

AUTOMATIC PLUME DETECTION FOR THE EUROPA IMAGING SYSTEM. G. Doran, K. L. Wagstaff, M. Cameron, I. Daubar, C. Phillips, Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA, 91109-8099, USA (gary.doran@jpl.nasa.gov).

Introduction: A key science goal of the Europa Clipper mission [1], which will make over 40 flybys of Jupiter’s moon Europa starting in the late 2020s, is to determine whether there any ongoing geologic processes such as plumes that release material from Europa’s subsurface ocean. Studying fresh material from Europa’s ocean can provide insights into its habitability both by Europa Clipper and potential future surface missions. To search for plumes, Europa Clipper will use its suite of instruments, including the Europa Imaging System (EIS) [2], which will look for direct visible evidence of plumes above the limb of the body illuminated by sunlight. Because the spacecraft will be so far from Earth, downlink constraints will limit the number of images that can be sent back to scientists, and the rate at which they can be transmitted. The ability to *automatically detect plumes within images onboard the spacecraft* could enable selective or reprioritized downlink so that the most scientifically relevant observations arrive back on Earth first to enable tactical planning of follow-up observations during subsequent flybys. We are exploring algorithms for automatic plume detection, which we describe here along with results on labeled analogue images.

Analogue Data: Active plumes have previously been observed on moons and other planetary bodies such as comets. Moons such as Io and Enceladus are more relevant to our study than highly non-spherical comet nuclei. Therefore, we include observations of these bodies with a number of different instruments: the Galileo Solid-State Imager (SSI), the Cassini Imaging Science Subsystem (ISS), and the New Horizons Long Range Reconnaissance Imager (LORRI). In addition to observations of bodies with active plumes, we include observations of Mercury by the MESSENGER spacecraft using the Mercury Dual Imaging System (MDIS), which is similar to EIS in design. Finally, we include observations of Europa

Table 1: The total number of images in our dataset of each body observed by the instruments listed, followed in parentheses by the number of images containing plumes.

Instrument	Body	Images (w/Plumes)
SSI	Io	19 (12)
	Europa	22 (0)
	Ganymede	13 (0)
	Callisto	7 (0)
LORRI	Io	84 (60)
MDIS	Mercury	123 (0)
ISS	Enceladus	40 (40)
Total:		308 (112)

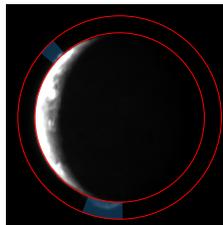


Figure 1: An example showing labeled plumes (light blue shaded regions) in the LORRI observation LOR_0035092817_0X630 of Io. The limb and annulus boundary are shown in red.

and other icy moons (Ganymede and Callisto) from SSI. Table 1 summarizes the data that we used in our study.

To provide ground truth for any algorithm that attempts to automatically detect plumes in the observations described above, we manually annotated each image in our dataset. Our annotations consist of both limb and plume labels, as shown in Figure 1. The limb is specified by a set of points in the image to which a circle is then fit. After the limb of the body is fit, the presence of a plume is indicated by labeling any portion of the annulus surrounding the body in which a plume is visible. We have made our labels available online,¹ and all of the images are available via the Planetary Data System (PDS).

Approach: To automatically detect plumes, we base our approach on prior work for onboard plume detection [3]. Whereas prior work investigated plume detection for potentially non-spherical bodies such as comets, we use a simplified algorithm for the imaging targets in this data set, which are approximately spherical. The algorithm begins by finding the high-contrast limb of the body using a Canny edge detector [4]. Then, a circle is fit to these edges to model the limb using Random Sample Consensus (RANSAC) [5], which ignores potential spurious edges detected due to hot pixels, stars, or the terminator. After the outline of the body is found, a plume search is conducted by looking within a ring-shaped, annular region surrounding the body from 101% to 120% of the body’s estimated radius (see Figure 2, left).

The plume search begins by averaging the pixel intensity, measured in raw digital numbers (DNs), within each of 1024 radial "sectors" of the annulus. Any sector not contained within the image (if part of the body falls outside the field of view) is ignored. Statistics are then computed using the pixel intensities of the remaining sectors to search for outliers. We use the inter-quartile range (IQR) as a robust measure of the typical variance due to noise in the intensities across sectors. Then, we detect plume candidates as those sectors with average intensity values that fall above the median sector intensity plus a multiple of the IQR. Thus, this approach uses an adaptive

¹<https://zenodo.org/record/2556063>

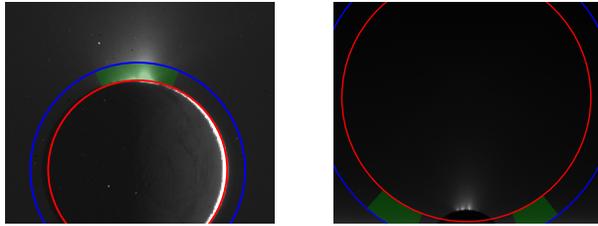


Figure 2: (left) An example limb and plume detection in observation N1635814521 of Enceladus by Cassini ISS. The detected limb is shown in red, with the search annulus edge shown in blue. Green indicates a region where a plume was detected. (right) An example limb-finding failure in a high phase angle observation (N1635781564) leading to false plume detections.

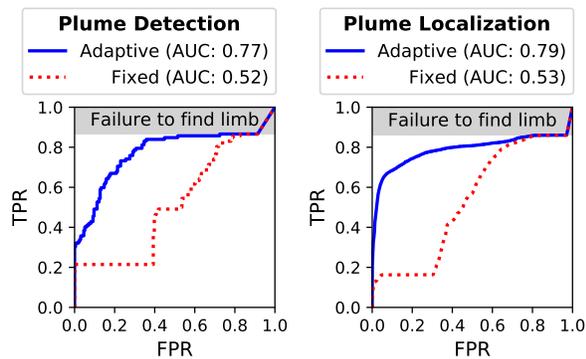


Figure 3: Receiver Operating Characteristic (ROC) curves comparing the performance of the adaptive- and fixed-threshold algorithms across all images, with area under the curve (AUC) shown in the legend. The plume detection results show performance on a binary plume/no plume decision for each image, whereas the plume localization results show performance on making a separate decision for each visible annulus sector.

threshold based on the observed IQR of the data. As a baseline, we also investigated a variant of this algorithm that instead uses a *fixed* threshold on the average pixel intensity within each sector.

Results: To evaluate the plume-finding algorithm, we ran the adaptive and fixed threshold approaches on the 308 images in our dataset. Instead of selecting a single threshold for each algorithm, we use receiver operating characteristic (ROC) curves to study how classification performance in terms of true positive rate (TPR) and false positive rate (FPR) varies as the threshold sweeps from high to low values, capturing fewer true positives but also generating fewer false positive detections (see Figure 3). We separately consider the problems of “plume detection,” deciding whether an image contains a plume or not, and “plume localization,” determining which of the visible annulus sectors contain plumes.

As an initial observation, we see that both curves are shifted down somewhat from the top of the graph (as

indicated by the gray shaded region), only rising to the top-right corner near the end of the curve. For a perfect classifier, ROC curves would rise as quickly as possible to a TPR of 1.0, producing an area under the curve (AUC) of 1.0. In this case, the reason for the downward shift is due to the set of images for which the limb-finding step fails, for which no plumes can be detected. The limb-finding fails if RANSAC cannot find a solution, or the solution it finds has a body radius that deviates by more than 50% from the expected value (hypothetically generated using telemetry predicts). Common causes of failure include high phase angle observations (Figure 2, right) and missing data that produces high-contrast regions within the body. Plumes are only detected in these cases if the threshold is at the lowest possible setting, so plumes are always reported (TPR and FPR of 1.0).

The results also show that, as expected, the adaptive approach significantly outperforms the fixed threshold. Because DNs can vary significantly across images due to difference in exposure times and instrument characteristics, using a single fixed threshold produces AUCs that are close to randomly guessing (0.5). The adaptive approach performs better at the localization task than at the detection task. This behavior is likely due to the fact that a single spurious plume detection leads to an incorrect classification for an entire image, so a higher TPR is only achieved at the expense of a higher FPR. Nonetheless, the AUC values indicate that the algorithm has a 77% chance of assigning a higher plume detection score to an image with a plume than to one without, so it can effectively prioritize images for downlink better than chance.

Conclusions: We have demonstrated an effective approach for detecting plumes within analogues of the images to be acquired with the EIS instrument onboard Europa Clipper. The algorithm is designed to run onboard the spacecraft to prioritize downlink, but it can also be deployed as part of ground processing to prioritize scientists’ focus of attention as soon as data reaches Earth. We are currently investigating similar approaches to detect anomalies or interesting phenomena within other Europa Clipper instruments [6, 7].

Acknowledgments: The authors thank Serina Diniega and Ashley Davies for providing feedback on this work, and the JPL R&TD program for funding support. This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. ©2019 California Institute of Technology. U.S. Government sponsorship acknowledged.

References: [1] Pappalardo, R.T. *et al.* (2017) *EPSC*. [2] Turtle, E. *et al.* (2018) *COSPAR*. [3] K.L. *et al.* (2014) *Astrophysical J.* [4] Canny, J. (1986) *TPAMI*. [5] Fischler, M.A. *et al.* (1981) *Comm. ACM*. [6] Doran, G. *et al.* (2018) *EPSC*. [7] Wagstaff, K.L. *et al.* (2019) *LPSC*.