Enabling Onboard Detection of Events of Scientific Interest for the Europa Clipper Spacecraft

Kiri L. Wagstaff
Gary Doran
Ashley Davies
kiri.l.wagstaff@jpl.caltech.edu
gary.b.doran.jr@jpl.caltech.edu
ashley.g.davies@jpl.caltech.edu
Jet Propulsion Laboratory,
California Institute of Technology
Pasadena, CA, USA

Saadat Anwar
Srija Chakraborty
sanwar@asu.edu
schakr34@asu.edu
Arizona State University
Tempe, AZ, USA

Marissa Cameron
Ingrid Daubar
Cynthia Phillips
marissa.e.cameron@jpl.caltech.edu
ingrid.daubar@jpl.caltech.edu
cynthia.b.phillips@jpl.caltech.edu
Jet Propulsion Laboratory,
California Institute of Technology
Pasadena, CA, USA

ABSTRACT

Data analysis and machine learning methods have great potential to aid in planetary exploration. Spacecraft often operate at great distances from the Earth, and the ability to autonomously detect features of interest onboard can enable content-sensitive downlink prioritization to increase mission science return. We describe algorithms that we designed to assist in three specific scientific investigations to be conducted during flybys of Jupiter's moon Europa: the detection of thermal anomalies, compositional anomalies, and plumes of icy matter from Europa's subsurface ocean. We also share the unique constraints imposed by the onboard computing environment and several lessons learned in our collaboration with planetary scientists and mission designers.

KEYWORDS

onboard data analysis, anomaly detection, image analysis, space exploration, resource-constrained computing

ACM Reference Format:
Kiri L. Wagstaff, Gary Doran, Ashley Davies, Saadat Anwar, Srija Chakraborty, Marissa Cameron, Ingrid Daubar, and Cynthia Phillips. 2019. Enabling Onboard Detection of Events of Scientific Interest for the Europa Clipper Spacecraft. In *The 25th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '19), August 4–8, 2019, Anchorage, AK, USA.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3292500.3330656

1 INTRODUCTION

Spacecraft operating at large distances from the Earth must conduct several activities in an entirely autonomous fashion, including navigation, data collection, and communication, due to the time it takes for commands to reach the spacecraft. A signal sent from the Earth takes up to 25 minutes to reach Mars and up to an hour to reach Jupiter. Therefore, spacecraft activities are typically transmitted in a pre-planned script rather than controlled through real-time commanding. This approach is very effective when the spacecraft position, environment, and other factors are known or can be predicted in advance. However, it precludes adaptation to unanticipated conditions and opportunities.

Our goal is to support planetary exploration by remote spacecraft using onboard data analysis methods when direct human oversight is not possible. This is particularly important when studying rare, transient events whose occurrence, location, and duration cannot be predicted. Examples include dust devils on Mars [16], meteorite impacts [10], lightning on Saturn [13], and icy plumes emitted by Enceladus and other bodies [17]. To enable fast detection and response to dynamic events, Chien et al. proposed an "agile science" approach to space exploration that combines onboard data analysis with a resource-aware planner to modify spacecraft activities in response to changing conditions and detections [7]. More generally, onboard data analysis enables responsive decision making such as how to prioritize observations based on their content and what follow-up observations to collect. These capabilities become progressively more valuable as spacecraft range to more distant destinations with correspondingly longer communication delays and smaller data downlink allocations.

In this work, we focus on the Europa Clipper mission, which plans to launch in the 2020's to explore, map, and characterize Jupiter's moon Europa [2, 21]. The mission's science objectives most relevant to this effort include the detection of thermal anomalies (hot spots), regions of unusual composition, and plumes, if any. The location, timing, and prevalence of these features is not known. In collaboration with Clipper mission and planetary scientists (including authors Davies, Cameron, Daubar, and Phillips), we have developed and evaluated data analysis methods to address each of these scientific priorities, subject to severe resource constraints. We describe these methods and their performance characteristics. The key lessons we learned about the constraints and opportunities involved in developing data analysis methods for operation onboard a spacecraft include:

- Simple methods with high interpretability and low resource consumption are most likely to be adopted by a mission.
- Collaboration with experts in spacecraft design and mission science objectives is vital to ensuring the effort is focused on meaningful problems with feasible implementations.
- Performance metrics should reflect audience priorities.
- Risk is an overriding concern in spacecraft development.

KDD '19, August 4–8, 2019, Anchorage, AK, USA

ACM ISBN 978-1-4503-6201-6/19/08…$15.00
https://doi.org/10.1145/3292500.3330656

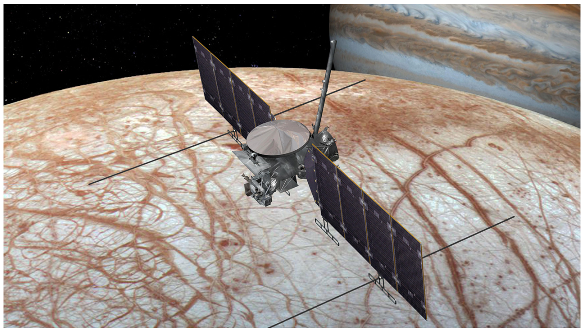

Figure 1: Europa Clipper spacecraft (artist's conception). Image credit: NASA/JPL-Caltech

2 MOTIVATION AND RELATED WORK

Jupiter's moon Europa, about the size of Earth's moon, is thought to contain an ocean of liquid water beneath its frozen, icy crust. Understanding the complex geophysical and geochemical processes on Europa is a high priority for NASA and the planetary science community. Europa is also a target of interest for astrobiology, as it may meet our current criteria for the existence of a biosphere: liquid water, the right chemical elements, and an energy source.

NASA's Europa Clipper mission (see Figure 1) has an overarching science goal to "Explore Europa to investigate its habitability" [21]. This goal is supported by three mission objectives: (1) Characterize the ice shell and any subsurface water; (2) Understand the habitability of Europa's ocean through composition and chemistry; and (3) Understand the formation of surface features, including sites of recent or current activity. The mission will employ a suite of nine science instruments to gather both remote sensing and in situ observations.

Europa Clipper will enter a highly elliptical orbit around Jupiter with a 14-day period punctuated by brief flybys of Europa to minimize the spacecraft's exposure to Jupiter's high radiation environment [12]. Most of the data will be collected during a four-hour period centered on the time of closest approach to Europa. The majority of each orbit (~10 days) will be devoted to gradually transmitting all of the data (~100 Gb per flyby) back to Earth. The instruments have the capability to collect far more data than can be transmitted; the onboard storage limit is 550 Gb. Prioritization of data to meet the downlink limits is vital to enable the planning process and to ensure that key scientific discoveries are not missed. We hypothesize that prioritizing downlink based on content-sensitive onboard data analysis will be more effective than the default priorities assigned during observation planning.

2.1 Challenges for Onboard Data Analysis

The high-radiation environment surrounding Europa creates a significant operational challenge. Radiation can cause the spacecraft CPU to reset, interrupting data collection or downlink activities. It can also corrupt data that is onboard the spacecraft waiting to be transmitted to Earth. Radiation-induced noise in the data can impact the quality of onboard data analysis [31]. Europa Clipper will therefore use a RAD750 processor, which is robust to radiation but operates at a maximum clock rate of 200 MHz, an order of magnitude slower than a typical desktop computer (or even a mobile phone). This severely limits the amount of onboard computing capability. Power constraints may impose further limits that preclude the use of otherwise idle periods.

A second important challenge is that data available onboard the spacecraft will be uncalibrated. In some cases, an approximate calibration can be performed [5], but it is necessary to demonstrate that the analysis algorithm performs sufficiently well on the pseudo-calibrated products. In other cases, it may be advantageous to develop an algorithm that is customized for operation directly on the uncalibrated data.

In addition to the technological and environmental challenges, there are also psychological and cultural factors that must be considered. Autonomous spacecraft analysis and decision making raises concerns about a corresponding increase in risk. The potential impact of an incorrect detection or flawed decision must be assessed and minimized. As much as possible, the onboard data analysis module must be isolated from the primary spacecraft activities so that it cannot interrupt basic health and safety operations.

2.2 Onboard Data Analysis Objectives and Related Work

Our goal is to develop data analysis methods that can run onboard the Europa Clipper spacecraft to detect and prioritize observations of high scientific value. Each observation will be assigned a default downlink priority prior to uplink. Onboard data analysis can decide whether to adjust the priority based on the observation's actual content. For example, an image that captures an erupting plume could receive a high priority to ensure that it is received quickly and can inform future targeting plans. An image that is dominated by radiation corruption or other artifacts could receive a low priority. Ultimately, this capability could enable more efficient exploration of Europa by (1) directing attention to data of most scientific value to inform the planning process and (2) increasing the "science return" of the mission (i.e., speed with which science objectives are met, crucial given the spacecraft's continual radiation bombardment).

In the onboard computing regime, with limited RAM and CPU resources, advanced machine learning methods such as deep convolutional networks cannot yet be employed due to their computational cost. Previous work with other spacecraft has likewise sought a balance between model performance and its computational cost or complexity. For example, the EO-1 Earth-orbiting spacecraft received a series of software updates that enabled it to detect volcanic eruptions [11] and to respond by automatically collecting follow-up images the next time the spacecraft passed over the same locations, or even to signal other spacecraft to observe the affected areas. The Swift spacecraft in Earth orbit detects gamma-ray bursts with one telescope and quickly points two other telescopes to observe the same source at x-ray and ultra-violet wavelengths [15]. The Mars Science Laboratory rover autonomously detects rock targets, ranks them according to current mission science priorities, and collects compositional spectra of the highest-priority targets using the ChemCam instrument, without any human intervention [14].

Several optimized machine learning models have been deployed for use onboard spacecraft. Support vector machines with linear and Gaussian kernels have been used onboard EO-1 to detect small sulfur deposits on glacial ice [19], a biosignature that could potentially be found on Europa. Random forests have been used in Earth orbit by a CubeSat [8] and EO-1 [30] to assess the fraction of an image that is covered by cloud, which can enable the prioritization of clear images ahead of those in which the surface is obscured.

Despite the use of machine learning onboard several spacecraft, in general onboard analysis and decision making is considered a novel, high-risk and high-payoff endeavor. Our target is the Europa Clipper spacecraft, which will travel much farther away than Earth or Mars orbit and for which each flyby will be unique. There may only be a single chance to capture a particular event or unusual surface feature, so content-sensitive prioritization could have a large scientific impact.

3 METHODS AND RESULTS

We have developed custom data analysis and event detection methods for three instruments on the Europa Clipper spacecraft. The instruments acquire different types of data and seek to achieve different scientific goals.

3.1 Thermal Anomaly Detection

Detecting thermal anomalies on the surface of Europa is essential for addressing a number of science goals set forth in the 2011 Planetary Science Decadal Survey [9] such as understanding cryovolcanism and plumes that can bring subsurface materials to the surface, enabling astrobiological studies of Europa's subsurface ocean [18, 27]. Previous studies of Europa's surface from spacecraft and ground-based telescopes have failed to definitively detect thermal anomalies caused by internal Europa activity (versus heating from the Sun). For example, the Photopolarimeter-Radiometer (PPR) instrument onboard the Galileo spacecraft did not detect any hot spots on Europa, but its detection sensitivity could permit anomalies up to 100 km^2 to go undetected [23]. Similarly, follow-up studies using the ground-based Atacama Large Millimeter Array (ALMA) observatory found potential anomalies in Europa's thermal inertia, but no definitive Europa heat sources [28].

The Europa Thermal Emission Imaging System (E-THEMIS) instrument on the Europa Clipper mission is designed to detect thermal anomalies on the surface of Europa. These anomalies could include geologically recent surface and sub-surface activities such as plumes, active vents, and resurfacing of material. Thermal characterization of the surface will also support landing site safety assessment for future Europa missions. E-THEMIS uses an uncooled microbolometer detector in a radiation-hardened design to withstand Jupiter's harsh environment. Filters attached to the microbolometer allow the instrument to detect thermal infra-red radiation in three discrete spectral bands: 7 to 14 μm, 14 to 28 μm, and 28 to 70 μm. E-THEMIS will collect Europa data at much higher spatial resolution than PPR.

3.1.1 Thermal Anomaly Detection Method. Because E-THEMIS is similar to the Thermal Emission Imaging System (THEMIS) instrument flown on the Mars Odyssey spacecraft, we build on an existing algorithm designed to detect thermal anomalies within Mars THEMIS data [5]. The algorithm operates by applying a rudimentary calibration to convert digital number (DN) values available onboard the spacecraft into approximate temperatures:

$$T = \alpha \ln \left[(\mathrm{DN} - o \times g) \times g \right] - \beta, \tag{1}$$

where o and g are the instrument offset and gain, and $\alpha = 101.85$ and $\beta = 505.69$ are empirically derived constants. Then, any pixel in the observation exceeding a pre-defined threshold temperature is flagged as an anomaly. Appropriate threshold temperatures will be much lower for Europa than for Mars. This threshold-based method will serve as a baseline for studying future approaches, such as incorporating spatial context around each pixel.

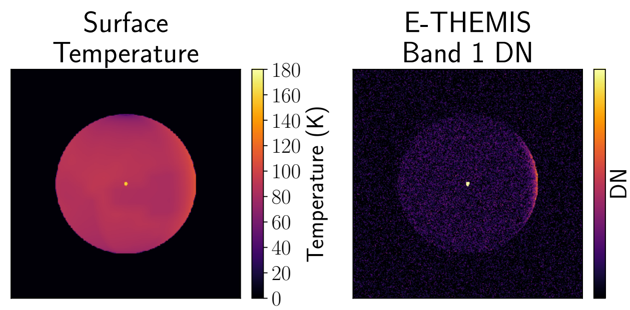

Figure 2: Example Europa thermal anomaly (temperature 160 K, area 5,000 km^2) as viewed by E-THEMIS at a distance of 50,000 km. Simulated surface temperature is converted into the digital number (DN) recorded by the instrument.

3.1.2 Results using Simulated E-THEMIS Data. No E-THEMIS data yet exists, so we generated 5000 simulated observations to characterize algorithm performance[1]. Each observation was generated by randomly selecting a spacecraft location (±50,000 km from closest approach to Europa) on one of 46 planned flybys. We used a thermal model to simulate the Europa surface temperature as a function of latitude, longitude, and local solar time (Sun position). We simulated the appearance of Europa within the E-THEMIS field of view given Europa's radius of 1591 km. We used ray tracing to determine, given the surface temperature model and the relative positions of Europa and the spacecraft, the amount of thermal radiation from Europa that a given pixel would collect. We used Monte Carlo integration over 100 uniformly random sub-pixel locations within each pixel.

We injected artificial thermal anomalies into the simulated Europa observations by selecting a random anomaly radius between 1 m and 25 km and random temperature between 130 K and 275 K, then modifying the modeled surface temperature (Figure 2, left). We added Gaussian noise with $\sigma = 1$ DN to each pixel, then used a lookup table derived from the E-THEMIS instrument model (included in the dataset online) to convert temperature to digital numbers and generate the final data product (Figure 2, right). The 5000 observations and 100 samples within each pixel provide sufficient coverage of the randomly varied parameters and allow resolving the pixel anomaly fraction to within 1%.

Figure 3 shows detection rates for anomalies of various sizes (in terms of the fraction of the pixel filled) and temperatures using a threshold of 140 K. This threshold is sufficiently high to guarantee

[1]Simulated E-THEMIS data set available at http://doi.org/10.5281/zenodo.2552814.

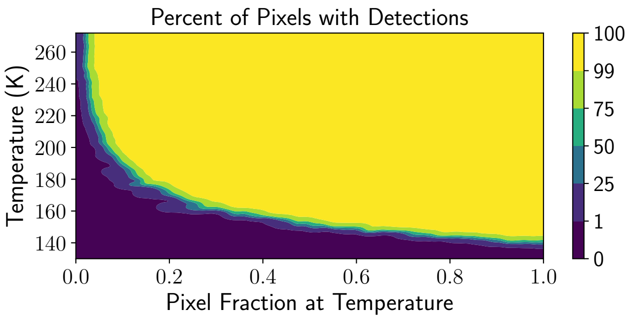

Figure 3: Thermal anomaly detection rates in simulated E-THEMIS Band 1 data with a threshold of 140 K, as a function of anomaly pixel fraction (size) and temperature.

that all detections are true positives, enabling a focus on recall. In this case, only Band 1 of the three-band simulated E-THEMIS data is used because it is the most sensitive to anomalies above temperatures of 140 K. The results show that, for example, anomalies comparable to the 170 K "tiger stripe" features detected by Cassini on Enceladus [22] could be detected by this approach even if they filled only ~20% of a pixel. At a 50,000 km offset, this translates to a 40 km^2 anomaly on the surface, with even smaller anomalies detectable at closer range, all of which are smaller and cooler than features detectable in PPR data. Thus, even this simple threshold-based approach has the potential to detect previously unknown thermal anomalies on the surface of Europa.

Future work will explore how performance can be improved with more sophisticated detection methods, and will improve the fidelity of the simulation, which currently does not include emission angle effects or radiation-induced, non-Gaussian noise. The use of simple, circular anomalies is sufficient for evaluating this threshold-based pixel-wise detection algorithm, but evaluating algorithms that take more context into account will require morphologically more realistic anomalies.

3.2 Compositional Anomaly Detection

The Europa surface is known to be composed primarily of water ice, and any significant anomalies can provide clues to the evolution of the surface and the types of surface alteration processes. These processes include modification of the surface due to cosmic ray bombardment, tidal heating from Jupiter, and the formation of crustal ridges and bands that may reveal material brought up from the subsurface ocean.

The Mapping Imaging Spectrometer for Europa (MISE) is a Dyson imaging spectrometer that collects three-dimensional "cubes" with a spatial dimension of 320 by 300 pixels and 480 spectral bands (0.8 to 5.0 μm) [3]. MISE will collect compositional observations of Europa with a spatial resolution of 25 m per pixel at a range of 100 km [3]. The peaks and troughs in each pixel's spectrum allow the identification of organics, salts, and radiation-altered materials. Fine-grained knowledge of the ice composition can inform our understanding of the habitability of Europa's ocean. MISE will also investigate the geological history of Europa's surface and search for areas that are actively being resurfaced. MISE is sensitive to different water ice structures that are indicators of age [4]. Fresh surface deposits can provide a window into Europa's interior.

MISE data products are large (527 Mb each). It is expected that only three to eight MISE cubes can be downlinked per flyby, so they must be carefully chosen. Anomaly detection can help ensure that materials with rare or unexpected composition are marked for priority downlink.

3.2.1 Compositional Anomaly Detection Methods. Spectral anomaly detection in hyperspectral images has been studied extensively. In this study, we employed two commonly used anomaly detection methods, one batch and one iterative.

The first algorithm is the widely used Reed Xiaoli spectral anomaly detector (RX) [6, 24]. RX builds a global model of the background spectral distribution from the spectral mean and covariance of the pixels in a given scene. RX computes an anomaly score for each pixel with respect to this model. For a hyperspectral image X with d spectral bands, the RX anomaly score A_{RX} of pixel x_i is

$$A_{RX}(x_i) = (x_i - \mu)^T \Sigma^{-1} (x_i - \mu), \tag{2}$$

where $x_i \in R^d$ is the spectrum of pixel x_i, μ is the mean spectrum across all pixels in X, and Σ is the background covariance matrix computed from X across all d bands. The inverse covariance matrix in equation (2) projects the spectra x_i along the components corresponding to the smallest eigenvalues [6]. A high A_{RX} value indicates that x_i is anomalous with respect to the current scene.

The second algorithm is DEMUD, an anomaly detection method that seeks to minimize redundancy in the top-ranked selections [32]. While RX computes a global, independent anomaly score with respect to the data background distribution, DEMUD instead uses an SVD model of its previous selections to select the item that is most different from everything previously seen. The DEMUD anomaly score for x_i is its reconstruction error:

$$A_D(x_i) = ||x_i - (UU^T(x_i - \mu) + \mu)||_2, \tag{3}$$

where U contains the top k principal components and $UU^T(x_i - \mu) + \mu$ is the reconstruction of x_i using U. U is initialized using the highest-scoring item based on an SVD of the full data set, or if that is too costly, a randomly chosen item can be used. DEMUD selects x' to maximize A_D, updates U to include x', and re-ranks the remaining items. The final ranking is the order of selection. DEMUD's anomaly score can also be expressed as

$$A_D(x_i) = ||\tilde{U}^T(x_i - \mu)|| \text{ or} \tag{4}$$

$$A_D^2(x_i) = (x_i - \mu)^T \tilde{U}\tilde{U}^T (x_i - \mu), \tag{5}$$

where $\tilde{U}$ contains the principal components $k + 1$ through n. In this form we see that RX and DEMUD both assess novelty by distance to the mean in the space defined by the principal components with least variance. The primary difference is that DEMUD's model changes with each new selection, while the RX model is fixed.

We compared the unsupervised RX and DEMUD methods to a supervised matched filter approach [20]. The Target Presence score T_{MF} of a given mineral s for pixel x_i is

$$T_{MF}(x_i) = \frac{s^T \Sigma^{-1} x_i}{\sqrt{s^T \Sigma^{-1} s}}. \tag{6}$$

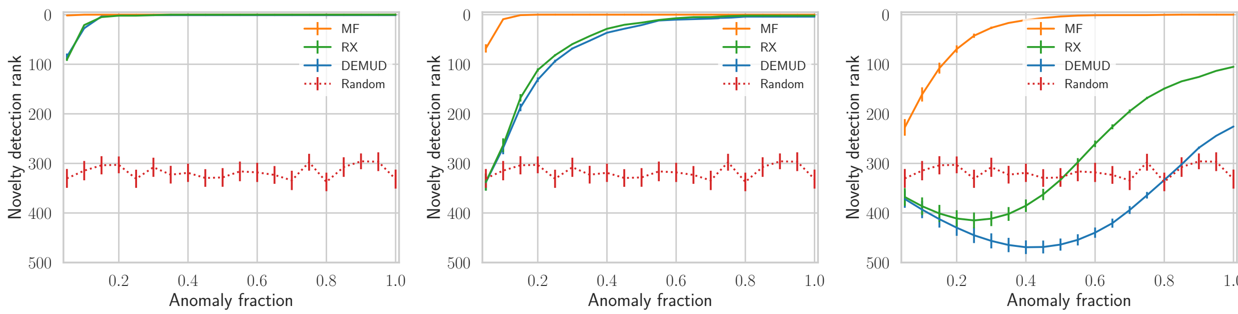

(a) Hyalite (b) Epsomite (c) Ulexite

Figure 4: Novelty detection curves for three minerals injected into NIMS observation 14ENEUR15H01B.

The matched filter provides an upper bound on the discoverability of mineral s within a given data set.

3.2.2 Results using NIMS Analogue Data. To evaluate the anomaly (RX and DEMUD) and target (MF) detectors, we injected artificial compositional anomalies into existing observations of Europa. Since MISE data does not yet exist, we used analogue data from the Galileo Near Infrared Mapping Spectrometer (NIMS) instrument, which collected observations of Europa at almost the same spectral range (0.8 to 5.2 μm), albeit lower spatial resolution. We randomly selected a pixel x_i and replaced it with $x_i{}'$, a linear mixture between x_i and the laboratory spectrum s for a selected mineral, given a specified anomaly fraction $0 < f_r \leq 1.0$:

$$x_i{}' = f_r s + (1 - f_r) x_i.$$

We used minerals (sulfates, silicates, and oxides) that would be of interest if discovered on Europa.

We evaluated the ability of the anomaly and target detectors to detect an injected anomaly by assessing its assigned novelty rank. For each algorithm, we sorted all of the pixels in the image in decreasing order of anomaly score and reported the rank at which the injected anomaly was discovered (lower is better). Results for NIMS Europa observation 14ENEUR15H01B[2] are shown for three minerals in Figure 4. At high anomaly fractions, the perturbed pixels are easily detected (novelty rank is 0). When less of the anomalous mineral is present, detection becomes more difficult, and the novelty rank increases. The supervised matched filter provides an upper bound. Minerals vary in terms of their novelty (and therefore detectability using RX or DEMUD). Hyalite is very different spectrally from data in this NIMS observation, as shown in Figure 5, so it is easily detected. Epsomite and ulexite are more challenging to detect because they are similar to the original data and therefore not inherently anomalous. Ulexite detection for small anomaly fractions performs worse than randomly selecting a pixel.

Some minerals, such as ulexite, exhibit non-monotonic novelty detection curves. We found that this occurs when the addition of a small amount of s (low f_r) causes x_i to initially become more similar to the background distribution and then, as f_r continues to rise, the perturbed pixel becomes more like s. Even 100% pure ulexite is difficult to distinguish from the original Europa observations, so it is unsurprising that small amounts could result in a spectrum that is more similar to the background. In contrast, since the matched filter directly computes the similarity between x_i and s, it is not influenced by the data distribution and it does not exhibit this inflection.

[2] https://pds-imaging.jpl.nasa.gov/data/go-j-nims-3-tube-v1.0/go_1115/europa/14e002ci.qub; see full description of the data set in Supplementary Materials.

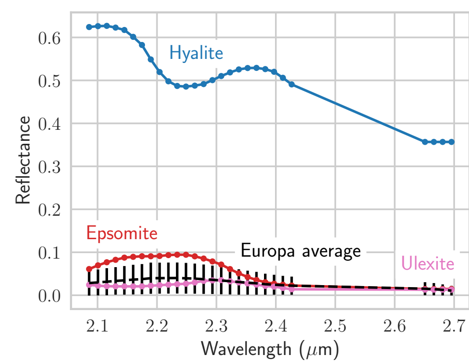

Figure 5: Spectra for minerals in Figure 4 and mean data spectrum for NIMS observation 14ENEUR15H01B.

We summarize the discoverability of a given mineral anomaly type by computing its average novelty rank across all anomaly fractions, for 100 trials (see Table 1). Overall, we found that RX generally out-performed DEMUD in discovering injected anomalies (exceptions are bloedite and opal). For this data set, across different mineral groups, the property that most determines discoverability is whether the mineral is hydrated. Hydrated minerals such as epsomite, opal, and ulexite were consistently the least anomalous, whereas anhydrous minerals like sulfur, olivine, and rutile were quickly discovered. This is consistent with our current understanding of Europa's surface composition (mostly water ice).

3.3 Plume Detection

A third important science objective of the Europa Clipper mission is to characterize any plume activity, which would provide key information about current processes at work beneath the frozen surface. Plumes have been observed emitting from Saturn's moon

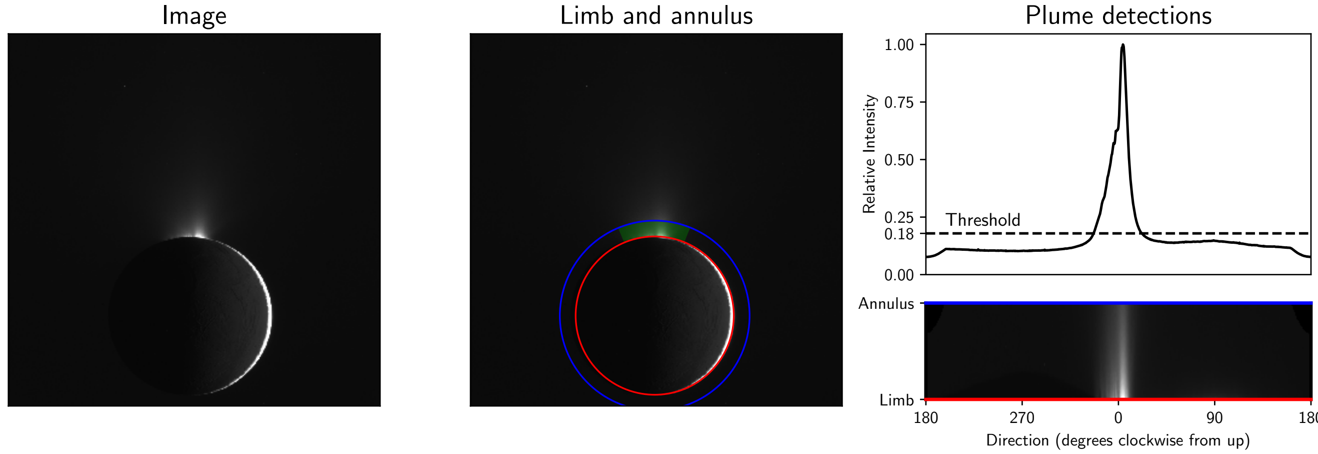

Figure 6: Plume detection within Cassini ISS observation `N1635814183_1` of Enceladus (left). Center: detected limb (red), annulus boundary (blue), and detected plume (green). Lower right: annulus unwrapped so that center corresponds to straight up. Upper right: average pixel intensity for each annulus column; dashed line indicates adaptively chosen detection threshold.

Table 1: Discoverability of different minerals injected into NIMS Europa observation `14ENEUR15H01B` in terms of mean (and standard error of) novelty rank (lower is better) across all anomaly fractions (100 trials). A random baseline achieves a novelty rank of 317.41 (4.16) for all minerals. Hydrated minerals are marked with (H).

Mineral	RX	DEMUD	MF
	Sulfates		
Sulfur	0.19 (0.01)	1.43 (0.04)	0.00 (0.00)
Polyhalite	1.63 (0.04)	3.55 (0.13)	0.03 (0.00)
Bloedite (H)	24.01 (0.92)	20.04 (0.83)	0.75 (0.18)
Epsomite (H)	58.52 (2.29)	64.52 (2.31)	3.91 (0.49)
	Silicates		
Olivine	0.45 (0.01)	1.89 (0.07)	0.02 (0.00)
Pyroxene	1.11 (0.02)	3.24 (0.06)	0.05 (0.00)
Hyalite (H)	6.07 (0.16)	6.63 (0.39)	0.08 (0.01)
Opal (H)	256.55 (8.21)	197.12 (7.39)	32.33 (2.91)
	Oxides		
Rutile	0.55 (0.01)	2.18 (0.08)	0.03 (0.00)
Cuprite	11.08 (0.34)	17.44 (0.74)	0.40 (0.10)
Chromite	66.22 (2.33)	92.23 (3.54)	5.65 (0.81)
Ulexite (H)	282.55 (8.37)	391.06 (9.60)	33.94 (2.88)

Enceladus [17], and there are hints from observations by the Hubble Space Telescope that similar activity could be present on Europa [25, 26], but as yet no definitive conclusions have been reached.

The Europa Imaging System (EIS) is composed of narrow- and wide-angle visible-wavelength cameras with the ability to image in color and in stereo [29]. EIS will produce global maps of Europa at 100 m per pixel, with high resolution imagery at 1 m per pixel or better in selected locations. Images from EIS will help constrain the formation of surface features, address small-scale regolith processes, and characterize potential landing sites for a possible future landed mission. More distant images from EIS will allow searches for plumes and other ongoing or recent geologic activity, as well as to characterize the shape of Europa, which has implications for subsurface structure.

3.3.1 Plume Detection Method. We adapted a method previously proposed for onboard detection of plumes from other bodies in the Solar System such as comets or Saturn's moon Enceladus [33]. This approach creates a model of the target using a convex hull to accommodate non-spherical small bodies. Europa is sufficiently spherical to not require this complexity, so we instead apply Canny edge detection followed by a Random Sample Consensus (RANSAC) circle-fitting algorithm to find the limb (apparent edge). Next, we conduct a plume search within a ring ("annulus") that spans from 101% to 120% of the estimated radius of the body within the image (see Figure 6). To increase robustness to pixel noise, plume detection operates on average intensity values for 1024 annular sectors surrounding the body. Any sectors whose average pixel values exceed an adaptively determined threshold (i.e., 1.5 times the inter-quartile range (IQR) across all sectors) are marked as plume candidates.

3.3.2 Results using Analogue Data. We used a variety of existing analogue datasets to evaluate the proposed plume detection approach. They comprise 308 observations of Europa and other small planets and moons collected by other spacecraft instruments (see the Supplementary Materials for a full listing). Many of these images exhibit effects likely to appear in EIS images as well, such as detector saturation, cosmic ray strikes, and low signal to noise ratio. For each image, we manually labeled the limb of the body and (if present) plumes to serve as ground truth[3].

We compared the adaptive-threshold plume detection method with a baseline method that employs a fixed threshold of 10 DN units. First, we computed area under receiver operating characteristic (ROC) curve (AUC) in terms of plume detection, which maps

[3] EIS analogue image limb and plume labels available at http://doi.org/10.5281/zenodo.2556063.

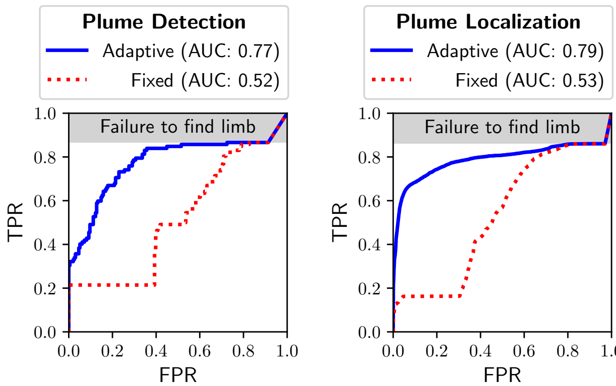

Figure 7: Plume detection and localization performance for fixed (solid, blue) and adaptive (dashed, red) thresholds.

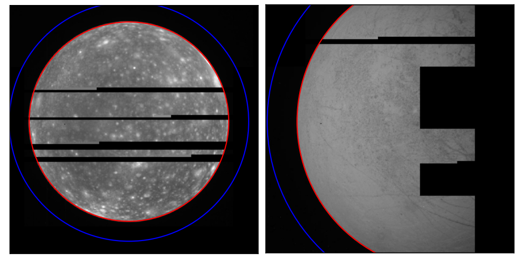

(a) Successful limb detections despite artifacts in images of Callisto (left) and Europa (right)

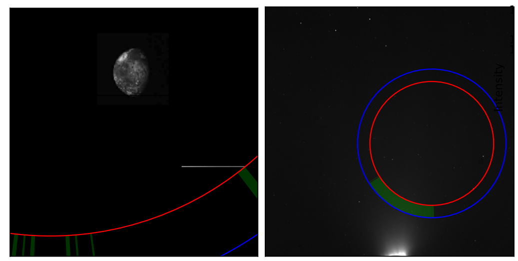

(b) Incorrect detections due to artifacts (Io, left) or body not in view (Enceladus, right)

Figure 8: Example successful and failed plume detections.

directly to prioritization decisions for each image (Figure 7, left). If plume activity in any annulus sector was detected, the image was marked as "contains plume." If limb-finding failed, then this was treated as "no plume." The reduction in AUC due to limb-finding failure appears in the plots as gray shaded regions. The adaptive threshold yielded significant improvement in AUC relative to using a fixed threshold, which performs poorly because DNs can vary significantly across images due to difference in instrument characteristics. We also evaluated plume localization, i.e., decisions made at the individual annulus sector level (Figure 7, right). This corresponds to a scenario in which the spacecraft transmits only the (cropped) active plume regions to reduce downlink consumption. Once again, the baseline approach performs near chance levels, whereas the adaptive threshold offers a significant advantage.

We investigated individual outputs to identify the strengths and weaknesses of this approach. We found that the limb detection process is robust to the presence of missing data or other artifacts (Figure 8(a)). In both cases, the limb was correctly modeled, and no spurious plumes were detected. However, limb detection also yielded incorrect results for some cases (Figure 8(b)). In the left image (Io), the circle was fit to a bright block of pixels in the lower right part of the image. In the right image (Enceladus), the body is not in view, only the plume, and a spurious circle is fit to stray pixels in the scene. We would therefore only employ this method for images with sufficient coverage of the body and true limb regions.

4 OPERATIONAL SCENARIOS

The preceding section described several methods for onboard analysis of scientific data collected by the Europa Clipper spacecraft and an evaluation of each component's performance on analogue or simulated data sets. In this section, we address the practical questions of how and where this capability could be deployed in the onboard spacecraft environment. We describe options and recommendations for an onboard science architecture. We also discuss the lessons we have learned to date in our collaboration with Clipper mission personnel and investigation scientists. These lessons are the key takeaways necessary to move this technology forward for deployment onboard a spacecraft.

4.1 Europa Clipper Operational Scenarios

We have designed an onboard science architecture for Europa Clipper that is customized to the existing spacecraft design. Figure 9 shows how Clipper's instruments will send data to the Bulk Data System (BDS) for temporary storage onboard. Each observation is stored in a pre-assigned priority bin. After the flyby, when downlink to Earth begins, data is selectively transferred from the priority bins to the the staging filesystem according to the downlink priority table. This table can be updated at any time by commands sent from Earth (e.g., prior to the next downlink). Data is packetized and transmitted to the ground (Earth).

The onboard science algorithms can be deployed in two possible locations. First, they can be integrated into each instrument's individual flight software. This permits the greatest flexibility, since priority decisions can be made according to individual instrument science objectives. However, the computational resources of individual instruments are even more constrained than those of the main spacecraft CPU; some instruments have only a special-purpose field-programmable gate array (FPGA) that contains code optimized for data collection and nothing more. In addition, this design choice would preclude any cross-coordination between instruments (e.g., for a thermal anomaly detection by E-THEMIS to increase the priority of a coincident spectral observation by MISE) without the addition of new cross-instrument communication channels.

The second option is for data analysis to take place on the main spacecraft CPU. It has the greatest computational resources available onboard and is the final decision point prior to any data being transmitted to the Earth. Updates to the downlink priority table can be made immediately by the spacecraft itself. The main CPU can

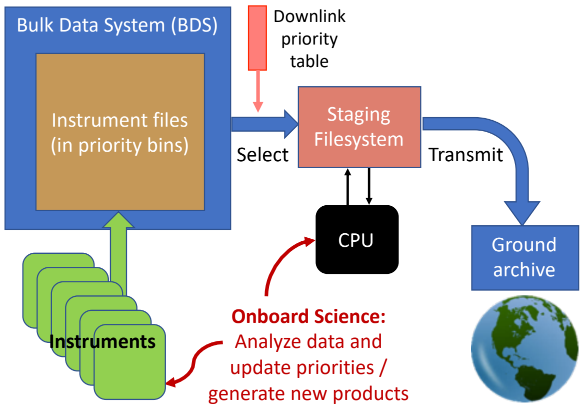

Figure 9: Onboard science architecture for Europa Clipper

synthesize detections from multiple instruments to enable cross-coordinated prioritization decisions. In addition, Europa Clipper will have the ability to partition CPU resources (memory and cycles) so that different software modules can be individually updated and will not interfere with other activities, which increases the robustness of the mission to faults and also permits strict control over resource consumption [2]. The drawback of running data analysis methods on the main CPU is that the instruments will already have stored all candidate data products on the BDS, in a fragmented format that was not designed for ease of onboard analysis. However, overall the benefits of this choice outweigh the drawbacks.

The Europa Clipper mission is planning 46 flybys of Europa over the period of several Earth years. We would likely not use onboard data analysis and prioritization on the first flyby but instead collect real Clipper observations of Europa with which to evaluate the algorithms described in this paper and then adapt them as needed. After demonstrating satisfactory performance on these early observations, we would activate the algorithms for operation onboard the spacecraft.

4.2 Lessons Learned

The Europa Clipper mission is still under development, with a launch date more than four years into the future. Our collaboration with the scientists and mission designers has been and will continue to be critical to tailoring onboard data analysis methods to the science objectives and the operational constraints of the mission. Our work to date has yielded the following lessons that can also inform future development of onboard science analysis for other missions.

Simpler analysis methods are often best. In the quickly evolving world of machine learning research, it is tempting to apply the latest, most sophisticated approaches to a given problem. However, spacecraft operations impose severe resource constraints (memory and CPU) that require the exploration of the tradeoff between runtime and accuracy. For example, while deep convolutional neural networks are commonly used today for image analysis tasks, they are infeasible for use on a single-core 200-MHz processor. Developing more capable radiation-hardened (e.g., multicore) processors is an ongoing area of investigation [1]. To date we have found good performance using basic computer vision and statistical methods that have low computational demands. One of our next steps will be to benchmark each algorithm on a RAD750 processor to obtain realistic estimates of memory and CPU consumption.

Another primary benefit of using simpler analysis methods is that they can be readily understood by experts from outside of computer science or artificial intelligence. Methods such as threshold application, outlier detection, and circle-fitting plus statistical testing are general enough to be accessible to the wider scientific community. Interpretability is essential for adoption and trust.

Performance metrics should be chosen to capture behavior of interest to the audience. For the benefit of the KDD audience, we have focused on traditional machine learning metrics such as probability of detection, precision, and recall. We are in the process of working with mission engineers to conduct mission simulations [12] in which we will quantify performance in ways that have immediate meaning to the spacecraft community, such as the increase in science return (e.g., fraction of downlinked images that contain a feature of interest such a a plume) and number of flybys required to meet the science objectives (fewer is better).

Risk is an overriding concern in spacecraft design and development. A new technology or capability must demonstrate large potential benefits for any increase in risk to be tolerated. Autonomous decision making is generally viewed as increasing risk: if a spacecraft's actions depend on unknown factors such as its environment or the frequency with which a feature of interest occurs, the outcome cannot be precisely modeled and checked for safe outcomes in advance. It is necessary to explore a wide range of possible scenarios and demonstrate no increased risk (or a tolerable level of risk) to the mission's operations. As noted above, processor partitioning is an excellent risk mitigation design choice.

5 CONCLUSIONS

We have developed algorithms to support three scientific objectives for the upcoming Europa Clipper mission: the detection of thermal anomalies, compositional anomalies, and plumes. Each algorithm demonstrates sufficiently high performance to satisfy the planetary science investigators. Our next steps will be to benchmark the algorithms on a RAD750 processor and quantify potential improvements in science return using mission simulations. The lessons learned from this collaboration can benefit future efforts to deploy data analysis for use onboard other spacecraft as well.

ACKNOWLEDGMENTS

We thank Steve Chien (spacecraft autonomy), Diana Blaney (MISE PI), Philip Christensen (E-THEMIS PI), and Serina Diniega (MISE Investigation Scientist) for their guidance and support for this investigation. We also thank the JPL R&TD program for funding support and the Europa Clipper project (including Robert Pappalardo, Barry Goldstein, L. Alberto Cangahuala, and others) for giving generously of their time and knowledge. The data sets used in this work were obtained from the Planetary Data System (PDS). This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

REFERENCES

[1] Jon Ballast, T. Amort, M. Cabanas-Holmen, E. H. Cannon, R. Brees, C. Neathery, S. Fischer, and W. Snapp. 2015. A method for efficient radiation hardening of multicore processors. In *Proceedings of the IEEE Aerospace Conference*. IEEE, Big Sky, MT.

[2] Todd Bayer, B. Buffington, J. F. Castet, M. Jackson, G. Lee, K. Lewis, J. Kastner, K. Schimmels, and K. Kirby. 2017. Europa mission update: Beyond payload selection. In *Proceedings of the IEEE Aerospace Conference*. IEEE, 1–12. https://doi.org/10.1109/AERO.2017.7943832

[3] Diana L. Blaney, R. N. Clark, J. B. Dalton, A. G. Davies, R. O. Green, M. M. Hedman, C. A. Hibbits, Y. Langevin, J. I. Lunine, T. B. McCord, C. Paranicas, S. L. Murchie, F. P. Seelos, and J. M. Soderblom. 2017. The Mapping Imaging Spectrometer for Europa (MISE) investigation: Exploring Europa's habitability using compositional mapping. In *Proceedings of the 48th Lunar and Planetary Science Conference*. Lunar and Planetary Institute, Houston, TX. Abstract #2244.

[4] Diana L. Blaney, C. Hibbitts, R. O. Green, R. N. Clark, J. B Dalton, A. G. Davies, Y. Langevin, J. I. Lunine, M. Hedman, T. B. McCord, S. L. Murchie, C. Paranicas, F. P. Seelos IV, J. M. Soderblom, S. Diniega, M. Cable, D. Thompson, C. Bruce, A. Santo, R. Redick, D. Hahn, H. Bender, B. Van Gorp, J. Rodriguez, P. Sullivan, T. Neville, S. Lundeen, M. Bowers, J. Hayes, B. Bryce, R. Hourani, J. Mize, F. Graham, E. Zarate, L. Moore, K. Maynard, I. M. McKinley, D. Johnson, P. Aubuchon, J. Fedosi, R. Wehbe, R. Calvet, and P. Mouroulis. 2019. The *Europa Clipper* Mapping Imaging Spectrometer for Europa (MISE): Using compositional mapping to understand Europa. In *Proceedings of the 50th Lunar and Planetary Science Conference*. Lunar and Planetary Institute, Houston, TX.

[5] Rebecca Castano, Kiri L. Wagstaff, Steve Chien, Timothy M. Stough, and Benyang Tang. 2007. On-board analysis of uncalibrated data for a spacecraft at Mars. In *13th Int'l Conf. on Knowledge Discovery and Data Mining*. 922–930.

[6] Chein-I Chang and Shao-Shan Chiang. 2002. Anomaly detection and classification for hyperspectral imagery. *IEEE Transactions on Geoscience and Remote Sensing* 40, 6 (2002), 1314–1325. https://doi.org/10.1109/TGRS.2002.800280

[7] Steve Chien, Brian Bue, Julie Castillo-Rogez, Dero Gharibian, Russell Knight, Steve Schaffer, David R. Thompson, and Kiri L. Wagstaff. 2014. Agile science: Using onboard autonomy for primitive bodies and deep space exploration. In *Proceedings of the International Symposium on Artificial Intelligence, Robotics, and Automation for Space (i-SAIRAS)*.

[8] Steve Chien, J. Doubleday, D. R. Thompson, K. L. Wagstaff, J. Bellardo, C. Francis, E. Baumgarten, A. Williams, E. Yee, E. Stanton, and J. Piug-Suari. 2017. Onboard autonomy on the Intelligent Payload EXperiment CubeSat mission. *Journal of Aerospace Information Systems* 14, 6 (2017), 307–315. https://doi.org/10.2514/1.I010386

[9] National Research Council. 2011. *Vision and Voyages for Planetary Science in the Decade 2013–2022*. National Academies Press.

[10] Ingrid J. Daubar, A. S. McEwen, S. Byrne, M. R. Kennedy, and B. Ivanov. 2013. The current martian cratering rate. *Icarus* 225, 1 (2013), 506–516. https://doi.org/10.1016/j.icarus.2013.04.009

[11] Ashley G. Davies, S. Chien, R. Wright, A. Miklius, P. R. Kyle, M. Welsh, J. B. Johnson, D. Tran, S. R. Schaffer, and R. Sherwood. 2006. Sensor web enables rapid response to volcanic activity. *Eos* 87, 1 (2006), 1–5. https://doi.org/10.1029/2006EO010002

[12] Eric W. Ferguson, Steve S. Wissler, Ben K. Bradley, Pierre F. Maldague, Jan M. Ludwinski, and Chistopher R. Lawler. 2018. Improving spacecraft design and operability for Europa Clipper through high-fidelity, mission-level modeling and simulation. In *Proceedings of the SpaceOps Conference*. https://doi.org/10.2514/6.2018-2469

[13] Georg Fischer, M. D. Desch, P. Zarka, M. L. Kaiser, D. A. Gurnett, W. S. Kurth, W. Macher, H. O. Rucker, A. Lecacheux, W. M. Farrell, and B. Cecconi. 2006. Saturn lightning recorded by Cassini/RPWS in 2004. *Icarus* 183, 1 (2006), 135–152. https://doi.org/10.1016/j.icarus.2006.02.010

[14] Raymond Francis, T. Estlin, G. Doran, S. Johnstone, D. Gaines, V. Verma, M. Burl, J. Frydenvang, S. Montaño, R. C. Wiens, S. Schaffer, O. Gasnault, L. DeFlores, D. Blaney, and B. Bornstein. 2017. AEGIS autonomous targeting for ChemCam on Mars Science Laboratory: Deployment and results of initial science team use. *Science Robotics* 2, 7 (2017). https://doi.org/10.1126/scirobotics.aan4582

[15] Neil Gehrels and Péter Mészáros. 2011. Gamma-ray bursts. *Science* 337, 6097 (2011), 932–936. https://doi.org/10.1126/science.1216793

[16] Ronald Greeley, P. L. Whelley, R. E. Arvidson, N. A. Cabrol, D. J. Foley, B. J. Franklin, P. G. Geissler, M. P. Golombek, R. O. Kuzmin, G. A. Landis, M. T. Lemmon, L. D. V. Neakrase, S. W. Squyres, and S. D. Thompson. 2006. Active dust devils in Gusev crater, Mars: Observations from the Mars Exploration Rover Spirit. *Journal of Geophysical Research: Planets* 111, E12 (2006). https://doi.org/10.1029/2006JE002743

[17] Candice J. Hansen, L. Esposito, A. I. F. Stewart, J. Colwell, A. Hendrix, W. Pryor, D. Shemansky, and R. West. 2006. Enceladus' water vapor plume. *Science* 311, 5766 (2006), 1422–1425. https://doi.org/10.1126/science.1121254

[18] Paul O. Hayne, P. R. Christensen, J. R. Spencer, O. Abramov, C. Howett, M. T. Mellon, F. Nimmo, S. Piqueux, and J. A. Rathbun. 2017. Possible nature and detectability of endogenic thermal anomalies on Europa. In *Lunar and Planetary Science Conference*, Vol. 48. Lunar and Planetary Institute, Houston, TX.

[19] Lukas Mandrake, U. Rebbapragada, K. L. Wagstaff, D. Thompson, S. Chien, D. Tran, R. T. Pappalardo, D. Gleeson, and R. Castano. 2012. Surface sulfur detection via remote sensing and onboard classification. *ACM Transactions on Intelligent Systems and Technology* 3, 4 (2012). https://doi.org/10.1145/2337542.2337562

[20] Nasser M. Nasrabadi. 2014. Hyperspectral target detection: An overview of current and future challenges. *IEEE Signal Processing Magazine* 31, 1 (2014), 34–44. https://doi.org/10.1109/MSP.2013.2278992

[21] Robert T. Pappalardo, D. A. Senske, H. Korth, D. Blaney, D. Blankenship, P. Christensen, S. Kempf, K. Retherford, E. P. Turtle, J. H. Waite, J. Westlake, G. C. Collins, K. Hand, J. Lunine, M. McGrath, F. Nimmo, C. Paty, J. Soderblom, J. R. Spencer, C. Paranicas, S. Solomon, and the Europa Science Team. 2017. The Europa multiple-flyby mission: Synergistic science to investigate habitability. In *Proceedings of the 48th Lunar and Planetary Science Conference*. Lunar and Planetary Institute, Houston, TX. Abstract #2732.

[22] Carolyn C. Porco, P. Helfenstein, P. C. Thomas, A. P. Ingersoll, J. Wisdom, R. West, G. Neukum, T. Denk, R. Wagner, T. Roatsch, S. Kieffer, E. Turtle, A. McEwen, T. V. Johnson, J. Rathbun, J. Veverka, D. Wilson, J. Perry, J. Spitale, A. Brahic, J. A. Burns, A. D. DelGenio, L. Dones, C. D. Murray, and S. Squyres. 2006. Cassini observes the active south pole of Enceladus. *Science* 311, 5766 (2006), 1393–1401. https://doi.org/10.1126/science.1123013

[23] Julie A. Rathbun, Nathaniel J. Rodriguez, and John R. Spencer. 2010. Galileo PPR observations of Europa: Hotspot detection limits and surface thermal properties. *Icarus* 210, 2 (2010), 763–769. https://doi.org/10.1016/j.icarus.2010.07.017

[24] Irving S. Reed and X. Yu. 1990. Adaptive multiple-band CFAR detection of an optical pattern with unknown spectral distribution. *IEEE Transactions on Acoustic Speech and Signal Processing* 38 (1990), 1760–1770. https://doi.org/10.1109/29.60107

[25] Lorenz Roth, Joachim Saur, Kurt D. Retherford, Darrell F. Strobel, Paul D. Feldman, Melissa A. McGrath, and Francis Nimmo. 2014. Transient water vapor at Europa's south pole. *Science* 343, 6167 (2014), 171–174. https://doi.org/10.1126/science.1247051

[26] William B. Sparks, K. P. Hand, M. A. McGrath, E. Bergeron, M. Cracraft, and S. E. Deustua. 2016. Probing for evidence of plumes on Europa with HST/STIS. *The Astrophysical Journal* 829, 2 (2016). https://doi.org/10.3847/0004-637X/829/2/121

[27] William B. Sparks, B. E. Schmidt, M. A. McGrath, K. P. Hand, J. R. Spencer, M. Cracraft, and S. E Deustua. 2017. Active cryovolcanism on Europa? *The Astrophysical Journal Letters* 839, 2 (2017), L18. https://doi.org/10.3847/2041-8213/AA67F8

[28] Samantha K. Trumbo, Michael E. Brown, and Bryan J. Butler. 2017. ALMA thermal observations of a proposed plume source region on Europa. *The Astronomical Journal* 154, 4 (2017), 148. https://doi.org/10.3847/1538-3881/AA8769

[29] Elizabeth Turtle and Alfred McEwen. 2018. Investigating Europa's geology, ice shell, and potential for current activity with the Europa Imaging System (EIS). In *42nd COSPAR Scientific Assembly*. Abstract B5.3-52-18.

[30] Kiri L. Wagstaff, Alphan Altinok, Steve Chien, Umaa Rebbapragada, Steve Schaffer, David R. Thompson, and Daniel Tran. 2017. Cloud filtering and novelty detection using onboard machine learning for the EO-1 spacecraft. In *Proceedings of the IJCAI Workshop on AI in the Oceans and Space*.

[31] Kiri L. Wagstaff and Benjamin Bornstein. 2009. K-means in Space: A Radiation Sensitivity Evaluation. In *Proceedings of the Twenty-Sixth International Conference on Machine Learning*. 1097–1104.

[32] Kiri L. Wagstaff, Nina L. Lanza, David R. Thompson, Thomas G. Dietterich, and Martha S. Gilmore. 2013. Guiding scientific discovery with explanations using DEMUD. In *Proceedings of the Twenty-Seventh Conference on Artificial Intelligence*. AAAI, Palo Alto, CA, 905–911.

[33] Kiri L. Wagstaff, David R. Thompson, Brian D. Bue, and Thomas J. Fuchs. 2014. Autonomous real-time detection of plumes and jets from moons and comets. *The Astrophysical Journal* 794, 1 (2014). https://doi.org/10.1088/0004-637X/794/1/43

SUPPLEMENTARY MATERIALS

This supplementary material contains the information necessary to reproduce the experiments described in the paper.

Section 3.1: Thermal Anomaly Detection

The simulated E-THEMIS data set (5000 images) and accompanying simulated thermal anomalies can be downloaded at http://doi.org/10.5281/zenodo.2552814.

Section 3.2: Compositional Anomaly Detection

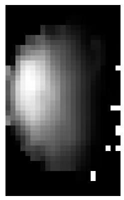

The data set used to evaluate compositional anomaly detection is Galileo NIMS observation 14ENEUR15H01B of Europa (see right). It consists of 874 pixels observed at 135 wavelengths from 2.085 to 3.988 μm. It can be downloaded at https://pds-imaging.jpl.nasa.gov/data/go-j-nims-3-tube-v1.0/go_1115/europa/14e002ci.qub.

NIMS data pre-processing. NIMS data may contain missing values, which are indicated with NaN values. For each NIMS cube, we first removed all pixels that were composed entirely of missing observations. For the remaining pixels, any missing values were imputed using linear interpolation from the remaining non-NaN spectral channels for that pixel.

NIMS spectral library. We injected known mineral spectra to create synthetic compositional anomalies. The mineral spectra were obtained from the USGS NIMS spectral library, which is available at https://archive.usgs.gov/archive/sites/speclab.cr.usgs.gov/spectral.lib04/lib04-NIMS.html. This library includes 498 spectra with reflectance reported from 0.8 to 2.7–3.0 μm.

For the experiments reported in Section 3.2, the intersection of the wavelength ranges for 14ENEUR15H01B and the spectral library restricts analysis to the range 2.085 to 2.7 μm, as shown in Figure 5.

DEMUD parameter. The DEMUD algorithm was run with the number of principal components, k, set to 10.

Section 3.3: Plume Detection

We compiled 308 images from previous missions to serve as analogues for the data EIS is expected to collect. The closest instrument analogue is the Mercury Dual Imaging System (MDIS) instrument flown on the MESSENGER spacecraft. While Mercury is a rocky body without any plume activity, the MDIS cameras had a similar design to that planned for EIS. Other predecessor instruments vary in terms of their imaging properties but are useful analogues because of their choice of targets. These include images of Jupiter's moons Callisto, Europa, Ganymede, and Io and Saturn's moon Enceladus collected by the Galileo Solid-State Imaging (SSI), Cassini Imaging Science Subsystem (ISS), and New Horizons Long Range Reconnaissance Imager (LORRI) instruments.

Images were selected to simulate the views of Europa that EIS will collect. Some contain the entire body of interest, often with partial illumination from the side so as to highlight plume activity. Others contain only a portion of the body to enable a close view of the limb where plume activity may be present. The ISS Enceladus images were chosen from Cassini's seventh flyby of Enceladus, with known plume activity present. Likewise, the Galileo SSI and LORRI Io images have known volcanic activity. The remaining targets do not exhibit plume activity and instead serve as control subjects for plume detection.

Our manual labels of the limb and plume(s) for each image can be downloaded at http://doi.org/10.5281/zenodo.2556063. All source images are available at https://pds-imaging.jpl.nasa.gov/search/. The following tables list the specific image product identifiers.

Cassini ISS narrow-angle camera (Enceladus), 40/40 contain plumes			
N1635781564_1	N1635781640_1	N1635781702_1	N1635781742_1
N1635781798_1	N1635781940_1	N1635781995_1	N1635782028_1
N1635782061_1	N1635782096_1	N1635804540_1	N1635804756_1
N1635804850_1	N1635804944_1	N1635805182_1	N1635805649_1
N1635805743_1	N1635805837_1	N1635808826_1	N1635809286_1
N1635809569_1	N1635809625_1	N1635809687_1	N1635809727_1
N1635809799_1	N1635813819_1	N1635813867_1	N1635813923_1
N1635814065_1	N1635814245_1	N1635814301_1	N1635814379_1
N1635814521_1	N1635815659_1	N1635815737_1	N1635815815_1
N1635815957_1	N1635816588_1	N1635816682_1	N1635816920_1

Galileo SSI (Europa), 0/22 contain plumes		
IMAGE_ID, File	IMAGE_ID, File	IMAGE_ID, File
G1E0004, 5139r	G7E0010, 4500r	10E0001, 2778r
12E0003, 4626r	12E0004, 4639r	12E0005, 4652r
14E0054, 4842r	14E0056, 4878r	14E0057, 4901r
14E0058, 4914r	14E0062, 4965r	14E0063, 4978r
14E0064, 5000r	14E0066, 5026r	14E0067, 5039r
14E0068, 5052r	15E0096, 4429r	19E0008, 4726r
25E0021, 6314r	25E0025, 6366r	25E0029, 6427r
25E0031, 6452r		
Galileo SSI (Io), 12/19 contain plumes		
C3I0040, 1300r	C9I0015, 3178r	C9I0016, 3200r
C9I0017, 3204r	C9I0018, 3207r	G2I0020, 6300r
G8I0010, 8645r	G8I0011, 8700r	G8I0012, 8723r
G8I0013, 5045r	G8I0014, 5100r	G8I0015, 5123r
G8I0019, 2445r	10I0028, 4204r	10I0029, 4207r
11I0012, 0085r	11I0015, 3485r	31I0001, 5146r
31I0003, 5547r		
Galileo SSI (Callisto), 0/7 contain plumes		
G2C0004, 2900r	G7C0001, 6200r	30C0027, 1600r
30C0028, 1901r	30C0029, 2201r	31C0001, 4300r
31C0001, 4301r		
Galileo SSI (Ganymede), 0/13 contain plumes		
C9G0010, 1000r	C9G0011, 1013r	C9G0012, 1026r
C9G0013, 1039r	C9G0016, 1078r	E6G0020, 1500r
E6G0021, 1513r	G1G0001, 2000r	14G0001, 3078r
20G0001, 6900r	30G0003, 7600r	30G0004, 7900r
30G0005, 8200r		

MESSENGER MDIS narrow-angle camera (Mercury), 0/26 contain plumes			
EN0108828545M	EN0108828592M	EN0108828634M	EN0108828799M
EN0108828804M	EN0108828855M	EN0108828906M	EN0108829090M
EN0108829172M	EN0108830151M	EN0108830222M	EN0108830339M
EN0108830513M	EN0108830518M	EN0108830711M	EN0131773890M
EN0131773947M	EN0131774207M	EN0131774306M	EN0131774343M
EN0131774348M	EN0131774578M	EN0131774583M	EN0131774634M
EN0131774685M	EN0162736761M		
MESSENGER MDIS wide-angle camera (Mercury), 0/97 contain plumes			
EW0108816478C	EW0108816480D	EW0108816482E	EW0108817678C
EW0108817680D	EW0108817682E	EW0108818878C	EW0108818880D
EW0108818882E	EW0108820007C	EW0108820012D	EW0108820017E

(cont.)

MESSENGER MDIS wide-angle camera (Mercury), cont.			
EW0108820022F	EW0108820027G	EW0108820032H	EW0108820037I
EW0108820042J	EW0108820047K	EW0108820052L	EW0108820057A
EW0108829678A	EW0108829683L	EW0108829688K	EW0108829693J
EW0108829698I	EW0108829703H	EW0108829708G	EW0108829713F
EW0108829718E	EW0108829723D	EW0108829728C	EW0131764500C
EW0131764505D	EW0131764510E	EW0131764515F	EW0131764520G
EW0131764525H	EW0131764530I	EW0131764535J	EW0131764540K
EW0131764545L	EW0131764550A	EW0131775228A	EW0131775232L
EW0131775236K	EW0131775240J	EW0131775244I	EW0131775248H
EW0131775252G	EW0131775256F	EW0131775260E	EW0131775264D
EW0131775268C	EW0131777088A	EW0131777092L	EW0131777096K
EW0131777100J	EW0131777104I	EW0131777108H	EW0131777112G
EW0131777116F	EW0131777120E	EW0131777124D	EW0131777128C
EW0162736786C	EW0162736790D	EW0162736794E	EW0162736798F
EW0162736802G	EW0162736806H	EW0162736810I	EW0162736814J
EW0162736818K	EW0162736822L	EW0162736826A	EW0162739786C
EW0162739790D	EW0162739794E	EW0162739798F	EW0162739802G
EW0162739806H	EW0162739810I	EW0162739814J	EW0162739818K
EW0162739822L	EW0162739826A	EW0162741039C	EW0162741043D
EW0162741047E	EW0162741051F	EW0162741055G	EW0162741059H
EW0162741063I	EW0162741067J	EW0162741071K	EW0162741075L
EW0162741079A			

New Horizons LORRI (Io), 60/84 contain plumes					
Full product id: lor_{id}_0x630_sci					
0034583519	0034583522	0034599119	0034599122	0034630919	0034630922
0034685519	0034685522	0034698119	0034698122	0034721219	0034721222
0034769759	0034785119	0034785122	0034821014	0034821017	0034821020
0034829594	0034829597	0034829600	0034844219	0034844222	0034860614
0034860617	0034860620	0034873619	0034873622	0034889879	0034940519
0034940539	0034940542	0034943534	0034943537	0034943540	0034966574
0034966577	0034966580	0034974374	0034974377	0034974380	0034981619
0034981639	0034981642	0035015234	0035015237	0035015240	0035022029
0035022049	0035022052	0035029094	0035029097	0035029100	0035040494
0035040497	0035040500	0035077919	0035077939	0035077942	0035092814
0035092817	0035092820	0035121554	0035121557	0035121560	0035129519
0035129539	0035129542	0035140199	0035140219	0035140222	0035149919
0035149939	0035149942	0035208179	0035215619	0035215639	0035215642
0035222819	0035222839	0035222842	0035230619	0035230639	0035230642