{-1}\mathbf{d}_0 - \mathbf{T}^{-1}\mathbf{G}_0\mathbf{m})^T (\mathbf{T}^{-1}\mathbf{d}_0 - \mathbf{T}^{-1}\mathbf{G}_0\mathbf{m})) \quad 20$$

which has the form of an unweighted least-squares problem with data vector $\mathbf{d} = \mathbf{T}^{-1}\mathbf{d}_0$ and design matrix $\mathbf{G} = \mathbf{T}^{-1}\mathbf{G}_0$.

We solve this set of equations using a QR decomposition, which lets us factor the $m \times n$ \mathbf{G} matrix into two matrices, \mathbf{Q} ($m \times n$) and \mathbf{R} ($n \times n$) where \mathbf{Q} has the property that $\mathbf{Q}^{-1} = \mathbf{Q}^T$, and \mathbf{R} is upper-triangular. The solution to the least-squares equation is then

$$\mathbf{m} = \mathbf{R}^{-1}\mathbf{Q}^T\mathbf{d} \quad 21$$

For our problem, \mathbf{R} is sparse, with less than a quarter of its elements populated, but \mathbf{Q} is generally dense. Calculating all of the elements of \mathbf{Q} would be prohibitively expensive in terms of memory and computational cost. To avoid this, we use an algorithm [Davis, 2011] from the *SuiteSparse* library of mathematical utilities that lets us efficiently calculate the Rz factorization, which gives \mathbf{R} and the product $\mathbf{z} = \mathbf{Q}^T\mathbf{d}$ without explicitly forming \mathbf{Q} .

The solution for \mathbf{m} is then just

$$\mathbf{m} = \mathbf{R}^{-1}\mathbf{z} \quad 22$$

and because \mathbf{R} is triangular, we can solve this equation by back-substitution in a number of operations equal to the number of non-zero entries in \mathbf{R} . In the iterative data editing procedure (3.4.7), we recalculate this solution at each iteration, removing data with excessively large scaled after each iteration.

The formal error in \mathbf{m} (calculated in section **Error! Reference source not found.**) is found

$$\sigma_m^2 = \text{diag}(\mathbf{R}^{-1T}\mathbf{R}^{-1}) \quad 23$$

Which we find by squaring the elements of \mathbf{R}^{-1} and calculating their row-wise sum. Likewise, the formal error for any linear function \mathbf{F} applied to \mathbf{m} is:

$$\sigma_{\mathbf{Fm}}^2 = \text{diag}((\mathbf{FR}^{-1})^T\mathbf{FR}^{-1}) \quad 24$$

which we calculate by squaring the elements of \mathbf{FR}^{-1} and taking their row-wise sum. This property is useful for letting us propagate the errors in the inversion for \mathbf{m} to any quantity linearly related to \mathbf{m} , such as its spatial average, or its derivatives, as calculated in 3.4.11.

If n_c is the number of columns in \mathbf{G} , and n_{nz} is the number of non-zeros in \mathbf{R} (where $n_{nz} \ll n_c^2$), calculating the Rz factorization of \mathbf{G} takes $O(n_c^3)$ arithmetical operations, calculating \mathbf{m} takes $O(n_{nz})$ operations, and calculating \mathbf{R}^{-1} for error propagation takes on the order of n_c^3 operations, or, for the grid sizes chosen for ATL14/15, around four times as long. Relative to these costs, none of the other steps require significant time or storage to compute. Because of this, solving for errors on the full-resolution grid is impractical. Instead, we use the data selected based on the full-resolution solution to re-generate the fitting matrices for a lower-resolution version of the grids, and solve for the errors at this lower resolution. We then interpolate the error values calculated with these coarse grids to the full-resolution grid.

Our solution for \mathbf{m} is then:

1. Calculate \mathbf{G}_d for all data points $h(x,y,t)$, and the smoothness-constraint matrix \mathbf{G}_c . Mark all points as *selected*.
2. Vertically concatenate the *selected* rows of \mathbf{G}_d with \mathbf{G}_c to form \mathbf{G}_0 . Form a matrix \mathbf{T}^{-1} with the inverse of the uncorrelated errors in h and the constraint-confidence estimates on its diagonal. Form a data vector \mathbf{d}_0 equal to \mathbf{h} padded with one zero value for each row in \mathbf{G}_c .
3. Multiply \mathbf{G}_0 and \mathbf{d}_0 by \mathbf{T}^{-1} to give \mathbf{G} and \mathbf{d} .
4. Calculate the RZ factorization of \mathbf{G} and \mathbf{d} , and solve for \mathbf{m} .
5. For overlapping 20-km subdomains, on 10-km centers, calculate *sigma_extra_sub*, the additional error needed to make the residual distribution consistent with the error estimate (see 3.4.8) within each 20-km subdomain. Assign each point a *sigma_extra* value equal to a weighted quadratic average of the values for each subdomain that contains it, with inverse-linear weighting based on the distance to the subdomain centers. Set any *sigma_extra* values greater than *sigma_extra_max* to *sigma_extra_max*.
6. Calculate the scaled model residual $\mathbf{r}_s = (\mathbf{h} - \mathbf{G}\mathbf{m}) / (\sigma_h LI^2 + \sigma_{extra}^2)^{1/2}$ and calculate the RDE of \mathbf{r}_s for all *selected* points, *sigma_d_s*.
7. If the current value of *sigma_d* is more than 1% different from the previous value, mark all points for which $|\mathbf{r}_s| < 3 \sigma_{d_s}$ as *selected*, and return to step 2.
8. Based on the selected points from step 6, generate a new set of fitting and constraint matrices with four times coarser resolution than the original set, and concatenate them to form \mathbf{G}_4 . Concatenate the height values from step 6 with zeros to form \mathbf{d}_4 .
9. Calculate the RZ factorization of \mathbf{G}_4 and \mathbf{d}_4 . Calculate \mathbf{R}_4^{-1} , the inverse of \mathbf{R}_4 .
10. Based on \mathbf{R}_4^{-1} , use Eq. 20 and Eq. 21, scaled by $\max(1, \sigma_{d_s})$ to calculate:
 - a. the errors in \mathbf{m} at each point (the errors in the differences between the height at each point and its reference-surface height)
 - b. the errors in the first temporal derivative in \mathbf{m} at each point
 - c. the errors in the year-on-year differences in \mathbf{m} at each point
 - d. the errors in the block means of \mathbf{m} at each point [d-f are calculated a scales of 10, 20, and 40 km]
 - e. the errors in the block means of the quarterly $d\mathbf{h}/dt$.
 - f. the errors in the block means of the year-on-year differences in \mathbf{m} .
11. Interpolate error values for the results of 9a-9c to the full grid resolution.

6.0 ALGORITHM IMPLEMENTATION

This section describes the specific implementation of the algorithm for program development.

Our implementation of the ATL14/15 algorithm is written in Python, and makes heavy use of the numeric python (numpy) library. We have developed a set of Python classes that represent finite-difference grids and operators on these grids, which are provided in the *LSsurf* package. Tide corrections are handled in the *pyTMD* package, and data organization, reading, and writing are handled in the *pointCollection* package. The code used this project is included in the following online repositories:

Main project repository: <https://github.com/SmithB/ATL1415>
Finite-difference fitting: <https://github.com/SmithB/LSsurf/>
Tide corrections: <https://github.com/tsutterley/pyTMD>
Data read/write and indexing: <https://github.com/SmithB/pointCollection>

6.1.1 Top-level module : ATL11_to_ATL15.py

The coordinating script that assembles the data and generates inputs for the functions in *LSsurf* is called *ATL11_to_ATL15.py*. This script takes input arguments specifying the location and extent of a subdomain, the resolution of the inversion, and the location of the data files. It reads and filters the data, calls the *smooth_xytb_fit_aug* module to calculate the height-change solution, and writes the data out to a file.

The *ATL11_to_ATL15.py* can also calculate the formal errors in a height-change-and-DEM solution. These errors are largely independent of the data heights and solution values, depending only on the data locations and timing, the error estimates, and the constraint values, which allows us to estimate model errors for a reduced-resolution solution using the edited data and misfits from a full-resolution model solution. This greatly improves the efficiency of the error calculation, because the error propagation steps are excessively computationally costly for a full-resolution solution.

The complete solution, including errors, is calculated by running *ATL11_to_ATL15.py* to completion for each tile at full resolution, then reloading the data and editing parameters from the full-resolution solution and calculating the errors at lower resolution (4 times coarser for the DEM, twice as coarse for the *dz* surfaces), then linearly interpolating these errors to the full-resolution grids. The DEM and height-change surfaces are derived from the full-resolution solution (the first time *ATL11_to_ATL15.py* is called), and the errors are derived from the lower-resolution error calculation (the second time *ATL11_to_ATL15.py* is called).

6.1.1.1 Inputs for each tile

The input arguments, and non-configurable input values specified for each ATL14/15 tile, are:

Table 9. Inputs to smooth_xytb_fit_aug.py

Arguments	ATBD section	Description	Default (Rel001 value)
$xy0$	3.4.5	Tile center, in projected coordinates	Different for each tile
W_{xy}	3.4.5	Tile width	81 km
Time_range	3.3.1	First and last time epochs for the solution	2018.75, 2021.25 CE
max_data_count	6.1.3	Maximum number of data in a tile	1.2×10^6
$E_RMS_d2z0dx2$	3.4.6.1	Constraint value applied to the second spatial derivative of the DEM height (corresponds to σ_{xx0})	Spatially dependent
E_RMS_d2zdt2	3.4.6.2.1	Constraint value applied to the second temporal derivative of the height change (corresponds to σ_{tt})	$200000 \text{ m}^2\text{yr}^{-3/2}$
$E_RMS_d3zdx2dt$	3.4.6.2.2	Constraint value applied to the second spatial derivative of the height change rate (corresponds to σ_{xxt})	$5 \times 10^{-5} \text{ yr}^{-1/2}$
Σ_{geo}	Error! Reference	Estimated spatially	6.5 m

	source not found.	correlated component of the geolocation error	
<i>data_gap_scale</i>	3.4.3	Assumed maximum gap between measurements	2500 m
<i>reference_epoch</i>	3.3.1	Temporal reference index corresponding to the DEM	5 (2020.0)
<i>ice_mask</i>	3.2.1, 3.4.6.3	Masks used in selecting data and choosing weighting parameters	See 3.2.1
<i>tide_mask</i>	3.2.4	Tide mask used to decide whether to apply the tide model and <i>dac</i> corrections	See 3.2.4
<i>geoid</i>	3.4.8	Geoid	EGM2008
<i>Tide_model</i>	3.2.4	Tide model	Arctic: AOTIM-5-2018 Antarctic: CATS_2008
<i>tide_scale_map</i>	3.2.4	Scaling applied for grounding-zone flexure	ATL11_0314_tide_adj_scale_200m.h5, provided in the masks/Antarctic directory in the ATL1415 repository
<i>W_pad</i>	3.4.10	Width at the edge of each tile that is not included in the final mosaic (weights are set to 0)	10 km
<i>W_taper</i>	3.4.10	Width inside of the distance	10 km

		specified by W_{pad} over which tile weights taper from 0 to 1	
--	--	--	--

6.1.2 Reading ATL11 data

6.1.2.1 *read_ATL11* function

For the initial group of tiles, designated *centers* in 3.4.8, The *read_ATL11* function *ATL11_to_ATL15* reads ATL11 data from individual granules for a specified region. It uses a pre-computed index of data locations (stored as GeoIndex.h5) to identify all reference points that fall within the tile boundaries, reads the ATL11 along-track data from the ATL11 files, and reads the corresponding crossover data for these points. Data from reference points with *fit_quality* values of 1 or 3 are eliminated from the datasets, as are data that fall outside the rock and ice masks (see 3.2.1).

6.1.3 Decimating data: *decimate_data* function

The *decimate_data* function avoids using vastly excessive numbers of data points for tiles near the South Pole (latitudes close to 88° S). The target number of data points for each tile is specified in the *max_data_count* parameter, and the current (release 001) value of 1.2×10^6 , gives strong solution constraints but relatively efficient calculations, implies a target density of data points, ρ_{target} , equal to 182 km^{-2} for an 81x81-km tile. The function is only used for tiles containing more than 1.2×10^6 data; for these tiles, the domain is divided into 5x5 km subdomains; the density of points within each subdomain, ρ_{sub} , is calculated, and its ratio to ρ_{target} is calculated. If $\frac{\rho_{sub}}{\rho_{target}} > 1$, then for each RGT, a sequence of target reference point numbers between the first reference-point number and the last reference-point number for each RGT is generated, with a mean spacing of $\frac{\rho_{sub}}{\rho_{target}}$. For each RGT, only those reference points that fall within this sequence are included in the solution.

6.1.4 Applying tides and atmosphere correction: *apply_tide_model* function

For each data point, we interpolate the tide mask (*tide_mask* parameter) to determine whether the point should be considered grounded, or floating. For the points identified as floating, the tide value is calculated based on the tide model (*tide_model* parameter) using the PyTMD package. The combined tide and *DAC* model is then scaled by the value interpolated from the tide-scale map at each data point.

6.2 Smooth-surface fitting: *smooth_xytb_fit_aug* function

The *smooth_xytb_fit_aug* function (smooth x, y, time, and bias fit, using an augmented-sigma error model) is a wrapper for classes within the *LSsurf* module. It is a general-purpose function that can be used to fit a variety of data, which are embodied in data structures that specify the location (x, y), time (t), and height (z) of the measurements, as well as error estimates (σ) and correlated errors (σ_{corr}). We map ATL11 data fields to *smooth_xytb_fit_aug* input fields as described in Table 10:

Table 10. Correspondence between ATL11 parameters and *smooth_xytb_fit_aug* variables

<i>Smooth_xytb_fit_aug</i> variable	ATL11 parameter
x, y	Latitude, longitude, projected into polar-stereographic coordinates (standard latitude of 71° S for Antarctica, standard latitude of 70° N for the Arctic)
<i>time</i>	ATL11 converted to years: $\text{delta_time}/24/3600/365.25 + 2018 + 0.5/365.25$
z	h_corr , corrected for ocean tide and <i>dac</i>
<i>sigma</i>	h_corr_sigma .
<i>sigma_corr</i>	Slope magnitude calculated from the RPT-median of the across-track and along-track slopes, times the vector sum of the across-track and along-track geolocation uncertainty (one value for each RPT)

smooth_xytb_fit_aug defines finite-difference grids object for the DEM (the $z0$ grid) and for the height-difference grids (the dz grid). These objects store the relationship between grid node locations and the model vector, which comprises all of the grid-node heights and bias values, concatenated together so that they can be multiplied by matrices. It also constructs the matrices embodying the finite-difference operators described in Eq. (10), and performs the iterative data-editing. We describe each step below.

6.2.1.1 Equation assembly

6.2.1.1.1 Interpolation operators

The *lin_op.interp_mtx* function calculates sparse matrices that describe linear interpolations in x , y , and time, as described in Eqns 1 and 2. The sum of these interpolations is described by a matrix that is the horizontal concatenation of the interpolation matrices for $z0$ and dz , which forms the total interpolation matrix, G_data . The expected misfit values for G_data are the σ values associated with the data points.

Bias operators are set up in the *setup_bias_fit* function. One bias parameter is established for each unique combination of RGTs and cycles within the data. Each of these combinations is assigned to a new model parameter, and a matrix is created, *G_bias* with a row every equation in *G_data* and a column for each bias parameter. For each equation in *G_data*, the column corresponding to the RGT and cycle in the matrix is set to 1, the others are set to unity. This matrix is appended to the right-hand side of *G_data*, so that the bias parameters appear in the model vector after the *z0* and *dz* parameters. A constraint equation matrix (equal to an identity matrix for each bias parameter) is created; its expected misfit values are equal to the median *sigma_correlated* value for the data from the corresponding RGT and cycle.

6.2.1.1.2 Smoothness constraints

The smoothness constraints are set up in the *setup_smoothness_constraints* function. It uses methods on the *lin_op* class to setup differential operators for the terms in Eq. 9. These methods generate sparse matrices describing finite-difference approximations of each derivative of functions sampled on the *dz* and *z0* grids. The vertical concatenation of these matrices forms the constraint matrix *G_c*. Each row these matrices, when multiplied by the model vector **m**, calculates the square root of one node's contribution to the integrals in Eq. 9. For those equations that include nodes that fall outside the model domain (i.e. edge nodes, or nodes whose neighbors overlap areas outside the ice mask), the equation is omitted. This leads to weaker constraints on the edge nodes of the domain, an effect that we take into account with our tiling scheme (see 3.4.5). The expected misfits for the constraint equations (**C_r** in Eq. 9) are set based on the input expected derivative values (see 3.4.6). These values are increased by a factor of 10 (decreasing the weight of the equations by the same factor) for equations that include at least one node that is outside the ice-only mask (see 3.4.6.3).

6.2.1.1.3 Assembling the equations

To ensure that the reference epoch (corresponding to Jan 1, 2020) is zero, for each operator, the columns that correspond to parameters in the *z0* matrix that fall within the reference epoch are removed. This implies that any operator that would have used a value from the reference epoch instead implicitly uses a value of zero. This is implemented by creating a sparse identity matrix with one row and column for each model parameter, and deleting the reference-epoch columns. This matrix (**Ip**, for Identity-parsing matrix) when multiplied by the full model matrices returns matrices that omit the reference epoch; the process can be reversed by multiplying edited matrices (or vectors) by the transpose of **Ip**, to yield matrices (or vectors) with zero values for the reference epoch.

6.2.1.2 Iterating the fit

The iterative fitting step (3.4.7) is carried out by the *iterate_fit* function. This function accepts arguments including the assembled fitting matrix, **G_coo** (where *_coo* indicates that the matrix is a sparse matrix in coordinate-list format), the *data* structure, and the inverse-square-root data covariance matrix, including the constraint-equation magnitudes. For up to six steps, it calculates a best-fitting model based on the data and the fitting matrix, calculates the robust spread of the data residuals for those points still included in the fit (*sigma_hat*), and repeats the fit using only those points whose residuals are more than $3 * \max(1, \sigma_hat)$, and are identified as *editable*

(see **Error! Reference source not found.**). Note that at each step, the full set of residuals data residuals are considered in the three-sigma-edit test, so data that were removed from the fit in a previous step may be included in a subsequent step if the model changes in such a way that their residual is once again smaller than the threshold. The fitting procedure ends after six iterations or if two subsequent steps result in the same set of edited data. The function returns a model estimate ($m0$), a robust spread estimate σ_{hat} , a vector indicating which points were included in the edit, in_TSE , and a set of scaled model residuals, rs_data .

Once the iterations are finished, the ice mask is updated to remove cells whose surface height is lower than the geoid height, and this update is reflected in the ice_area output variables.

6.2.1.3 Calculating averaging operators: the `setup_averaging_ops` function

This function calculates a set of operators that, when multiplied against the model vector, calculate derived parameters including the height-change rate, and reduced-resolution maps of height change and change rates. It first calculates a set of matrices that calculate the height-change rate on a per-pixel basin, at quarterly ($dzdt_lag_1$), annual ($dzdt_lag_4$) and full-mission ($dzdt_lag_N$, where N is four times the number of complete years in the mission) resolution.

It then calculates operators that average dz and dz/dt over 10 km, 20 km, and 40 km. For each scale, the area-weighted average of the parameter is calculated: For each grid cell in the output, a tapered rectangular window is calculated spanning the dz cells within half the averaging width of the output cell in each direction (i.e. from -5 km to 5 km for cells in the 10-km grid, inclusive). Cells on the edges of the window are assigned half weight, and cells on the corner of the window are assigned quarter weight, so that the weights for adjacent cells sum to unity. The weight values are then multiplied by the ice mask (downscaled to the dz grid), and by the area of each output grid cell, accounting for the area distortion implicit in a polar-stereographic grid and for ice-front advance or retreat. The sum of all weight values for an output grid cell is saved as its ice_area parameter (indicating the total area averaged to produce the output value) and the weights, divided by ice_area , are assigned to a row in the averaging matrix. To calculate the lagged dz/dt matrices, two sets of weights are calculated; one for the first epoch in height-change calculation, one for the second, and the two sets of weights are subtracted to form each row in the averaging matrix. This function also calculates the downsampled ice mask.

Each averaging matrix is assigned to an output grid that records the position of the averaged-cell centers and the time values associated with each. For the dz/dt matrices, the time values are set to the midpoint between their two time epochs.

6.2.1.4 Parsing the fit output: the `parse_model` function

The `parse_model` function translates model parameters and data residuals into quantities with units that can be saved into intermediate fitting files, prior to mosaicking into a final product. This function translates the model vector into the proper dimensions for the output grids, and evaluates each term in equation 9 to indicate how each model term contributes to the total residual metric. It also applies the averaging matrices (6.2.1.3) to compute reduced-resolution estimates of the height-change parameters.

The data count per cell is calculated by summing the interpolation matrix over all data points:

$$\begin{aligned} N_{\delta z} &= \mathbf{I}_3 \mathbf{1} \\ N_{z_0} &= \mathbf{I}_2 \mathbf{1} \end{aligned} \quad 1$$

Here \mathbf{I}_2 and \mathbf{I}_3 are 2D and 3D interpolation matrices for the DEM and the dz surfaces, and $\mathbf{1}$ is a vector of ones. The sum of the weights associated with each data point gives an effective number of data points for each grid cell, but because more than one cell is involved in interpolating the model to the data, the count for each data value is spread over several grid points, the largest contribution going to the node closest to the point. This summation gives data counts for both the z_0 and for the dz surfaces.

The model RMS misfit is mapped to the z_0 and dz nodes by calculating:

$$\begin{aligned} misfit_{\delta z} &= (\mathbf{I}_3 \mathbf{r}^2)^{\frac{1}{2}} \\ misfit_{z_0} &= (\mathbf{I}_2 \mathbf{r}^2)^{\frac{1}{2}} \end{aligned} \quad 2$$

Here \mathbf{r} are the edited model residuals. The scaled misfit is calculated in the same way, but replacing \mathbf{r} with \mathbf{r}/σ .

6.2.1.5 Evaluating the error model: the `calc_and_parse_errors` function

The `calc_and_parse_errors` function performs a similar function to the `parse_model` function, but calculating equations 20 and 21 to estimate the uncertainties in the model parameters. It evaluates the RZ decomposition of the fitting matrix, \mathbf{G} (as described in 5.1, equations 18 and 19), and invert the \mathbf{R} matrix to form \mathbf{R}^{-1} . The errors in the model parameters are then found by applying equation 18 to \mathbf{R}^{-1} to find the errors in the DEM and the height surfaces, and by applying equation 21 for each averaging matrix to find the errors in the height change rates and the averaged model quantities.

6.2.1.6 Saving tile data: the `save_fit_to_file` function

The `save_fit_to_file` function writes the input data and output fields to an hdf-5 file. It creates a `/data` group to hold the input ATL11 data, and writes each field read (or derived) from the ATL11 files into this group. It saves the components of the model residual listed in Table 3. ATL14/15 `tile_stats` group parameters **Error! Reference source not found.** (without the `R2_` prefix) in the `/RMS` group, and saves their expected values in the `/E_RMS` group. It also saves the estimated biases for each RGT and cycle in the `/bias` group.

All gridded outputs from the algorithm (the DEM, the height-change surfaces, the error estimates, and the averages, the gridded misfit values, and the data counts) are written into separate groups, with x , y , and, where applicable *time* variables to indicate the grid locations of the measurements.

6.3 Mosaicking the tiles: the `regen_mosaics` script

The `regen_mosaics` script reads the tiles saved by `ATL11_to_ATL15` and combines them into seamless mosaics. For each of a list of variables, it reads the tile files, and writes data into two

grids: a sum-of-weights grid, and a weighted-sum grid. Each point in the weighted-sum grid contains the sum of the variable values multiplied by a tile-specific weighting function, from every tile that overlaps that point. Each point in the sum-of-weights grid contains the sum of the weights from all tiles for each point. The final product is calculated by dividing the weighted-sum grid by the sum-of-weights grid, filling in *invalid* for all points that have zero sum-of-weights. Note that error parameters are calculated based on the *prelim* tiles, because error estimates derived from the *matched* tiles have proven to be unrealistically small.

6.4 Final product generation: ATL14_write_nc and ATL15_write_nc

These scripts rewrite the outputs from *regen_mosaics* to the final product to rename algorithm-internal fields to final product variables, and add metadata fields.

7.0 TEST DATA AND RESULTS

7.1.1.1 Verification of algorithm amplitude response.

To check that our inversion code produces results that match the analytical prediction for a sinusoidal input (13), we used data that included the errors, positions, and timings of measurements from two sites in Greenland, replaced the height values with harmonic functions, and calculated the model response to these functions. First, we generated a set of data with only temporal variations and a period of 0.5 yr, for two different phases of the seasonal signal:

$$\begin{aligned} z_d &= \cos(2\pi t/0.5 \text{ yr}) \\ z_d &= \sin(2\pi t/0.5 \text{ yr}) \end{aligned} \tag{25}$$

Figure 7 shows the input signals, and the area-mean model responses to these signals. the input signals, and the area-mean model responses to these signals.




Figure 7. Model time response to simulated data.

Specified signal (dots), signal values at seasonal nodes (stars), and recovered model response (lines) for data sampling representative of central Greenland.

For the cosine signal, the peak-to-peak amplitude of the recovered signal is slightly larger than 2, because the least-squares approximation of the sinusoidal temporal signal slightly overshoots the signal itself. For the sine signal, the nodes of the signal correspond exactly with the model

temporal nodes, so the recovered signal is zero, except at the edges where the seasonal cycle is not fully sampled, and the solution diverges from zero. We can estimate the total model response at a period of 0.5 yr to be the root sum of squares of the sine and cosine responses, averaged over the time period that is well sampled by the measurements (2019.0-2020.5).

We performed this calculation for a range of periods, and for different constraint values (0.25, 1, and 4 times the default value of 200000) to show how the model amplitudes relate to the constraint value and period, and for two different 80x80 areas in Greenland. Both areas were in central Greenland, at around 72°N latitude, one area included the coast, where slopes are higher resulting in decimeter-scale correlated errors, the other was in the interior, where the flat surface produced small, cm-scale errors. Figure 8 shows the recovered amplitudes as a function of period.




Figure 8. Recovered long-wavelength amplitude as a function of period.

Solid lines indicate the inland response, dashed lines indicate the coastal response.

At large times, the response ratio approaches unity, and at small times, it tends towards zero. As expected, the smaller σ_{tt} suppress short periods strongly. All three σ_{tt} values allow a peak at $\tau \approx 0.08$ yr, which reflects aliasing of the input signal to a longer period because of imperfect temporal sampling of the signal by the tracks. Similarly, the amplitudes around $\tau=0.6$ yr in the weakly constrained curves reflects the multilinear representation of the signal in the model. The amplitudes are weaker in the coastal region, especially at short periods, largely because the error estimates are larger.

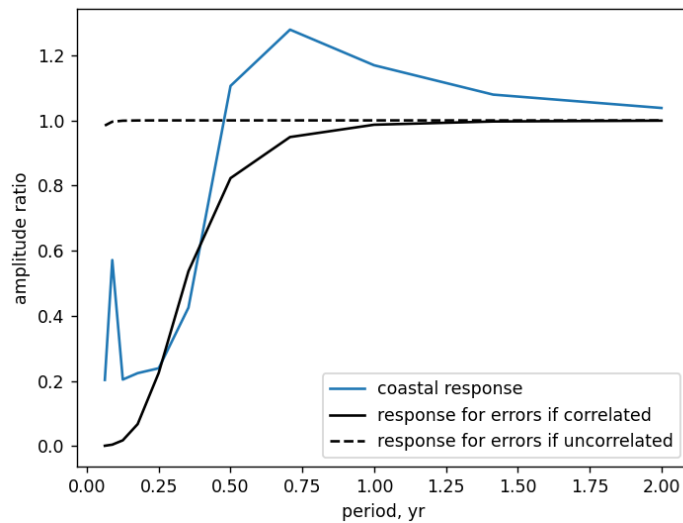


Figure 9. Analytic and computed amplitude response to sinusoidal signals.

The blue curve shows the attenuation as a function of period based on the distribution of errors and measurements from the coastal location. The predicted attenuation curve considering only errors that are uncorrelated from point to point is shown with a dashed black line, the predicted attenuation considering only uncorrelated errors is shown with a solid black line.

In Figure 9, we compare the analytic representation of the attenuation curve with the inversion results for sinusoidal inputs for the coastal location. We used the analytic attenuation formula (13) to predict the attenuation-vs-period curve assuming that all errors were uncorrelated, with magnitude equal to the weighted average of the ATL11 h_corr_sigma values, and to predict the curve assuming that all errors were correlated by RGT and cycle, with magnitude equal to the weighted average of the ATL11 $h_corr_sigma_systematic$ values. The curve based on the assumption of correlated errors falls off sharply in amplitude for signals shorter than around 0.5 yr, which roughly matches the inversion response. We assume that the differences between the analytic and inverted curves is determined by the temporal sampling of the RGTs, which is coarser than the uniform sampling assumed in the analytic calculation.

GLOSSARY/ACRONYMS

ATLAS	ATLAS Advance Topographic Laser Altimeter System
GSFC	Goddard Space Flight Center
ICESat-2 MIS	ICESat-2 Management Information System
PSO	ICESat-2 Project Support Office
SIPS	ICESat-2 Science Investigator-led Processing System

- Baumhoer, C. A., A. J. Dietz, C. Kneisel, H. Paeth, and C. Kuenzer (2021), Environmental drivers of circum-Antarctic glacier and ice shelf front retreat over the last two decades, *Cryosphere*, 15(5), 2357-2381.
- Consortium, R. (2017), Randolph Glacier Inventory – A Dataset of Global Glacier Outlines: Version 6.0: Technical Report, Global Land Ice Measurements from Space., in *Digital Media.*, edited, Colorado, USA.
- Davis, T. A. (2011), Algorithm 915, SuiteSparseQR: Multifrontal Multithreaded Rank-Revealing Sparse QR Factorization, *Acm Transactions on Mathematical Software*, 38(1).
- Greene, C. A., A. S. Gardner, N.-J. Schlegel, and A. D. Fraser (2022), Antarctic calving loss rivals ice-shelf thinning, *Nature*.
- Morlighem, M. (2020), MEaSURES BedMachine Antarctica, Version 2, edited, Boulder, Colorado USA. NASA National Snow and Ice Data Center Distributed Active Archive Center.
- Pfeffer, W. T., et al. (2014), The Randolph Glacier Inventory: a globally complete inventory of glaciers, *Journal of Glaciology*, 60(221), 537-552.
- Wahba, G. (1990), 2. More Splines, in *Spline Models for Observational Data*, edited, pp. 21-39, SIAM.