Supporting Information for "Tracking eruption column thermal evolution and source unsteadiness in ground-based thermal imagery using spectral-clustering"

C. R. Rowell[1], A. M. Jellinek[1], Y.T. Gilchrist1

[1]Department of Earth, Ocean, and Atmospheric Sciences, University of British Columbia, Vancouver , British Columbia , Canada

Contents of this file

Text S1. Methods Supplement. Here we described additional details for the methods used in the main manuscript related to thermal image data processing. The code repository and further documentation are available at https://github.com/colinrr/locust.git.

1. Thermal Imagery Pre-processing

1.1. Data Format Conversion and Frame Registration

The thermal imagery, initially recorded in the propriety InfraTec *.irb* format, were first converted to an ASCII file format using one of two methods: via the Infratec IRB2ASCII executable, or the program IRT_Analyzer. The resulting text files are then converted into MATLAB data format. For a given sequence of thermal image frames, a single *.mat* file

is generated for each frame, as a well as a file containing a table of header information with rows corresponding to each frame, titled *frameHeads.mat*.

Following conversion to MATLAB file format, frame registration is performed as needed when the camera viewfield was not perfectly steady. For instance, wind periodically caused minor shaking of the camera during video recording, which must be corrected to avoid large spatial displacements of imaged objects between frames. Figure S1 shows an example of temperature differences between frames for Event 3 before and after registration. A brief period of strong wind gusts caused significant shaking of the camera, such that temperature differences become large between frames for a given pixel due to the displacement of imaged objects. Re-alignment of these images using the MATLAB image registration package prevents erroneous plume segmentation and velocity analysis during later processing steps.

1.2. Plume segmentation using plumeTracker

To obtain image masks identifying pixels as belonging to an eruption plume, we use the **plumeTracker** software of Bombrun, Jessop, Harris, and Barra (2018), modified to facilitate efficient data input and output in our workflow (the core image segmentation algorithm is unchanged). In most instances, clear skies provided excellent thermal contrast between plumes and the background of thermal images, and **plumeTracker** segmentation was very effective in these instances for producing binary image masks. In instances where clouds are present in the vent region (either atmospheric vapor or outgassing from the vent itself), the **plumeTracker** software struggled to differentiate between cloud and plume motions. In these cases and as necessary, we applied **plumeTracker** to a pre-processed

version of the thermal images, in which we enhanced contrast prior to segmentation and/or perform manual masking of image regions in which no plume was present.

1.3. Mapping Pixels to Physical Coordinates

To calculate the positions, dimensions, and velocities of image features, it is necessary to obtain estimates of pixel dimensions (in meters) at the target (vent or plume) distance, accounting for the viewing angle of the camera. Figure S2 shows the view geometry and relevant parameters. Here we generate a mapping function that converts pixel coordinates to distances relative to the camera reference frame. Using a modified form of the geometric corrections in Harris (2013, Chapter 9), we first assume that (a) the target plume lies in a 2D vertical plane, centered above the vent and horizontally perpendicular to the camera view azimuth (as in Figure S2a), and (b) that only vertical tilt of the thermal images is present, with no horizontal tilt (the camera was leveled during field observations). We further assume that lens distortion is negligible (a correction is applied internally in the camera), such that the thermal images represent an ideal camera perspective projection.

We first need the true azimuth, ϕ_0, and elevation angle relative to horizontal, θ_0, of the camera view centerline. Using the geodetic coordinates of the camera observation point, a visible reference point on the Sabancaya edifice, and the estimated vent location (see Figure S2a,c), we calculate the azimuth, elevation angle, and slant distance from the camera observation point to the reference point, $(\phi'_{ref}, \theta'_{ref}, D_{ref})$, and to the Sabancaya vent, $(\phi'_v, \theta'_v, D_v)$. Here the $'$ indicates that angles are *apparent* angles with respect to the camera reference frame (i.e. less than the true angle relative to camera centerline, see Figure S2b). The coordinates of the reference point and vent location are estimated from the ALOS 12 digital elevation model (ASF-DAAC, 2015) and Google Earth satellite imagery.

The angular distance in azimuth and elevation angle between the camera centerline and the reference point are, respectively:

$$d\theta_{ref} = \left[i_{ref} - (\tfrac{n_i}{2} + 0.5)\right] \beta_i \,, \tag{S.1}$$

$$d\phi_{ref} = \left[j_{ref} - (\tfrac{n_j}{2} + 0.5)\right] \beta_j \,, \tag{S.2}$$

where $d\theta_{ref}$ and $d\phi_{ref}$ are the angular distances within the camera field of view between the reference point and camera centerline; (i_{ref}, j_{ref}) are the vertical and horizontal pixel coordinates of the reference point in the image, respectively; (n_i, n_j) are the vertical and horizontal image resolution in pixels, respectively; and (β_i, β_j) are the vertical and horizontal angular fields of view for a single pixel. To retrieve the geodetic elevation angle and azimuth, we then apply an apparent angle correction to account for the camera tilt to obtain the camera elevation angle θ_0 and azimuth ϕ_0 (in geodetic reference frame):

$$\theta_0 = \theta_{ref} - d\theta_{ref} = \tan^{-1}\left(\frac{\tan(\theta'_{ref})}{\cos(d\phi_{ref})}\right) - d\theta_{ref} \,, \tag{S.3}$$

$$\phi_0 = \phi'_{ref} - d\phi'_{ref} = \phi'_{ref} - \tan^{-1}\left(\cos(\theta_0)\tan(d\phi_{ref})\right) \,. \tag{S.4}$$

Using a similar apparent angle correction, the horizontal angular distance from camera centerline $d\phi_v$ and elevation angle θ_v to the vent location are given by:

$$d\phi_v = \tan^{-1}\left(\frac{\tan(\phi'_v - \phi_0)}{\cos(\theta_0)}\right) \,, \tag{S.5}$$

$$\theta_v = \tan^{-1}\left(\frac{\tan(\theta'_v)}{\cos(d\phi_v)}\right) \,. \tag{S.6}$$

We then calculate horizontal distance from camera to vent D_x and slant distance from camera to the centerline pixel D_{los} (in the projected vertical plane) according to:

$$D_x = D_v \cos(d\phi_v)\cos(\theta_v) \,, \tag{S.7}$$

$$D_{los} = \frac{D_x}{\cos(\theta_0)} \,. \tag{S.8}$$

Once the above calculations are performed for a given camera location, image coordinates i, j in units of pixels may be readily converted to coordinates in meters in the projected image plane (x, z) by:

$$z = D_x \tan\left[\theta_0 - (i - 0.5 - \tfrac{n_i}{2})\right] \beta_i \,, \tag{S.9}$$

$$x = (D_{los} + (D_x \tan(\theta_0) - z)\sin(\theta_0)) \tan\left[(j - 0.5 - \tfrac{n_j}{2})\right] \beta_j \,. \tag{S.10}$$

We numerically propagate uncertainty in the positions of the vent and reference feature to estimate uncertainty in pixel dimensions and absolute position. From these calculations for Events 1 and 2 (Observation Site 1), the image center pixel dimensions are $(dx, dz) = (3.3 \pm 0.06, 3.5 \pm 0.06)$ m, the camera vertical field of view at the vent distance is 2720 m, and estimated pixel positional errors for features near the center image are $\pm(1, 30)$ m in x and z coordinates, respectively, and increase approximately linearly towards the image edges to a maximum of $\pm(30, 75)$ m. For Event 3 (Observation Site 2), pixel dimensions are $(dx, dz) = (2.7 \pm 0.06, 2.9 \pm 0.06)$ m, the vertical field of view at vent distance is 2270 m, and estimated positional errors are similar to those for Events 1 and 2. The analyses presented below rely almost exclusively on pixels from features near image centers with relatively lower uncertainty values, and therefore absolute positional error are generally $\leq \pm(15, 60)$ m for all three events. Note that these error values are estimated for the geometrical projection alone, assuming objects lie in the projection plane (Figure S2), and result primarily from uncertainty in the absolute positions of the reference feature and vent.

1.4. Image re-gridding

Image stretching during the above image projection step results in pixel positions that are not regularly spaced in (x, z) (see Figure S2c,d). The resampling of thermal images and masks into regularly gridded spatiotemporal coordinates is done in two steps. First, individual frames are resampled from projected positions onto regularly sampled grid coordinates (x, z) (Figure S2e). Re-gridding is performed using a linear interpolation in the MATLAB 'scatteredInterpolant' function. Next, because the thermal camera frame rate is slightly variable, we used the MATLAB 'resmaple' function with piecewise cubic Hermite interpolation. This function outputs a MATLAB structure which contains a regularly sampled data cube $\Delta T(x, z, t)$.

2. Two-dimensional (2D) Velocity Field Retrieval Using Optical Flow Analysis

Here we describe the details of obtaining 2D velocity fields using the Optical Flow toolbox of Sun, Roth, and Black (2014). Their code was used unaltered, and we do not include it in the our available code, but driver scripts and additional functions built for the filtering process described below are included. There are two key controls that govern the extent to which we resolve the velocity field with Optical Flow. First, the sampling interval must fully resolve the turnover time of the coherent structure we are tracking (generally $\gtrsim 5$ s). Second, the rise speed of any feature must be similar or greater than the "grid speed" dx/dt set by the image spatial resolution dx and sampling rate dt, such that the classical Courant number is of order 1.

$$C = \overline{||u, v||}\frac{dt}{dx} \gtrsim 1, \tag{S.11}$$

where $\overline{||u,v||}$ is the mean of the velocity field magnitude in the plume. For pixel sizes of about 3 m and frame intervals of 0.1 to 0.2 s, velocities of 15 to 30 m/s are comparable to a displacement of about 1 pixel per frame. Motions significantly faster than this are not fully resolved because displacements are very large between frames, but tracking motions much slower than this introduces unphysical noise to the velocity field estimation, as shown in Figure S3. Therefore the 10 Hz sampling rate of the camera is higher than is required to resolve coherent structures rising at typical speeds of about $\sim$ 5 to 25 m/s or turning over with periods of $\gtrsim$ 10s, and we first downsample the video to every second frame before applying the Optical Flow analysis. The choice of every second frame balances resolving slower plume motions while also minimizing the number of frames for which velocity fields must be interpolated. To eliminate remaining high-frequency noise after obtaining optical flow fields, we then apply a low-pass filter to the velocity fields over the time dimension t, and finally interpolate velocity fields for frames not included in downsampling. The frequency cutoff for low-pass filtering is chosen based on the Courant condition such that it is about half the ratio of dominant plume flow velocities to the pixel grid spacing,

$$f \approx (0.5)\frac{\overline{||u,v||}}{dx}, \tag{S.12}$$

This choice of frequency cutoff (about 0.5 to 1 Hz) produces velocity fields that are spatially coherent and robustly capture plume motions with minimal noise over the timescales of plume rise or the overturn of large eddies ($\gg$ 1 s). Obtaining optical flow velocities in this manner produces velocities that are accurate compared with manual calculations to within a few percent.

3. Atmospheric Profile Removal Methods Supplement

Here we present additional details on the atmospheric profile removal steps, including height uncertainty estimation, late-time filtering, and additional checks on the choice of ΔT_{mode}.

3.1. Height uncertainty estimation

We estimate the height uncertainty of pixels in the projected images, so that pixels with large height uncertainty can be filtered out when fitting the atmospheric profile to the data. Figure S4 shows the two main sources of pixel height uncertainty, and the estimate of the two uncertainties for a sample frame of Event 3. For a given column feature, error in the height estimation arises if the feature is farther from or closer to the camera than the assumed projection plane above the vent. Because of the upward camera viewing angle, an object that is farther than the projection plane will likely have an underestimated height for features at the edge of the column, and an object that is closer will likely have an overestimated height. Therefore for a perfectly vertical, conical plume centered above the vent, the 3D shape of the plume (i.e. the plume radius) will generally cause the projected height to overestimate the true height since the outer visible portions of the plume will lie outside the projection plane and relatively closer to the camera. We refer to this height uncertainty as the radial error in z (Figure S4a). Additional (and greater) height uncertainty arises where the plume center axis is not centered above the vent (e.g. due to wind transport), such that the position of features may be ambiguously closer to or farther from the camera than implied by the assumed projection plane. We refer to this uncertainty as the axial uncertainty in z (Figure S4b).

For all pixels in all frames as a function of height, we estimate the combined height uncertainty using the geometrical relationships shown in Figure S4, combined with simplifying assumptions on column geometry. We use the shape of the plume masks derived from the plumeTracker algorithm (Bombrun et al., 2018, yellow line in Figure S4c) to estimate both radial and axial height uncertainties. First, we use the mask width as a proxy for column diameter. Second, we use distance in x between the mask centerline and the vent as a proxy for the magnitude of uncertainty in distance to the camera. This choice is based on the assumption that plume advection along the visible x coordinate is of similar magnitude to advection into or out of the projection plane. This is broadly true for our three Events based on field observations, in which columns were generally shifted north-northwest by prevailing winds at altitude. Figure S4d shows the resulting estimated maximum combined height uncertainty (generally occurring near the plume centerline) for pixels in the plume mask as a function of altitude z, where the solid blue line gives the estimated height radial error, and the shaded blue region gives the axial uncertainty resulting from the column axial position. In this example, the plume is approximately centered above the vent for $z \lesssim 6800$ m a.s.l., and the estimated height error increases linearly with both increasing plume radius and elevation angle from the camera, but remains less than about 50 m. Above 6800 m where the plume is bent away from the vent, the uncertainty in height quickly increases in both magnitude and sign, ranging from about -50 to +150 m.

3.2. Late Time Filtering

The transient Events 2 and 3 are conducive to atmospheric profile fitting by virtue of rapid decay in source heat flux. At large time after the Event onset, remnants of

the column become increasingly quiescent and approach thermal equilibrium with the background atmosphere. Figure S5 compares ΔT for two frames of Event 3 at $t = 60$ and $t = 245$ s. ΔT is both much lower and much more uniform throughout the column for the late-time frame, and we take advantage of this to apply the atmospheric filtering. Figure S5b shows a heat map of $T_b - T_a$ as a function of time for Event 3, where each column of data is the normalized probability density function (PDF) for pixels in a single frame. Initially, during the peak heat flux from the vent, the PDF of ΔT_i has a mode of about 60 K and wide standard deviation of about 30 K. At late times as column heat flux decays, however, the values converge towards an approximately constant mode of about -12.4 K with a standard deviation of only about 1 K (red box), indicating that the plume is near thermal equilibrium with the atmosphere at this relative temperature difference between image and satellite data. This constant value corresponds to peak of filtered PDFs as shown in Figure 4(g-i) of the main manuscript, from which we estimate $\Delta T_{mode} = -13.6$ K.

3.3. Apparent Minimum Temperatures and Atmospheric Profile Fit Residuals

To add confidence in our estimate of ΔT_{mode} based on filtered PDF half-maxima, here we discuss an additional filtering step. As shown in Figure 4, the action of column rise and turbulent entrainment and mixing causes column temperatures to approach those of atmospheric background approximately asymptotically. At any given height in the column, turbulent entrainment and thermal mixing act to combine two packages of fluid: the relatively hot plume mixture, and the cold and comparatively uniform temperature of the ambient air. As a result, temperature distributions in the data are frequently bi-modal, and the statistical distribution of the coldest temperatures can be used to

find pixels that represent the coldest background temperature being entrained into the column and towards which temperatures are approaching asymptotically. To this end, for all filtered pixel *rows* within the plume (see Figure 4e) we find the minimum value of ΔT. Since the initial masking process should eliminate the majority of partially transparent pixels, we expect that these apparent ΔT minima are statistically representative of the coldest atmosphere being drawn into the plume.

Figure S6(a-c) shows the resulting PDFs of row-wise apparent minima for Events 1-3, respectively, as purple dashed lines. The filtered PDFs of Figure 4(g-i) are shown in blue and orange for comparison. For the Event 3 example in Figure S6c, this procedure gives a population of about 20,000 apparent minima values (about 400 pixel rows within the range of good height estimates over about 50 frames). The apparent minima peaks approximately coincide with the half-maximum value of the initially filtered pixels (vertical dashed grey lines), except in the case of Event 1. Because Event 1 is a steady plume and lacks a well-equilibrated 'late-time' population of pixels, the apparent minima peak is more diffuse. We take the long tails in apparent minima PDFs at colder temperatures as representative of partially transparent pixels incorporating information from cold sky background. That apparent minima peaks for Events 2 and 3 coincide with peak half-maxima of the filtered pixels (orange lines) adds confidence that these values correspond approximately to the atmospheric background temperatures, motivating our choice of ΔT_{mode}. Since this is not the case for Event 1, however, for consistency we skip this additional apparent minima step in our final results, and retain only the half-maxima definition of ΔT_{mode} as described in the main manuscript.

Panels (d), (e), and (f) in Figure S6 show the residual errors in the modal value of $T_b - T_a$ as a function of height for Events 1, 2, and 3, respectively. The blue and orange curves correspond to the mode for all plume pixels and filtered plume pixels, respectively, as in panel (e). As expected for all events, the fit for all pixels is generally poor near the vent where temperatures are high and above about 7500 to 8000 m a.s.l., where plumes becomes increasingly transparent and pixel height estimates are increasingly uncertain. In between these heights for Events 2 and 3, fit of filtered pixels with the atmospheric profile is generally very good. This is in part facilitated by the transient nature of these events, which allows a better fit for thermally well-mixed plumes at late time. For Event 1, which is both steady in time (and thus lacks "late-time" data to use) and increasingly bent by winds with altitude above the vent (see Figure 3a), there is a more narrow height range of about 6700 to 7300 m a.s.l. where the profile fit is good. Below this height the steady plume is not thermally well mixed, and above this height the plume is likely both increasingly transparent and strongly transported out of the projection plane, resulting in large height uncertainty.

3.4. Effects on Power Law Fits of ΔT_{mode} Approximation

The linear approximation using ΔT_{mode} in Equation 6 (main manuscript) is valid so long as the maximum values of ΔT are not more than 100 to 200 K. For example, assuming that the satellite atmospheric profile gives the true temperature (about 267 to 275 K between 6500 and 7500 m a.s.l.), then for the largest estimated magnitude of $\Delta T_{mode} = -12.4$ K (Event 3, Figure S6c), Equation 5 implies a combined emission and transmission loss $(\varepsilon\xi) \approx 0.83$. In this case, the maximum error introduced to our ΔT approximation for the hottest (unsaturated) pixels is about 7 K, and typically less than about 2 K. This

approach is further aided in our case by the combination of high altitude (observation sites about 5200 m a.s.l.), a dry environment, and generally clear-sky background, which are well-suited to minimal (though non-zero) transmission loss of thermal emissions at our viewing distances (Harris, 2013).

Figure S7 compares power law fit results using the scalar temperature correction ΔT_{mode} (Equation 6) against the results of applying a multiplicative correction to brightness temperature data (Equation 5). Panel (a) shows absolute temperatures after applying four different multiplicative corrections ($\epsilon\xi$) ranging from 0.6 to 1.0 to brightness temperature data. A value of 0.83 gives the best fit with the satellite-derived atmospheric profiles. In panel (b), we calculate the excess temperature 95th percentile. For cases besides ($\epsilon\xi$) = 0.83, multiplicative corrections must be adjusted by an additive constant to align the decay curves for comparison with the scalar-adjusted ΔT employed in the main manuscript (black line in panel (b)). In panels (c) and (d), the power law fit results show that resulting B estimates are insensitive to the choice of correction, though RMS fit error increases for low correction values of ($\epsilon\xi$) = 0.6.

References

ASF-DAAC. (2015). Dataset: ALOS PALSAR_Radiometric_Terrain_Corrected_high_res; Includes Material © JAXA/METI 2007.
doi: 10.5067/Z97HFCNKR6VA

Bombrun, M., Jessop, D., Harris, A., & Barra, V. (2018, February). An algorithm for the detection and characterisation of volcanic plumes using thermal camera imagery. *Journal of Volcanology and Geothermal Research*, *352*, 26–37. doi: 10.1016/j.jvolgeores.2018.01.006

Harris, A. (2013). *Thermal Remote Sensing of Active Volcanoes: A User's Manual*. Cambridge University Press.

Sun, D., Roth, S., & Black, M. J. (2014, January). A Quantitative Analysis of Current Practices in Optical Flow Estimation and the Principles Behind Them. *Int J Comput Vis*, *106*(2), 115–137. doi: 10.1007/s11263-013-0644-x

Table S1. Table of variables for supporting information.

Variable	Description	Units
i_{ref}	Vertical image coordinate of reference feature	pixels
j_{ref}	Horizontal image coordinate of reference feature	pixels
n_i	Height of camera field of view	pixels
n_j	Width of camera field of view	pixels
β_i	Vertical angular field of view for a single pixel	radians
β_j	Horizontal angular field of view for a single pixel	radians
ϵ	Emissivity of target plume	-
θ_0	Camera view centerline elevation angle to horizontal	radians
$d\theta_{ref}$	Angular elevation to reference point along camera centerline	radians
ξ	Transmissivity of the atmospheric path between camera and plume	-
ϕ_0	Camera view centerline azimuth	rad
$d\phi_{ref}$	Angular distance: camera centerline to reference azimuth	rad
Notation:		
$'$	E.g. θ': apparent angle relative to a camera-centered frame of reference	

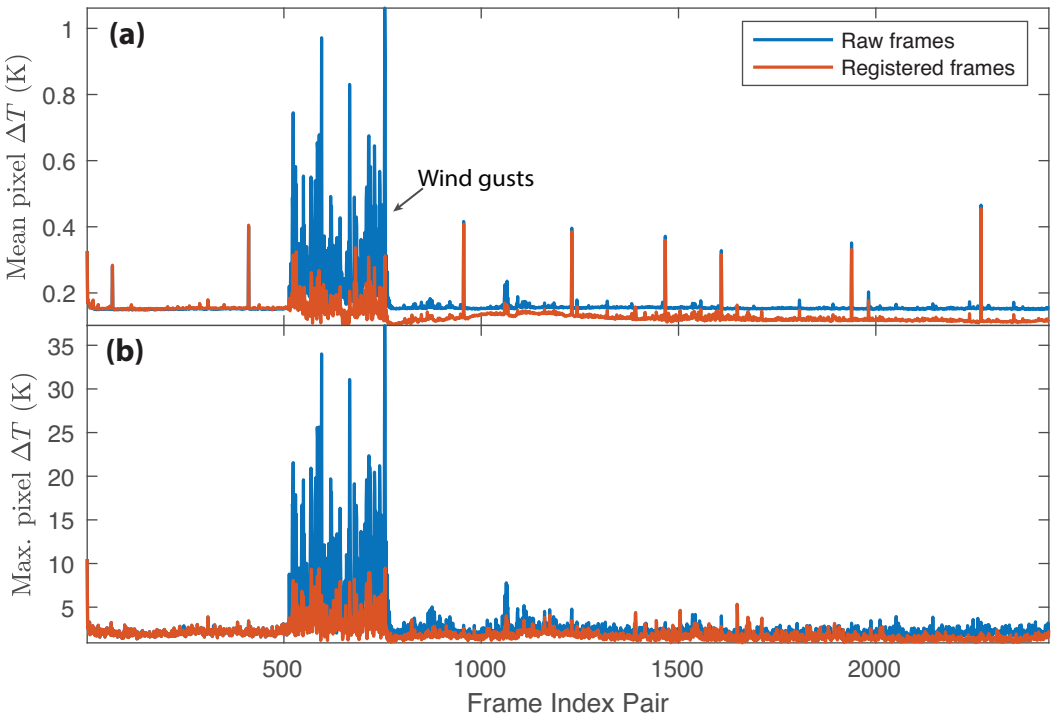

Figure S1. Mean (a) and maximum (b) temperature difference between subsequent frames for Event 3. Blue and red lines show temperature differences for raw and registered frames, respectively. Re-alignment of images during a period of strong wind-gusting (frames 500-750) reduced the frame-to-frame pixel error.

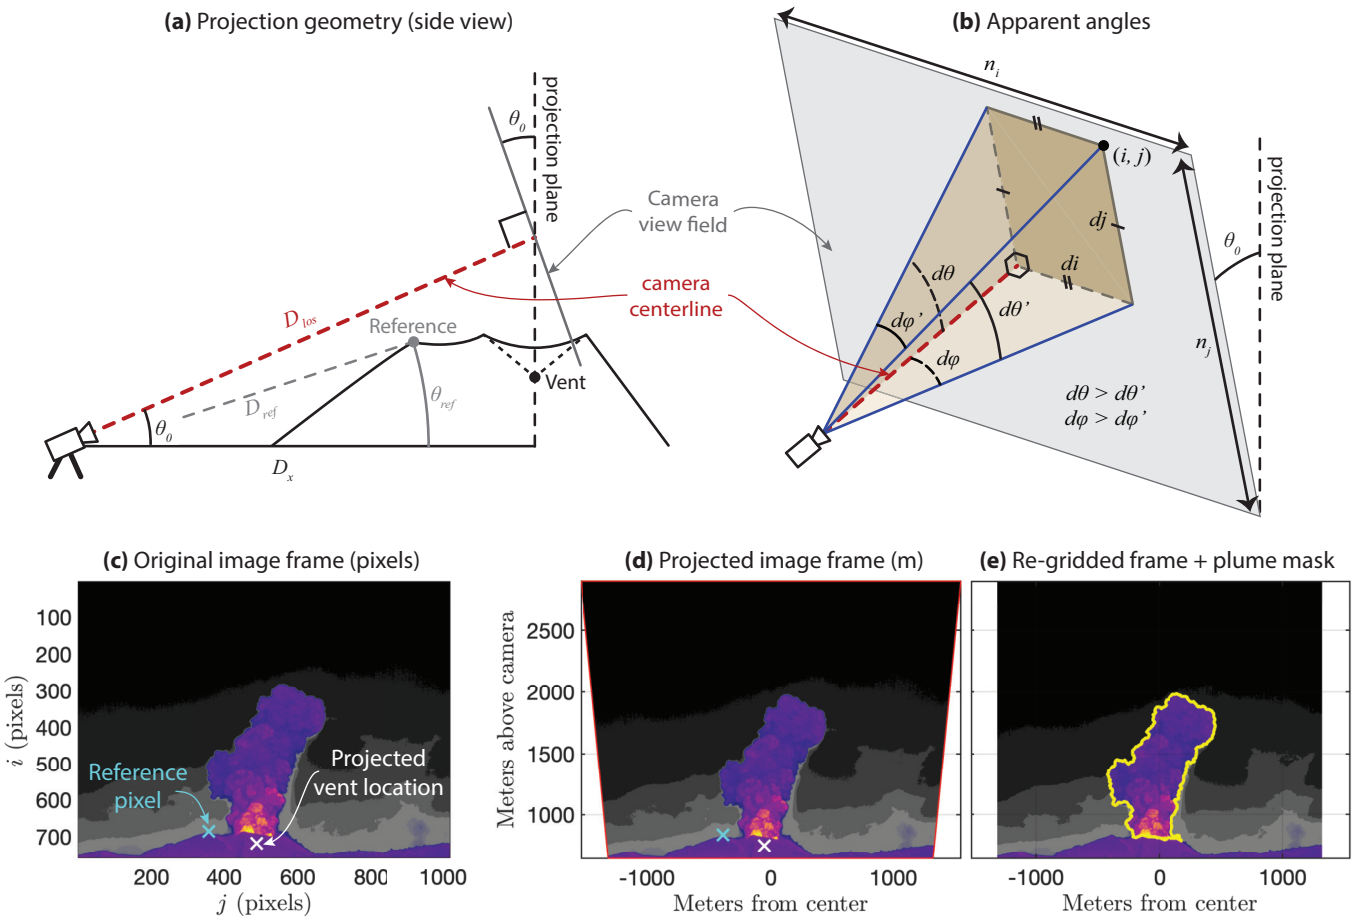

Figure S2. Field observation geometry for mapping image coordinates to spatial coordinates. (a) Geometry of field setup, perpendicular to camera view, showing the camera elevation angle θ_0, the camera field of view, and the projected image plan centered above the vent. (b) Schematic of angles and apparent angles (denoted with $'$) relative to a frame-of-reference centered on the camera image centerline for an arbitrary pixel coordinate (i, j). (c) Example raw image frame, showing the pixel coordinates of the reference pixel (i_{ref}, j_{ref}) and projected vent location (i_t, j_t). (d) The image from panel (c) after projection onto the vertical plane centered above the inferred vent location. (e) Projected thermal frame after the image re-gridding step.

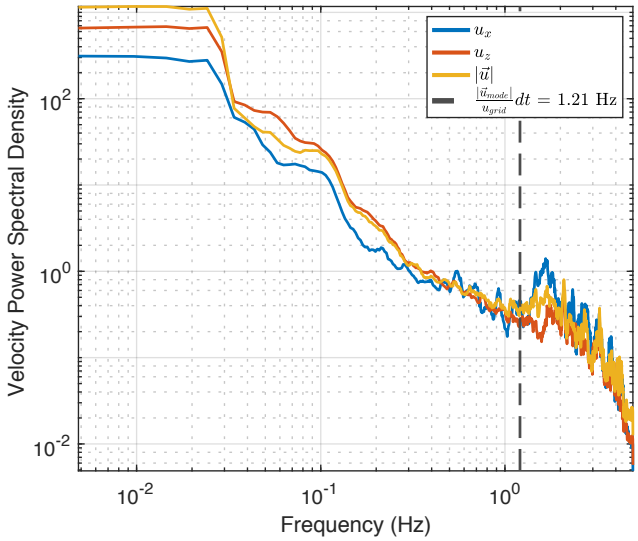

Figure S3. Optical flow velocity power spectral density for a sequence of thermal images from Event 3, using the maximum possible time resolution (about 10 frames per second). The Courant frequency of 1.21 Hz, obtained from Equation S.11 is marked by the vertical grey dashed line. Significant noise occurs at frequencies above this cutoff, which corresponds to sub-pixel displacements between frames.

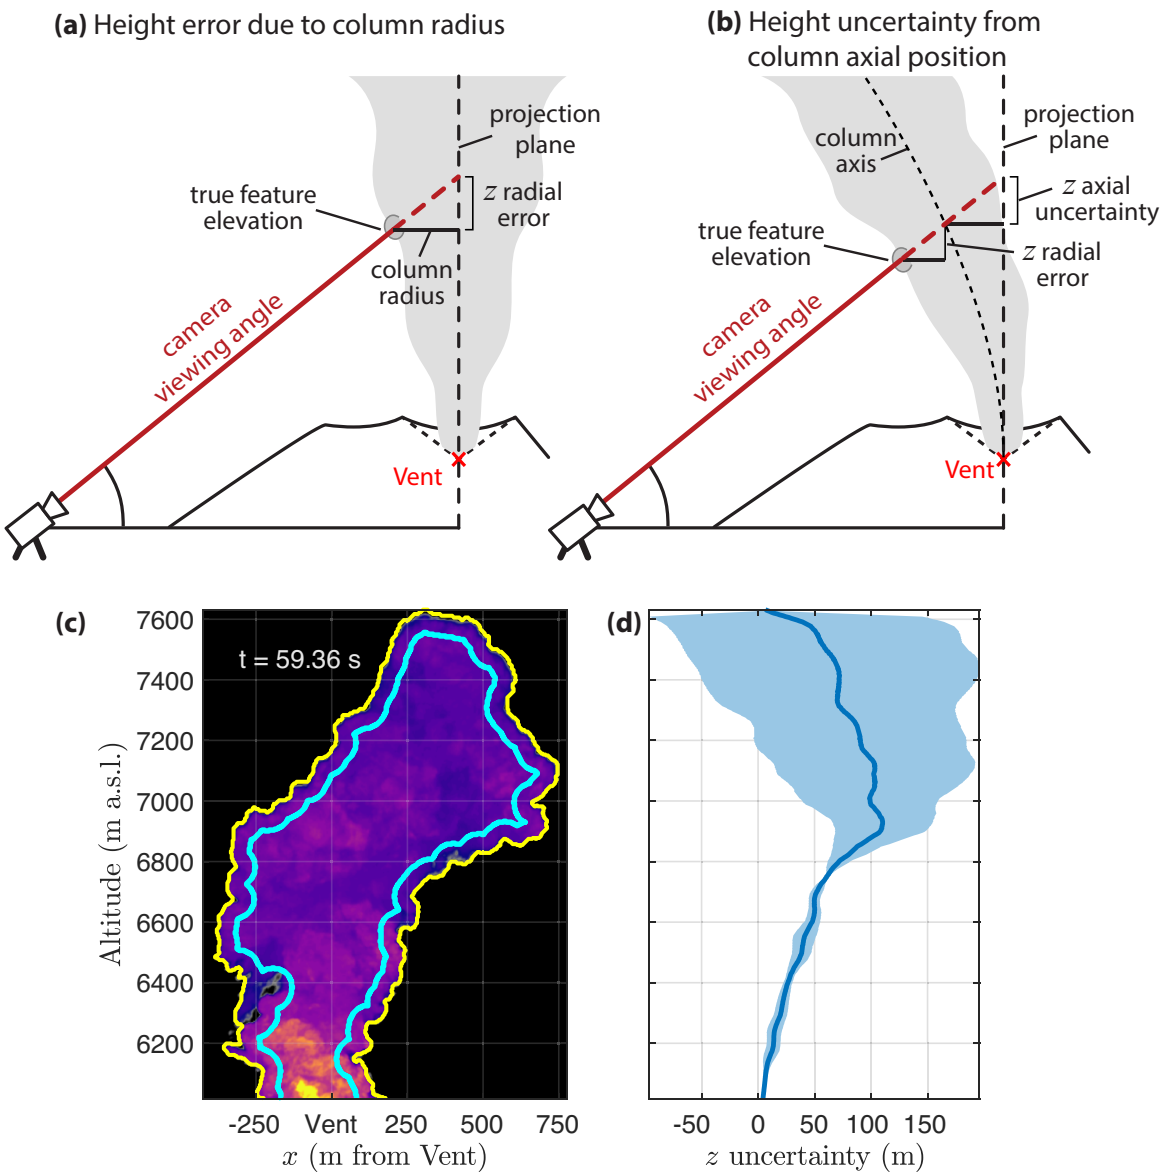

Figure S4. Sources of pixel height uncertainty. (a) Projection height error resulting from radial thickness of the column predictably causes overestimation of the true height of features at the visible column exterior. (b) Projection height uncertainty arises from uncertainty in the axial position of a bent column, is additive to the radial error and may produce an under- or overestimation of the height of features, depending on whether the axial position is away from or towards the camera, respectively. (c) Sample frame for estimating height uncertainty for $t = 60$ s in Event 3. (d) Column centerline uncertainty for the frame in (c). Solid blue line gives the estimated radial z error, and the shaded region gives the uncertainty from axial position.

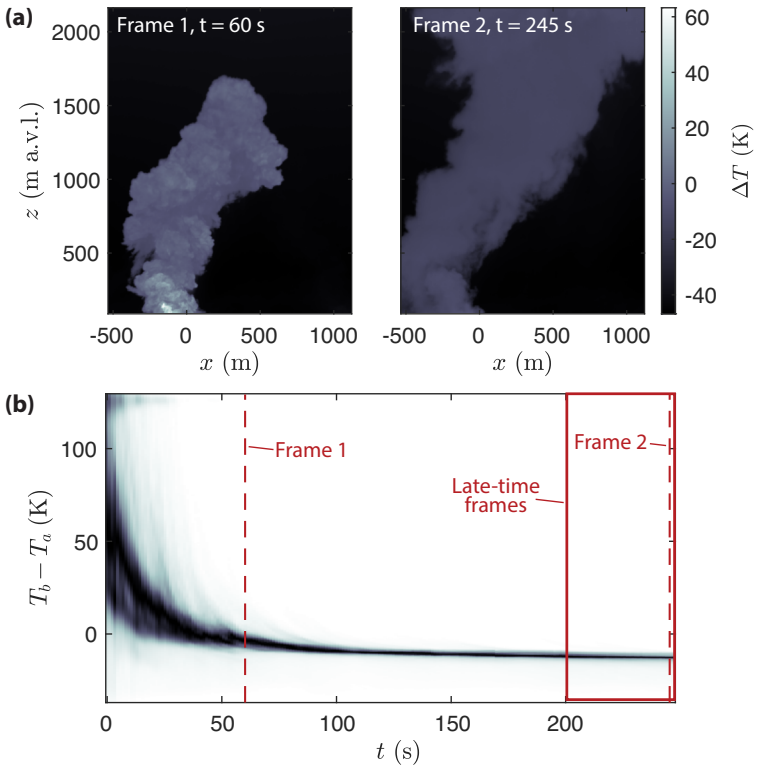

Figure S5. Example of late-time filtering for Event 3 atmospheric profile removal. (a) Comparing an 'early-time' frame where the column ΔT values are elevated well above background with a 'late-time' frame where column ΔT values are relatively homogeneous relative to background and column pixels remain largely opaque to background radiation. (c) Probability density map of initial pixel excess temperature $\Delta T_i = T_b - T_a$) versus time in the Event 3 plume, for all plume pixels. The PDF for each frame (pixel column here) is normalized to the maximum value in the frame, and blue dots indicate the mode value for each frame. Initially high ΔT at the start of the explosion decays with time towards a flat value corresponding to the background atmospheric profile (≈ -12.4 K). The red box highlights the late-time frames (standard deviation $<$ 1 to 2 K) used to obtain PDF functions for atmospheric profile fitting (Figure 4, main manuscript).

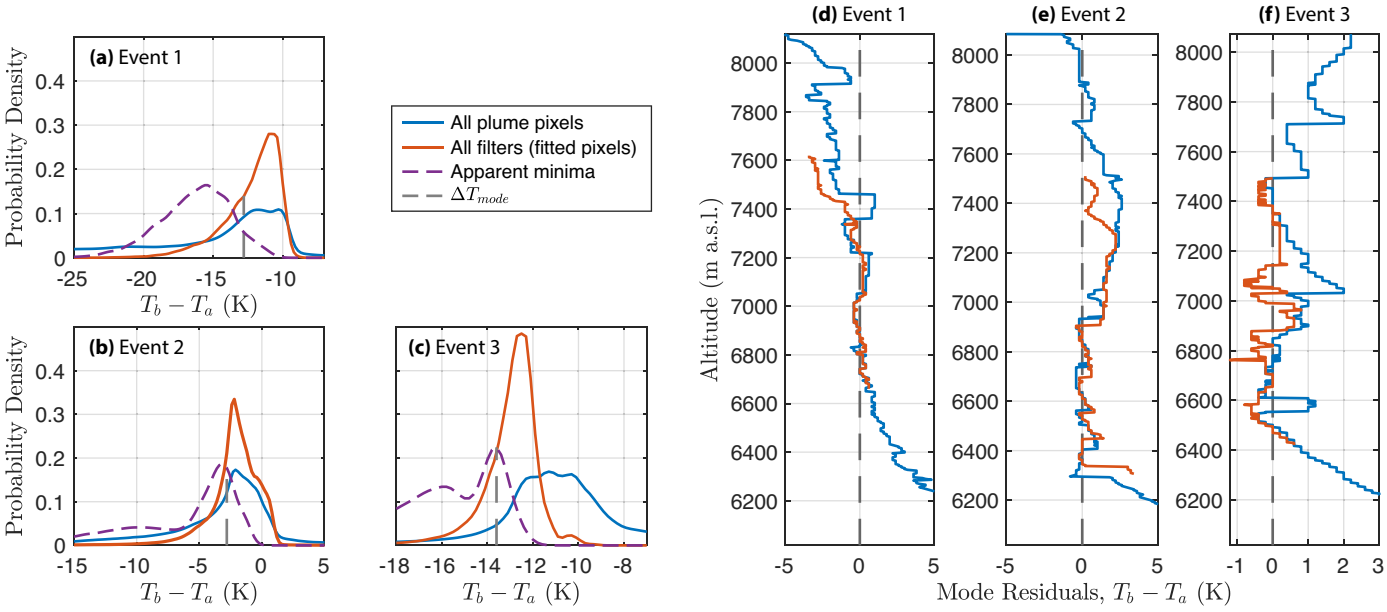

Figure S6. Apparent minima PDFs are similar to filtered PDF half-maxima. (a)-(c) PDFs of $T_b - T_a$ for Events 1-3. Blue and orange lines are identical to those in Figure 4, where blue lines show the distribution of all plume pixels, and the orange line results after filtering pixels for vent proximity, large height uncertainty, and late time. The purple line shows the distribution of apparent row-wise minima. (f) Mode $_b - T_a$ as a function of height for Event 1. As in (a)-(c), blue line shows all pixels and orange line shows filtered pixels. (g) As (f), for Event 2. (h) As (f), for Event 3.

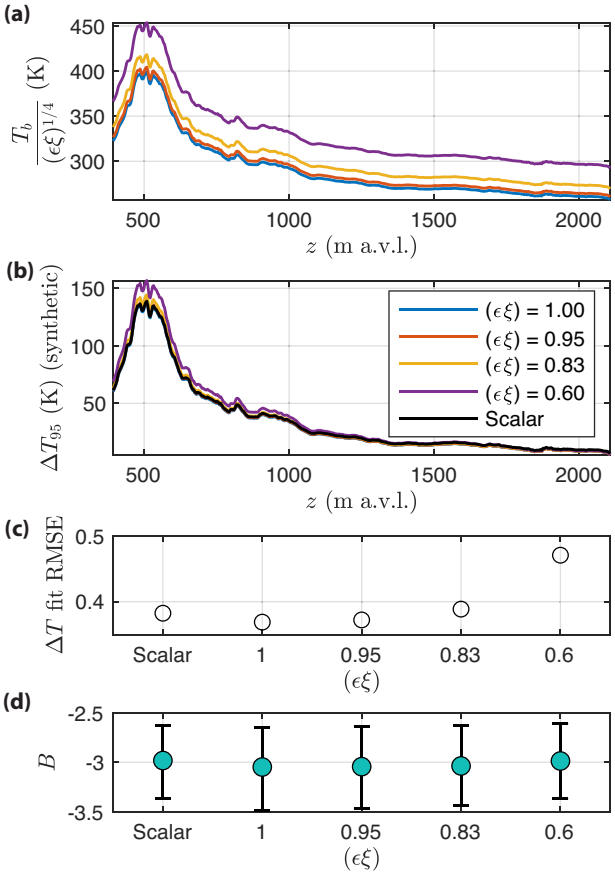

Figure S7. Comparison of power law fit results using the additive ΔT_{mode} atmospheric profile correction versus multiplicative corrections from assumed emission and transmission loss. (a) Absolute temperature for the tracked leading front of Event 3 after applying emission and transmission corrections ($\epsilon\xi$). (b) ΔT_{95} after apply a scalar shift to align the decay curves about 1800 m. (c) RMS error (K) of power law fit for each correction. (d) Estimated B values for each correction.

Table S2. Table of lower (min.) and upper (max.) height limits for linear R fit and power law B fit. Bold/black denotes manually chosen limits, cyan denotes limits imposed by automatic pixel filters, and light gray indicates that the full extent of the data set was used. See also Figures S8-S15 below for comparison.

Event No.	Track No.	**Fit height limits (m)**			
		R min.	R max.	ΔT min.	ΔT max.
1	1	116	**377**	**112**	616
1	2	99	**420**	**98**	581
1	3	**149**	**300**	**105**	850
1	4	**146**	**212**	78	809
1	5	102	**325**	**112**	571
1	6	**173**	**392**	**95**	758
1	7	599	**510**	78	755
1	8	**146**	**323**	**109**	602
1	9	122	**358**	**132**	874
1	10	122	**300**	**109**	632
1	11	200	364	**112**	442
2	1	139	**540**	**145**	1080
2	2	142	342	**190**	425
2	3	103	425	**107**	**426**
2	4	107	**182**	**121**	415
2	5	**149**	**375**	75	545
2	6	**175**	**312**	**103**	506
2	7	**163**	**353**	**89**	489
2	8	132	**384**	**89**	555
3	1	**540**	**1172**	**558**	1518
3	2	**924**	**1157**	**393**	1518
3	3	176	**321**	**214**	495
3	4	129	**305**	**233**	622
3	5	**300**	**597**	**274**	**720**
3	6	118	634	**198**	88
3	7	129	640	98	837

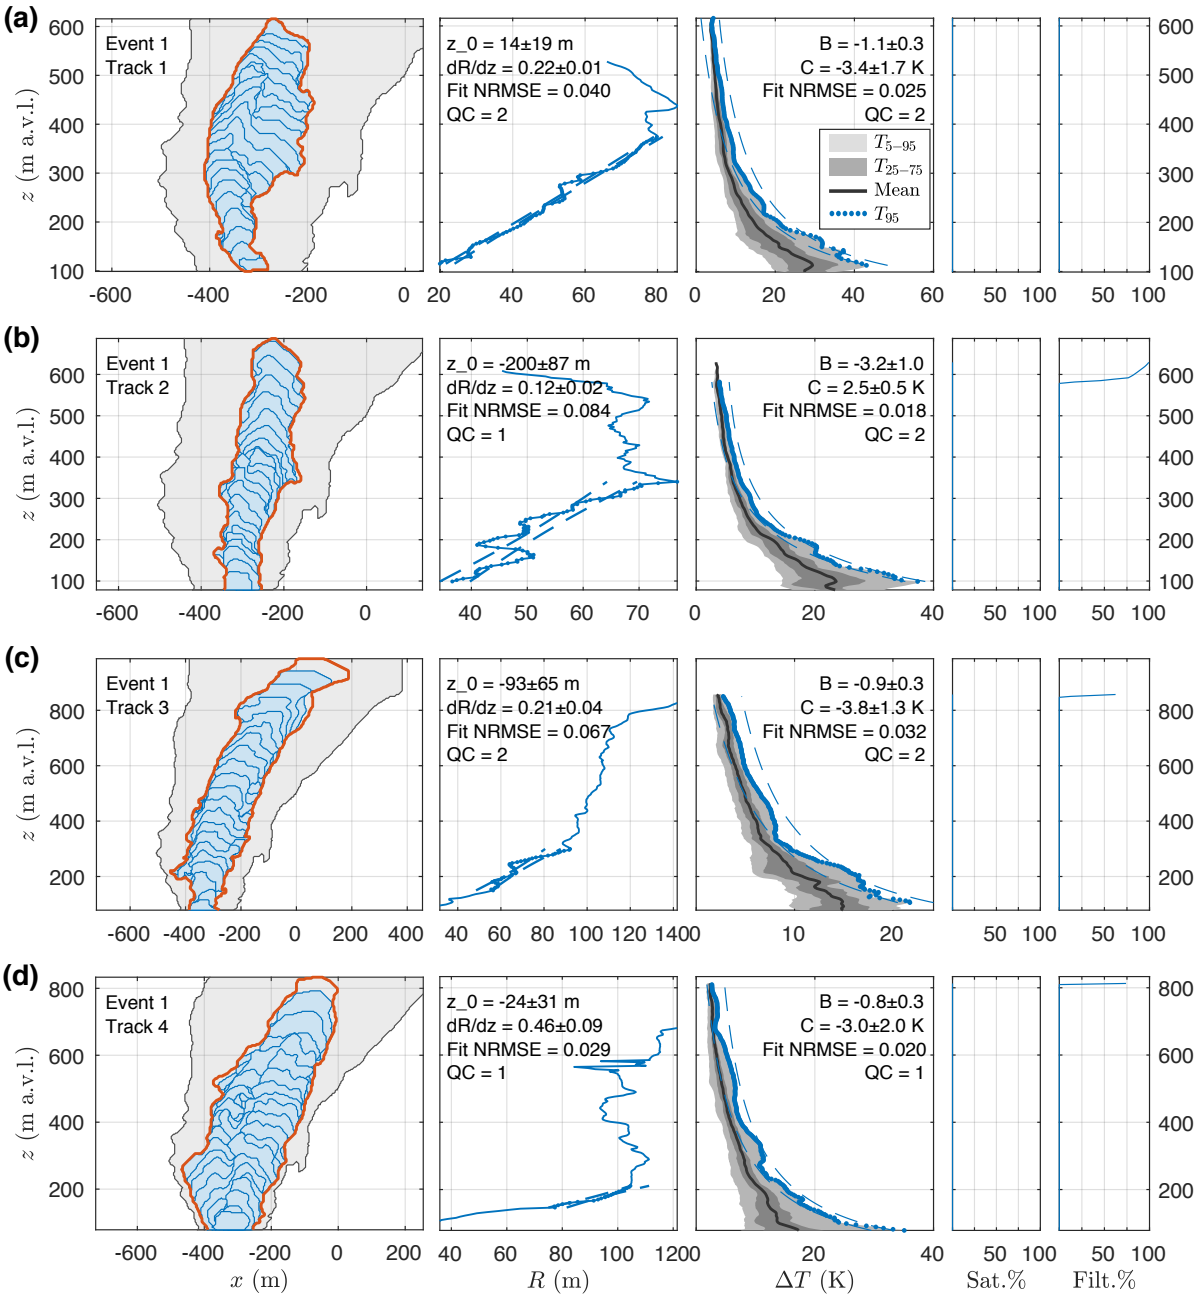

Caption on next page.

Figure S8. Tracking and curve-fitting results for Event 1, tracks 1-4. Left-most column shows tracked feature outlines at various time steps in blue, and the total outline (enclosing all tracked pixels for all times) in orange. The gray outline shows the time-averaged image mask as shown in Figure 7 of the main manuscript. The second column from the left shows the tracked feature effective radius (taken as the square root of total pixel area) versus height (of pixel values at each time step) above the vent. The linear model 95% confidence bounds for $R(z)$ are shown with blue dashed lines. The third column from the left shows evolution with height of the reconstructed excess temperature distribution. Light gray and dark gray fields show the 5-95 and 25-75 percentile range, respectively, and the dark gray line gives the mean. The curve-fitted ΔT_{95} is highlighted with blue dots, and the fit confidence 95% bounds are again shown with blue dashed lines. Note that the plotted power law fit confidence bounds shown are for the central z_0 estimate only, but the printed uncertainty in B includes for the range of uncertainty in z_0. We also report the manually identified quality check (QC) value, where 0 = poor, 1 = ok, 2 = excellent (see text for description). The second-from-right and right-most columns show, respectively, the percentage of pixels at each height that are removed from curve-fitting due to either temperature saturation or large height error. Saturated and height-filtered pixels are removed from the temperature fit, while only height-filtered pixels are removed from the radius fit.

276

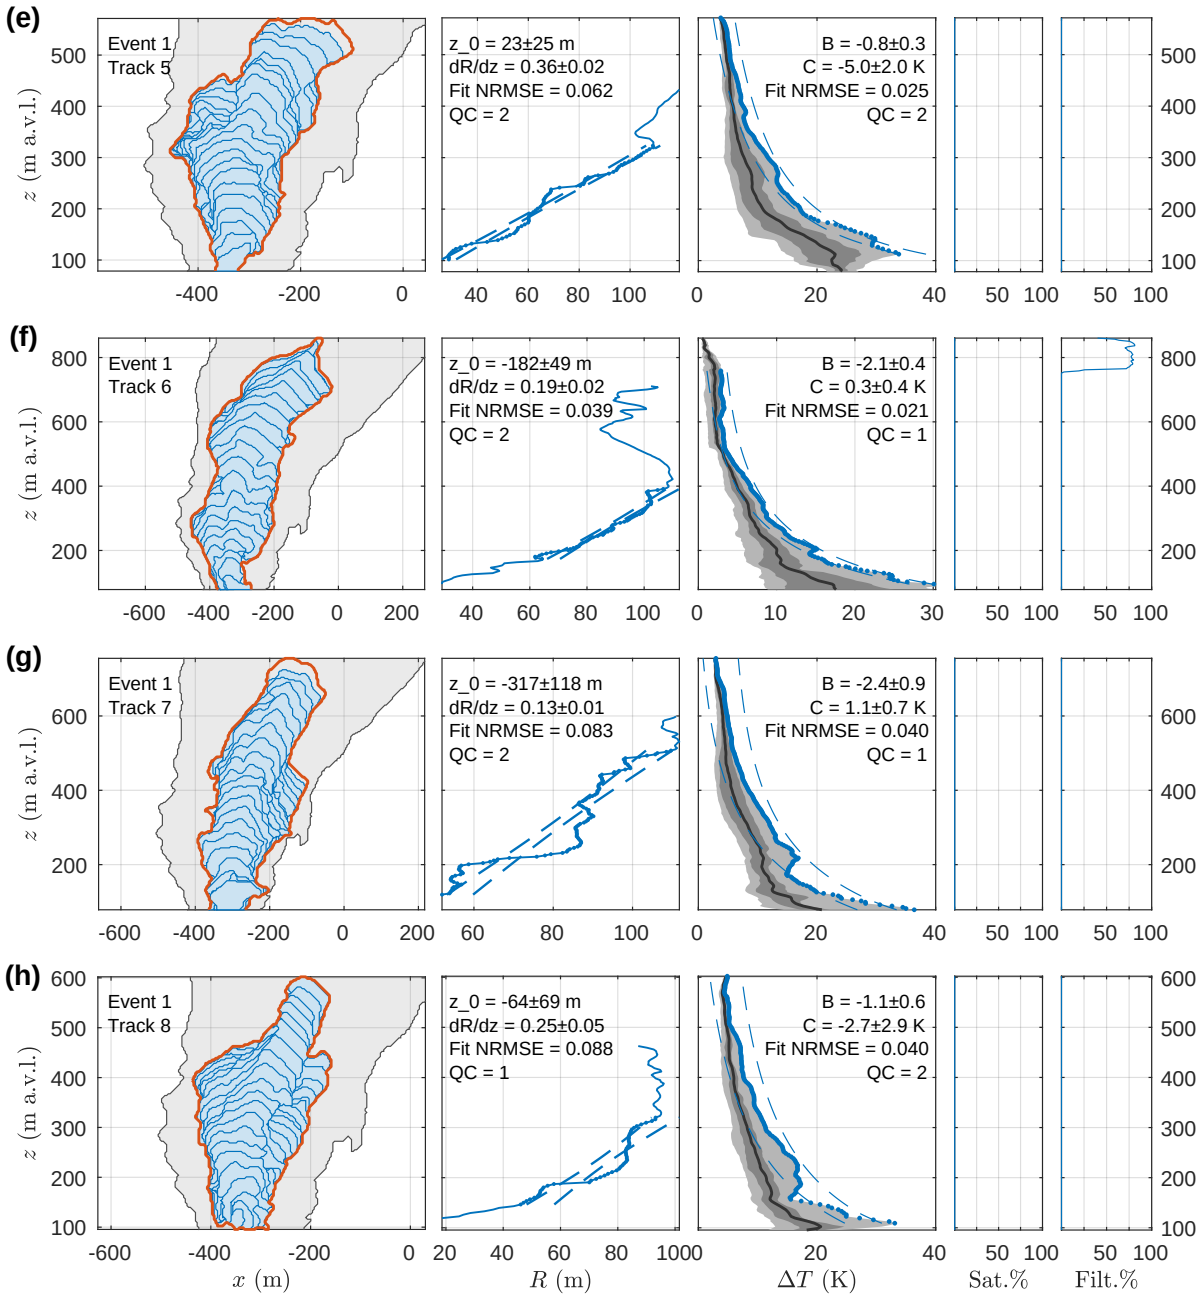

Figure S9. Tracking and curve-fitting results for Event 1, tracks 5-8. See the caption of Figure S8 for full description.

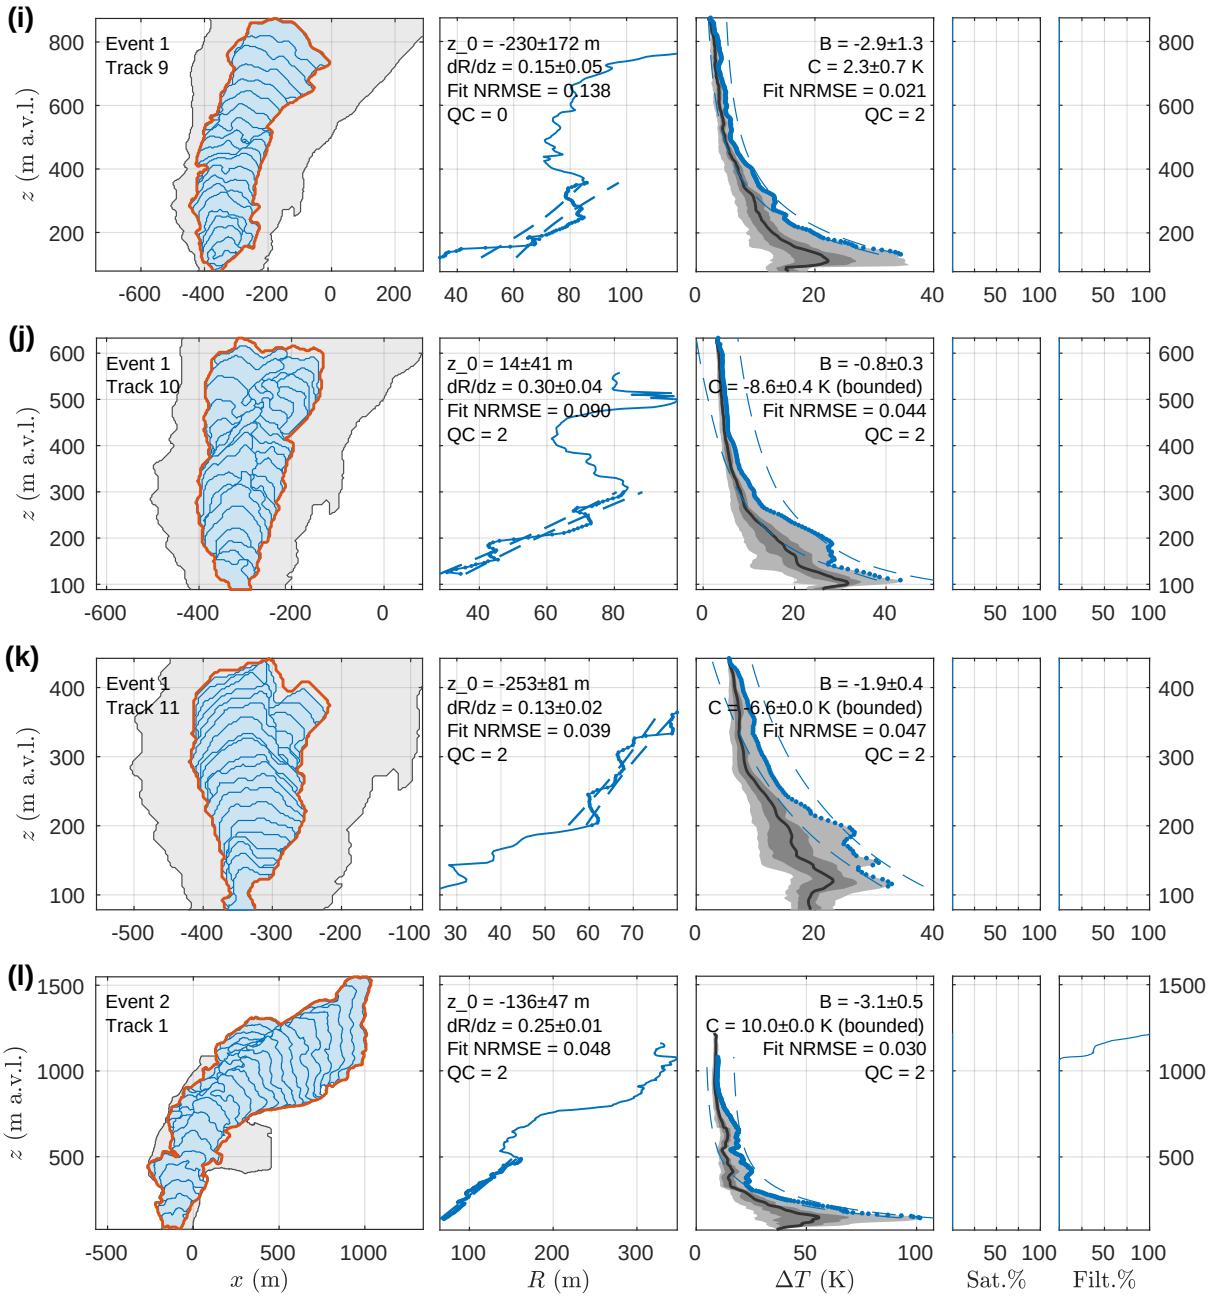

Figure S10. Tracking and curve-fitting results for Event 1, tracks 9-11 and Event 2, track 1.

See the caption of Supplementary Figure S8 for full description.

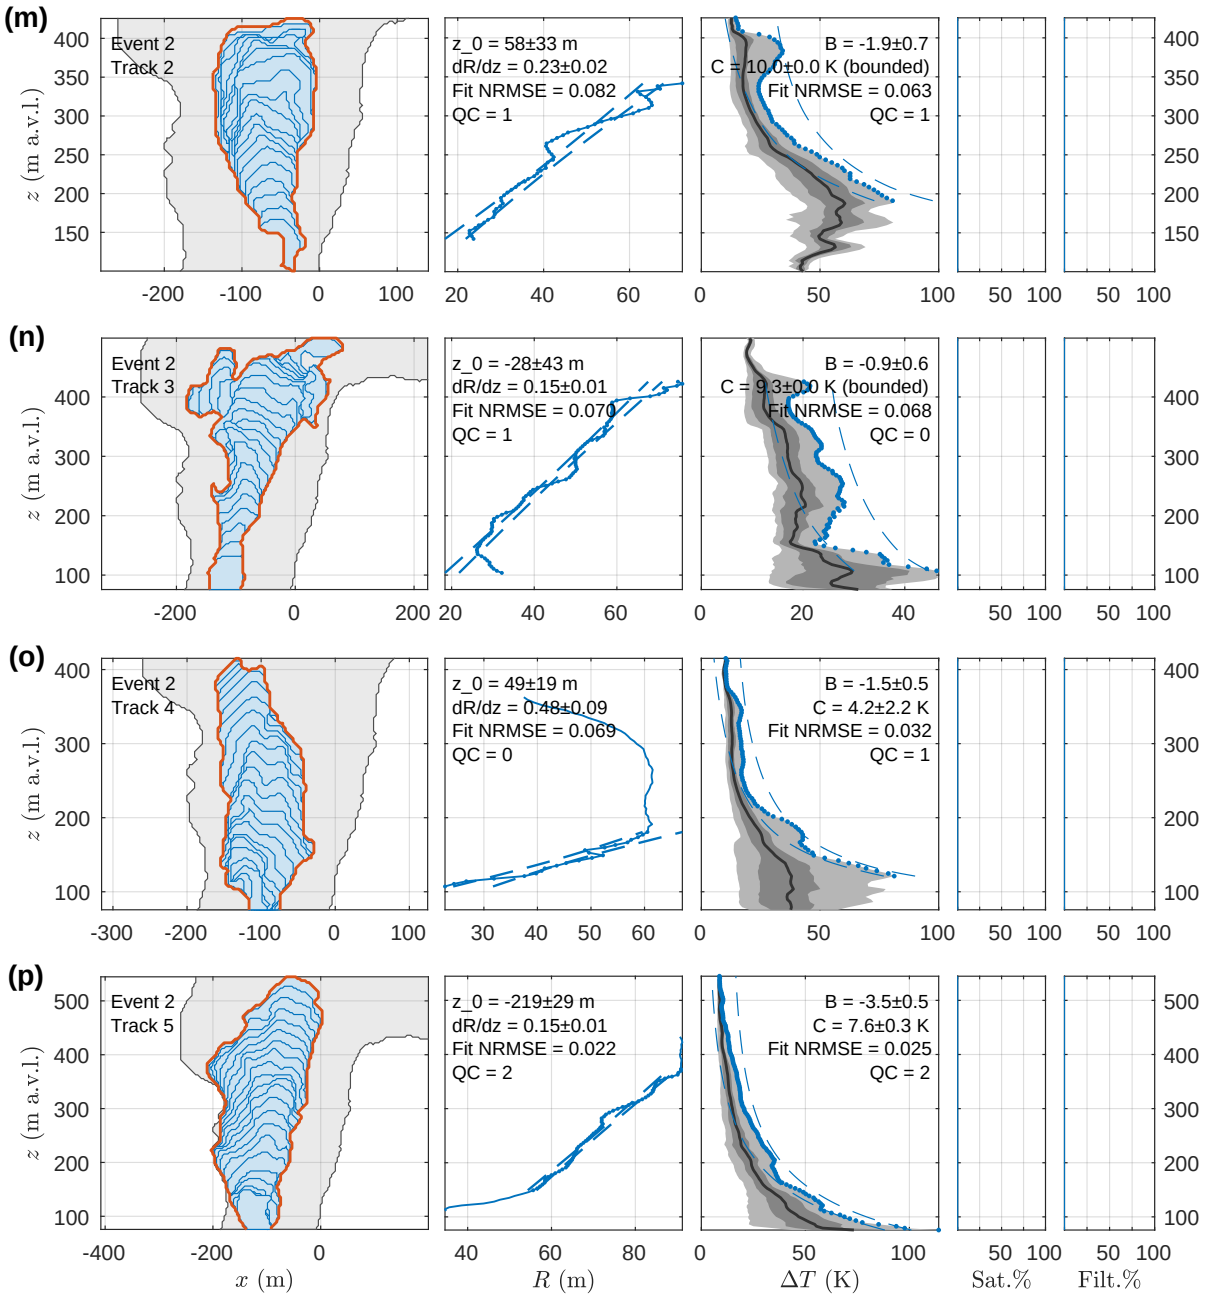

Figure S11. Tracking and curve-fitting results for Event 2, tracks 2-5. See the caption of Supplementary Figure S8 for full description.

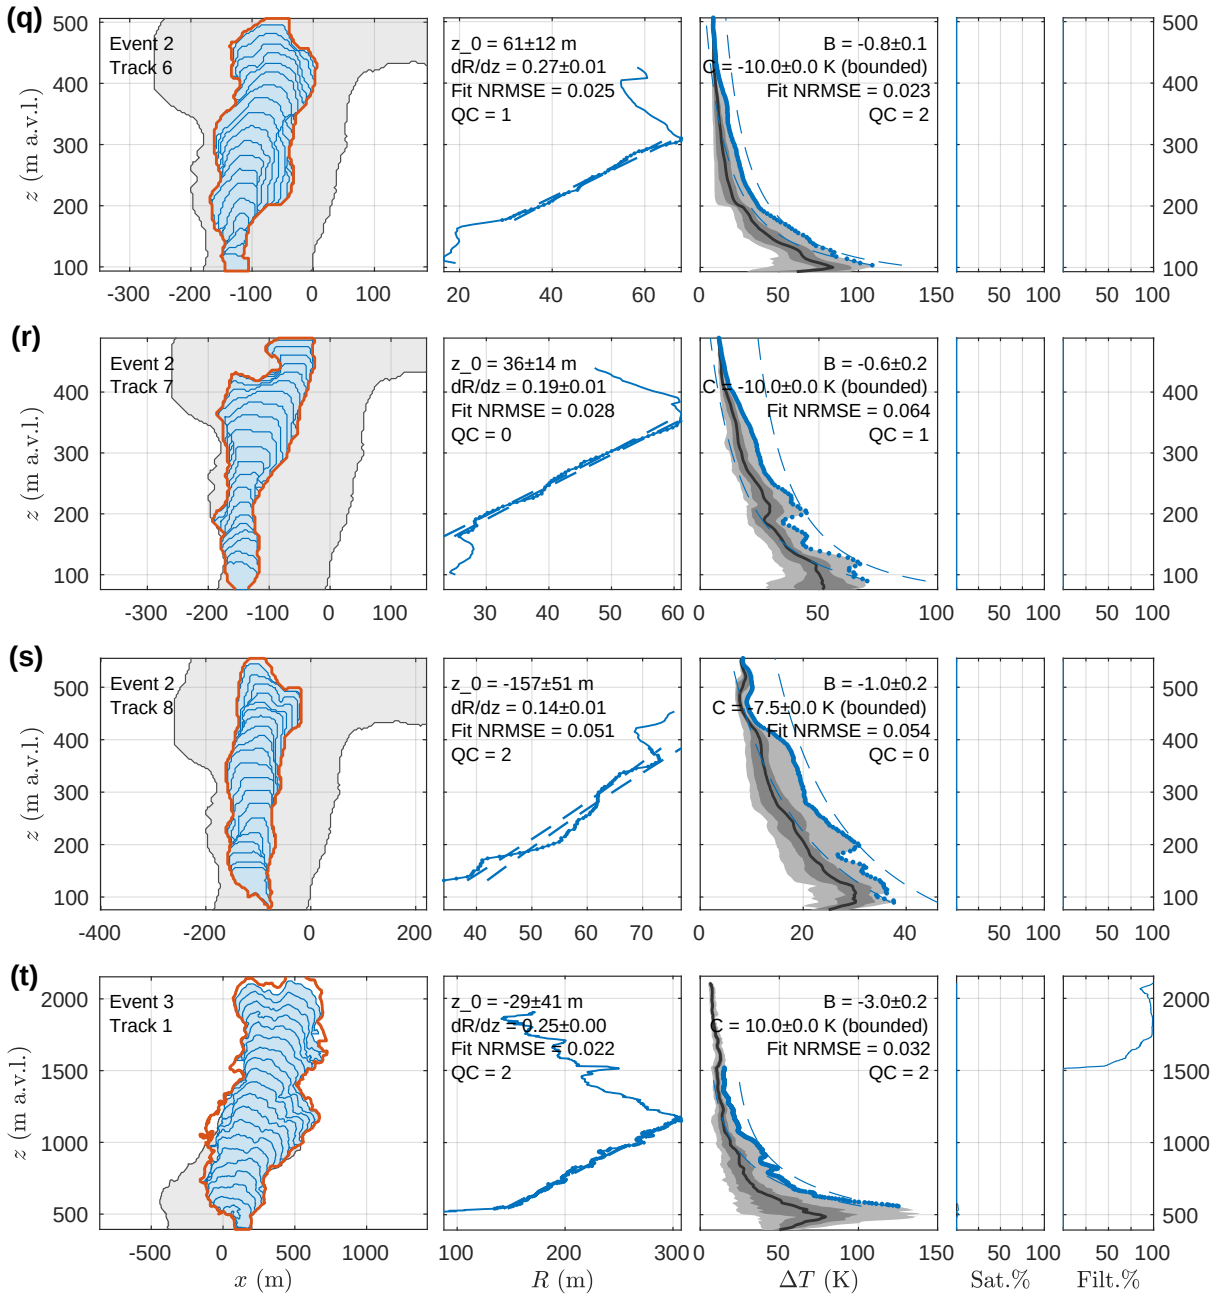

Figure S12. Tracking and curve-fitting results for Event 2, tracks 6-8 and Event 3, track 1.

See the caption of Supplementary Figure S8 for full description.

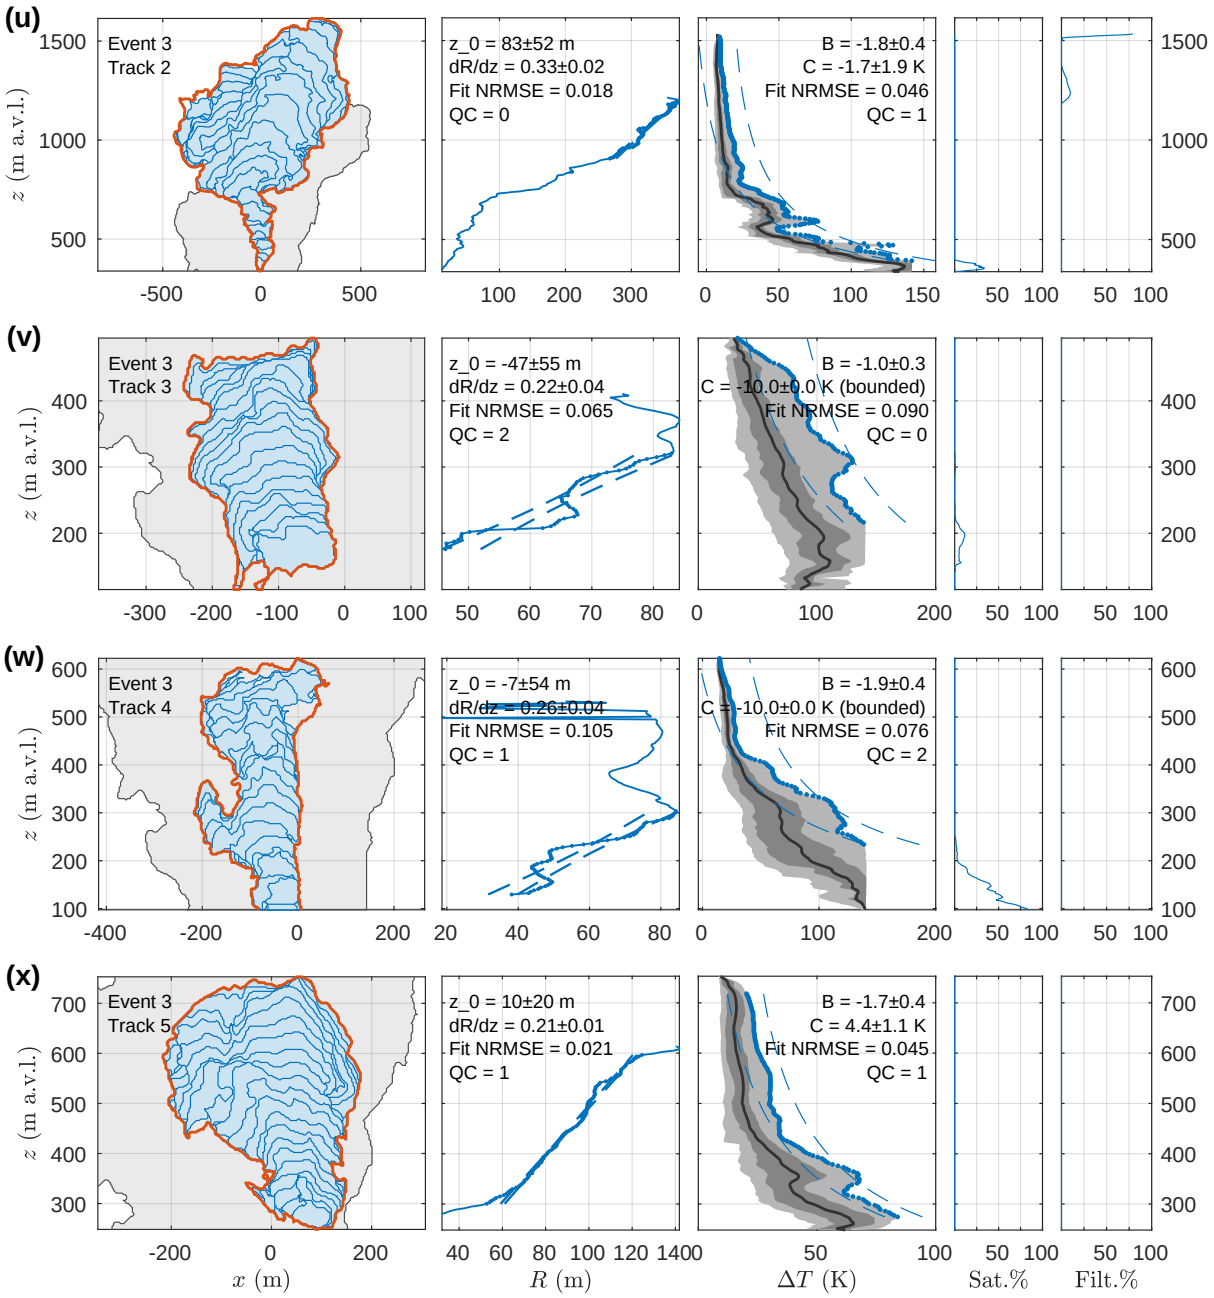

Figure S13. Tracking and curve-fitting results for Event 3, tracks 2-5. See the caption of Supplementary Figure S8 for full description.

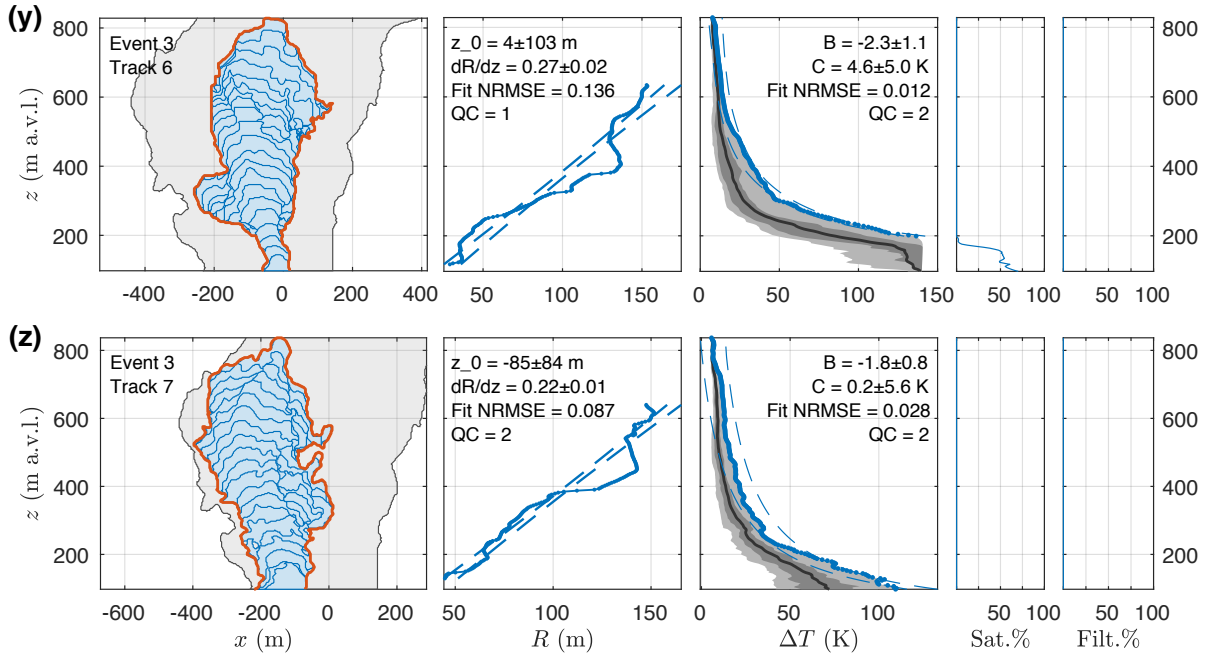

Figure S14. Tracking and curve-fitting results for Event 3, tracks 6-7. See the caption of Supplementary Figure S8 for full description.

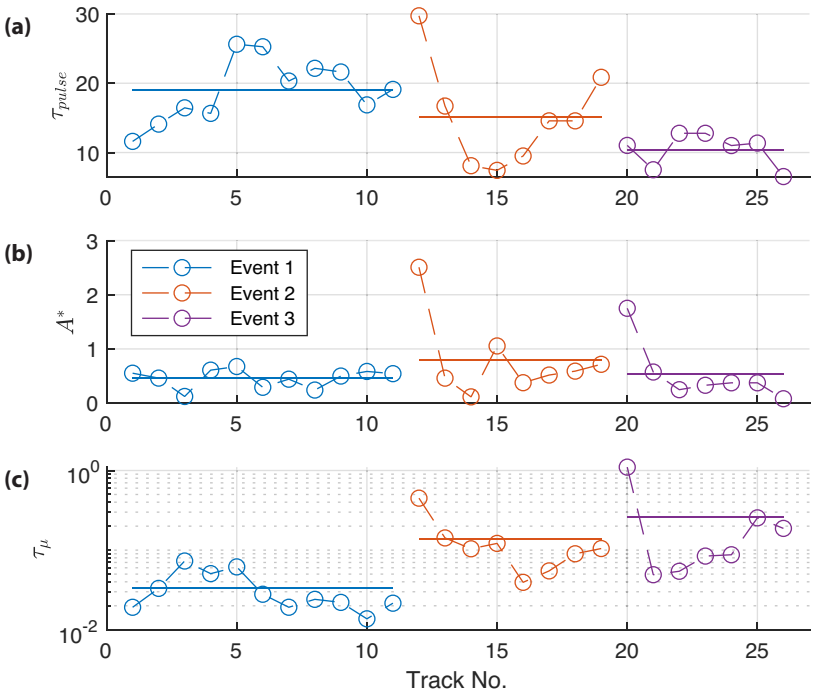

Figure S15. Components of Pu_μ for all tracks of Events 1-3. Horizontal solid lines show unweighted track averages. Track numbers are in chronological order for each event. (a) A^*, (b) τ_{pulse}, (c) τ_μ.