Appears in Proc. AI in the Oceans and Space Workshop, International Joint Conference on Artificial Intelligence, Melbourne, Australia, August 2017.

Automatic detection of plumes from 67P/Churyumov-Gerasimenko in OSIRIS/Rosetta image sequences: A preliminary report

David Brown and **William Huffman** and **David R. Thompson** and **Steve Chien***
Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109, USA

Holger Sierks
Max-Planck-Institüt fur Sonnensystemforschung (MPS), Justus-von-Liebig-Weg 3, 37077 Göttingen, Germany

Abstract

We describe a computer vision based methodology for detecting plumes of the comet 67P/Churyumov-Gerasimenko from imagery acquired by the OSIRIS visible camera during escort. In our approach, we use a 3D model of the comet nucleus, Rosetta orbiter spacecraft position and pointing information, and instrument boresight information to formulate an initial guess at the outline of the nucleus in the image. We apply a series of image processing steps to handle the intensity variation between the sunlit and shadowed regions of the nucleus. We use an active contours method to fit an outline to the nucleus and dilate the outline to produce a range of constant altitude contours above the surface. In each of these 1D contours, we then detect plumes by finding regions of local intensity maxima. Plumes are indicated by a series of maxima in proximity at successively higher altitudes. This technique has utility in both ground-based analysis of plume sequences as well as onboard applications, such as isolating short sequences of high activity for priority downloading or triggering follow-up observations with additional instruments.

Introduction

Plumes and outbursts from small bodies in our solar system are of great science interest. The timing and makeup of such activity provides valuable insight into the processes, structure, and composition of the bodies under observation (Lin, Z.-Y. et al. 2015). Furthermore, understanding the properties of these small bodies - particularly comets, which are increasingly being shown to be remnants of the early solar system (Davidsson, B. J. R. et al. 2016) - enhances our understanding of the origin and evolution of the solar system.

The landmark Rosetta mission (Glassmeier, Boehnhardt, and Koschny et al 2007), which escorted the comet 67P/Churyumov-Gerasimenko (67P/C-G) from August 2014 through September 2016, provided a treasure trove of data to contribute to humanity's knowledge of comets and solar system formation. In particular, we are considering images from the OSIRIS visible camera (Keller, Barbieri, and Lamy et al 2007) that were designed to detect and track plume activity. These and other OSIRIS data have already revealed tremendous knowledge about the structure, composition, and activity of 67P/C-G (Sierks et al. 2015).

*Corresponding author. Email: steve.a.chien@jpl.nasa.gov

However, to date, analysis of plumes and outbursts on 67P/C-G has relied on having manually labeled plumes (Schmitt et al. 2017), precluding the possibility of doing such analysis automatically. In this article, we describe an automatic plume detection pipeline focused on linear plumes in extended duration image sequences that could address this limitation, enabling higher throughput small body plume analysis and new responsive, autonomous science efforts.

Motivation and Related Work

Agile Science

One particularly exciting application of plume detection and tracking is its potential onboard usage to enhance science in future missions to comets and other bodies. Onboard knowledge of plumes could enable applications ranging from data reduction to autonomous response (Chien 2016; Chien and Wagstaff 2017). Detection of plume activity could be used to triage data from extended duration plume tracking campaigns. As an example, software on the Mars Exploration Rover Opportunity (Castano et al. 2008) enables the rover to take long sequences of images, search for dust devils on the surface of Mars, and return only portions of the images with dust devils. In the plume case, extended duration plume campaigns could take long sequences of images and only return images containing plumes or even image portions containing plumes. Alternatively, the software could return statistics on the number of plumes detected, such as location, intensity and other summarizations.

With increased autonomous control, onboard analysis could trigger additional imaging. For example, upon detection of plumes, a higher cadence of images or additional images using different lens filters could be acquired. With even greater autonomy, other measurements could be targeted using narrow field of view instruments. For the Rosetta case, the narrow field of view MIRO submillimeter instrument (Gulkis, Frerking, and Crovisier et al 2007) could be used to image the plume source regions guided by detections. Alternatively, the MIRO instrument could trace the outline of the plume and across the plume at various heights. Given the dynamic, unpredictable nature of the plume and high pointing accuracy required, these types of interactive measurements are only possible with a high level of autonomy. This type of autonomy is very similar to that of the AEGIS software

currently operational on the Mars Exploration Rover (Estlin et al. 2012) and Mars Science Laboratory (Francis et al. 2017) in which the rover acquires wide field of view imagery to identify targets and then follows up with narrow field of view imagery (zoomed in and/or additional filter wheels or using the ChemCam laser).

Challenges

67P/C-G has a highly irregularly shaped body, which yields additional challenges over plume detection on other bodies: low fidelity edge tracing risks missing out on plumes near the surface of the body. Previous works on autonomous plume detection have made varying degrees of assumptions about the shape of the body of interest, with the earliest work from Wagstaff et al. working with spherically-shaped bodies (Wagstaff et al. 2006), Bue et al. assuming elliptically-shaped bodies (Bue et al. 2007), and more recently, Wagstaff et al.'s work on Hartley 2 assuming convex bodies (Wagstaff et al. 2014). Although these methods are applicable to a wide variety of solar system bodies, they are not designed to consider intricately shaped bodies like 67P/C-G, with crevices that would be excluded even using more advanced methods like convex hull. Additionally, even convex hull methods would miss out on short plumes that don't extend beyond the convex hull outline.

In addition to shape, the intensity variation between illuminated and dark parts of the comet is extreme - on the order of 100-fold - with the dark regions of the comet sometimes being even darker than the surrounding space background. Combined with the irregularly shaped body, these differentially illuminated regions of the comet can yield dark regions that appear to have outgassing activity because they are occluding illuminated regions behind them. Additionally, research on 67P/C-G using Rosetta's MIRO instrument has shown that illumination is only a partial contributor to outgassing activity, and residual outgassing activity does occur after sunset (Lee and von Allmen 2015). Thus, plume detection algorithms would benefit from being able to detect borders and classify plumes in both light and dark extremes, necessitating dynamic operations that do not depend on global properties of the image.

For this work, we focus on plumes that are set against a background of space as opposed to the 67P/C-G nucleus. Detecting plumes on the nucleus presents additional challenges due to the extremely low signal-to-noise ratio of plumes vs the body of the comet. While these plumes are also of interest, we leave these for future work.

Methods

In order to address the various challenges of plume detection on irregularly shaped and illuminated bodies, we developed a multi-phase plume detection pipeline. First, we use the known geography of 67P/C-G and the pointing information of OSIRIS to derive a nucleus mask estimate. Due to errors in the pointing and geometry information, the nucleus mask derived from the shape model and spacecraft kernels did not give the fidelity needed to reliably detect small plumes near the surface. Thus, we use an image processing pipeline to extract the edge outline of both the light and dark regions of the nucleus. The geometry-based mask is then used as a rough initial estimate of the outline, and this mask and image-based edge detections are fed into an active contours algorithm to generate a high-fidelity outline of the nucleus of 67P/C-G. Finally, this outline is used as a starting point from which the plume detection processes begin their search.

Nucleus Projection Mask

In order to generate an initial outline of the 67P/C-G nucleus for a particular image taken, we used the SHAP5 (Jorda et al. 2016) shape model developed by the OSIRIS team. This model consists of $\sim 3x10^6$ triangles specified as points offset from the estimated center of mass. In order to determine whether each pixel in a given OSIRIS image is expected to contain the nucleus, we project each pixel into the plane orthogonal to the boresight of OSIRIS at the time the image was taken, and look at whether a ray with its origin at OSIRIS and intersecting this projected pixel intersects with the shape model (Fig 1).

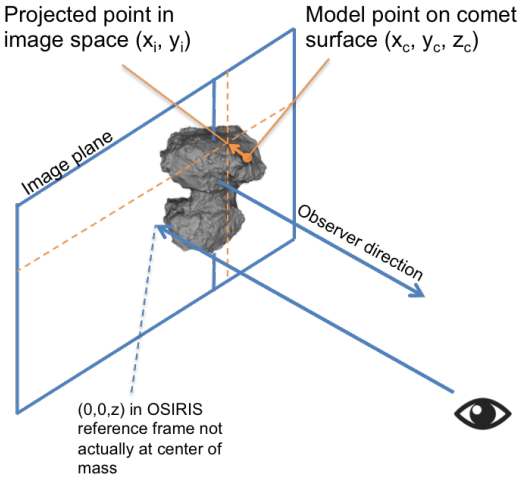

Figure 1: An example point projection onto the observer orthogonal plane.

Doing this individually for each pixel-triangle pair would require $\mathcal{O}(n * m)$ operations, where n is the number of pixels and m is the number of triangles, or on the order of 10^{13} operations for our shape model and a 2048x2048 image. Thus, we used two alternatives for generating the initial model estimate. The first was to project each triangle of the shape model into the orthogonal plane. While this does result in minor projection errors, these errors are negligible at the typical distances between Rosetta and 67P/C-G. For close-up images where projection error would be larger, an alternative approach was to downsample the nucleus mask estimate. Depending on computational capability available, the downsample factor can be chosen accordingly. Either of these methods, and a variety of downsample factors for the second method, are acceptable to pass along the pipeline, since the high-fidelity edge detection only requires a rough estimate of where the nucleus might be.

Image Processing for Edge Accentuation

Popular image segmentation techniques such as snakes (Kass, Witkin, and Terzopoulos 1988) and active contours without edges (Chan and Vese 2001) use a variety of image properties such as gradients and mean intensities to guide the energy functions that result in the optimal borders around single objects. However, regardless of the image statistics used, they rely on some regularity existing within a particular object. In our initial attempts at applying these algorithms to OSIRIS images of 67P/C-G, none of them could reliably detect both the dark and light regions of the nucleus as the same object. This is not surprising, as most applications of image segmentation are looking at everyday objects and scenes where the difference between light and shadow are not as extreme as they are in space applications. To accommodate this, devised an image processing pipeline that could accentuate the boundaries of the nucleus and reduce the difference between the light and dark regions, generating an image that is more amenable to image segmentation techniques.

Before proceeding with the dataset-agnostic image processing, we performed two dataset-specific corrections on the images based on known image corruption or image artifacts. First, in the pixel region [1646:2026, 38:57], there is a visible streak of very low intensity pixels relative to the surrounding space that could have been caused by an instrument defect or data corruption. We removed this image artifact by applying a 7x7 median filter to this region. We also discovered a radial artifact along the top and bottom of each image. In order to prevent these regions from biasing the intensity range, we trimmed all borders that contained a substantial portion (50%) of pure black pixels after normalization.

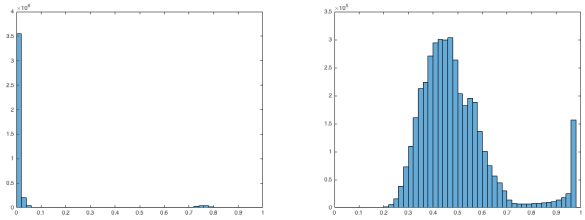

(a) Histogram of normalized intensities

(b) Histogram of log scale intensities

Figure 2: Before and after log scale transformation

Log scale Looking at the histogram of pixel intensities after removal of image artifacts and normalization (Fig 2a) reveals two distinct and narrow regions: 1. the low intensity pixels representing the dark parts of the comet and the surrounding space, and 2. the high intensity pixels representing the illuminated parts of the nucleus. Looking at the image directly at this point (Fig 3a) reveals an almost indistinguishable boundary between the dark regions of the comet and the surrounding space. In order to disambiguate the regions of the dark nucleus and the surrounding space, the differences

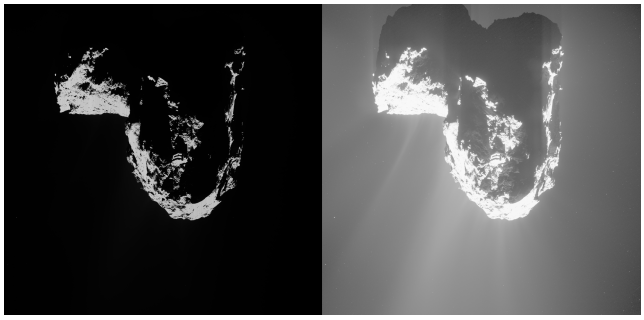

(a) Normalized image

(b) After log scale transformation

Figure 3: Before and after log scale transformation

in intensities in the darker regions needs to be exaggerated. To do this, we transformed the pixels into log scale by taking the $\log_e$ of each pixel value. For pixel values that are at 0, we set them to e^{-1} times the minimum non-zero pixel. After this log scale transformation, the histogram shows a much broader distribution of dark intensities (Fig 2b), and the borders become clearly distinguishable (Fig 3b).

Edge Accentuation At this point, the dark regions of the comet are more distinguishable from the surrounding space, but from our testing, the relative intensity difference is still not sufficient for reliable image segmentation. In order to accentuate the dark and light borders equally, we then applied a series of image processing steps:

1. Apply a 7x7 median filter
2. Apply a 5x5 relative variance filter (variance * (1-mean))
3. Apply a Sobel gradient magnitude filter
4. Combine the relative variance and gradient magnitude weighted 1:2
5. Apply a 7x7 median filter
6. Apply additional double Sobel gradient magnitude filter

Active Contours for Nucleus Tracing

Active contours is an image segmentation method that starts with an estimate of an image outline, and uses an energy function iteratively to determine where and how far to move each point in the contour, theoretically converging on the optimal outline of the object. Active contours was an attractive algorithm for segmenting the 67P/C-G nucleus, since we can be confident in the assumption that the nucleus is an isolated, continuous, object.

The active contour method we used was one developed by Chan and Vese, as implemented in MATLAB R2015B (Chan and Vese 2001). This active contour method is an improvement on previous active contour methods in that it works on images with a noisy background and less discrete edge transitions, a use case that coincides with our imaging environment. Since the inside of the nucleus can be a very noisy environment with many edge transitions between light and dark, we want to ensure that the first edge that active

contours encounters is the outermost edge. So, we start by dilating the nucleus mask by a number of pixels greater than our expected projection error. In our case, this was found experimentally to be 85 pixels. Then, we run the active contour method with a Smoothing Factor of 1.0 to avoid the contour fixating on some noisy outliers, and a contraction bias of 0 to prevent the contour from being pulled inside the nucleus in areas where there are very ambiguous boundaries. For our dataset, we found that active contours tends to converge at around 5000 iterations (Fig 4).

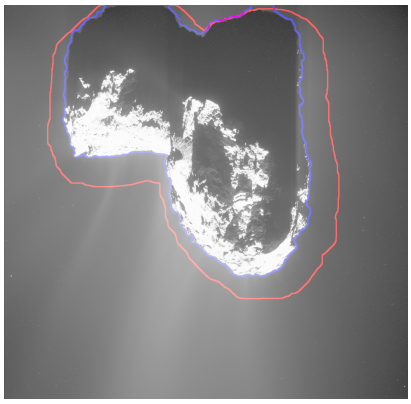

Figure 4: The dilated geometry-based projection (red) and the results of active contours after 5000 iterations (blue).

Plume Detection

Our plume detection technique is divided into three steps: search along altitudes above the nucleus for bright sections that may indicate plumes, join the bright sections across altitudes into two-dimensional plume regions, and fit one or more line segments to each plume region representing the plume (while filtering out implausible detections). The goal of this approach is to disambiguate between neighboring plumes early in the process by matching possible region components together, rather than producing a large number of detections immediately and then fitting consistent line segments across all detections.

Candidate Detections Along Altitudes We typically expect plumes to be linear features extending from the nucleus and somewhat near normal in orientation to the nucleus boundary, since the comet still has mostly convex parts, and plumes are emitted outwards from the comet body. Therefore, when tracing along an altitude above the nucleus boundary, we expect plumes to appear as a continuous segment of abnormally bright pixels surrounded by dark space. In practice, detecting these regions is complicated by high-frequency noise, cosmic rays, and image artifacts, as well as overlapping and intersecting plumes and systematic differences in brightness across the image. In order to properly ignore the high-frequency confusers, we use a narrow median filter to make detections. Similarly, we avoid low-frequency confusers by using a wide median filter to characterize the background signal along the altitude. A detection is therefore an area where the value of the narrow median filter reaches some threshold above the value of the wide median filter.

The sizes of the narrow and wide median filters are parameter choices that determine the scale of plumes the algorithm is looking (approximately: not more than twice as narrow as the narrow filter, and not larger than half the width of the wide filter). The threshold determines how aggressive the algorithm is in making detections and should be dependent on the amount of noise and non-plume-related variation in the image - we set this as a constant multiple of the standard deviation of the filter differences. While this naively produces a constant number of detections per altitude, these are later filtered on connectivity requirements within and across neighboring altitudes.

It is not necessary to use the same filter sizes across all altitudes in an image. Many plume regions begin as narrow and sharp features near the nucleus contour and transition into wider and more diffuse features far from the nucleus contour (this is not universally true, e.g. plumes whose sources are occluded by the nucleus may only be seen when they are already relatively diffuse). We found that using three different parameter regimes and using filter sizes increasing with distance from the nucleus contour worked well on a small number of test images, and therefore used this breakdown across our dataset. However, other choices of regimes (or continuously varying parameters, etc.) may be more suitable in other applications.

Plume Region Segmentation True plumes regions should not only appear continuous along an altitude, but also be continuous along the normal to the altitudes. In order to join the altitude detections into plume regions, we employed connected components segmentation. To determine connectivity, we defined a neighborhood region of a detection A. To begin, fit a line segment to the points in A (we used a simple linear regression, then projected all of the points in A along that line and used the smallest segment containing all of the projections). Use that line segment as a side of rectangle, with the other side length h, and then mirror that rectangle across the segment. The space covered by both rectangles is the neighborhood region of A. Then, a detection A and a detection B are connected if and only if their neighborhood regions overlap.

This definition of connectedness provides for a very aggressive matching process, and therefore must be traded with the aggressiveness of the detection threshold in the previous section to ensure suitable plume regions. This is a design choice; plumes which begin distinct near the nucleus can blend into a larger plume further away, in which case there may be different situations in which they should be considered a single plume or separate plumes.

Plume Segment Fitting With candidate plume regions identified they can now be fit by a line segment best describing the plume. An ideal fit would trace along the brightness maxima of the plume from its origin near the nucleus contour out to the visible extent of the plume. In order to generate such a fit, we use the RANSAC (Fischler and Bolles 1981) line-fitting algorithm to find the best line describing the points in the plume region, weighted by pixel intensity

of the points. By leveraging the finite range of the RANSAC fitting cost function, we can ensure that the shape of the region (especially near the extremities where it may be diffuse) does not bias the fit away from the brightness maxima. Further, we discard fits which are not close to orthogonal to the plume contour. While this may discard a few true plumes which originate at atypical angles, we have found that in our dataset this has a net benefit in avoiding local cost function minima.

Finally, we employ a small number of filters on the resultant plumes to decrease the number of false detections. We extend the plume detections to intersect with the nucleus contour, and any plumes which fail to intersect are deemed implausible and discarded. Further, we remove small plumes which are intersected by larger plumes, reasoning that these are detections of smaller pieces of the same plume.

Results

OSIRIS Image Dataset

Images for testing plume detection came from a set of images designated for dust monitoring by the OSIRIS team. The dataset contains 27 images taken by the OSIRIS wide-angle camera (WAC) on November 22, 2014 at 20-minute intervals over a 9-hour period. Images from OSIRIS WAC are 2048x2048 16-bit greyscale images and encoded using the NASA PDS .IMG file format. Images were calibrated and corrected via the OSIRIS calibration pipeline level 3 (Tubiana, C. et al. 2015).

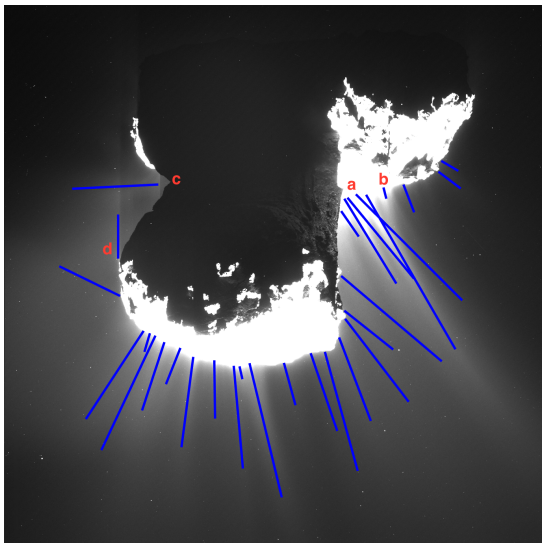

Figure 5: Example image of labeled plumes (blue) **a.** Plume detection successfully capturing plumes in a concave portion of the image. **b.** Short plumes close to the surface can be captured due to tightly-fitting nucleus contour. **c.** Plume doesn't extend to nucleus due to nucleus projection mask error and difficulty of active contours in tightly fitting dark regions. **d.** Illumination artifacts can still lead to false positives.

Statistics

For plume detection on the WAC dataset, we looked at three partially overlapping altitude regions, from 5-100, 10-200, and 10-500 pixels above the surface. For each of these regions, we found an average of 122 ± 18, 72.2 ± 14, and 165 ± 32 distinct clusters of potential plume candidate pixels. After applying connected components segmentation and fitting and plume segment extension and filtering, this resulted in a mean 23.5 ± 4.7 plumes per image, with an average plume length of 245 ± 156 pixels, which represents $1.89x10^3 \pm 1.4x10^3$ m if the plumes were projected in the plane of the image. These plume quantities are within the expected average plume counts of previously characterized plume sequences ($p = 0.20$) (Schmitt et al. 2017).

The average runtime for the nucleus projection on a modern laptop was 101 ± 3 s. It then took 703 ± 20 s for image processing and 5000 iterations of active contours. Finally, the average runtime for the plume detection was 196 ± 26 s.

Discussion

Our initial results show promise for detecting a diversity of plume shapes under variable lighting conditions. Qualitatively, the plume detection algorithm reliably finds the most prominent plumes, and captures many of the smaller plumes near the surface (Fig 5). However, further work will need to be done to characterize the precision and specificity of the plume detection algorithm by comparing to manually labeled plumes.

Despite the efforts of the image processing component of the pipeline to accentuate edges, there were still regions of the dark part of the comet body that the active contours failed to capture. There is a tradeoff between increasing true positives (capturing dark regions of the comet) and increasing false positives (capturing space as part of the comet), and there was no selection of parameters that yielded a perfect result. Using additional imaging modalities that are not subject to the limitations of albedo, such as microwave detection or depth sensing, could work to alleviate these errors. Other methods of improvement might include putting lower bounds on the active contours algorithm to discourage inward shrinking, using information about the expected area of the body as projected in the image frame, or separately classifying the dark and light regions of the comet and merging the results.

Additionally, the image processing and plume detection pipeline has difficulty dealing with different kinds of sensor noise, such as extremely grainy images or image artifacts. For this dataset, we had to apply custom artifact removal components based on finding the artifacts that exist across all images. Future versions of this work could use methods to find artifacts that are consistent across datasets and remove them, or use existing common image processing techniques to help with dealing with noise.

Although not yet optimized for onboard analysis, the flexibility in each of the components of the pipeline allows for a variety of potential analysis applications. As the active contours approach is an iterative algorithm, the number of iterations can be tailored to the requirements for a particu-

lar mission. For onboard analysis of novel, distant objects where a high fidelity shape model has not yet been established, more iterations could be done to get a higher resolution fit on the object body. Alternatively, for applications where a high fidelity shape model and pointing information exist, fewer iterations might be necessary. Similarly, the number of iterations can be modified to meet particular processing requirements, at the expense of contour fidelity.

Our multi-stage nucleus projection, plume detection, and plume correlation pipeline describes a robust and flexible method for finding plumes in a noisy, high-contrast environment. With additional work in reducing the runtime, generalization to work on additional body types and sensing modalities, and migration to on-board computing environments, this method could be an invaluable tool in science missions that seek to detect and respond to emissions in distant solar bodies where the light time to ground operations is greater than the expected lifetime of the events.

References

Bue, B.; Wagstaff, K.; Castano, R.; and Davies, A. 2007. Automatic onboard detection of planetary volcanism from images. *Lunar and Planetary Science XXXVIII*.

Castano, A.; Fukunaga, A.; Biesiadecki, J.; Neakrase, L.; Whelley, P.; Greeley, R.; Lemmon, M.; Castano, R.; and Chien, S. 2008. Automatic detection of dust devils and clouds on mars. *Machine Vision and Applications* 19(5):467–482.

Chan, T., and Vese, L. 2001. Active contours without edges. *IEEE Transactions on Image Processing* 10.

Chien, S., and Wagstaff, K. L. 2017. Robotic space exploration agents. *Science Robotics* 2(7).

Chien, S. 2016. Agile science flight software . a new paradigm for space missions. Workshop on Spacecraft Flight Software.

Davidsson, B. J. R.; Sierks, H.; Gttler, C.; Marzari, F.; Pajola, M.; Rickman, H.; AHearn, M. F.; Auger, A.-T.; El-Maarry, M. R.; Fornasier, S.; Gutirrez, P. J.; Keller, H. U.; Massironi, M.; Snodgrass, C.; Vincent, J.-B.; Barbieri, C.; Lamy, P. L.; Rodrigo, R.; Koschny, D.; Barucci, M. A.; Bertaux, J.-L.; Bertini, I.; Cremonese, G.; Da Deppo, V.; Debei, S.; De Cecco, M.; Feller, C.; Fulle, M.; Groussin, O.; Hviid, S. F.; Hfner, S.; Ip, W.-H.; Jorda, L.; Knollenberg, J.; Kovacs, G.; Kramm, J.-R.; Khrt, E.; Kppers, M.; La Forgia, F.; Lara, L. M.; Lazzarin, M.; Lopez Moreno, J. J.; Moissl-Fraund, R.; Mottola, S.; Naletto, G.; Oklay, N.; Thomas, N.; and Tubiana, C. 2016. The primordial nucleus of comet 67p/churyumov-gerasimenko. *Astronomy and Astrophysics* 592:A63.

Estlin, T.; Bornstein, B.; Gaines, D.; Anderson, R.; Thompson, D.; Burl, M.; and Castano, R. 2012. Aegis automated science targeting for the mer opportunity rover. *ACM Transactions on Intelligent Systems and Technology* 3:50.

Fischler, M. A., and Bolles, R. C. 1981. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM* 24(6):381–395.

Francis, R.; Estlin, T.; Doran, G.; Johnstone, S.; Gaines, D.; Verma, V.; Burl, M.; Frydenvang, J.; Montaño, S.; Wiens, R. C.; Schaffer, S.; Gasnault, O.; DeFlores, L.; Blaney, D.; and Bornstein, B. 2017. Aegis autonomous targeting for chemcam on mars science laboratory: Deployment and results of initial science team use. *Science Robotics* 2(7).

Glassmeier, K.; Boehnhardt, H.; and Koschny et al, D. 2007. The rosetta mission: Flying towards the origin of the solar system. *Space Science Reviews* 128:1–21.

Gulkis, S.; Frerking, M.; and Crovisier et al, J. 2007. Miro: Microwave instrument for rosetta orbiter. *Space Science Reviews* 128:561–597.

Jorda, L.; Gaskell, R.; Capanna, C.; Hviid, S.; Lamy, P.; urech, J.; Faury, G.; Groussin, O.; Gutirrez, P.; Jackman, C.; Keihm, S.; Keller, H.; Knollenberg, J.; Khrt, E.; Marchi, S.; Mottola, S.; Palmer, E.; Schloerb, F.; Sierks, H.; Vincent, J.-B.; AHearn, M.; Barbieri, C.; Rodrigo, R.; Koschny, D.; Rickman, H.; Barucci, M.; Bertaux, J.; Bertini, I.; Cremonese, G.; Deppo, V. D.; Davidsson, B.; Debei, S.; Cecco, M. D.; Fornasier, S.; Fulle, M.; Gttler, C.; Ip, W.-H.; Kramm, J.; Kppers, M.; Lara, L.; Lazzarin, M.; Moreno, J. L.; Marzari, F.; Naletto, G.; Oklay, N.; Thomas, N.; Tubiana, C.; and Wenzel, K.-P. 2016. The global shape, density and rotation of comet 67p/churyumov-gerasimenko from preperihelion rosetta/osiris observations. *Icarus* 277:257 – 278.

Kass, M.; Witkin, A.; and Terzopoulos, D. 1988. Snakes: Active contour models. *International Journal of Computer Vision* 1:321–331.

Keller, H.; Barbieri, C.; and Lamy et al, P. 2007. Osiris - the scientific camera system onboard rosetta. *Space Science Reviews* 128:433–506.

Lee, S., and von Allmen, P. 2015. Spatial and diurnal variation of water outgassing on comet 67p/churyumov-gerasimenko observed from rosetta/miro in august 2014. *Astronomy and Astrophysics*.

Lin, Z.-Y.; Ip, W.-H.; Lai, I.-L.; Lee, J.-C.; Vincent, J.-B.; Lara, L. M.; Bodewits, D.; Sierks, H.; Barbieri, C.; Lamy, P. L.; Rodrigo, R.; Koschny, D.; Rickman, H.; Keller, H. U.; Agarwal, J.; AHearn, M. F.; Barucci, M. A.; Bertaux, J.-L.; Bertini, I.; Cremonese, G.; Da Deppo, V.; Davidsson, B.; Debei, S.; De Cecco, M.; Fornasier, S.; Fulle, M.; Groussin, O.; Gutirrez, P. J.; Gttler, C.; Hviid, S. F.; Jorda, L.; Knollenberg, J.; Kovacs, G.; Kramm, J.-R.; Khrt, E.; Kppers, M.; La Forgia, F.; Lazzarin, M.; Lpez-Moreno, J. J.; Lowry, S.; Marzari, F.; Michalik, H.; Mottola, S.; Naletto, G.; Oklay, N.; Pajola, M.; Roek, A.; Thomas, N.; Liao, Y.; and Tubiana, C. 2015. Morphology and dynamics of the jets of comet 67p/churyumov-gerasimenko: Early-phase development. *Astronomy and Astrophysics* 583:A11.

Schmitt, M.; Tubiana, C.; Gttler, C.; Sierks, H.; Vincent, J.-B.; El-Maarry, M. R.; Bodewits, D.; Mottola, S.; Fornasier, S.; Hofmann, M.; Barbieri, C.; Lamy, P. L.; Rodrigo, R.; Koschny, D.; Rickman, H.; AHearn, M. F.; Agarwal, J.; Barucci, M. A.; Bertaux, J.-L.; Bertini, I.; Cremonese, G.; Da Deppo, V.; Davidsson, B.; Debei, S.; De Cecco, M.; Deller, J.; Fulle, M.; Gicquel, A.; Groussin, O.; Gutirrez,

P. J.; Hviid, S. F.; Ip, W.-H.; Jorda, L.; Keller, H. U.; Knollenberg, J.; Kramm, J. R.; Khrt, E.; Kppers, M.; Lara, L. M.; Lazzarin, M.; Lpez-Moreno, J. J.; Marzari, F.; Naletto, G.; Oklay, N.; Pajola, M.; Prasanna, D.; Shi, X.; Scholten, F.; Toth, I.; and Thomas, N. 2017. Long-term monitoring of comet 67p/churyumov-gerasimenkos jets with osiris onboard rosetta. *Monthly Notices of the Royal Astronomical Society*.

Sierks, H.; Barbieri, C.; Lamy, P. L.; Rodrigo, R.; Koschny, D.; Rickman, H.; Keller, H. U.; Agarwal, J.; A'Hearn, M. F.; Angrilli, F.; Auger, A.-T.; Barucci, M. A.; Bertaux, J.-L.; Bertini, I.; Besse, S.; Bodewits, D.; Capanna, C.; Cremonese, G.; Da Deppo, V.; Davidsson, B.; Debei, S.; De Cecco, M.; Ferri, F.; Fornasier, S.; Fulle, M.; Gaskell, R.; Giacomini, L.; Groussin, O.; Gutierrez-Marques, P.; Gutiérrez, P. J.; Güttler, C.; Hoekzema, N.; Hviid, S. F.; Ip, W.-H.; Jorda, L.; Knollenberg, J.; Kovacs, G.; Kramm, J. R.; Kührt, E.; Küppers, M.; La Forgia, F.; Lara, L. M.; Lazzarin, M.; Leyrat, C.; Lopez Moreno, J. J.; Magrin, S.; Marchi, S.; Marzari, F.; Massironi, M.; Michalik, H.; Moissl, R.; Mottola, S.; Naletto, G.; Oklay, N.; Pajola, M.; Pertile, M.; Preusker, F.; Sabau, L.; Scholten, F.; Snodgrass, C.; Thomas, N.; Tubiana, C.; Vincent, J.-B.; Wenzel, K.-P.; Zaccariotto, M.; and Pätzold, M. 2015. On the nucleus structure and activity of comet 67p/churyumov-gerasimenko. *Science* 347(6220).

Tubiana, C.; Gttler, C.; Kovacs, G.; Bertini, I.; Bodewits, D.; Fornasier, S.; Lara, L.; La Forgia, F.; Magrin, S.; Pajola, M.; Sierks, H.; Barbieri, C.; Lamy, P. L.; Rodrigo, R.; Koschny, D.; Rickman, H.; Keller, H. U.; Agarwal, J.; AHearn, M. F.; Barucci, M. A.; Bertaux, J.-L.; Besse, S.; Boudreault, S.; Cremonese, G.; Da Deppo, V.; Davidsson, B.; Debei, S.; De Cecco, M.; El-Maarry, M. R.; Fulle, M.; Groussin, O.; Gutirrez-Marques, P.; Gutirrez, P. J.; Hoekzema, N.; Hofmann, M.; Hviid, S. F.; Ip, W.-H.; Jorda, L.; Knollenberg, J.; Kramm, J.-R.; Khrt, E.; Kppers, M.; Lazzarin, M.; Lopez Moreno, J. J.; Marzari, F.; Massironi, M.; Michalik, H.; Moissl, R.; Naletto, G.; Oklay, N.; Scholten, F.; Shi, X.; Thomas, N.; and Vincent, J.-B. 2015. Scientific assessment of the quality of osiris images. *Astronomy and Astrophysics* 583:A46.

Wagstaff, K.; Castano, R.; Davies, A.; and Bue, B. 2006. Automating the detection of enceladus-style plumes. *Bulletin of the American Astronomical Society*.

Wagstaff, K. L.; Thompson, D. R.; Bue, B. D.; and Fuchs, T. J. 2014. Autonomous real-time detection of plumes and jets from moons and comets. *The Astrophysical Journal* 794(1):43.

Acknowledgments: This work was conducted, in part, at the Jet Propulsion Laboratory, California Institute of Technology, under a contract from the National Aeronautics and Space Administration.