Remote Sensing of Environment xxx (xxxx) xxxx

Contents lists available at ScienceDirect

Remote Sensing of Environment

journal homepage: www.elsevier.com/locate/rse

Land ice height-retrieval algorithm for NASA's ICESat-2 photon-counting laser altimeter

Benjamin Smith[a], Helen A. Fricker[b,*], Nicholas Holschuh[c], Alex S. Gardner[d], Susheel Adusumilli[b], Kelly M. Brunt[e,f], Beata Csatho[g], Kaitlin Harbeck[e,h], Alex Huth[a,c], Thomas Neumann[e], Johan Nilsson[d], Matthew R. Siegfried[i]

[a] University of Washington, Applied Physics Laboratory, Polar Science Center, United States of America
[b] Scripps Institution of Oceanography, University of California, San Diego, United States of America
[c] Department of Earth and Space Sciences, University of Washington, United States of America
[d] Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109, United States of America
[e] Cryospheric Science Lab, NASA Goddard Space Flight Center, Greenbelt, MD 20771, United States of America
[f] Earth System Science Interdisciplinary Center, University of Maryland, College Park, MD 20740, United States of America
[g] Department of Geological Sciences, University at Buffalo, Buffalo, NY 14260, United States of America
[h] SGT, Inc/KBR Wyle
[i] Department of Geophysics, Colorado School of Mines, Golden, CO 80401, United States of America

ARTICLE INFO

Keywords:
ICESat-2
Laser altimeter
Ice sheets
Glaciers
Land ice
Cryosphere

ABSTRACT

The Ice, Cloud, and land Elevation Satellite-2 (ICESat-2) and its sole scientific instrument, the Advanced Topographic Laser Altimeter System (ATLAS), was launched on 15 September 2018 with a primary goal of measuring changes in the surface of the Earth's land ice (glaciers and ice sheets). ATLAS is a photon-counting laser altimeter, which records the transit time of individual photons in order to reconstruct surface height along track. The ground-track pattern repeats every 91 days such that changes in ice sheet surface height can be estimated through time. In this paper, we describe the set of algorithms that have been developed for ICESat-2 to retrieve ice sheet surface height from the geolocated photons for the Land Ice Along-Track Height Product (ATL06), and demonstrate their output and performance using a synthetic dataset over various land-ice surfaces and under different cloud conditions. We show that the ATL06 algorithm is expected to perform at the level required to meet the ICESat-2 science objectives for land ice.

1. Introduction

On 15 September 2018, the National Aeronautics and Space Administration (NASA) launched ICESat-2 (Ice, Cloud, and land Elevation Satellite-2) into a near-polar orbit (92° inclination). ICESat-2 is the follow-on mission to ICESat (in operation from 2003 to 2009), which was the first laser altimeter in a near-polar orbit around Earth (Schutz et al., 2005). ICESat's Geoscience Laser Altimeter System (GLAS) was a profiling laser altimeter operating at 1064 nm (infrared) to produce non-overlapping ~72 m footprints every 172 m along track (Schutz et al., 2005). For each footprint, the time delay between each laser pulse from the satellite and the detected reflection from the surface was used to determine the range, which was then converted to an estimate of height above a reference ellipsoid (Schutz et al., 2005; Sun et al., 2017).

ICESat (Abdalati et al., 2010) had only a single beam with a broad (30-70 m) footprint, which limited its performance when measuring sloping surfaces in two ways: it reduced the accuracy of single shot heights due to waveform broadening (Yi et al., 2005), and it introduced uncertainties in derived height-change estimates from repeat-track analysis, due to cross-track slope effects from non-exact repeat tracks (Smith et al., 2009; Moholdt et al., 2010). These limitations complicated height-change detection and had a significant impact in the outlet glaciers of Greenland and Antarctica, where both the surface slopes and the true changes are significant (Pritchard et al., 2009). ICESat-2 (Markus et al., 2017) carries a single instrument, the Advanced Topographic Laser Altimeter System (ATLAS). To improve spatial coverage and enable determination of local cross-track slope, ATLAS splits the transmit laser pulse into six beams configured in a 2 × 3 array (three pairs; Fig. 1), with spacing within and between pairs determined by

* Corresponding author.
E-mail address: hafricker@ucsd.edu (H.A. Fricker).

https://doi.org/10.1016/j.rse.2019.111352
Received 5 November 2018; Received in revised form 23 July 2019; Accepted 26 July 2019

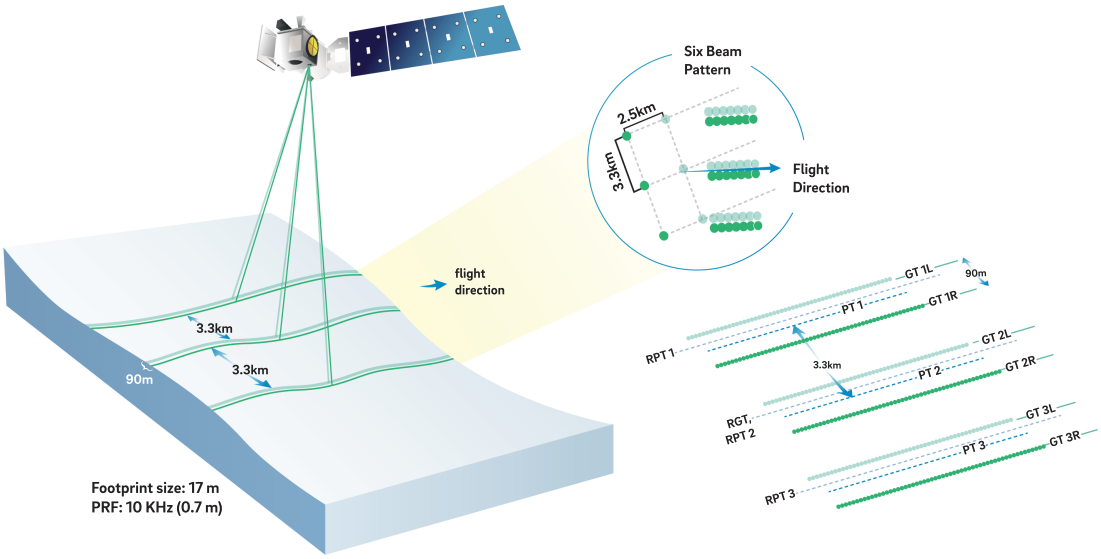

Fig. 1. Schematic of the ATLAS six-beam pattern. The central RPT (Reference Pair Tracks) follows the RGT (Reference Ground Track, which matches the nadir track of the predicted orbit). GTNX (where N is the beam pair number (1–3) and X is L (left) or R (right)) are the ground tracks generated by successive ATLAS spots (green circles). ATLAS is shown here in the forward direction, with the weak beams on the left side of the beam pair. The weak and strong beams are pitched relative to each other such that the weak beams lead the strong beams by ~2.5 km. When ATLAS is oriented in the backward orientation, the relative positions of weak and strong beams change; the strong beams are on the left side of the GT pairs and lead the weak beams. Measured Pair Tracks (PTs) are defined by line that bisects the pairs of GTs, and deviate slightly from the RPTs because of inaccuracies in repeat-track pointing. The separation of GTs within each pair is greatly exaggerated compared to the separation between RPTs. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

spacecraft orientation. Extensive processing of these data, combined with knowledge of the spacecraft position and orientation, provide estimates of the height of the surface beneath the satellite with centimeter accuracy.

ICESat-2/ATLAS's measurement strategy is new to spaceborne laser altimeters. ICESat's GLAS used an infrared laser that emitted pulses of light 40 times per second (40 Hz), with a per-pulse energy of between 5 and 70 mJ, depending on laser age (Sun et al., 2017). ATLAS uses green laser light (532 nm), pulsing at 10,000 times per second (10,000 Hz), with an expected per-pulse energy of around 0.12 mJ in the three strongest beams. Unlike GLAS, which collected many millions of return photons from each laser pulse using an analog detector, each of ATLAS's detectors receive, under ideal conditions, fewer than twelve return photons from each laser pulse. Each of these photon "events" (i.e., a photon hitting a detector) is recorded individually, a technique known as "photon counting". The result of this new observational strategy is that the downlinked ATLAS data are fundamentally different from GLAS data. GLAS reported the power of the return as a function of time whereas ATLAS provides the arrival time of individual photons. When combined with the pointing angle of individual beams and the position of the observatory in space, these photons can be assigned latitudes, longitudes and heights relative to an ellipsoid (Neumann et al., this issue). This novel measurement approach requires a new set of algorithms to retrieve precise surface heights from ATLAS data for all applications (land ice, sea ice, land and vegetation heights, inland water heights, sea surface heights, and cloud layering and optical thickness).

The ICESat-2 Science Definition Team has developed a set of processing techniques to obtain accurate surface-height measurements from ATLAS data for these applications (e.g., geolocated photon clouds (ATL03; Neumann et al., 2019; sea ice (ATL07; Kwok et al., 2019); vegetation (ATL08; Neuenschwander and Pitts, 2019). These techniques take advantage of the large number of pulses that ATLAS emits every second by aggregating the height data for all photons returned from consecutive pulses, such that surface height estimates can be derived from collections of tens or hundreds of photons, as opposed to the approximately 6 to 12 photons returned from a single laser pulse. In this paper, we describe the set of algorithms that we have developed to retrieve height estimates for ICESat-2's land ice along-track data product (ATL06) and demonstrate their output and performance using a synthetic dataset. We present a summary of the ATL06 algorithms to ensure that they are reviewed by and available to the broader scientific community; detailed descriptions are in the ATLAS ATL06 Algorithm Theoretical Basis Document (ATBD; Smith et al., 2019).

2. ATLAS specifications

2.1. ICESat-2 specifications

ICESat-2's 2 × 3 beam array results in a ground pattern consisting of three pairs of beams collecting observations in the along-track direction (Fig. 1). The separation between beam pairs and between beams within each pair is controlled by the spacecraft's yaw angle; ICESat-2 will orbit with a yaw angle of 2° during nominal operations, setting beam-pair separation at ~3.3 km and beams within a pair by ~90 m (Fig. 1). Each laser pulse has a ~17 m diameter spot on the terrestrial surface and an along-track sampling interval of ~0.7 m, determined by the spacecraft velocity (~7 km/s) and laser repetition rate (10 kHz). Each beam pair has a "weak" beam and a "strong" beam (where strong is defined as approximately four times brighter than weak); this is designed to provide the necessary dynamic range to capture enough return photons from both bright (e.g., ice) and dark (e.g., ocean) surfaces. The ATLAS instrument emits a series of green laser pulses that are received and counted by a micro-pulse single-photon sensitive detector; this type of detector is also known as a "photon-counting" detector. The onboard Laser Reference System (LRS) combines observatory star trackers with inertial navigation data to provide absolute laser pointing direction, while GPS receivers provide data on the observatory position in space (Neumann et al., this issue).

2.2. ATLAS laser and reflected pulses

ATLAS's laser pulse has a temporal distribution of energy that is approximately Gaussian. The full width at half maximum of the pulse is 1.7 ns, which corresponds to a standard deviation of 0.68 ns. Converted

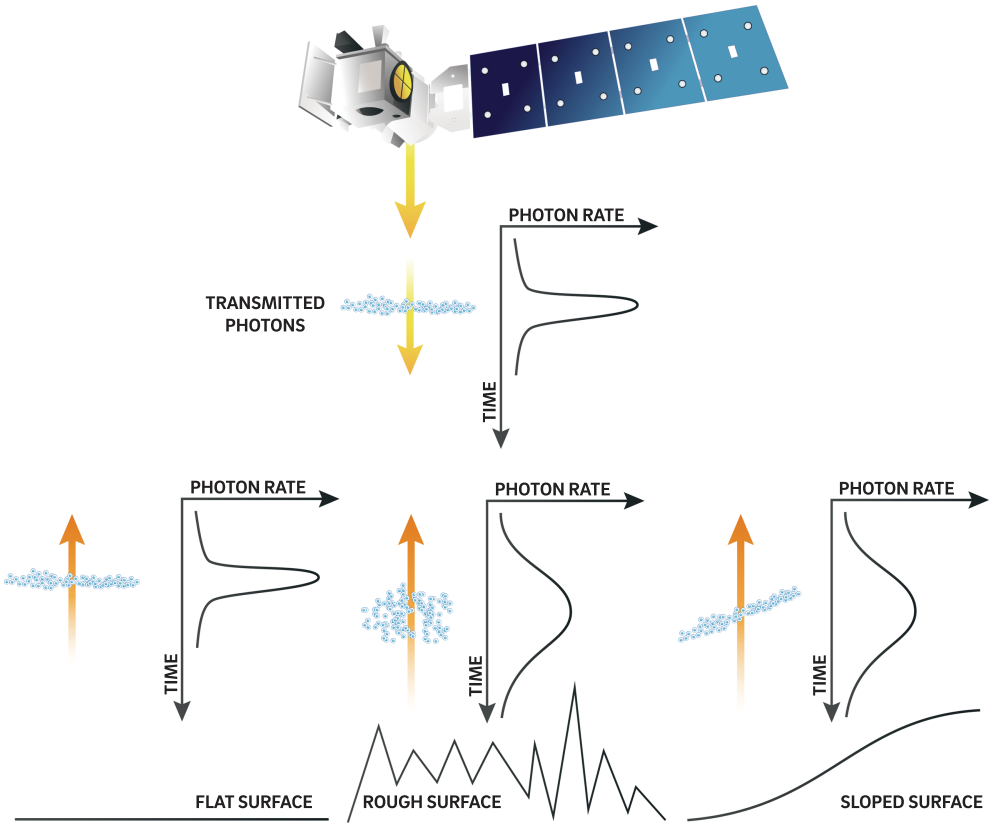

Fig. 2. Schematic of expected ATLAS returns from different surface types: flat surface, rough surface, and sloped surface, with three corresponding expected photon cloud distributions and return waveforms. In both the rough and sloped cases, the spread of path lengths of the returned photons is increased and the return waveform is broadened.

to ranges, these values correspond to distances of 0.26 m and 0.1 m, respectively. The character of the ATLAS return pulse will be determined by the shape of the transmit pulse and the nature of the reflecting surface. For many regions of ice sheets and glaciers, the surface may be approximated as a Lambertian plane. It is anticipated that each reflected laser pulse will have a Gaussian distribution in space, with intensity falling to $1/e^2$ of its peak value over a distance $W/2$ from the center of the beam, where W is the spot diameter. The reflected laser pulse is also anticipated to have an approximate Gaussian distribution in time, with a standard deviation σ_{rx}.

If the surface is sloped, photons from the edge of the spot farthest from the satellite will be delayed relative to photons from the edge nearest the satellite (Fig. 2). Similarly, a rough surface will yield early photons (from peaks) and late photons (from troughs), further spreading the returned photons. If the angle between the beam and the surface normal is ϕ, and the surface height within the spot has a Gaussian distribution with RMS deviation R relative to the plane of the surface, then the measured temporal distribution of the returned photons will be Gaussian as well (Yi et al., 2005), with a temporal standard deviation σ_{rx} equal to the quadratic sum of the spreading effects of the transmit pulse, the surface slope, and the surface roughness:

$$\sigma_{rx} = \left[\sigma_{tx}^2 + \left(\frac{W}{8c}tan\varphi\right)^2 + \left(\frac{R}{2c}\right)^2\right]^{1/2} \quad (1)$$

where c is the speed of light. For ATLAS, W is ~17 m, and the transmit pulse duration (σ_{tx}) is around 0.68 ns, so spreading due to sloping surfaces will be smaller than σ_{tx} for slopes up to about 2.7°.

2.3. ATL03 algorithm

Photons entering the ATLAS telescope include a small number of photons emitted by the ATLAS laser that are returning from the surface, and, during daylight, a large number of "background" photons from the sun. Most of these background photons will be removed by tunable filters that allow only photons with wavelengths close to the laser's central wavelength to pass through (Neumann et al., this issue). However, some photons pass these filters, and their arrival times are also recorded. In addition to these, the ATLAS detectors will contribute a small number of false triggers that will be recorded as photons. The rate of false triggers is expected to be between two and three orders of magnitude smaller than the solar photon rate over sunlit surfaces. The rate of solar photons depends on the solar elevation angle and top of atmosphere reflectance. In general, to distinguish signal photons from background photons, we search for clusters of photons with respect to height (i.e., photons with a similar range). Sunlight scattered from a bright surface (e.g., snow) for a solar elevation of 45° will produce detected photons at rates around 12 MHz. For comparison, a surface return with as few as three photons distributed over 0.5 m of height (i.e., range) produces an instantaneous return rate of 900 MHz. Signal returns are also distinct from the background because signal photons will be clustered at similar heights between one pulse and the next (i.e., information from many pulses can be used to identify signal photons from a single pulse).

The first step in ATLAS data processing is the generation of ATLAS L2A Global Geolocated Photon Data (ATL03). In order to reduce the volume of photon data telemetered from ATLAS for ground processing, on-orbit algorithms select a window of photons around the likely surface. The ATL03 algorithm provides an absolute time and position (latitude, longitude, and height with respect to the WGS-84 ellipsoid) for these photons and identifies the subset that likely contains the surface. The window height varies from tens to thousands of meters over the ice sheets and includes both signal photons (from the ice surface) and background photons (from the sun and instrument). The ATL03 algorithm applies at least one classification for each photon (land, ocean, sea ice, land ice, and inland water) based on predefined overlapping masks. The ATL03 photon classification algorithms identify each photon as either likely signal or likely background, with a confidence value for each likely signal photon based on the ratio of the

photon density around each photon to that expected for background noise only. ATL03 also applies geophysical corrections (e.g., solid Earth tide, atmospheric delay). ATL03 data are arranged by beam in the along-track direction and segmented into granules that are several minutes long, according to predefined regions separated by latitude. The segmentation of granules is consistent between repeat orbits. The basic steps of the ATL03 algorithm over the ice sheets are described in Neumann et al. (this issue). ATL03 is the common input for all other ATLAS algorithms over land surfaces: land ice (ATL06 algorithm; this paper), sea ice (ATL07; Kwok et al., 2019), vegetation (ATL08, Popescu et al., 2019), and ocean (ATL12, Morison et al., 2019).

3. ATL06 algorithm

The ATL06 algorithm was developed to derive estimates of land-ice height that capture details of the surface at small (< 50 m) spatial scales, while retaining along-track surface-slope information that can help distinguish accurate surface-height estimates from false surface detections that can result from background noise and clouds (blunders). Over the vast majority of the Earth's land ice, the surface is smooth, with low (< 1°) variations in surface slopes at scales less than a few hundred meters (Markus et al., 2017). This geometry allows us to approximate surface profiles using a segmentation technique, where we divide the along-track photon data from each beam into short (40 m) overlapping segments. For each segment, we fit the photon data as a function of along-track distance with a linear model. The surface height is then taken to be the value of the model at the center of the segment, and the slope of the line (in the along-track direction) is recorded (Fig. 3).

The segment length of 40 m is about twice the width of the horizontal distribution of photons from a single pulse, and is thus the shortest distance over which the along-track slope of each segment can be determined with a precision of a few percent or better. We space the segment centers by 20 m along the ground tracks, so that consecutive segments overlap by 50%. This means that errors in the heights and slopes estimated for consecutive segments are not independent. When both beams in a pair are able to resolve the surface height, offset by 90 m across track, we also estimate the cross-track surface slope.

The ATL06 algorithm performs the following functions using the ATL03 data as input: (i) selects photons for the along-track location of each segment (Section 3.1); (ii) fits line segments to those photons with an iterative rejection of background photons (Section 3.2); (iii) calculates instrument bias corrections (Section 3.3); and (iv) determines output parameters for each segment (Section 3.4). We describe each of these steps in detail below.

3.1. Initial photon selection

The ATL03 calculation geolocates photons and assigns each photon a land-ice specific classification (*signal_conf_ph*) which has three categories [five values]: (i) surface or "likely signal" (with high [4], medium [3] or low [2] confidence); (ii) within 10 vertical meters of the surface [1]; or (iii) background [0]. ATL06 processing refines this classification by decreasing the height of the surface window to produce a well-characterized selection of photon heights relative to the ATL06 40-m segment, and, in rare cases when the ATL03 classification has failed to detect the surface, by using a backup algorithm (defined below) to search for weak surface returns.

Estimating surface height based on the mean height of a collection of signal and background photon heights leads to errors in the surface height that depend on the relative abundance of each type of photon and on the vertical extent of the surface window within which photons are distributed. Selecting photons within a narrower window results in both a smaller fraction of background photons and in a smaller vertical distribution of their height, each of which leads to a smaller contribution from the background photons to the error in the surface height. We use a two-stage process to identify the smallest vertical range of photons that can be used to define the surface: (i) the *photon-classification* stage, where a subset of the photons are classified as likely containing the surface (i.e., "likely signal" photons), and (ii) the *photon-selection-refinement* stage, where we decrease the height of the window around the segment that best fits the selected photons to exclude background photons while retaining the signal photons.

We established a *photon-distribution criterion* to determine whether

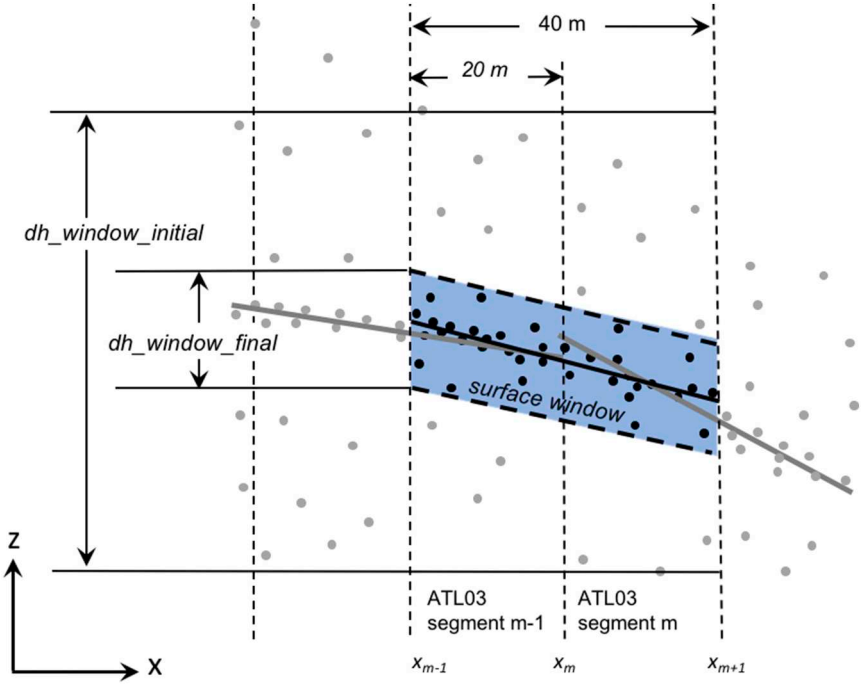

Fig. 3. ATL06 segment fitting: Along-track segments fit to photon heights (gray circles) as a function of along-track distance (x) for segment *m* (black line) and neighboring segments *m + 1* and *m-1* (gray lines). Photons selected by the ATL06 photon-refinement algorithm for reference point *m* are shown in black.

the number of photons in a 40-m segment is adequate to define its height and slope: the minimum number of photons per segment must be at least 10, and the horizontal distance between the first and last photon must be at least 20 m. Segments with smaller photon counts, or that have small along-segment separation between photons, often do not have a well constrained height or slope.

For each segment, our algorithm first attempts to use the ATL03 photon classification to find a collection of photons that define the surface height. If the photons flagged by ATL03 as possible signal photons meet our criterion, they are used to define an initial segment. These flagged photons are then fit with a sloping line segment, and the robust spread (see Appendix A) of the residuals to this line, σ_r, is calculated. The flagged photons, plus any unflagged photons that fall within a window whose height is either ± 1.5 m or ± $3\sigma_r$ (whichever is larger), are passed to the segment-fitting and selection refinement algorithm.

In some cases where the surface returns are weak or the surface is complex, the ATL03 signal selection algorithm may not identify sufficient photons to meet our criterion. In these cases the ATL06 algorithm attempts to use a backup algorithm to find a group of photons that meet these criteria. The backup algorithm doubles the size of the along-track segment from 40 m to 80 m without changing the segment center. All photons within the 80 m segment are collected in 10-meter vertical bins. The vertical bin with the largest photon count (N_{max}) is identified as a likely location of surface signal, as well as all bins whose photon counts are greater than N_{max}-$N_{max}^{1/2}$. This threshold ensures that we capture all bins whose photon counts are not significantly different from the maximum, assuming that all counts were drawn from a Poisson distribution. The photons within these vertical bins, and within ± 20 m along track of the segment center, are then passed to the segment-fitting and selection-refinement algorithm. Information about the data used in the initial photon classification is provided in the *signal_selection_source* parameter of the ATL06 data product.

3.2. Segment fitting and photon selection refinement

The basis of our segment-fitting and photon selection refinement process is the statistical "three-sigma editing" algorithm. The simplest version of this process would fit a collection of photons from an individual 40 m segment of ground track with a linear model, then calculate the residuals between the photons and the model. It would then calculate the standard deviation of the residuals, σ, and then edit the collection of photons to contain only those that lie within ± 3σ of the line. The process would be repeated until two subsequent iterations produced the same collection of photons within a predefined threshold. Our version of the process makes three improvements to the simplest version: (i) instead of the standard deviation, we estimate σ using a robust estimator (Appendix A) that takes into account the likely number of background photons in the collection. This choice of statistic allows our iterations to converge even in the presence of background photons that would otherwise dominate the calculation; (ii) instead of using 3σ as the half height of the window in the next iteration, we use the maximum of 3σ and 0.75 times the half height of the current window. This option prevents large changes in the window between subsequent steps from missing the surface entirely; (iii) we limit the minimum window height to ± 1.5 m. This hard limit prevents the algorithm from converging on a window that excludes too many potential signal photons.

The ATL06 fitting process begins with the likely signal photons from the photon-classification stage of ATL03 algorithm, or by the ATL06 backup algorithm in cases where ATL03 was unsuccessful in identifying a likely surface. We refer to this initial collection of likely signal photons using the subscript *likely*, and refer to the collection of photons identified at each step of the iterative routine with the subscript *i*. At each iteration of the fitting process, we:

1. Perform a least-squares linear fit to the currently selected photons, as a function of (x_{photon}-x_{ctr}), where x_{photon} are the photon along-track coordinates, x_{ctr} is the current segment center. This gives the segment mean height and surface slope [*h_mean*, *dh/dx*] and residuals to the model, r_i.
2. Calculate the median, r_{med}, and background-corrected spread (Appendix A) of the distribution, σ_o, of the residuals for the selected photons. If σ_o is > 5 m, it is set to 5 m. Calculate the expected spread of the return photons, *h_expected_RMS* based on the current *dh/dx* estimate using Eq. (1), assuming zero roughness.
3. Calculate the residuals of all of the likely signal photons to the current model estimate, r_{likely}.
4. Select photons from among the initially selected photons for which $|r_{likely}$-$r_{med}| < H\_window/2$, where *H_window* represents the height of the window and *H_window* = max(6 σ_o, 6 *h_expected_RMS*, 0.75 *H_window_last*, 3 m). Here *H_window_last* is the height of the window from the previous iteration.
5. If step *d* has changed the photon selection, and if fewer than twenty iterations have happened, return to (a).
6. Propagate the error of the least-squares fit for the last iteration to give error estimates for the segment height and slope, assuming that the errors in the individual photon heights are *max*(σ_o, *h_expected_RMS*).

The least-squares fitting step estimates the surface height at the center of the segment, x_{ctr}. Because the least-squares calculation effectively gives a weighted mean of the selected photon heights, we report this value as *h_mean*.

3.2.1. Evaluating the quality of surface-height estimates

Under strong-signal conditions where surfaces will typically be well approximated by a linear along-track model over 40 m length scales, our algorithm reliably converges in less than five iterations. For weak signals and high background rates, the algorithm may fail to converge on a narrow (3 m vertical) window that includes the surface. Most often, these blunders result in large values of H_{window} and small photon counts within the window. We tested the refinement process for three cases using simulated data (Fig. 4). In Case A (Fig. 4a), the surface location from ATL03 was unambiguous, and the algorithm converged on a narrow (3 m) window after one iteration. In Case B (Fig. 4b), there were fewer likely signal photons classified by ATL03, and they were distributed over a wider window. The ATL06 algorithm converged on a 3-m window after six iterations. In Case C (Fig. 4c), there were no photons flagged by ATL03, and the ATL06 algorithm went through 14 iterations before converging on a wide, sloping window, centered 5 m below the surface. Cases A and B represent successful window refinement, while case C represents a blunder.

We characterize the signal strength by calculating the signal-to-noise ratio (*SNR*):

$$SNR = \frac{N_{tot} - N_{BG}}{N_{BG}} \tag{2}$$

hwhere N_{tot} is the total number of photons in the window, and N_{BG} is the number of background photons in the window, calculated as:

$$N_{BG} = BGR\ \frac{2\,H_{window}}{c} \tag{3}$$

where *BGR* is the background photon rate averaged over 50 shots as determined in the ATL03 algorithm (Neumann et al. (this issue); H_{window} is the window height; and *c* is the speed of light.

It is unlikely that we will observe large *SNR* values for segments fit exclusively from noise photons, which lets us use this parameter to reject segments resulting from blunders. To quantify the probability of a blunder for each segment, we calculate the probability that the algorithm would converge to a segment with a signal-to-noise ratio at least as large as the observed SNR_{obs} if the input were pure random noise at

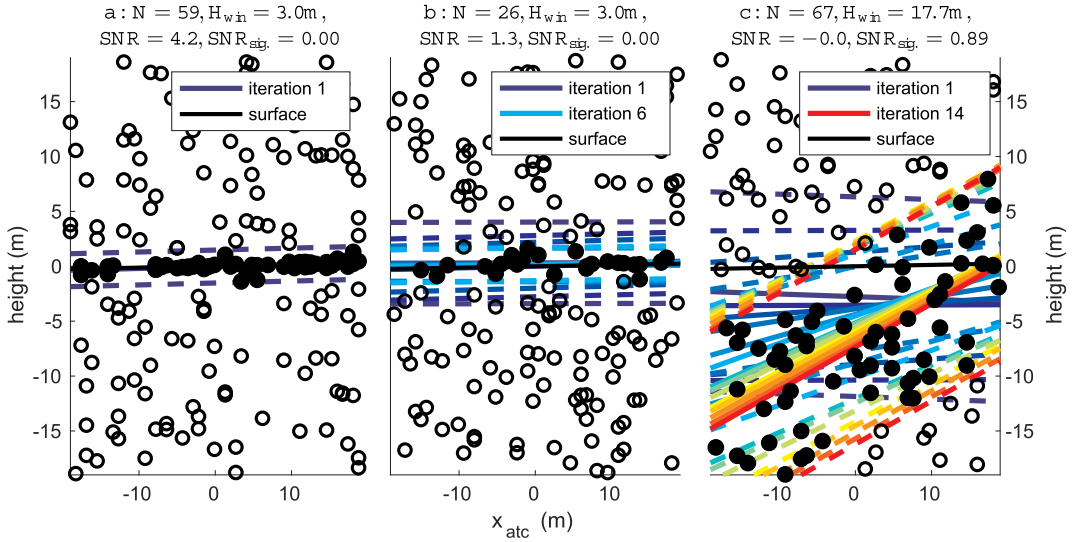

Fig. 4. ATL06 photon selection refinement process. Photon selection refinement process shown for three cases based on simulated ATL03 photons (open circles: all photons; filled circles: final ATL06-selected photons): a) segment where the signal has moderate-strength, where ATL03-flagged photons (not shown) defined a narrow surface window; b) segment with a weaker signal, where six iterations were required to converge on the surface; c) segment with no ATL03-flagged photons, where, after 14 iterations, the algorithm converged to a wide, sloping window centered 5 m below the surface. In all plots, solid colored lines show the segment height for different iterations, proceeding in spectrum order from blue to red. Dashed lines show the bounds of the surface window, with colors matching the surface-fit color. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

rate *BGR*, distributed over the initial height range, $H_{initial}$, from the photon-selection algorithm (Section 3.1). Using a lookup table that we generated by applying the ATL06 algorithm to simulated noise-only ATL03 data (Smith et al., 2018) we calculate:

$$snr\_\text{significance} = \text{P}(SNR > \text{SNR}_{\text{obs}} \mid \text{BGR}, \text{H}_{\text{initial}}) \quad (4)$$

Small values of *snr_significance* imply a low likelihood that an observed segment results from a blunder. ATL06 processing provides valid surface heights only when *snr_significance* < 0.05, and segments with *snr_significance* > 0.02 are flagged as potentially unusable. Note that in Fig. 4, the non-blunder cases (A and B) both had *snr_significance* < 0.005, while the blunder (C) had a *snr_significance* of 0.89. In each case, the *snr_significance* test resulted in the correct decision.

To allow users to make a simple decision about whether to use a particular segment or reject it, we defined a quality parameter *ATL06_quality_summary* which is calculated for each segment. This parameter combines information about the signal-selection process, the signal-to-noise ratio, the spread of residuals, and the estimated error, and has values of 0 and 1. A value of 0 indicates that the quality is high. The criteria for this selection are described in Section 3.4.

3.3. Instrumental biases and their corrections

After data segmentation, we improve our estimate of the surface height and estimate the error in our recovered height based on the number and distribution of photons within each 40 m segment around the linear fit. Although the results of the iterative least-squares fitting algorithm provide an estimate of the ice-sheet surface height, three major sources of bias remain:

1. **Sensitivity to outliers**: The mean-based height estimate used in the iterative fit is sensitive to outlying residuals, including those associated with background photons, and is also sensitive to delays associated with volume scattering of laser light in the snowpack (Smith et al., 2018).
2. **Transmit-pulse shape**: The transmit pulse is not symmetric in time around its mean; therefore, truncation of the return pulse can lead to biased estimates of the surface height.
3. **First-photon bias**: The detector electronics are less likely to detect photons that arrive on the late (long-range) side of each return pulse. This occurs because the detectors require a finite amount of time to recover from a photon-detection event, and can lead to upward biases in the surface height estimates.

Each of these biases could be influenced by changes in the transmit-pulse shape, by changes in surface roughness, or, in the case of the first-photon bias, by changes in return strength and width; changes in any of these biases would lead to time-varying biases in the apparent surface height. We have developed strategies to mitigate the impact of each of these types of errors by analyzing the distribution of the residuals for each segment. These “residual-based corrections”, each of which is described below, are either applied to the final surface-height estimate or included in the ATL06 product to account for these biases.

3.3.1. Sensitivity to outliers

Under ideal conditions, the slope-corrected mean of the selected photons (*h_mean* from Section 3.2) provides an accurate estimate of the surface height. The accuracy of this estimate declines as atmospheric scattering (Yang et al., 2011) and volume scattering (Smith et al., 2018) introduce delays in a portion of the surface photons, and also when the surface window fails to converge completely. Median-based estimates of surface height are more robust than mean-based estimates because they are less sensitive to the tails of the distributions caused by scattering processes. For subsurface scattering, using a median-based estimate of the surface height reduces these biases by a factor of two or better compared to a mean-based estimate (Smith et al., 2018), and we expect similar improvements in biases due to blowing snow (Yang et al., 2010).

The difference between the mean-based surface height and the median-based surface height is equal to the median of the residuals between the selected photons and the linear segment model, and is given in the *med_r_fit* parameter, which is added to *h_mean* to give the median-based height. Note that *h_mean* + *med_r_fit* gives the median-based height, uncorrected for the first-photon bias. As we describe in Section 3.3.3, the mean-to-median correction in the ATL06 land-height parameter is calculated during the first-photon-bias correction. The *med_r_fit* parameter is not used in the calculation of *h_li*, but is included as part of the land-ice product so that users wishing to calculate surface

heights without a first-photon bias (or wishing to implement a different first-photon bias model) can do so using only ATL06 parameters.

3.3.2. Transmit pulse shape bias

3.3.2.1. Cause of the transmit pulse bias. The surface ranges used in ATL06 height estimates are calculated based on the difference between the time that each pulse was transmitted and the time estimated for a collection of received photons (photon events). The times of the transmit pulses are calculated based on an estimate of the centroid of all transmit photons for each pulse. Height estimates based on the mean of *all* received photons should give an unbiased estimate of the surface height. However, the asymmetry of the transmit pulses leads to three potential problems:

1. Height estimates based on windowed means of received photons will have biases whose magnitudes depends on the height of the window, with larger windows associated with smaller biases. This dependence comes about because larger windows capture a sample of photons whose statistics better approximate those of an un-windowed sample;
2. Height estimates based on the median time of the received photons will be biased, because the median is less sensitive to the photons in the long-time tail of the received-photon distribution; and
3. The magnitude of biases in windowed means and medians depend on the roughness and slope of the surface, with larger roughness/slopes associated with smaller biases. This dependence comes about because larger slopes and roughnesses mix photons between the leading and trailing edge of the return, to produce increasingly random samples, less influenced by the shape of the transmit pulse.

We developed the ATL06 algorithm based on two sources of information about the expected pulse shape. The earliest source was a set of waveforms measured from an engineering prototype laser, which had a severely skewed pulse shape. Later, we obtained samples of data from the ATLAS Transmitter Echo Path (TEP), which monitors the shape of the transmit pulse by directing a small number of photons from each pulse through a delay fiber and back to the detectors (for additional TEP details, see Neumann et al., this issue). The TEP observations show a narrower, more symmetric pulse shape compared to the prototype laser, but still are slightly asymmetric in time with a relatively slow decline in power on the trailing edge of the pulse compared to the leading edge (Fig. 5a). While we expect the TEP measurements to be representative of on-orbit ATLAS pulse shape, the prototype laser measurements provide a clearer (i.e., higher magnitude) representation of the pulse-asymmetry bias, and we feel that it is helpful to present both.

Fig. 5a and b show pulses from ATLAS, and from the prototype laser, respectively. Skew in the ATLAS pulse leads to a mean-median bias of 14 mm (with the median higher (earlier) than the mean) for a return from a flat surface; from a rough surface and the same pulse, the mean-median bias is only 5 mm, because the roughness reduces the skew of the received pulse. Because the ATLAS pulse is narrow compared to the ± 1.5 m surface window, the bias on the mean is approximately zero. The prototype-laser pulse shows more significant biases (which is, again, not expected for the ATLAS flight lasers), with a flat-surface median bias of 120 mm, a rough-surface median bias of 90 mm, and a rough-surface mean bias of 24 mm. This last value is nonzero because a small amount of the energy from the tail of the prototype pulse is lost from the edge of the surface window.

To demonstrate how the magnitudes of the biases depend on the transmit pulse shape, the window height, and the surface roughness, we generated the equivalent of 50,000 segments of simulated photons from a strong beam over a bright surface of variable roughness (Appendix B). For each segment, we calculated the windowed mean and median return time for window heights between ± 1.5 and ± 4 m for an expected ATLAS pulse (Fig. 6a, b) and for the prototype pulse (Fig. 6e, f). The

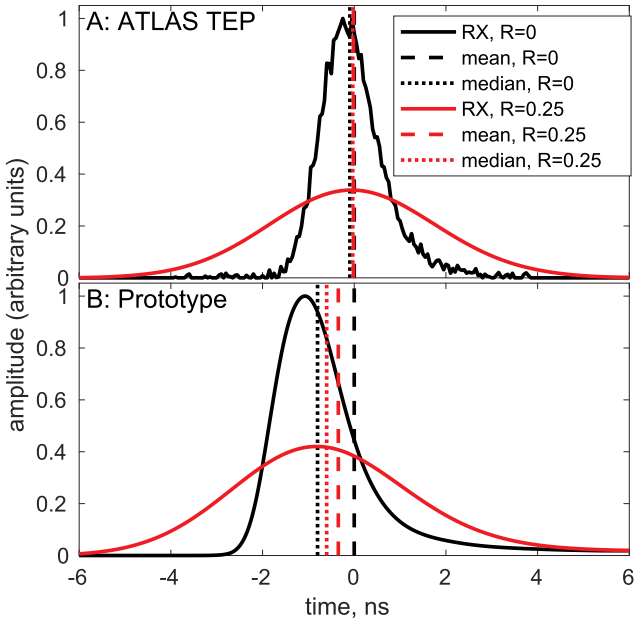

Fig. 5. Transmit pulse shape bias. a) The shape of the return laser pulse from a smooth, flat surface (equivalent to the transmit pulse), as predicted based on the ATLAS TEP, and the modeled return shape after reflection from a rough (0.25 m RMS) surface. b) The shape of a skewed return pulse from a smooth, flat surface, as predicted from measurements on a prototype laser, and the modeled return shape after reflection from a rough (0.25 m RMS) surface. For both plots we show the mean (dashed line) and median (dotted line) photon heights for the transmit and return pulses.

only significant bias for the ATLAS pulse is in the median at near-zero roughness, where bias magnitude reaches up to 15 mm. The prototype-pulse biases are larger, with magnitudes as high as 120 mm for the median over smooth surfaces, and 64 mm for the mean over smooth surfaces with a narrow window (again, not expected for the ATLAS flight lasers). For consistent transmit pulse shapes and surface conditions, we expect these biases to remain constant. However, if the transmit pulse shape changes throughout the mission (as the lasers age), and as surface roughness varies seasonally, we expect the biases to change with time.

3.3.2.2. Transmit pulse shape bias correction. To estimate the transmit laser-pulse shape bias, we calculate a simulated return using TEP measurements from the ATL03 data product (Neumann et al., this issue), the measured variance of the simulated surface return, σ_{rx}^2, and the height of the surface window, as follows:

1. We artificially broaden the transmit pulse waveform by convolving it with a Gaussian function whose sigma parameter is equal to $(\sigma_{rx}^2 - \sigma_{tx}^2)^{1/2}$.
2. We iteratively window the simulated waveform around its centroid, using a window whose height matches that used for the surface return, to obtain a mean-based measurement consistent with that described in Section 3.2. The iterations ensure that the window is centered on the centroid of the photons contained within it.
3. Once the iterations are complete, we calculate the median and mean of the windowed broadened waveform to obtain *tx_med_corr* and *tx_mean_corr*. *These variables become corrections which must be added to an uncorrected value to correct for the bias*. Thus, *h_mean* + *tx_mean_corr* gives a corrected mean-based estimate of the surface height and *h_mean* + *med_r_fit* + *tx_med_corr* gives a transmit-pulse-shape-bias-corrected, median-based estimate of the surface height.

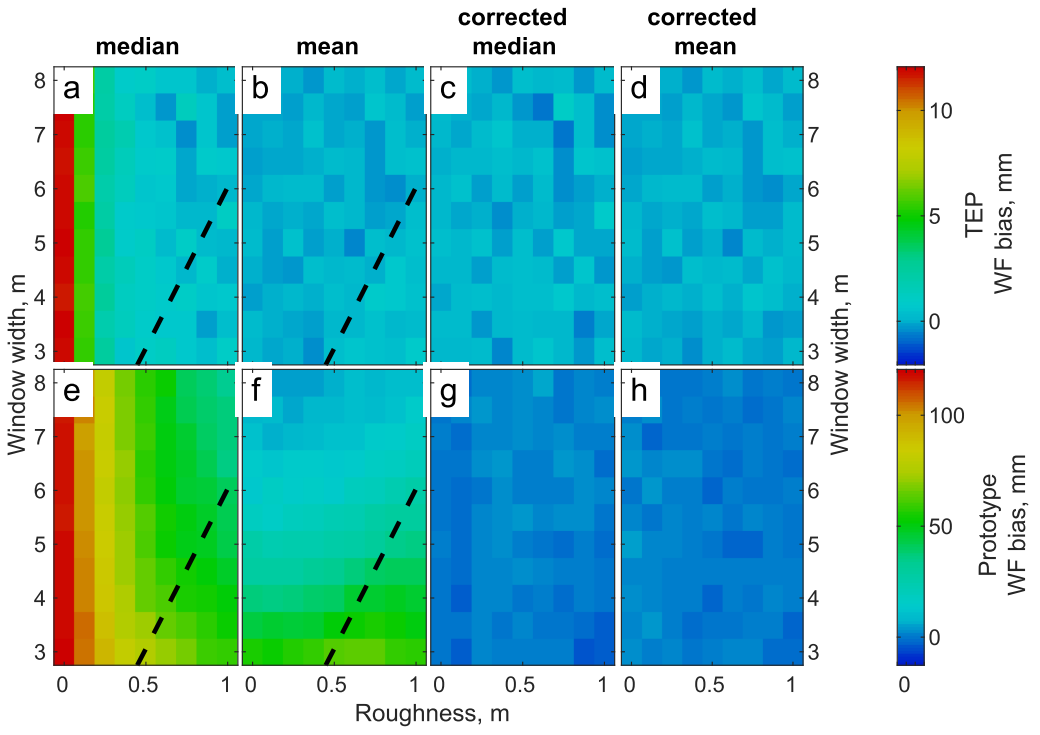

Fig. 6. Transmit pulse-shape bias correction. The transmit pulse-shape bias, and its correction, as a function of surface roughness and window height. The top panels (a–d) show transmit pulse-shape biases and corrected values for simulated data based on the expected ATLAS waveform, as measured by the instrument's TEP. The bottom panels (e–h) show results calculated for an early prototype laser that had a more severely skewed pulse. Panels a and e show the bias in the uncorrected median, b and f show the bias in the uncorrected mean, c and g show the bias in the corrected median, and d and h show the bias in the corrected mean. Note that the color scales are different between the prototype and TEP results, with much larger biases for the prototype laser. The dashed line in 6a, 6b, 6e, and 6f shows the truncation expected based on the three-sigma editing applied to photon return times during photon-selection refinement. We do not expect to see window heights in the region to the lower right of this line. (For interpretation of the references to color in this figure, the reader is referred to the web version of this article.)

These calculations are performed for each segment, using the TEP measurement provided in the source ATL03 file. Fig. 6c and d show the magnitudes of the corrected median and mean for the ATLAS pulse, and Fig. 6g and h show the same values for the prototype-laser pulse. In both cases, the correction reduces the magnitude of the biases to < 1 mm. Most variations in the corrected values result from random scatter that are visible in Fig. 6c and d because of the color scale's small range of values.

3.3.3. First-photon bias

3.3.3.1. Cause of the first-photon bias. The ATLAS receiver uses solid-state multi-pixel photomultiplier elements to detect individual photons. Each pixel (i.e., one element of a photomultiplier detector) can sense only one photon arrival every ~1 ns. If a photon reaches a pixel within ~1 ns of an earlier arrival, the second photon's arrival will not be recorded. The period of time after a detection event is referred to as the "analog-stage deadtime". Each pixel remains inactive (i.e., is insensitive to another photon) until one analog deadtime interval or greater has passed with no arriving photons. This stage of the detector is potentially paralyzable, in that if the input rate of photons into a given pixel exceeds 1 GHz, the pixel will produce no output after the first photon arrival, because each new photon will overlap and extend the deadtime.

The amplified output of the pixels feeds the input stage of a set of timing modules. This input stage imposes a "digital-stage deadtime" of 3.2 ns on the inputs to the timing modules, but, unlike the analog-stage deadtime, this stage is nonparalyzable: after detecting a pulse, this stage is inactive for the duration of the digital-stage deadtime (3.2 ns), then becomes active again.

The combination of the effects of the analog-stage and digital-stage deadtimes implies that photons arriving earlier from a given pulse are more likely to be detected than those arriving later, potentially resulting in a positive bias in estimates of surface height. This problem, which we term the first-photon bias (FPB), is most severe over bright or highly specular surfaces, especially if the combination of surface slope and roughness produce return durations approximately equal to the deadtime. Detector deadtime also reduces the number of photons detected, resulting in apparent reductions in the surface reflectance. ATLAS's detectors reduce the effects of the FPB by splitting each surface return between four pixels (for a weak beam) or 16 pixels (for a strong beam) so that during the brightest expected returns from snow surfaces, each pixel will receive, on average, 0.75 photons for each transmit pulse. This ratio is the same for both weak and strong beams because the number of detectors is proportional to the beam's output power. This greatly reduces the FPB magnitude relative to a single-channel instrument with the same pulse energy, but does not eliminate it in the case of high photon-rate returns. The number of received photons can be adjusted on orbit by modifying the ATLAS laser energy level.

Fig. 7 shows a simulated photon-return histogram for a strong-beam ATL06 segment (i.e., all photons accumulated from 57 pulses) over a flat surface. The simulated distribution of arriving photons is drawn from a symmetric normal distribution with a standard deviation of 1 ns (0.15 m), representing an ATLAS pulse with a small amount of broadening due to surface roughness. The photon count incident on the detector is set to 16 per pulse, about 30% higher than that expected for a strong beam over a bright surface. Because of the FPB, the histogram of detected photons is skewed to earlier arrival times, such that its mean and centroid are each about 40 mm too high, and only about 63% of all photons are detected.

We simulated the FPB for 50,000 segments of strong-beam data, for a range of surface reflectance values (0–1.5; the value of 1.5 accounts for the enhanced reflectance in the backscatter direction that has been observed for snow surfaces (Kaasalainen et al., 2006)) and a range of effective surface roughness values. The surface reflectance values are used to determine the number of photons entering the detector: If the effective reflectance is 1, N_{signal} is equal to the predicted white-surface value of 0.75 times the number of pixels in the detector; smaller reflectances result in proportionally smaller signal counts. We then calculated the average magnitude of the FPB for mean (Fig. 8a) and median (Fig. 8b) photon return times, as well as the error in the reflectance that would be estimated based on the detected photon count (Fig. 8c). The largest fraction of photons is lost for smooth, bright surfaces, where as few as 55% of total incident photons may be

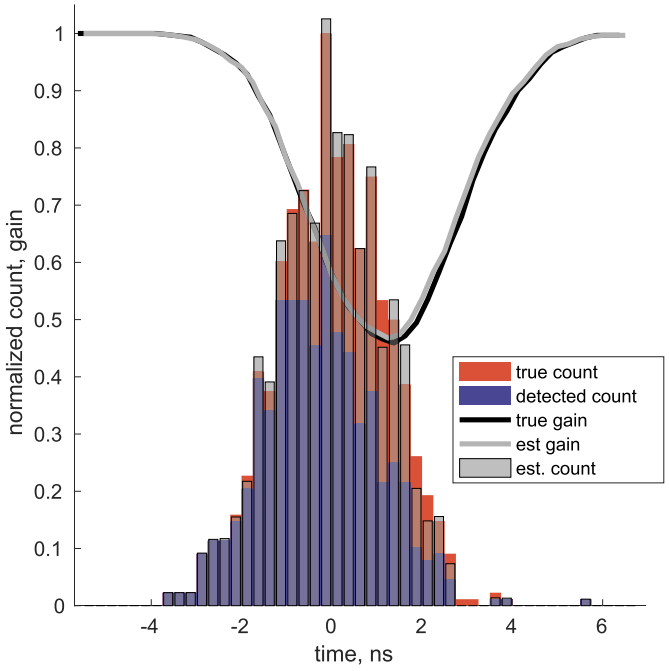

Fig. 7. The effect of first-photon bias on return shapes. The total (red) and detected (blue) distribution of photons from a simulated collection of photons from 57 pulses from a moderately rough (standard deviation of 15 cm), bright (16 photons/pulse) segment, as measured by a strong beam. The black curve shows the (true) effective gain calculated during the simulation. The gray curve (nearly overlying the black curve) shows the gain estimated as part of the correction. The gray bars show the recovered estimate of the photon distribution. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

detected. The resulting bias in these cases is relatively small, because the reflected pulse is narrow. We found the largest biases for roughness values to be around 0.2 m, where the reflected-pulse duration is comparable to the deadtime; in these cases, for unit surface reflectance, the bias in the mean photon height is around +29 mm, and the bias in the median is close to +43 mm.

We can describe how the FPB modifies the returned pulse by calculating the ratio between the histogram of detected photons for a collection of pulses, $N_d(\tau)$, and the histogram of all photons (detected and undetected) for that collection of pulses, $N(\tau)$, where τ is the difference between the arrival time for each photon and the round-trip travel time to the surface:

$$N_d(\tau) = G(\tau)N(\tau) \tag{5}$$

If we know whether the pixels of the detector were active at a particular time for a particular pulse, we can calculate the gain $G(\tau)$for a detector with N_{pixels} pixels, as a function of time for a collection of photons from N_{pulses} pulses:

$$G(\tau) = \frac{1}{N_{pulses}N_{pixels}} \sum_{pulses,pixels} A(\tau, px, p) \tag{6}$$

where $A(\tau,px,p)$ is a function that is equal to 1 when pixel px is active for pulse p and zero otherwise. Averages of A (over multiple pulses and pixels) tend toward zero when the rate of photon arrivals is large, and remain low for a period of around one deadtime after the rate of photon arrivals falls back to zero. The gain for the simulated return is plotted in Fig. 7; it falls from unity at the start of the return, to a minimum of around 0.47 shortly after the peak of the return, then recovers back to unity over the subsequent 3 ns.

3.3.3.2. First-photon bias correction. We correct for the FPB for 40 m segments containing N_{pulses} return pulses. For each pulse within this segment, we assume that the effective gain can be represented by the same function $G'(\tau)$. We estimate $G'(\tau)$ based on the residuals between the best-fitting linear model (Section 3.2) and the photon ranges, and then use this function to correct the measured histogram of photon arrival times to an estimate of the histogram of all arrival times, detected or undetected. To estimate the gain function, we convert the photon height residuals to a histogram of time residuals, $N_{detected}(\tau)$. We then calculate an estimate of the gain, $G'(\tau)$:

$$G'(\tau) = 1 - \frac{1}{N_{pulses}N_{pixels}} \sum_{\tau - t_{dead} < \tau' < \tau} N_{detected}(\tau') \tag{7}$$

Here, we assume that every detected photon marks the start of an interval during which one pixel was inactive for one pulse and average over all pulses and pixels to calculate the mean number of dead pixels over the segment. We consider only the nonparalyzable component of the deadtime due to the digital-stage input. Except at unusually high photon-return rates (> 1 GHz/pixel), the 1 ns analog-stage deadtime should have little effect on the return. We then use $G'(\tau)$ and the detected photon count, $N_{det}(\tau)$, to estimate the number of photons that were incident on the detector: $N_{est}(\tau) = N_{det}(\tau)/G'(\tau)$.

In Fig. 7 we show the estimated gain for the collection of pulses and the resulting $N_{est}(\tau)$. The true and estimated gain are identical within a few percent, and the estimate of $N_{est}(\tau)$, matches $N(\tau)$ to within the precision expected based on the Poisson-random sampling that produced the detected photons (gray curve in Fig. 7). Errors in corrected mean (Fig. 8d) and median (Fig. 8e) height and error in the reflectance calculated from the corrected photon count (Fig. 8f) are consistently small (within a few mm of zero for heights, 1–2% for reflectantance), with fluctuations driven largely by the statistical noise in the calculation. These statistics were generated for a strong beam; similar calculations for a weak beam show around twice as much scatter, but no net bias. The error estimate in the FPB correction is calculated from the statistical uncertainty in the median of the photon distribution, weighted by the inverse gain; in this calculation we assume that there is no significant error in the calculated gain.

The ATL06 land-ice height parameter, *h_li*, includes median-based corrections for biases due to the first-photon bias (*fpb_med_corr*) and due to asymmetry of the transmit pulse (*tx_med_corr*). The sum of *h_fit* and these two corrections gives the median-based surface height.

3.4. Summary of ATL06 parameters

In addition to basic altimetry information (latitude, longitude, height, uncertainty, and timing of segments), the ATL06 product contains a variety of parameters that are intended to allow users to distinguish high-quality segments from segments with potentially large errors. ATL06 files will be distributed by the National Snow and Ice Data Center (NSIDC) as HDF5 (Hierarchical Data Format, version-5) files that contain tiered groups of parameters. Each file contains one top-level group for each of the six beams. Within each beam, a top-level group contains the basic altimetry information, and subgroups contain the quality-and-correction parameters. Detailed descriptions of all parameters are available in the ATL06 ATBD (Smith et al., 2018), and a full summary of all parameters is provided here: https://nsidc.org/sites/nsidc.org/files/technical-references/ATL06-data-dictionary-v001.pdf. We describe the top-level parameters below:

3.4.1. Land-ice height

The land-ice height estimate in the ATL06 product (*h_li*) is the sum of the segment-center height from the linear fit (*h_mean*), the transmit-lpulse-shape correction (*tx_med_corr*), and the first-photon-bias correction (*fpb_med_corr*):

$$h\_li = h\_mean + tx\_med\_corr + fpb\_med\_corr \tag{8}$$

Note that both the transmit-pulse shape correction and the first-photon-bias correction are median-based corrections, so that *h_li* gives

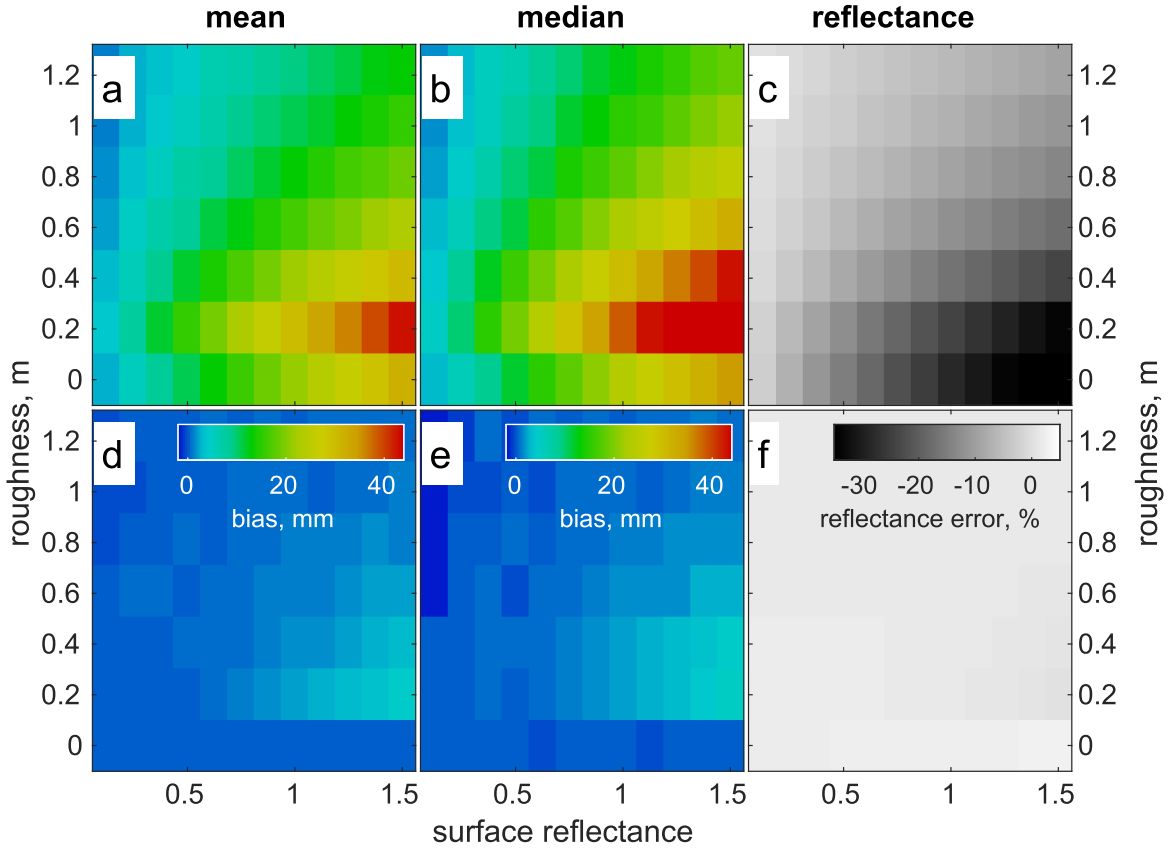

Fig. 8. First-photon bias correction. The first-photon bias, and its correction, as a function of surface reflectance and surface roughness. Panels a–c show the error due to the first-photon bias in a: the mean height, b: the median height, and c: the estimated surface reflectance. Each pixel in each panel gives the mean error over 50,000 segments of simulated data. Panels d–f show the errors in the corresponding corrected values.

an estimate of the median photon height.

The two correction values are provided so that users can assess the influence of their variability on height-change estimates, or can convert height estimates to mean-based estimates (using the *tx_mean_corr* and *fpb_mean_corr* parameters).

3.4.2. Error estimates

The uncertainty in the land-ice height is a function of both ranging errors and errors in geolocation of the return. The land-ice height uncertainty (*h_li_sigma*) gives the surface height error due to ranging errors only and is calculated as the maximum of the propagated error in the least-squares fit of the selected photons and the statistical uncertainty in the first-photon bias correction. We use the maximum of these error estimates rather than their root sum square because the two are partially redundant; each calculates the sampling error in the height estimate in a different way, and we choose the more pessimistic of the two to provide a conservative error estimate. The error component due to the combination of errors in the satellite position, errors in the geolocation of the laser spot, and errors in the along-track coordinate of the surface slope is given by *sigma_geo_h*. In estimating the mean error of a collection of ATL06 measurements, *h_li_sigma* should be treated as uncorrelated between non-overlapping segments (i.e. with centers separated by 40 m or more), while *sigma_geo_h* should be treated as a systematic error that is correlated at the scale of a few kilometers or more.

3.4.3. Segment geolocation

The *latitude*, *longitude*, and *delta_time* parameters give the geographic coordinates of the centers of the 40 m segments and the segment timing. These are calculated from the latitude, longitude, and time values for the ATL03 photons selected for each segment by evaluating the least-square regression at the along-track coordinate of the segment center. The errors in the segment geolocation parameters are given in the *geolocation* subgroup.

3.4.4. Quality flag

The *atl06_quality_summary* flag is intended to give a high-level indication of whether a particular segment should be used for height-change calculations, where 0 indicates highest quality and 1 lowest quality. The flag is set to 0 only if all of the following are true:

1. The signal selection (Section 3.1) was made based on ATL03 photons identified as low or better confidence (*signal_conf_ph* ≥ 2).
2. The spread of photon residuals around the linear model (*h_robust_spread*), calculated as the background-corrected robust dispersion estimate, is < 1 m.
3. The estimated error in the surface height (*h_li_sigma*) is < 1 m.
4. The segment is unlikely to have resulted from a blunder in signal selection (*snr_significance* < 0.02).

4. ATL06 algorithm testing

The ICESat-2 project collected data with airborne simulators (e.g. Multiple Altimeter Beam Experimental Lidar, MABEL; McGill et al., 2013) that were used extensively in the development of the ATL03 algorithm (Neumann et al., this issue). These data, however, do not provide accurate estimates of the true surface height at the sub-centimeter precision that is required of ATL06 surface-height estimates. We tested our ATL06 surface-finding algorithm against synthetic data generated using a workflow described in Appendix B. By testing ATL06 algorithms with synthetic data, we have been able to quantify the ATL06 algorithm's performance against a known (true) surface height. Since ICESat-2 data are now available, we also demonstrate the ATL06

algorithm performance with real ATL03 data from ATLAS.

4.1. Photon selection and linear regression

We expect the signal-refinement strategy that we have implemented to produce a mixture of valid surface heights and blunders. Blunders will remain because we expect the *snr_significance* < 0.02 threshold to identify most, but not all of the blunders (errors of omission), and in some cases the algorithm will misidentify some of the valid segments as blunders (errors of commission). To help quantify the expected performance of the ATL06 algorithm under a range of surface roughness, signal strength, and background-noise conditions, we generated segments of random synthetic data representative of a weak beam for background photon rates between 0.25 and 10 MHz, for effective surface reflectance values between 2% and 100%, and for effective surface roughness values of 0 and 2 m. The variations in effective surface reflectance incorporate the combined effects of reduced albedo surfaces (e.g., due to impurities in the ice) and attenuation due to clouds. The effective surface roughness captures the combined effects of surface roughness and slope on the received waveform. In this experiment, for each value of the background rate, effective surface reflectance, and roughness, we generated 1600 segments of random synthetic ATL03-like data. In each segment there were N_{sig} signal photons, whose heights have a Gaussian temporal distribution calculated based on the effective surface roughness and a transmit pulse with a sigma = 0.68 ns using Eq. (1), and N_{noise} noise photons, whose heights have a uniform random distribution over a 200 m total window height. N_{sig} and N_{noise} were generated as Poisson random variables, with expected values calculated:

$$N_{sig,expected} = 3\, N_{pulses} R_{eff} \tag{9}$$

$$N_{noise,expected} = \frac{BGR\, H_{window}}{c/2} \tag{10}$$

where: 3 is the number of signal photons expected for a weak beam under clear-sky conditions; N_{pulses} is the number of pulses in a typical 40 m segment (57); R_{eff} is the uncorrected surface reflectance; *BGR* is the background rate; H_{window} is the initial window height (200 m); and c is the speed of light. We then used the ATL06-only (backup) signal-finding routine and the segment fitting and refinement procedure to shrink the surface window from its initial height of 200 m. The result of this process for each segment was a surface-height estimate, h_{seg}, a surface-slope estimate dh_{seg}/dx, an error estimate for the surface height, $\sigma_{h_{seg}}$, a refined window height, W_{final}, and a count of the photons within the window, N_{final}. Based on the input background rate, W_{final}, and N_{final}, we calculated *snr_significance* for each segment for which the window contained > 10 photons.

To evaluate the effectiveness of the *snr_significance* criterion, we first evaluated whether each refined segment represented a valid surface detection or a blunder, using the arbitrary criterion that valid detections should have $|h_{seg}\text{-}h_{true}| < 1$ m and $|dh_{seg}/dx| < 0.1$. We tested two criteria for accepting or rejecting segments: one where segments would be accepted if *snr_significance* < 0.05, the other with *snr_significance* < 0.02; in either case, segments were only considered if the propagated error in the segment center height was < 1 m. Fig. 9a and b show the fraction of segment fits resulting in a blunder for the low- and high-roughness surfaces, as a function of *BGR* and $N_{sig,\ expected}$. Blunder rates increased with *BGR* and surface roughness, while they decreased with signal strength. For signal strengths greater than around 1 photon/pulse, the algorithm found the surface virtually 100% of the time. Fig. 8b and f show the fraction of blunders that remain in the dataset after all of those segments with *snr_significance* > 0.05 were eliminated. Over most cases considered, this criterion removed most blunders for smooth surfaces, and reduced blunders for rough surfaces to < 30% of the segments for which any surface height was reported. Reducing the acceptance threshold to 0.02 (Fig. 9c and g) improved the unidentified blunder rate in both cases. The highest concentration of unidentified blunders is found for rough surfaces, for weak returns (0.1–0.3 photons/pulse) and low (0.25–2 MHz) background rates. We expect that most users will choose not to use segments with such low signal strengths, which can be readily identified based on the low number of photons included in the segment fit. The converse problem, of valid segments eliminated by the *snr_significance* < 0.02 criterion, is illustrated in Fig. 9d and h. The highest concentration of valid segments eliminated by the criterion were for rough surfaces with low-to-moderate signal strength (0.3–1 photon/pulse) and large (2–8 MHz) background rates. These results show that testing segments against the *snr_significance* statistic is an effective way to remove signal-selection blunders, and that applying this test at our default level of 0.02 is unlikely to result in an unacceptable rate of rejection of segments that otherwise might be useful.

The overall accuracy of the ATL06 segment-fitting algorithm and the *snr_significance* < 0.02 criterion is illustrated in Fig. 10; here, three types of error estimates for all accepted segments are shown as a function of *BGR* and the number of signal photons per pulse, for smooth (Fig. 10a–c) and rough (Fig. 10d–f) surfaces. When the signal is strong (> 1 photons/pulse), particularly over smooth surfaces, the RMS values for surface height estimates are typically < 1 cm (left panels; Figs. 10a and d); however, for weaker signal levels the RMS value increases (i.e., the distribution of height estimate widens) indicating that blunders make up a significant fraction of the accepted segments. The robust estimator of the standard deviation, σ_R, (calculated as the interquartile range of the distribution, scaled to match that of a Gaussian distribution) has a similar pattern (center panels; Fig. 10b and e). This measure of spread increases more gradually with decreasing signal levels, indicating that a main driver of the spread as estimated by the RMS is the minority of outlying measurements within the distribution of the errors. The relationship between the estimated segment errors (σ_{est}, the propagated error in the least-squares segment fit) and the spread of the errors σ_R, is shown by mapping the robust spread of the ratio σ_R/σ_{est} over the 1600 segments that were simulated for each pixel in the map (right panels; Fig. 10c and f). For this statistic, values near unity indicate that the error estimates predict the data errors well; this is true in the in the moderate-to-strong signal domain (0.5–3 photons/pulse), but in the low-signal region, the prevalence of outlying values tends to produce data errors that are significantly larger than the estimated errors. Combined, these results show that the estimated errors generally give a reliable approximation of the expected spread in the surface height values in areas where the surface roughness is low and when signal strengths are adequate; however, under more marginal conditions, outlying values in the error distribution may produce substantial variation in surface heights.

4.2. ATL06 algorithm performance using simulated data over real topography

To demonstrate the performance of the ATL06 algorithm under realistic ice sheet conditions, we generated a set of test data for a region on the west coast of Greenland (Fig. 11). We generated synthetic ATL03 data along the ICESat-2 reference tracks, using the 8 m ArcticDEM (Porter et al., 2018). While in smoother regions of the ice sheet this digital elevation model (DEM) should be a reasonable approximation of the true ice-sheet surface, it likely underestimates the roughness of the most heavily crevassed outlet-glacier surfaces. Nevertheless, this simulation experiment allows us to assess how well the ATL03 surface-detection and ATL06 selection-refinement algorithms perform over complex surfaces, and how well the combined transmit-pulse -shape and first-photon bias corrections perform over smooth surfaces.

To demonstrate the effects of surface roughness on height recovery, we simulated ATL03 data along a track that encounters rough topography near Jakobshavn Glacier (Fig. 11b) and one for a smoother region inland (Fig. 11c). We then ran the ATL06 algorithm for both of these cases. The ATL03 photon heights and ATL06 segments are both

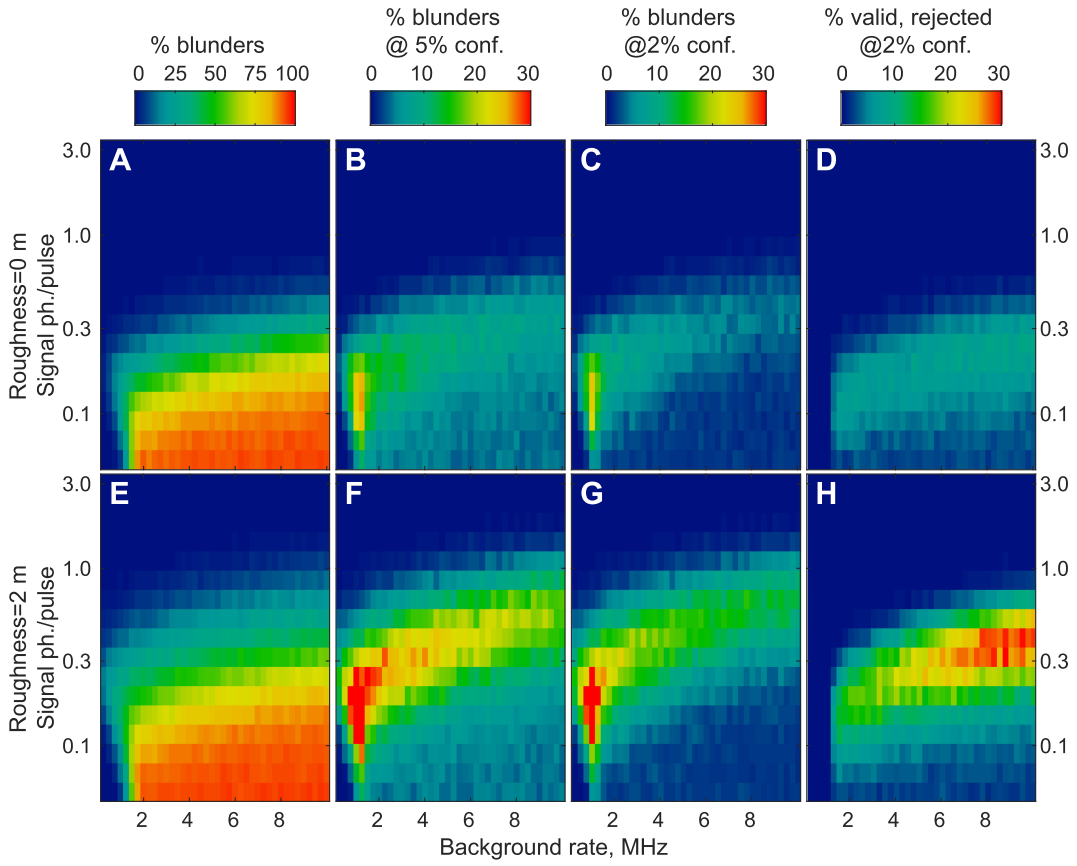

Fig. 9. Accuracy of the SNR-significance test. Test cases demonstrating the success of the signal-to-noise-significance test for identifying blunders in photon selection. The top row (a–d) shows the test performance for a smooth, flat surface (roughness = 0 m, slope = 0). The bottom row (e–h) shows the test performance for a rough (2 m RMS) or strongly sloping (26°) surface as a function of the background photon rate and the number of detected photons per transmit pulse. Panels a and e show the total fraction of blunders for all input segments if no test is applied. Panels b and f show the fraction of blunders that remain if the test is applied with a 5% confidence threshold. Panels c and g show the fraction of blunders that remain if the test is applied with a 2% threshold. Panels d and h show the fraction of valid segments that are (erroneously) rejected if the test is applied with a 2% threshold.

shown, as is the distribution of the height errors around the "true" surface height sampled at the segment centers. Neither set of errors shows a significant bias, and the smooth-region errors are smaller than those in the rough area. The high-roughness error histogram reflects the complex structure of the surface in this area: the robust spread, σ_R (Section 3.3.3) is four times smaller than the RMS error. We suggest that the difference between these two metrics is due to a minority of segments that cross the edges of crevasses or seracs that have large height errors: the relatively coarse (40 m) resolution (compared to the length scale of roughness) cannot match the true structure of the surface, while the majority of relatively flat segments are more consistent with the DEM. Users of real ATL06 data will have the option of using the misfit and surface-slope parameters on ATL06 to select segments from the smoother regions, which will have smaller errors, at the expense of lower spatial coverage.

We present the bias statistics for the smooth region in Fig. 11c (Fig. 12). The combination of the first-photon and transmit-pulse-shape errors produce a bias in ATL06's *h_li* parameter with a mean of 0.1 m and an RMS of 0.14 m. Applying the corrections for outliers, transmit-laser-pulse shape, and first-photon bias reduces the mean bias to the sub-centimeter level, and reduces the RMS error to 0.09 m (Fig. 12a). In reality, the accuracy of surface heights will be limited by the geolocation accuracy of the photons, resulting in additional errors on the order of the ATL06 algorithm error for our example data set. Applying a realistic geolocation error (6.5 m RMS) to the data results in a small increase in the recovered vertical error (Fig. 12b; Markus et al., 2017); in a region with a larger mean surface slope, geolocation error would likely be the dominant source of error.

For the smooth region, we investigated the effect of clouds on the precision of the recovered errors (Fig. 13). We repeated the simulation with two-way cloud optical depths of 3, 2, 1, and 0, which correspond to transmittances of 5%, 15%, 36%, and 100% respectively for cases I to IV. In these simulations, both clouds and detector deadtime reduced the number of detected photons, resulting in average per-segment photon counts that ranged from ~9 detected photons per pulse (for case IV with no clouds) to ~0.6 photons per pulse (for case I with the thickest clouds). For the two cloudiest cases (I and II), errors increased significantly with optical thickness. Errors were relatively consistent between the moderately cloudy (III) and clear (IV) case. This is likely because the first-photon bias correction introduces a small amount of random error in the recovered surface heights, and this effect is larger for the strongest returns. Over smooth surfaces, the ATL06 algorithm can retrieve accurate surface heights through moderately thick clouds, with errors increasing significantly only when clouds attenuate more than about 85% of the input photons (cloud optical thickness > 2).

4.3. Preliminary ATL06 algorithm performance with real ATLAS data

By testing ATL06 algorithms with synthetic ATLAS data, we have been able to exactly quantify the performance against a known (true) surface height at each step. During the commissioning phase of ICESat-2, we have begun to test the ATL06 algorithm with real ATL03 data from ATLAS, to investigate how it performs over different surfaces. As an example, we show granule ATL03_20181109172947_06430103_001_01.h5

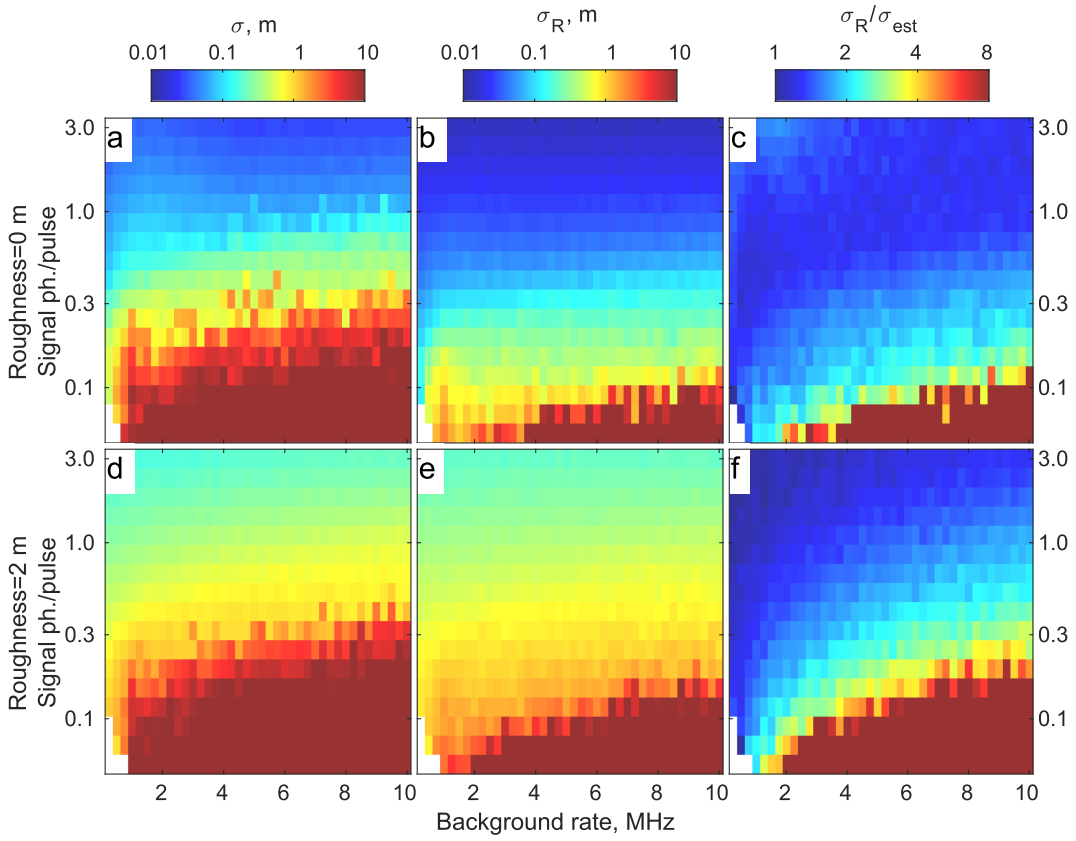

Fig. 10. Accuracy of selected segments. The accuracy of heights from segments selected using the SNR-significance test, and their associated error estimates. Panels a–c were modeled for a smooth (0 m RMS), flat surface, and panels d-f were modeled for a rough (2 m RMS) or strongly sloping (26°) surface. Error estimates are shown in three ways: (i) left panels (a and d) RMS of recovered surface heights; (ii) center panels (b and e) robust spread (σ_R) of recovered surface heights; and (iii) right panels (c and f) robust spread of the recovered surface heights divided by the estimated errors in the surface-height estimates. Each statistic is shown as a function of background-photon rate and signal strength (detected photons per transmit pulse).

along a ground track in West Greenland, 9 November 2018, which spans sea ice, exposed rock, and ice (Fig. 14). Preliminary analyses show that the algorithm correctly selects the photons that correspond to the surface, and produces height estimates even over rough and complex surfaces. We note that the quality of the ATL06 segment fits scales with the number of photons, and the signal strength here suggests that the granule was collected under clear (cloud-free) conditions. This is the same track used in the companion ATL03 paper (Neumann et al., this issue; their Fig. 9). This

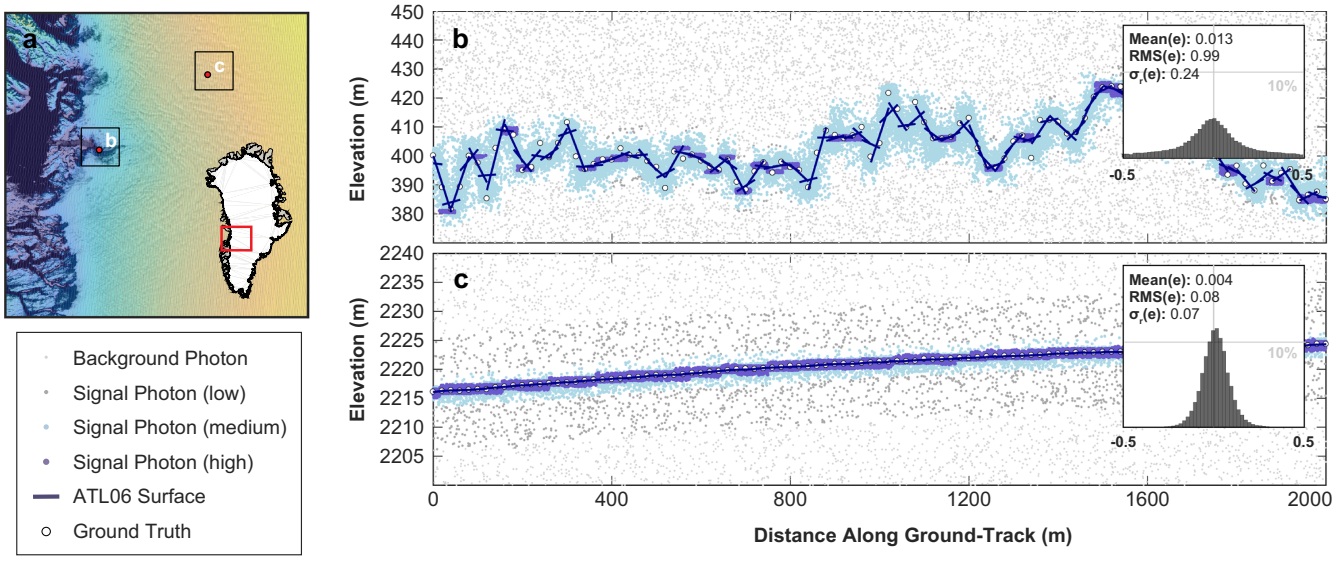

Fig. 11. ATL06 performance over West Greenland. Performance of ATL06 surface height retrieval algorithm for synthetic examples in the Jakobshavn Glacier catchment. a) Location of synthetic ICESat-2 coverage (gray lines) spanning a range of surface types, including: (b) rough surfaces in streaming flow; and (c) smooth surfaces in the ice sheet interior. Inset histograms show the errors in retrieved ATL06 surface height (computed minus true height), with their mean, root mean squared deviation (RMS), and robust spread (σ_r), in meters.

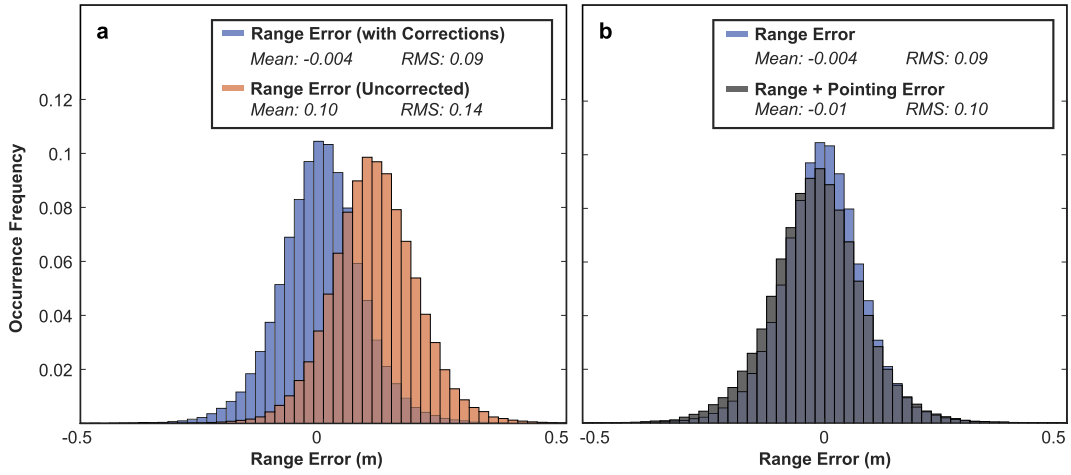

Fig. 12. Range errors in ATL06 retrieval. Error distributions for computed ATL06 surface heights (computed minus true height) over the smooth surface of Fig. 11c. a) Components of the error related to range only. Applying the transmit pulse shape and first-photon bias corrections reduced the mean error in the computed surface height by ~10 cm. b) The total accuracy of surface-height estimates, including the error due to mislocation of the spot on the ground.

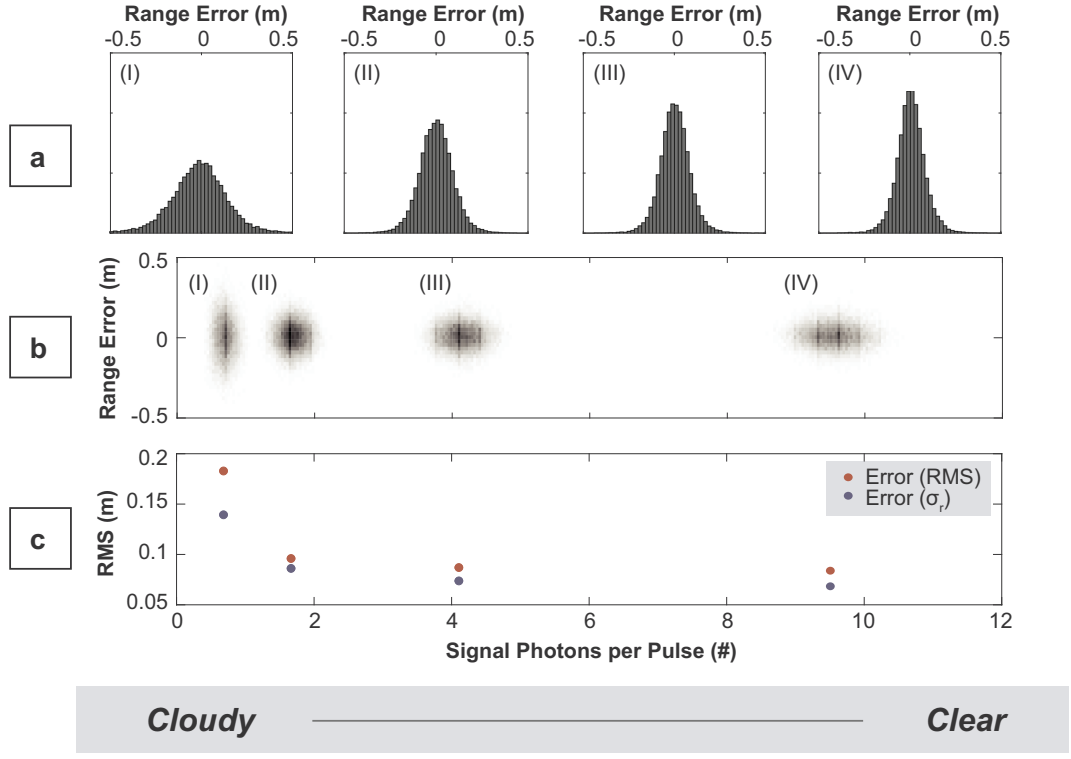

Fig. 13. Errors in ATL06 retrieval due to clouds. Distribution of surface fit errors for a variety of cloud conditions for the smooth area in Fig. 11c. a) Histograms (I-IV) correspond to regional simulations with a range of cloud conditions where I is the cloudiest and IV is the least cloudy; (b) relation between the signal strength (recovered signal photons per pulse) and the distribution of per-segment height errors; (c) relationship between the spread of errors (measured with the RMS and the robust spread and the signal strength.

gives us confidence that ICESat-2 will be able to recover heights over rough topography, even in the ice sheet margins.

5. Summary

We have presented the algorithms that will be used to generate surface height estimates over land ice for the ICESat-2 mission (ATL06) from lower level geolocated photon data (ATL03). We described the rationale and application of the surface fitting algorithm and corresponding corrections for transmit-pulse shape and first-photon bias. We also demonstrated how the ATL06 algorithms perform under various ice-sheet surface and cloud conditions using synthetic ATL03 test data. We make the following observations about the ATL06 algorithm performance:

1. The correction for the transmit-pulse-shape bias reduces the expected bias from > 10 mm to < 1 mm over a wide range of received-pulse shapes and can produce accurate results for a strongly skewed transmit-pulse shape.
2. The correction for the first-photon bias reduces the expected bias from > 40 mm to < 1 mm for the worst expected surface roughness and surface reflectance conditions.

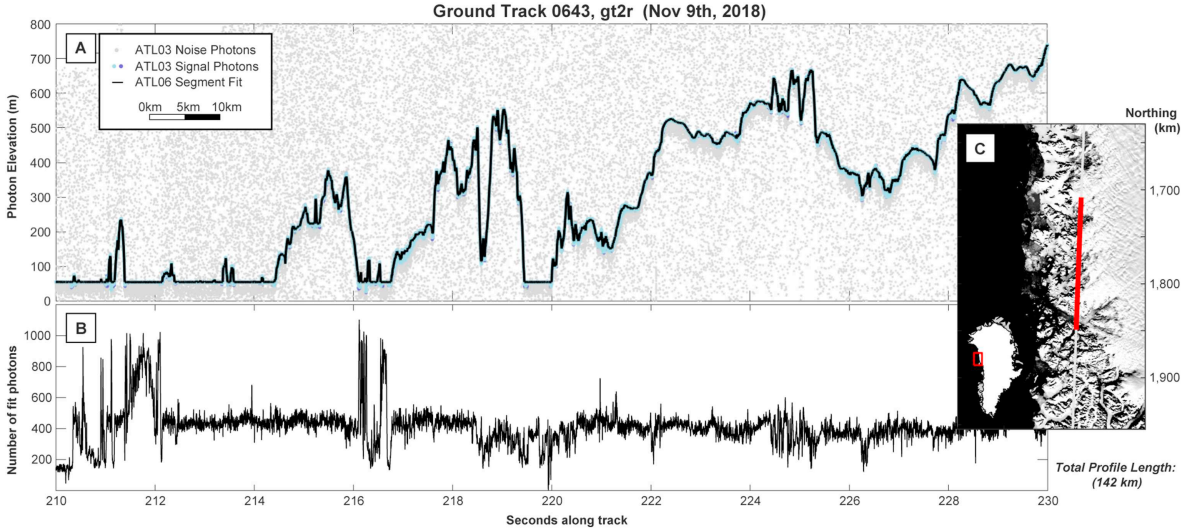

Fig. 14. ATL06 algorithm applied to real ATLAS data. a) ATL06 segment fits overlying ATL03 geolocated photons collected by ATLAS over West Greenland, 9 November 2018, spanning sea ice, exposed rock, and ice (see inset map c). b) number of identified signal photons in the fit. Note that the quality of the ATL06 segment fits scales with the number of photons. This is the same track used in the companion ATL03 paper (Neumann et al., this issue; their Fig. 9).

3. Using simulated ATL03 data derived from realistic surface topography, ATL06 algorithm performance is largely consistent with that inferred from simpler test data.
4. For smooth surfaces:
   a. the ATL06 algorithm can accurately estimate surface height even under moderately (optical thickness of < 2) cloudy conditions.
   b. evaluation of segment-fit quality allows rejection of essentially all surface-detection blunders as long as the mean signal strength exceeds 0.5 photons per pulse. For rough (> 2 m RMS) surfaces, > 80% of potential blunders are rejected.
   c. segments accepted by the quality filter have high (< 0.1 m) accuracy for signal strengths > 0.3 photons/pulse. For rough (> 2 m RMS) surfaces, the accuracy is lower, but generally < 1 m.

Based on these observations, we expect the ATL06 algorithm to perform at the level required to meet the ICESat-2 Science Objectives for land ice. Specifically, we anticipate ATL06 will meet the ice-sheet-wide requirement, that: *ICESat-2 shall produce an ice surface elevation product that enables determination of ice-sheet elevation change rates to an accuracy of better than or equal to 0.4 cm/yr on an annual basis* (Markus et al., 2017). This requirement can only be met if there is no instrumental change that produces a bias trend larger than 0.4 cm/yr. The corrections for the biases we have addressed here reduce their maximum values to millimeter scales, well below 0.4 cm/yr threshold, unless there is a drastic change in the system impulse response or the detector deadtime. Ultimately, under clear-sky conditions, the uncorrelated error magnitude is expected to be on the order of 0.1 m. Thus, we expect ICESat-2 will meet its regional mission requirement: *produce an ice surface elevation product that enables determination of annual surface elevation change rates on outlet glaciers to an accuracy of better than or equal to 0.25 m/yr over areas of 100 km^2 for year-to-year averages* (Markus et al., 2017).

Acknowledgements

This work was funded by NASA Cryospheric Sciences Program in support of the ICESat-2 mission under the following awards: NNX15AE15G; NNX15AC80G; NNX16AM01G. We are grateful to Denis Felikson and an anonymous reviewer for their extremely thorough, thoughtful and constructive reviews of this paper.

Appendix A. Description of robust spread estimator used in ATL06 algorithm

Our strategy to shrink the surface window to improve the ratio of signal photons to noise photons relies on accurate estimates of the spread of the signal photons. Commonly used measures of signal spread generally overestimate the spread of signal photons under low signal-to-noise conditions, because the range of heights of the background photons will generally be much larger than that of the signal photons. To obtain valid estimates of spread under these conditions, we estimate the standard deviation of a collection of points based on the difference between its percentiles, taking into account the contribution of the background rate (*BGR*) to the total distribution of photons. The uncorrected statistic of this type is the robust spread (σ_r), so we refer to the corrected statistic as noise-corrected robust spread, $\sigma_{r,corr}$.

In this calculation, we begin with a collection of N_{tot} photons, with heights $\boldsymbol{z}$, distributed between z_{min} *and* z_{max}. If the rate of background photons is estimated to be R_{BG}, then we estimate that the collection is made up of noise photons and $N_{sig} = N_{tot}$ - N_{BG} signal photons, where $N_{BG} = 2(z_{max} - z_{min})R_{BG}/c$. In general, we expect to obtain less accurate estimates of the spread of the signal photons for large windows than for small because the number of background photons is higher for large windows and it is less obvious which are the signal photons.

Within the collection of photons whose spread we are estimating, we make an estimate of heights of the photons that represent the 25th and 75th percentiles of the signal photon heights: We sort the photons and assign each an index value, i, equal to its position in the list minus 0.5. Then,

z_{25} = the height of the highest photon for which $i < (0.25\ N_{sig} + (z - z_{min})2R_{\_BG}/c)$

z_{75} = the height of the lowest photon for which $i > (0.75\ N_{\mathrm{sig}} + (z - z_{min})2R_{BG}/c)$

$\sigma_{r,corr}$ is then found:

$\sigma_{r,corr.} = (z_{75}\_z_{25})/scale\_factor$

where *scale_factor* is chosen to be equal to the difference between the 75th and 25th percentiles of a normalized Gaussian distribution, or 0.6745. This choice means that if **z** has a Gaussian distribution, $\sigma_{r,corr}$ will equal its standard deviation. The uncorrected robust spread is calculated in the same way, except that R_{BG} is set to zero.

Because R_{BG} may not accurately reflect the background rate, and because even if R_{BG} were accurate, there is statistical noise in the distribution of background-photon heights, it sometimes happens that $z_{25} > z_{75}$. In these cases, we report the mean photon-to-photon distance, $\sigma_{r,corr} = (z_{max}\text{-}z_{min})/N_{tot.}$

We demonstrate the effectiveness of this measure of spread using a series of synthetic datasets that simulate a range of surface conditions and background-photon rates for 40-m segments measured with a weak beam. For each simulated segment, we generate N_{sig} signal photons, whose heights have a Gaussian distribution of standard deviation σ_{signal}, and N_{noise} background photons. The background photons are uniformly distributed over a 40-m vertical window, and the number of each set of photons is selected based on Poisson random variable with expected values N_{signal} and N_{noise} such that: $N_{signal} + N_{noise} = 125$.

The total number of photons is set based on the expected number of photons in a 40-m horizontal by 40-m vertical window for a background rate of 8 MHz. This represents a summer day at high latitude over a bright, ice or cloud surface. We generate data for 10,000 such segments for value of N_{signal} between zero and 125, and for σ_{signal} values between 0.1 and 1.25 m. For each segment, we calculate the standard deviation, the uncorrected robust spread, and the corrected robust spread of the photon heights.

Fig. A1 shows the statistics of the recovered spread estimates. Panel A1a shows the median value of the standard deviations of the 10,000 realizations of the photon distribution for each value of N_{signal} and σ_{signal}. Because the standard deviation is strongly sensitive to the distribution of outlying photons, its values are largely determined by the proportion of signal photons, achieving accurate results only for the zero-noise case (the right-hand side of the plot). Panel A1b shows the median values of the uncorrected spread; these strongly overestimate the spread of the signal photons for $N_{signal} < 70$. Panel A1c shows the median value of $\sigma_{r,corr.}$, with a much smaller color range than that used in panels A1a and A1b. The corrected robust spread values match the input values well for N_{signal} values greater than around 30. Panel A1d shows the spread (the σ_r) of the $\sigma_{r,corr}$ values. These spreads are small for signal levels greater than around 30 photons, but diverge for weaker signals.

These results show that even under adverse conditions, with large (40 m) initial windows and high background levels, we can expect the background-corrected spread estimate $\sigma_{r,corr.}$to give a reliable estimate of the spread of the signal photons for windows that include at least 25–30 signal photons, corresponding to an initial signal-to-noise ratio of around 0.25. By contrast, the two other measures of spread considered here, the standard deviation and the uncorrected spread, σ_r, both overestimate the spread of the surface photons. For this reason, we use $\sigma_{r,corr}$, as part of an iterative ground-finding scheme of the type we describe in Section 3.1 of this paper rather than either of these two other metrics.

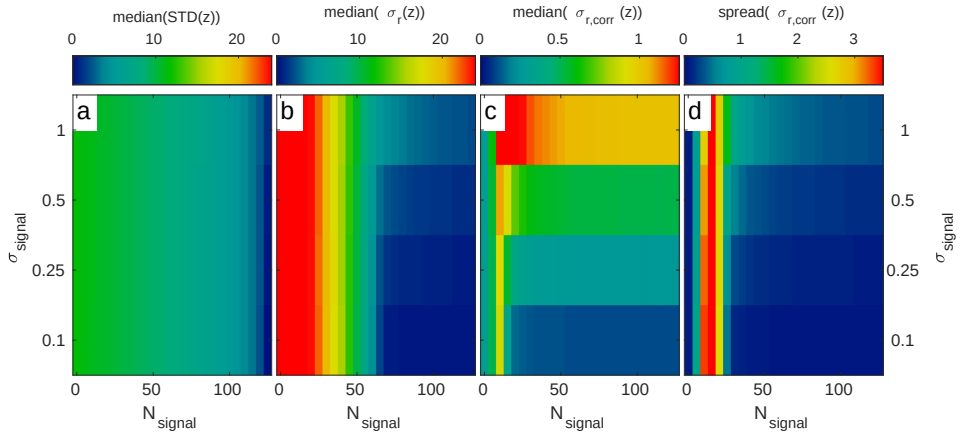

Fig. A1. Performance of different measures of the spread of a distribution in the presence of noise. Each measure is plotted a function of the input and values. Panel a shows the median of the (10,000) standard-deviation values for each value of σ_{signal} and N_{signal}. Panel b shows the median of the uncorrected σ_r values, and Panel c shows the median of $\sigma_{r,corr}$ values. Note that the color scale for a and b is much wider than that of c, and that the values on the right-hand side of c are (to within statistical noise) correct. Panel d shows the spread (calculated with σ_r) of the $\sigma_{r,corr}$ values.

Appendix B. Generation of test ATL03 data

We simulated ATLAS ATL03 data based on a transmit-pulse shape $TX(t)$, a value for the expected number of photons returned per pulse, spot width, the number of pixels in the detector, n_{ch}, and a DEM giving the shape of the surface measured. For the location of each pulse on the DEM, x_{pulse}, we generated a random number sampled from a Poisson distribution with a mean of N_{signal} to give the number of photons in each pulse. We then generated a location, x_{photon}, for each of these photons, sampled from a two-dimensional Gaussian distribution centered on x_{pulse} with width σ_{beam}. We sampled the DEM height at each of these locations to give a height value, z_p, for each of these photons. Modeling this interaction at the individual photon scale inherently simulates the pulse broadening, as described in Eq. (1). We converted these heights to two-way travel time values and added a random value sampled from the transmit-pulse shape, $TX(t)$, to each. This gave the time distribution of photons arriving at the receiver for each pulse. In the simulation output, the location for the photons in each pulse is reported as x_{pulse} (rather than x_{photon}), because with real ICESat-2 data, the only information available about the location of the photons is the calculated centroid location of the laser spot. No information remains about the location of individual photons relative to that spot.

For those simulations that included the effects of detector dead time, we then assigned each photon to a random detector pixel. For each pulse, and for each pixel, we checked each photon to see whether a previous photon had arrived in that pixel less than one analog deadtime earlier (< 1 ns). Next, we checked whether a previous photon that had passed through the analog detector had inactivated the digitizer-input stage of the detector (with a deadtime of 3.2 ns). In either case, the photon within the deadtime was removed from the analysis.

We assumed that an ICESat-2 pulse, traveling through a clear polar atmosphere and reflecting off a clean snow surface, would result in 12 photons entering the detector per pulse for a strong beam and 3 photons per pulse for a weak beam. Because of the first-photon bias effect, some of these photons would go undetected, resulting in as much as 25% fewer photon detections than signal photons entering the detector over very flat, very bright surfaces. Clouds and aerosols between ATLAS and the surface, or measurements over a surface with a reflectance < 98% at 532 nm wavelength can result in a smaller input photon count, which, for the purposes of this paper, are modeled simply by reducing the number of incident photons per pulse from 12 (or 3) to a smaller number.

References

Abdalati, W., Zwally, H.J., Bindschadler, R., Csatho, B., Farrell, S.L., Fricker, H.A., Harding, D., Kwok, R., Lefsky, M., Markus, T., Marshak, A., Neumann, T., Palm, S., Schutz, B., Smith, B., Spinhirne, J., Webb, C., 2010. The ICESat-2 laser altimetry mission. Proc. IEEE 98 (5), 735–751. https://doi.org/10.1109/JPROC.2009.2034765.

Kaasalainen, S., Kaasalainen, M., Mielonen, T., Suomalainen, J., Peltoniemi, J.I., Näränen, J., 2006. Optical properties of snow in backscatter. J. Glaciol. 52 (179), 574–584.

Kwok, R., Cunningham, G., Hancock, D., Ivanoff, A., Wimert, J., 2019. ICESat-2 Algorithm Theoretical Basis Document (ATBD) for Sea Ice Products (ATL07/10). NASA Jet Propulsion Laboratory, Pasadena, CA available at. https://icesat-2.gsfc.nasa.gov/science/data-products.

Markus, T., Neumann, T., Martino, A., Abdalati, W., Brunt, K., Csatho, B., Farrell, S., Fricker, H., Gardner, A., Harding, D., Jasinski, M., Kwok, R., Magruder, L., Lubin, D., Luthcke, S., Morison, J., Nelson, R., Neuenschwander, A., Palm, S., Popescu, S., Shum, C.K., Schutz, B.E., Smith, B., Yang, Y., Zwally, J., 2017. The ice, cloud, and land elevation satellite-2 (ICESat-2): science requirements, concept, and implementation. Remote Sens. Environ. 190, 260–273. https://doi.org/10.1016/j.rse.2016.12.029.

McGill, M., Markus, T., Scott, V.S., Neumann, T., 2013. The multiple altimeter beam experimental Lidar (MABEL): an airborne simulator for the ICESat-2 mission. J. Atmos. Ocean. Technol. 30 (2), 345–352. https://doi.org/10.1175/JTECH-D-12-00076.1.

Moholdt, G., Nuth, C., Hagen, J.O., Kohler, J., 2010. Recent elevation changes of Svalbard glaciers derived from ICESat laser altimetry. Remote Sens. Environ. 114 (11), 2756–2767. https://doi.org/10.1016/j.rse.2010.06.008.

Morison, J., Kwok, R., Smith, B., Jasinski, M., Hancock, D., Robbins, J., Dickinson, S., Urban, T., 2019. ICESat-2 Algorithm Theoretical Basis Document (ATBD) for Ocean Surface Height (ATL12). University of Washington, Seattle, WA, School of Oceanography available at. https://icesat-2.gsfc.nasa.gov/science/data-products.

Neuenschwander, A., Pitts, K., 2019. ICESat-2 Algorithm Theoretical Basis Document (ATBD) for Land and Vegetation Products (ATL08). Applied Research Laboratory, University of Texas, Austin, TX available at. https://icesat-2.gsfc.nasa.gov/science/data-products.

Neumann, T., Brenner, A., Hancock, D., Robbins, J., Saba, J., Harbeck, K., Gibbons, A., 2019. ICESat-2 Algorithm Theoretical Basis Document (ATBD) for Global Geolocated Photons (ATL03). NASA Goddard Space Flight Center, Greenbelt, MD available at. https://icesat-2.gsfc.nasa.gov/science/data-products.

Neumann, T.A., A. Martino, T. Markus, S. Bae, M.R. Bock, A.C. Brenner, K.M. Brunt, J. Cavanaugh, S.T. Fernandes, D.W. Hancock, K. Harbeck, J. Lee, N.T. Kurtz, P.J. Luers, S.B. Luthcke, L. Magruder, T.A. Pennington, L. Ramons-Izquierdo, T. Rebold, J. Skoog, T.C. Thomas (this issue), The ice, cloud, and land elevation satellite - 2 mission: a global geolocated photon product derived from the advanced topographic laser altimeter system. Remote Sens. Environ.

Popescu, S.C., Zhou, T., Nelson, R., Neuenschwander, A., Sheridan, R., Narine, L., Walsh, K.M., 2019. Photon counting LiDAR: an adaptive ground and canopy height retrieval algorithm for ICESat-2 data. Remote Sens. Environ. 208, 154–170. https://doi.org/10.1016/j.rse.2018.02.019.

Porter, C., Morin, P., Howat, I., Noh, M.-J., Bates, B., Peterman, K., Keesey, S., Schlenk, M., Gardiner, J., Tomko, K., Willis, M., Cloutier, M., Husby, E., Foga, S., Nakamura, H., Platson, M., Wethington Jr., M., Williamson, C., Bauer, G., Enos, J., Arnold, G., Kramer, W., Becker, P., Doshi, A., D'Souza, C., Cummens, P., Laurier, F., Bojesen, M., 2018. ArcticDEM. V1 Harvard Dataversehttps://doi.org/10.7910/DVN/OHHUKH.

Pritchard, H.D., Arthern, R.J., Vaughan, D.G., Edwards, L.A., 2009. Extensive dynamic thinning on the margins of the Greenland and Antarctic ice sheets. Nature 461 (7266), 971. https://doi.org/10.1038/nature08471.

Schutz, B.E., Zwally, H.J., Shuman, C.A., Hancock, D., DiMarzio, J.P., 2005. Overview of the ICESat mission. Geophys. Res. Lett. 32 (21). https://doi.org/10.1029/2005GL024009.

Smith, B.E., Fricker, H.A., Joughin, I.R., Tulaczyk, S., 2009. An inventory of active subglacial lakes in Antarctica detected by ICESat (2003–2008). J. Glaciol. 55 (192), 573–595. https://doi.org/10.3189/002214309789470879.

Smith, B.E., Gardner, A., Schneider, A., Flanner, M., 2018. Modeling biases in laser-altimetry measurements caused by scattering of green light in snow. Remote Sens. Environ. 215, 398–410. https://doi.org/10.1016/j.rse.2018.06.012.

Smith, B.E., Harbeck, K., Roberts, L., Neumann, T., Brunt, K., Fricker, H.A., Gardner, A., Seigfried, M.R., Adusumilli, S., Csatho, B.M., Holschuh, N., Nilsson, J., Paolo, F.S., 2019. ICESat-2 Algorithm Theoretical Basis Document (ATBD) for Land Ice Along-Track Height (ATL06). Applied Physics Laboratory, University of Washington, Seattle, WA available at. https://icesat-2.gsfc.nasa.gov/science/data-products.

Sun, X., Abshire, J.B., Borsa, A.A., Fricker, H.A., Yi, D., DiMarzio, J.P., Paolo, F.S., Brunt, K.M., Harding, D.J., Neumann, G.A., 2017. ICESat/GLAS altimetry measurements: received signal dynamic range and saturation correction. IEEE Trans. Geosci. Remote Sens. 55 (10), 5440–5454. https://doi.org/10.1109/TGRS.2017.2702126.

Yang, Y., Marshak, A., Varnai, T., Wiscombe, W., Yang, P., 2010. Uncertainties in ice-sheet altimetry from a spaceborne 1064-nm single-channel Lidar due to undetected thin clouds. IEEE Trans. Geosci. Remote Sens. 48 (1), 250–259. https://doi.org/10.1109/TGRS.2009.2028335.

Yang, Y., Marshak, A., Palm, S.P., Varnai, T., Wiscombe, W.J., 2011. Cloud impact on surface altimetry from a spaceborne 532-nm micropulse photon-counting Lidar: system modeling for cloudy and clear atmospheres. IEEE Trans. Geosci. Remote Sens. 49 (12), 4910–4919. https://doi.org/10.1109/TGRS.2011.2153860.

Yi, D., Zwally, H.J., Sun, X., 2005. ICESat measurement of Greenland ice sheet surface slope and roughness. Ann. Glaciol. 42, 83–89. https://doi.org/10.3189/172756405781812691.