

Image Processing Onboard Spacecraft for Autonomous Plume Detection[☆]

David R. Thompson*, Melissa Bunte¹, Rebecca Castaño², Steve Chien²,
Ronald Greeley³

Abstract

Previous missions have imaged active plumes at Io and Enceladus, as well as outgassing by cometary nuclei. It is often difficult to predict where and when these transient events will occur, so characterizing them requires collecting long image sequences with many redundant frames. This demands a prohibitive fraction of the spacecraft's limited cache and bandwidth, and precludes sustained surveys of plume activity. Onboard processing could enable long-term plume monitoring campaigns with high imaging rates. Specifically, spacecraft can analyze image sequences onboard to identify plumes, with events triggering preferential storage, prioritized transmission, or followup with coincident observations by Thermal or Visible Near-Infrared imagers. We propose a detection method based on horizon identification with Random Sample Consensus (RANSAC). The approach evidences reliable performance on a test set of plume images from Enceladus and Io.

Keywords:

Europa, Io, Enceladus, Plume Detection, Spacecraft Autonomy, Instruments and Data Processing

*Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr., Pasadena, CA 91109.

Email address: david.r.thompson@jpl.nasa.gov (David R. Thompson)

URL: <http://ml.jpl.nasa.gov/people/thompson.shtml> (David R. Thompson)

¹School of Earth and Space Exploration, Arizona State University, Tempe, AZ.

²Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr., Pasadena, CA 91109.

³Deceased 27 Oct 2011, formerly of the School of Earth and Space Exploration, Arizona State University, Tempe, AZ.

Previous missions have imaged active volcanic plumes at Io and Enceladus, as well as outgassing by cometary nuclei. These phenomena provide key constraints on subsurface processes and composition [1]. To date, over 70 distinct images containing plumes have been collected. In the case of Io alone, Galileo NIMS images of the Jovian system identified 37 hot spots that could be centers of volcanic activity associated with plumes [2]. However, plumes are difficult to detect since one cannot predict their precise timing or location in advance. The existing record consists mainly of disconnected events. Exceptions include a 10 hour sequence from the Voyager 2 observations of Io [3] which imaged multiple plumes on a trajectory through the Jovian system. Currently such sequences yield many redundant or uninformative frames. They demand a prohibitive fraction of the spacecraft's limited cache and bandwidth, which discourages long-term monitoring of plume activity.

This work investigates an onboard processing technique to enable sustained plume monitoring campaigns. A spacecraft in tour mode could collect surplus images at a high rate, analyze them onboard to detect plumes, and save the fraction containing key events for downlink. This could change the nature of the plume catalog. A sustained survey would provide a temporally-continuous record of plume activity, giving a new dimension to the previous dataset and permitting time series analysis to constrain surface and subsurface dynamics. Plume monitoring is especially appropriate for early mission or tour phases when narrow-field cameras could otherwise be idle. It could also be used in small bodies exploration, to characterize outgassing activity of cometary nuclei during rendezvous [4].

Onboard image processing has been used in various terrestrial and planetary missions. Similar technologies have been demonstrated for change detection on the WATCH system aboard the Mars Exploration Rovers [5]. Here, automated processing identified images containing dust devils. Candidate image regions were enclosed in subframe images for preferential transmission. Similar selective downlink strategies could enable plume surveys, with event detections triggering preferential storage and transmission of coincident observations by Thermal or Visible Near-Infrared imagers.

In prior research, Bue et al. [6] demonstrate an efficient detection algorithm based on finding ejected material beyond the planetary limb. We expand on their work with a new method that evidences strong performance on a test set of images from Enceladus and Io. We begin by describing the approach and then evaluate its performance on a set of full-disk images from Galileo and Cassini spacecraft. We briefly discuss the coverage and sensitiv-

ity offered by horizon-based plume detectors.

1. Approach

Desirable characteristics of a plume detection algorithm include: *sensitivity*, the ability to find faint plumes no more than a few pixels in height; *robustness* to noise that is pervasive in high-radiation environments such as Jupiter; and *graceful degradation* such that failure modes do not result in unnecessary false detections. Sensitivity over large areas of the planetary surface is desirable for accumulating statistics on total plume activity. However, it is equally important that performance be consistent over the sensitive areas.

With these factors in mind we propose an algorithm based on identifying the planetary horizon as in Bue et al. [6]. This approach can detect ejected plume material that extends beyond the sunlit limb. We assume that any significant lens distortion or off-axis effects have been corrected prior to running the algorithm. The basic method has three sequential stages. The first step uses a classical edge detector to find pixels on boundary between the planetary disk and the dark background. Second, we fit an ellipse to match this horizon. Finally, if an ellipse is found we identify illuminated plume material by analyzing pixel intensities in an annular region around the planetary limb.

1.1. Edge Detection

We first find candidate horizon pixels with the edge detection algorithm of Canny [7]. We favor the Canny algorithm because it is well studied and offers robust performance on our test set. The Canny detector initially smoothes the image with a low-pass operator to reduce noise. It then convolves a bidirectional Sobel operator to find pixels that have high spatial derivatives. A subsequent step known as nonmaximal suppression forces the final detected edge to be a narrow line. The algorithm follows the resulting traces along their dominant directions and suppresses any pixel not at the maximum. Finally, it performs a threshold operation to find the final detected edges. It accounts for edges fluctuating just over or under threshold by using a hysteresis method; it suppresses borderline weak detections that are not supported by neighboring strong edge pixels. The result is a set of edge pixels $E = \{(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)\}$.




Figure 1: Edge detection with the Canny edge detector. The horizon contour is the strongest edge in the scene, permitting a strict detection threshold. False “noise pixels” corresponding to surface features or imaging artifacts are expected at this stage. The data product is Galileo 3178r, Image credit: NASA/JPL/University of Arizona.

An example of edge detection input and output appears in the Galileo image of Io shown in Figure 1. It is not yet necessary to detect the entire planetary disk or to exclude all other features. For example, the edges in Figure 1 include several high-contrast surface features. Additionally, the unlit parts of the disk are undetected. Neither issue impacts performance. It is only important at this stage that a significant fraction of the limb is sampled to constrain horizon location and curvature, and that the limb comprises a significant fraction of the detected edge pixels. This is not difficult to achieve in practice since the sunlit horizon generally has the highest contrast of all the linear features in the scene. We use strict hysteresis thresholds of 0.5 and 0.1 for strong and weak edges respectively. This combination offers consistent performance across all contrast and illumination levels of the images considered in this study.

1.2. *Ellipse Fitting*

The next stage finds the ellipse that intersects as many edge pixels as possible. We presume the horizon line has the form of a conic $F(\mathbf{x}, \mathbf{a})$ where $\mathbf{x} = [x^2, xy, y^2, x, y, 1]$ and $\mathbf{a} = [a, b, c, d, e, f]$. we represent the conic using the second-order polynomial [8]:

$$\begin{aligned}
 F(\mathbf{a}, \mathbf{x}) &= \mathbf{a}\mathbf{x}^T = 0 \\
 ax^2 + bxy + cy^2 + dx + ey &= -f
 \end{aligned} \tag{1}$$

There are five degrees of freedom in the expression since the entire vector \mathbf{a} can be rescaled without changing the conic. In this work, we use the constraint $f = -1$ following Rosin [9], and solve the resulting set of linear equations with five candidate horizon points to define a single unique contour. For the data matrix given by

$$X = \begin{bmatrix} x_1^2, x_1 y_1, y_1^2, x_1, y_1 \\ \dots \\ x_5^2, x_5 y_5, y_5^2, x_5, y_5 \end{bmatrix} \quad (2)$$

the optimal conic is

$$\mathbf{a} = (X^T X)^{-1} X (-I)^T \quad (3)$$

Here $(-I)^T$ signifies a column vector of six -1 values. We count only true ellipses, categorically rejecting all fits corresponding to parabola (for which $ac = 0$) and hyperbola (for which $ac < 0$).

We search for the best ellipse using the Random Sample Consensus (RANSAC) method of Fischler and Bolles [10]. RANSAC is a robust estimator for fitting parametric models in the presence of outliers; it chooses random subsets of edge points and computes a closed-form fit for each. In our case, we choose subsets of size five and find the optimal interpolating ellipse using the closed-form solution. The RANSAC algorithm is the slowest stage of the plume detection procedure but one can trade some detection sensitivity for additional speed by limiting the number of random attempts. The actual number required varies by image quality. For example, consider a very poor quality image where only 25% of the edge points lie on the horizon. All five random pixels are horizon with probability $0.25^5 \approx 0.097\%$. However, with 10^4 random subsets the probability is greater than 99.99% that this will happen at least once. Designers can significantly improve speed by prefiltering the edge list with ephemeris information or with prior fits from earlier images in a sequence.

We use a series of simple rules to score the candidate ellipses. We automatically reject ellipses whose shape alone suggests they are obviously incorrect. These include ellipses that are greater than the size of the image or whose aspect ratio exceeds a user threshold. We use a value of 1.05 for our tests, which requires that valid ellipses be nearly circular. A large number

of candidate horizons may still survive this filter. We find the best horizon by counting the number of edge points that are located within 1 pixel of the ellipse. One can require a certain minimum fraction of edge pixels to fall within this boundary. This stage can easily recognize when an ellipse is non-sensical or the fit has failed, since in these cases a significant portion of the image’s edges or bright pixels will lie outside the horizon. For each candidate we compute the geometric distance from every edge point to the ellipse using the algorithm of Eberly [11]. This is an iterative approach based on Newton’s method. If faster approximations are required one can substitute the closed form algebraic error $F(\mathbf{x}, \mathbf{a})$ applied to normalized data. The total count is used to select the best elliptical shape which is then taken to be the horizon. Figure 2 shows several examples (large green ellipsoids).

1.3. Plume Detection

Having identified the planetary contour, we search an annular region beyond this horizon to find bright pixels corresponding to ejected material. We look for plumes in pixels at altitudes up to a user defined range of the limb (in this work we use altitudes up to 10% of the planetary diameter). The set of all pixels in this annular region are taken to be the “background.” The mean and standard deviation of these pixels form an estimate of the noise distribution in the vicinity of the planetary horizon. Any pixels with intensity greater than τ standard deviations above the mean are grouped into contiguous features. For this work, we favor $\tau = 5.0$ as our pixel intensity threshold; more lenient values would result in higher detection rates at the possible cost of false positive detections.

A final rule-based filtering step applies a series of simple tests on size and aspect ratio. Valid detections must be at least 5 pixels in area, and no larger along its longest dimension than a single planetary radius. We also filter detections touching the image borders since these are also likely to be artifacts. These tests are seemingly obvious but help to exclude several imaging artifacts, like those noise and stripe features of Figure 2 Lower Left, that could otherwise trigger a detection. Any remaining regions are estimated as plumes, and their centroids taken as a list of detection locations $D = \{(d_{xi}, d_{yi})\}_{i=1}^n, n \geq 0$.

Figure 2 shows an example of typical detections using this method. A green line follows the best-fit ellipse, and a plus shows the size and location of each detected plume. Image 3178r illustrates a small plume on the limb of Io. Image PIA09761 is shown as an example Enceladus detection, but




Figure 2: Typical detection results. The Green ellipsoids signify the best-fitting ellipse, and a plus sign shows the location and size of each detection. Clockwise from upper left: Galileo 3178r; PIA09761 shows an Enceladus scene (excluded from performance analysis); Voyager C2066857 shows multiple simultaneous plumes, with the smallest less than 3 pixels in height; Galileo 5147r evidences robust detection despite high noise and imaging artifacts. Image credit: NASA/JPL/University of Arizona (Galileo), NASA/JPL (Voyager), NASA/JPL/Space Science Institute (Cassini).




Figure 3: Typical failure modes. Left: Cassini N00028218, where glare and low SNR causes an incorrect disk fit. Right: Cassini W00065133, a widefield image with insufficient pixels to identify the limb. Neither failure results in a false detection. Image Credit: NASA/JPL/Space Science Institute.

its contrast is improved by manual postprocessing so we exclude it from our performance analysis. Image C2066854 contains multiple plumes which are successfully detected as independent events. Image 5147r succeeds despite significant noise and image artifacts.

Figure 3 shows typical failure cases. Errors in estimating the planetary disk can occur if image quality is poor or the limb consists of few pixels. Poorly-illuminated plumes occasionally fail to exceed the 5σ detection threshold. In neither case does the failure generate a “false positive” detection. The system could still recover and detect the plume on a subsequent image of the sequence.

2. Spatial Coverage

The proposed detection approach is only sensitive to a subset of the potential plume activity on the surface. Detected plumes must extend beyond the planetary horizon, meaning they must reach a minimum altitude before becoming visible. This section characterizes the fractional coverage of the planetary surface over which detection is possible, assuming that plumes must achieve a minimum apparent pixel distance beyond the image of the horizon. We derive the projected pixel altitude of the plume from known parameters such as focal length, planetary diameter, and the physical altitudes of plume and spacecraft.

We assume that the camera has been corrected for radial distortion so that we can treat it as a true linear device [12]. Without loss of generality we consider a single linear camera viewing a 2D scene. Figure 4 shows a schematic diagram of the camera and scene geometry. Point p_0 corresponds to the camera center. It is located at a distance d from the planetary surface and directly facing disk so that the camera z (depth) axis points directly from the camera center toward the planetary center. We represent the planet itself a circle for simplicity. A plume of altitude a and offset from centerline by an angle of θ_1 extends to a point $p_1 = (p_{1x}, p_{1z})$ where:

$$p_{1x} = (r + a) \sin \theta_1 \quad (4)$$

$$p_{1z} = r + d - (r + a) \cos \theta_1 \quad (5)$$

Its projection in the image plane for a focal length f is inversely proportional to the range [13]:

$$u_1 = f \frac{p_{1x}}{p_{1z}} = f \frac{(r + a) \sin \theta_1}{r + d - (r + a) \cos \theta_1} \quad (6)$$

The visible horizon corresponds to a point on the planetary surface with the largest-magnitude image projection. We have labeled this point p_2 . It lies on the tangent line from the camera center to the surface, so that its offset angle is given by $\theta_2 = \arccos(r/(r + d))$. These expressions allow us to relate pixel height in the image, spacecraft camera resolutions, focal lengths, and altitude to the height of the smallest detectable plume.

We model an example based on the Voyager 2 flyby of Io on June 1979 which included a 10-hour “volcano watch” image sequence. The Voyager spacecraft approached to within about 10^6 km of the moon, and imaged several small plumes on the surface. For this sequence Io’s apparent radius was approximately 170 pixels. [3] estimate the detected plumes to be approximately 70km high. This is consistent with our model, assuming the known Io planetary diameter and a range of 10^6 km. Figure 5 shows the apparent pixel altitude of hypothetical plumes that are 50km and 100km as a function of their longitudinal offset. In test analysis of the Io data the algorithm successfully detected plumes down to approximately 2 pixels in altitude. This suggests that for plumes of 50km and 100km height, an automated monitoring campaign using a similar trajectory and camera should be sensitive across about 15 and 30 degrees of longitude, respectively.

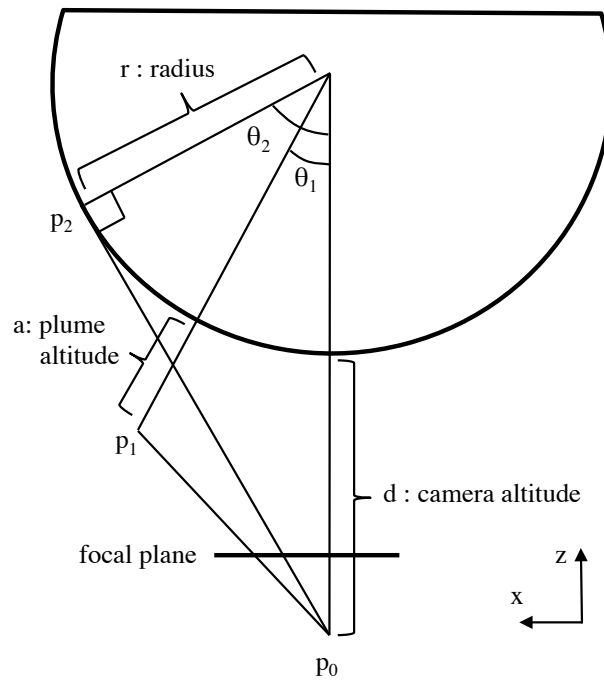


Figure 4: The projection of a plume reaching altitude a above the planet is determined by the relative sizes of plume and planet, as well as the distance and focal length of the camera. Sensitivity is determined by the apparent altitude in the image, that is, the difference in image projections between the plume tip p_1 and the planetary surface at the tangent line p_2 .




Figure 5: Apparent plume height in the image (in pixels) as a function of longitudinal offset on the planet’s surface. Unless noted otherwise, these curves use the geometry of Voyager 2’s Io observations on June 1979. The proposed algorithm regularly detects plumes at heights of 2-3 pixels in test data, suggesting sensitivity over at least 15 degrees of planetary longitude.

The solid line shows the projection that could be achieved with a closer approach to 2×10^5 km range. Note that the geometry of close approaches favors plumes on the moon’s camera-facing side, and this curve is shifted slightly to the left of the 90-degree offset.

Another factor affecting detectability (which we will not address here) is the apparent brightness of the plume material. This hinges on orbit-specific illumination conditions as well as unconstrained optical properties of the plume itself and the detector sensitivity of the instrument used.

3. Evaluation

We evaluated the retrieval and false positive rates of the plume detection algorithm on 19 sequences of single-band images of Io and Enceladus. The images were taken by Voyager, Galileo and Cassini spacecraft; each sequence contained one or more raw images that had been acquired in close succession. We limit our test set to images containing at least half of the planetary disk. A tour phase monitoring campaign would likely favor full-disk imagery

Sequence	Target	Images	Plume	Alt. (px)	Detected	#Detections	Notes
0085r	Io	1	A	1	N	0/1	
			B	3	Y	1/1	
2445r	Io	1	-	-	-	-	
3178r	Io	1	A	10	Y	1/1	
3200r	Io	3	A	9	Y	3/3	
3485r	Io	1	A	8	Y	1/1	
4204r	Io	2	-	-	-	-	
5045r-5123r	Io	3	A	9	Y	3/3	
5147	Io	1	A	14	Y	1/1	Artifacts
5407	Io	1	A	4	Y	1/1	
554x	Io	2	A	14	Y	2/2	Artifacts
6300r	Io	1	-	-	-	-	Artifacts
8XXr	Io	3	-	-	-	-	Plume ambiguous
C2064X	Io	2	-	-	-	-	
C2065X	Io	2	-	-	-	-	
C2066X	Io	18	A	3	Y	2/18	
			B	2	Y	2/18	
N00015-N00016X	Enc	3	A	6	Y	1/3	Faint plume
N148733X	Enc	19	A	>20	N	0/2	Low contrast
W0065133	Enc	2	A	7	N	0/2	$r < 20\text{px}$
W0065148	Enc	2	A	5	N	0/2	Artifacts
			B	5	N	0/2	
			C	5	N	0/2	

Table 1: Performance results. Columns show (Left to Right): a sequence identifier; target planetary body; the number of images in the sequence; a character identifying the plume event (there may be more than one physical plume per image sequence); the approximate altitude of the plume in pixels above the horizon based on visual inspection of the images; whether the plume was detected in at least one image; the number of images in which the plume was detected; and specific notes about image quality. There were no false detections in any of the 68 images.

to provide sensitivity across the entire limb, and a long horizon contour provides the best possible constraint on the position and curvature. The image set includes representatives from the Galileo *SSI* instrument, Voyager 2 and Cassini narrow- and wide-angle cameras. Many sequences consist of single plume events, but there are also longer sequences such as the *C2066X* series or the *N148733X* images of Enceladus. We ran the algorithm without alteration on all images using identical parameters.

Table 1 gives performance details. Several images in the test set did not contain plumes; we included these to test reliability against false positive detections. In fact, no false positives were generated with this dataset. All failures were “safe” in that missed plumes or incorrect horizon fits resulted in obvious failure conditions rather than mistaken detections. Of all sequences containing plumes, 76.9% generated at least one detection. Those sequences for which detection failed entirely generally suffered from artifacts or poor image quality (Figure 3). 11 distinct plumes were detected out of 17 total. These tests suggest that image quality and the pixel intensity of the plumes are primary constraints on plume detection performance.

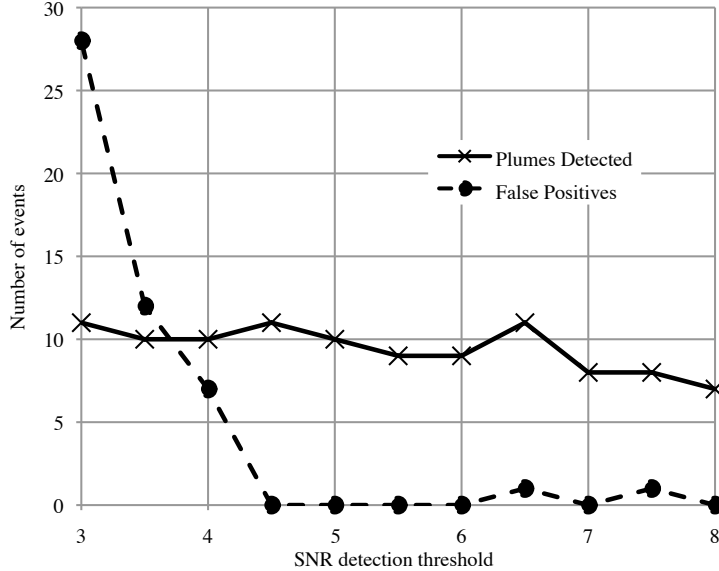


Figure 6: The fraction of plumes detected as a function of the SNR threshold τ . The tests in this work use $\tau = 5.0$ but more lenient settings may yield higher retrieval rates.

We found performance changed little with respect to edge detection thresholds or ellipse-scoring criteria. These parameters deal with horizon identification, a task that is straightforward for images with good contrast. After identifying the horizon, distinguishing weak plume features from image noise and planet glow is more challenging. This makes the detection intensity threshold the free parameter most affecting performance. Designers could tune this threshold to achieve desired system behavior. Figure 6 shows how accuracy and retrieval rates vary for different intensity thresholds; two curves indicate the total number of distinct plume events detected and the number of false positives returned. Due to the expense of transmitting data from spacecraft in an outer-planets mission we assume designers would prefer a low false positive rate. For this reason we favor a strict value of $\tau = 5.0$. However, if scientists wished to capture fainter events they could reduce this threshold at the risk of incurring false detections.



Figure 7: Left: Horizon models need not be ellipsoidal. Left: Nucleus of comet Hartley 2 as seen from the EPOXI flyby [4]. Center: Edge detection result. Right: Horizon identification using a convex hull, and subsequent detections of plume activity. Image credit: NASA/JPL-Caltech/UMD.

4. Discussion

We have described a robust system for detection of ejected plumes that could enable monitoring campaigns during characterization, flybys, and long tour phase operations when spacecraft telescopic imagers would otherwise be idle. Our basic strategy is to find the extent of the target body in the image and then mask it to reveal solar-illuminated plume material. These horizon-based approaches could provide a temporally continuous picture of plume activity, with significant benefit for our understanding of planetary volcanic processes. They are capable of detecting most plumes and return no false detections on our test set of plume images. A spacecraft could incorporate onboard ephemeris information for further improvements in both efficiency and detection accuracy. Higher-resolution cameras, that could simultaneously localize small features while imaging a large portion of the planetary disk, could also improve sensitivity to the smallest plumes. The horizon masking approach requires the ejected material to lie near the sunlit limb but it is possible to characterize the sensitive areas. Thus even limited monitoring coverage can still constrain plume activity rates. The algorithm should apply with very few changes to “partial disk” images where the horizon is nearly flat.

Plume detection software could be incorporated into a mission at little extra mass or power cost. The computation involved is generally tractable for spacecraft processors. The primary factors influencing computational cost are the number of RANSAC iterations and the edge contrast threshold

(controlling the number of edge points used by RANSAC). These values can be changed as needed to achieve the optimal sensitivity/time balance for specific operational requirements. Note that there may be other constraints to consider apart from computational resources, such as the limited lifetime of mechanical items such as doors and camera shutters.

The general horizon-masking approach could also support monitoring of outgassing from cometary nuclei. Simple variants to the horizon model can accommodate these non-ellipsoidal targets. Figure 7 shows one example from the EPOXI flyby of Comet Hartley 2 [4]. Edge detection finds surface contours as before, but instead of fitting an ellipse we use the enclosing polygon of edge points to estimate a conservative horizon. It is straightforward to compute these *convex hull* polygons [14]. This algorithm variant does not require RANSAC estimation and the resulting surface mask is effective for finding diffuse plumes. Future missions will augment the catalog of small bodies plume images, permitting more thorough performance studies for this class of targets.

Acknowledgments

We thank Brian Bue and Kiri Wagstaff for help with data, algorithms, and general advice. Richard Doyle and Tara Estlin provided additional ideas and counsel. This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration and funded through the internal Research and Technology Development program. Copyright 2011 California Institute of Technology. All Rights Reserved. US Government Support Acknowledged.

References

- [1] C. Porco, P. Helfenstein, P. Thomas, et al., Cassini observes the active south pole of Enceladus, *Science* 311 (2006) 1393–1401.
- [2] R. Lopes-Gautier, A. McEwen, W. Smythe, P. Geissler, L. Kamp, A. Davies, J. Spencer, L. Keszthelyi, R. Carlson, F. Leader, et al., Active volcanism on Io: Global distribution and variations in activity, *Icarus* 140 (1999) 243–264.
- [3] B. Smith, L. Soderblom, R. Beebe, J. Boyce, G. Briggs, M. Carr, S. Collins, A. Cook, et al., The Galilean satellites and Jupiter: Voyager 2 imaging science results, *Science* 206 (1979) 927.

- [4] M. A'Hearn, M. Belton, W. Delamere, L. Feaga, D. Hampton, J. Kissel, K. Klaasen, L. McFadden, K. Meech, H. Melosh, et al., Epoxi at comet hartley 2, *Science* 332 (2011) 1396.
- [5] A. Castaño, et al., Automatic detection of dust devils and clouds on mars, *Machine Vision and Applications* 19 (2008).
- [6] B. Bue, K. Wagstaff, R. Castaño, A. Davies, Automatic onboard detection of planetary volcanism from images, *Lunar and Planetary Sciences Conference Abstract* 1717 (2007).
- [7] J. Canny, *A Computational Approach to Edge Detection*.
- [8] A. Fitzgibbon, M. Pilu, R. Fisher, Direct least squares fitting of ellipses, *Pattern Analysis and Machine Intelligence* 21 (1999).
- [9] R. L. Rosin, A note on the least squares fitting of ellipses, *Pattern Recognition Letters* (1993) 799–808.
- [10] M. A. Fischler, R. C. Bolles, Random sample consensus, *Communications of the ACM* 24 (1981) 381–395.
- [11] D. Eberly, Distance from a point to an ellipse in 2d, *Magic Software* (Internet publication with modifications by Hui Ma) (2004).
- [12] R. Hartley, A. Zisserman, *Multiple view geometry in computer vision*, Cambridge Univ Pr, 2003.
- [13] D. Forsyth, J. Ponce, *Computer vision: a modern approach*, Prentice Hall Professional Technical Reference, 2002.
- [14] W. Press, *Numerical recipes: the art of scientific computing*, Cambridge Univ Pr, 2007.