NASA/TM-2019-220241



Fusion of visible and thermal-infrared imagery for SLAM for landing on icy moons

M. Gonzalez Felício Instituto Superior Técnico, University of Lisbon, 1049-001 Lisbon, Portugal

P. Michael Furlong SGT Inc/NASA Ames' Intelligent Robotics Group Michael Dille SGT Inc/NASA Ames' Intelligent Robotics Group Uland Wong SGT Inc/NASA Ames' Intelligent Robotics Group

NASA STI Program... in Profile

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA scientific and technical information (STI) program plays a key part in helping NASA maintain this important role.

The NASA STI Program operates under the auspices of the Agency Chief Information Officer. It collects, organizes, provides for archiving, and disseminates NASA's STI. The NASA STI Program provides access to the NASA Aeronautics and Space Database and its public interface, the NASA Technical Report Server, thus providing one of the largest collection of aeronautical and space science STI in the world. Results are published in both non-NASA channels and by NASA in the NASA STI Report Series, which includes the following report types:

- TECHNICAL PUBLICATION. Reports of completed research or a major significant phase of research that present the results of NASA programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA counterpart of peer-reviewed formal professional papers, but having less stringent limitations on manuscript length and extent of graphic presentations.
- TECHNICAL MEMORANDUM. Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.
- CONTRACTOR REPORT. Scientific and technical findings by NASA-sponsored contractors and grantees.

- CONFERENCE PUBLICATION. Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or co-sponsored by NASA.
- SPECIAL PUBLICATION. Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.
- TECHNICAL TRANSLATION. Englishlanguage translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services also include organizing and publishing research results, distributing specialized research announcements and feeds, providing information desk and personal search support, and enabling data exchange services.

For more information about the NASA STI Program, see the following:

- Access the NASA STI program home page at http://www.sti.nasa.gov
- E-mail your question to help@sti.nasa.gov
- Phone the NASA STI Information Desk at 757-864-9658
- Write to: NASA STI Information Desk Mail Stop 148 NASA Langley Research Center Hampton, VA 23681-2199

NASA/TM-2019-220241



Fusion of visible and thermal-infrared imagery for SLAM for landing on icy moons

M. Gonzalez Felício Instituto Superior Técnico, University of Lisbon, 1049-001 Lisbon, Portugal

P. Michael Furlong SGT Inc/NASA Ames' Intelligent Robotics Group Michael Dille SGT Inc/NASA Ames' Intelligent Robotics Group Uland Wong SGT Inc/NASA Ames' Intelligent Robotics Group

National Aeronautics and Space Administration

Ames Research Center

February 2019

The use of trademarks or names of manufacturers in this report is for accurate reporting and does not constitute an offical endorsement, either expressed or implied, of such products or manufacturers by the National Aeronautics and Space Administration.

Available from:

NASA STI Program / Mail Stop 148 NASA Langley Research Center Hampton, VA 23681-2199 Fax: 757-864-6500

Abstract

This paper addresses the problem of localization for landing on the surface of icy moons, like Europa or Enceladus. Due to the possibility of specular reflection as well as high bulk albedo, icy surfaces present new challenges that make traditional vision-based navigation systems relying on visible imagery unreliable. We propose augmenting visible light cameras with a thermal-infrared camera using inverse-depth parameterized monocular EKF-SLAM to address problems arising from the appearance of icy moons. Results were obtained from a novel procedural Europa surface simulation which models the appearance and the thermal properties simultaneously from physically-based methods. In this framework, we show that thermal features improve localization by 23% on average when compared to a visible camera. Moreover, fusing both sensing modalities increases the improvement in localization to 31% on average, compared to using a visible light camera alone.

1 Introduction

The icy moons of the gas giants Jupiter and Saturn are today considered among the most likely locations for finding other forms of life within our solar system. The presence of active plumes in Europa, first suggested by the Hubble Space Telescope (HST) and later sustained by Galileo's flybys (Roth et al., 2014; Sparks et al., 2016), and Enceladus evidenced by Cassini's flybys (Porco et al., 2006; Spencer et al., 2006; Hansen et al., 2006) make these two moons particularly interesting. The plumes provide i) strong evidence of the existence of liquid oceans underneath their icy surfaces and ii) a unique opportunity to make *in situ* measurements with no need of heavy and power intensive drilling or excavating mechanisms. For this reason, there has been a growing interest in dedicated missions to icy moons, e.g. (Hendrix et al., 2019; Phillips and Pappalardo, 2014; Eigenbrode et al., 2018; Hand et al., 2017).

However, entry, descent and landing (EDL) on these moons brings up unique challenges. Due to the long distance to Earth, human intervention is confined to high-level strategies, and onboard perception and decision making will be essential to missions success. For instance, NASA InSight, the latest probe to successfully land on mars, spent nearly seven months traveling trough space before going through the so-called "seven minutes of terror", i.e. the seven minute period between when the EDL procedure was initiated and when mission control could be aware of whether or not landing was successful due to the communications delay between Mars and Earth. Juno, the latest spacecraft launched to Jupiter needed five years to arrive to its destination, and the communication delay would be at least thirty five minutes. To the date, in EDL missions to the Moon and Mars, the landing site is selected *a pri*- ori by scientists based on orbital data. Icy surface on Europa and Enceladus may have unique dangers and geological features such as crevasses, jagged penitentes, liquids and ice thickness may not be visible from orbit due to resolution and perspective. Thus, onboard navigation will be crucial for landing safely and in the presence of control and sensor uncertainty.

While visible cameras sense light that is primarily reflected from a scene (in the planetary context), thermal-infrared cameras sense light emitted by objects due to their natural temperature. Therefore, unlike visible cameras, thermalinfrared cameras are more robust to ambient illumination conditions. This is especially relevant in the case of icy surfaces, where the high reflectance makes vision-based navigation systems relying solely on visible imagery unsuitable for localization (Fig. 1). While the nature of icy moon surfaces is not wellknown at this time, visible reflectance models predict the possibility of oblique (Fresnel) specularities, spectral shifts, and general loss of contrast affecting vision-based features Cook and Torrance (1982). The main contribution of this paper is in the exploration of using thermal-infrared sensing to overcome these difficulties specific to icy surfaces. We first demonstrate use of thermalonly navigation in lieu of a visible light camera and then explore thermal-visual fusion, through EKF-SLAM, to produce a further improved solution.

The remainder of this paper is organized as follows: section II summarizes the related work. The proposed methodology is described in section III. Simulation results illustrating the performance obtained are presented in section IV and discussed in section V. Section VI summarizes the contents of the paper and proposes future work.



Figure 1: Sequence of (non-consecutive) visible(top) and thermal (bottom) images captured from our simulator. The basis of this work is that visible images provide superior textural features in some areas while the thermal is better in others.

2 Related Work

As commodity thermal imaging sensors have only recently been made available at resolutions amenable to computer vision, thermal navigation has only been sparsely studied for robotic navigation, particularly for planetary applications. In one of the earliest applications of thermal-infrared cameras for motion estimation, Schön and Roll (2009) take advantage of already existing cameras in cars for human detection, to fuse this data with proprioceptive sensors through an extended Kalman filter (EKF). Similarly, Nilsson et al. (2011) compute an offline least mean squares motion estimate using a similar set of sensors.

Thermal-infrared cameras have also been investigated for perception through obscurants (e.g. smoke, airborn dust, fog) in GPS-denied conditions. Brunner et al. (2011) show that thermal cameras are less affected by the presence of smoke and dust, and can be used either alone or combined with visible cameras to improve SLAM performance and robustness of unmanned ground vehicles in outdoor environments. The novel contribution of their work is a metric to detect degradation and select which camera to rely in, before running feature detection, thereby saving computational time.

Vidas and Sridharan (2012) implemented a real-time hand-held SLAM using a thermal-infrared camera that is able to handle mechanic shutter interruptions (non-uniform conformity) common to these cameras. Emilsson and Rydell (2014) implemented a hand-held SLAM system based on a pair of stereo thermal-infrared cameras and an inertial measurement unit (IMU) primarily for use firefighters in smokey environments. More recently, Papachristos et al. (2018) employs a thermal-infrared camera fused with inertial sensor cues through an EKF framework to obtain localization of an aerial robot in an indoor environment with strong presence of smoke and fog.

Lastly, LIDAR technologies have proliferated in terrestrial applications to ostensibly handle many of the imaging challenges presented in this paper Kohlbrecher et al. (2011). However, LIDAR has many similar drawbacks to visual imaging in highly-reflective and specular surfaces since it is based on the principle of reflected light. Even in terrestrial snow and ice conditions, there is increased uncertainty and known failure conditions from icy geometry Kerekes et al. (2012), Deems et al. (2013). Space LIDAR is further limited in providing voluminous data suitable for SLAM by technological maturity, sensing range, mass, and power.

3 Approach

We have implemented a multi-modal SLAM algorithm using a framework similar to Brunner et al. (2011) which this work builds upon. We use one visible light camera and one thermal-infrared camera and fused data from the cameras to estimate vehicle egomotion.

To simplify the problem, we considered that the two cameras are coregistered (i.e. they observe the same portion of the scene at every instant) and synchronized (i.e. images are taken simultaneously) at 1 fps. The cameras point downward, along the negative z-axis in the robot frame. Images from both cameras are resized to a resolution of 640×480 pixels. Note however that these simplifying assumptions are not required to use our algorithm.

Our algorithm does not try to find common features between the two video streams, but tracks landmarks in each video stream over time. Landmarks in this case are SIFT features identified in the video streams. The current estimated vehicle velocity is used to predict where the tracked landmarks should appear in the next frame of the video. The differences between the observed landmark positions and the predicted ones are the correction signal which is fed into the Kalman filter that estimates the vehicle pose.

We produced ten unique trajectories over simulated terrain. The vehicle flew thirty trials for each trajectory. Errors and perturbations in actuation were modeled as additive Gaussian noise. The trajectories include two linear motions as well as an S-shape, half-ellipse and half-circle trajectories. The curved trajectories were included to test the performance of the algorithms, but do not necessarily represent feasible descent trajectories as we do not focus on propulsive control in this work.

Our work is distinguished from that of Brunner et al. (2011) by three primary differences. First, their application was operations in an environment with roads, plants, and some man-made objects such as vehicles, whereas we are operating on an icy surface that may not have such features to track. We instead comprehensively analyze performance in the context of icy planetary application.

A second differentiating factor is that the vehicle in (Brunner et al., 2011) is a ground-vehicle while ours is a flying vehicle some distance above the surface being tracked. When the vehicle is close to the ground there should always be features relatively close to the camera, meaning that smaller motions will produce greater motion of landmarks, making it easier to estimate vehicle motion. There is the distinct possibility in our problem setting that if the vehicle is at sufficient altitude, the parallax may not be sufficient to estimate vehicle motion. During the development of this approach we found that keeping track of relative motion and to slowing down the frame rate of our simulated cameras was needed in order to observe egomotion in a proper range. We also introduce a ranging altimeter into the estimation framework to aid in identifying range ambiguities and propagating state through feature gaps.

Finally, the third difference between our algorithm and that proposed in (Brunner et al., 2011) is how we replace landmarks they exit the field of view of the camera. Brunner et al. (2011) select from the available landmarks to fill the list of tracked landmarks with a uniform random distribution. We propose a sampling algorithm which uses a sampling distribution which attempts to ensure that the landmarks are spatially diverse through the field of view of the camera. Spatial diversity should improve the estimate of vehicle motion, and is discussed in detail in Section 3.3.

3.1 SLAM framework

Due to the projective nature of a camera, the 3-D position of a newly detected feature is not available for initialization because the uncertainty in the unobserved degree of freedom (the depth) is infinite, and because the projection to the camera (the measurement equation) is nonlinear, the extended Kalman filter cannot deal with it. Several delayed initialization methods such as (Kim and Sukkarieh, 2003; Bailey, 2003; Bryson and Sukkarieh, 2005; Lemaire et al.,

2005) treat newly detected features separately from the main probabilistic map, accumulating information in a special process over several frames to reduce depth uncertainty before insertion into the full filter.

However, these methods have the drawback that the bearing information from new features does not contribute to the estimation of the camera pose until included in the map. Moreover, distant features that retain low parallax over the many frames are usually rejected completely because their depth uncertainty is never significantly reduced. More recent methods (Kwok and Dissanayake, 2004; Sola et al., 2005; Sola, 2007) have been proposed that are immediately able to benefit from new features to improve camera motion estimates. These works still have the drawback of treating new features in a special way unlike the works of (Montiel et al., 2006; Civera et al., 2008; Marzorati et al., 2008; Sola, 2010) that exploit the inverse-depth to encode uncertainty up to infinity within a single Gaussian, keeping quasi-linearity of the observation function within all this uncertainty range, enabling the use of the EKF. For further discussion on these parametrizations the reader is referred to (Sola et al., 2012).

We use the inverse-depth parametrization from Montiel et al. (2006); Civera et al. (2008). In order to recover from scale, the vehicle's state, $X = (r, q, v, \omega)$, is propagated using the inertial measurements as follows:

$$r_{t+\Delta t} = r_t + v_t \Delta t + a_t^m \Delta t^2 / 2 \tag{1}$$

$$q_{t+\Delta t} = q_t + 1/2\,\omega_t^m \otimes q_t \Delta t \tag{2}$$

$$v_{t+\Delta t} = v_t + a_t^m \Delta t \tag{3}$$

$$\omega_{t+\Delta t} = \omega_t^m \tag{4}$$

where r and q are the position and the orientation quaternion, v and ω are the linear and angular velocities of the vehicle, a^m and ω^m are the linear acceleration and the angular velocity measured by the IMU, and Δt is the time step.

For each tracked feature, the ray from the first camera position from where the feature was observed encoded using the inverse-depth parameter as follows

$$L_i = (r_i^T, \theta_i, \phi_i, \rho_i)^T \tag{5}$$

where r_i is the camera optical center, θ_i and ϕ_i are the azimuth and elevation in world frame and ρ_i is the inverse of the feature depth along the ray.

In this work, we implement and compare SLAM using three different sensing configurations: only the visible camera, only the thermal-infrared camera, and both cameras fused ("dual modality"). In the *dual* configuration, the approach for fusing the information from two cameras is based on Sola et al. (2008), and consists of treating features from both cameras in the same map. For the single camera modalities a maximum of N = 32 features are tracked, as shown in Eq. (6), while an appended vector of N = 64 is considered in the *dual* modality case, as seen in Eq. (7).

$$M_Z = (L_{1,Z}^T, ..., L_{N,Z})^T, \quad Z \in \{V, T\}$$
(6)

$$M_D = (L_{1,V}^T, ..., L_{N,V}, L_{1,T}, ..., L_{N,T})^T$$
(7)

3.2 Feature detection and matching

The surface of icy moons are unstructured environments. Consequently we should not expect that primitive edge or corner detectors would identify stable, reliable features (Harris et al., 1988; Takacs et al., 2010).

Blob-like features are commonly observed in our data set, as seen in the sequence of frames in Fig. 1, so feature detectors which can key in on them are desirable. It would also be desirable for the feature detection method to be invariant to scale, as a landing vehicle will observe changes in scale of features during descent. Likewise, rotational robustness would be valuable as the vehicle may experience perturbations or change heading during descent.



Figure 2: SIFT Features in the visible light channel. The points marked with an cross in the image represent landmarks in the current frame. The blue dots are landmarks which are being track from the previous frame. The green line between them shows which current landmarks have been associated with currently tracked landmarks. The red crosses are landmarks which are not being tracked, and can be added to the tracked landmarks, should new ones be needed.

In these experiments we used SIFT features (Lowe, 2004) to identify stable features in the scene, illustrated in Fig. 2. SIFT was chosen because high quality implementations are widely available. However, any feature detector which meets the above criteria would be suitable for this application.

Note that in this work we do not attempt to match features between the two modalities to produce a measurement from multiple views. It may not be possible to use SIFT features to match across sensing modalities, as appearance descriptors may be very different. No further pre-processing techniques are applied to the images other than simple conversion to grayscale.

3.3 Feature sampling

This algorithm adds new features to the tracked landmarks as old features pass out of the field of view of the camera. In contrast, Brunner et al. (2011)

randomly select landmarks from the available features.

However, given the limited number of landmarks tracked relative to the number of SIFT features in a scene, there is the possibility that the tracked landmarks could be biased towards one region of the scene. This is particularly important as large regions of an icy surface with the same orientation could have severely reduced contrast at critical angles, e.g. Fresnel reflections Cook and Torrance (1982). A biased sampling of the scene would make for a less accurate estimate of egomotion.

Features well-spread over the image space are highly desirable as it increases the observability of the vehicle's state. We weighted the points in the sampling algorithm by the spatial diversity of the candidate points in order to increase the likelihood of collecting a stable set of points.

In Fig. 3 we computed the diversity score for every pixel in the image (Fig. 3b). We computed the diversity score using a kernel density estimator modeling the spatial distribution of the currently tracked landmarks.

We divide the landmarks into two sets. First is the set of tracked landmarks, $M_{Z,t}$, where Z is the SLAM modality currently being used, (*visual*, *thermal*, or *dual*), held over from the previous frame. There is also the landmarks identified in the current frame, \mathcal{L}_{t+1} . Through a matching process we remove points from \mathcal{L}_{t+1} that are matched to points in $M_{Z,t}$.

Should any of the landmarks in $M_{Z,t}$ not be matched to \mathcal{L}_{t+1} , then new landmarks need to be selected. The landmarks in \mathcal{L}_{t+1} are weighted by their Shannon surprise (Shannon, 1948), computed using a kernel density estimator. The kernel density estimator approximates a probability distribution of where currently tracked landmarks are located in the frame.

The kernel density estimator is generated from the locations of landmarks from the previous frame \mathcal{L}_t , which we denote $f_{M_{Z,t}}(l)$. The diversity score is given in Eq. (8). Here $f_{M_{Z,t}}(\cdot)$ is the kernel density estimator created from the currently tracked features.

$$D(l) = -log f_{M_{Z,t}}(l) \quad \forall \quad l \in \mathcal{L}_{t+1}$$

$$\tag{8}$$

By using the spatial distribution over $M_{Z,t}$, candidate landmarks that are further away from the currently tracked landmarks will be scored higher, as seen in Fig. 3b. Using these weights for sampling, as opposed to using a uniform weighting, it should increase the accuracy of the egomotion estimated from the tracked landmarks.



(a) Original frame

(b) Spatial Diversity (D)

Figure 3: The original video frame and the sampling weight based on the diversity score. In the frame there are few tracked landmarks (blue dots) in the bottom right hand corner, consequently the spatial diversity score is higher in this region. For illustrative purposes in the above images we have computed the values at every pixel. In the algorithm we only compute these values at candidate landmark sites.

4 Modeling Europa's Surface

To date, there have been no missions to the surface of Europa, or other icy moons. The current image resolution of the surface is no better than 10 meters per pixel. This is well below the resolution needed to make safe landing decisions. In future missions we might anticipate orbital image resolution of the landing site on par with the Mars Reconnaissance Orbiter, which has a resolution of approximately 1 meter per pixel. Until we have that data, however, the only way to develop and test vision-based algorithms is to model icy surfaces using current best-guess models.

4.1 Modeling Terrain Geometry

In this paper we model geometry and thermal qualities based on a Europan equatorial surface. The geometry of the terrain is procedurally generated using features that are randomly sampled but governed by statistical size-frequency and spatial distributions. Geometric features modeled at the time of this work include hills, valleys, craters, blocks and broken surfacing. The current version of the simulator also adds penitentes and crevasses. The simulated terrain has an area of $400 \times 400 m^2$, and has changes in elevation of 31.7m. The simulation was based on the simulator used in Lunar terrain modeling described by Allan et al. (2019), but with additional work to make the terrain better resemble icy environments.

4.2 Modeling Visual Appearance

To create the visual appearance of the scene we used illumination conditions that would be expected at Europa's equator. The contributing light sources include direct rays from sun and reflected light from Jupiter. We introduce a quasi-Phong bi-directional reflectance (BRDF) model (see Eq. (9)) that represents a regolith-ice mixture surface. This model is a compromise between physical accuracy and rendering efficiency. The irradiance E_i of each scene point is given by:

$$E_i = \rho_a + \alpha \times \rho_d(\hat{N}_i \cdot \hat{L}) + (1 - \alpha) \times \rho_s(\hat{R}_i \cdot \hat{V})^k$$
(9)

where \hat{N}_i , \hat{R}_i are the normal and reflection vectors of surface facet i; \hat{L} , \hat{V} are the lighting and view vectors; and α is a weighting function of the inverse phase angle that is zero when the lighting and view vectors oppose. The material properties of the surface include ρ_a , ρ_d and ρ_s which are the ambient, diffuse, and specular surface albedos respectively, and k is the Phong sharpness, which controls the width of the specular highlight. For the most part, we set $\rho_a = \rho_d$. Randomized albedo maps centered on the true value are generated to include "splotches" of variability on the terrain and striae, which are darker linear features (without corresponding geometry).

4.3 Modeling Thermal Appearance

An approximate thermal model was developed for Europa's surface conducive to real-time graphics. We first render the radiance of the Sun and Jupiter considering their apparent size (6.14 arcmin and 717.55 arcmin, respectively), Solar and Jupiter's irradiance $(50 W/m^2 \text{ and } 0.0897 W/m^2)$ and the reflection of the Sun in Jupiter $(0.0001 W/m^2)$. Shades are determined through the gradient of the procedurally-generated digital elevation model. The total absorbed irradiance is then corrected taking bolometric bond albedo (0.68) into account, and the temperature of Europa's surface is obtained through Stephen-Boltzman law considering the surface as a gray body with emissivity e = 0.94. We then use several smoothing passes on the temperature gradients in order to approximate the thermal diffusion at first order.

5 Results

In order to evaluate the SLAM performance we use the Euclidean norm of the position error per distance traveled E as a metric. For each trajectory, we obtained this metric through a linear regression of the position error versus the distance travelled. The results for the different trajectories are presented in Table 1.

Trajectory	$E\left(\% ight)$					
ITAJECTOLY	Vis	Vis T-IR				
1	12.31	6.56	6.60			
2	14.12	10.87	9.93			
3	8.04	3.11	3.36			
4	4.71	5.20	3.42			
5	3.62	4.32	3.71			
6	4.07	1.46	1.83			
7	7.15	5.18	4.76			
8	12.72	13.85	12.32			
9	3.85	2.49	1.61			
10	3.83	4.04	3.74			

Table 1: Euclidean norm of the position error per distance travelled E for ten trajectories. For each trajectory the position error accumulation rate was averaged over thirty trials

To compare the effectiveness of the algorithms, we compare the percent improvement in the error rate. The definition of the percent improvement is given in Eq. (10).

$$I_{\rm VisvsT-IR} = \frac{E_{\rm Vis} - E_{\rm T-IR}}{E_{\rm Vis}} \times 100\%$$
(10)

 E_V and E_T are the Euclidean norm of the position error of visible and thermal cameras, respectively, and similarly for the other pairings. A positive value means that an improvement was obtained by using the second SLAM sensing modality (e.g. *thermal*), and a negative value means the first SLAM sensing modality was better (e.g. *visible*) and vice-verse. In addition we compute Cohen's d to evaluate the effect size.

Table 2 summarizes the performance and the effect size. Table 3 reports the statistical confidence that the performance of the competing algorithms are different, using a ROPE analysis.

Table 2 shows the average improvement in visible and thermal-infrared (Vis vs T-IR), visible and dual modality (Vis vs Dual), and thermal-infrared and dual modality (T-IR vs Dual). Overall, the average improvement of *thermal* over *visible* is approximately 21.5%, and the fusion of the two modalities increases the average improvement to 31.2% over *visible* and 8.5% over *thermal*.

Comparing only the visible and thermal SLAM modalities, Table 3 shows that the thermal modality is superior to the visible modality for trajectories 1, 2, 3, 6, 7, and 9. In trajectory 10 thermal and visible are statistically undistinguishable. On trajectories 1, 7, and 9 the improvement from using thermal is very large (d > 1.20), and on trajectories 3 and 6 the improvement is huge (d > 2.0). For trajectories where *thermal* is not the clear winner (trajectories 4, 5, 8, and 10) in all but trajectory 5, the difference in performance is small (d < 0.5), as per Table 2. In the case of trajectory 5, there is a significant difference in performance, and the effect size is large (d > 0.8). However, in all other cases it is safe to say that the *thermal* modality is as good as or better than the *visible* modality. This result may be useful in informing sensor selection where Size, Weight, and Power requirements drive a choice between one sensor or the other.

However, if the possibility of using both instruments is available, the *dual* SLAM modality becomes the dominating approach. The benefits of fusing both modalities is expected, as there is twice as much information available than in one modality alone.

In Table 1 the visible camera SLAM has the best performance for trajectory 5, the thermal-IR SLAM has the best performance for trajectories 1, 3, and 6. In all other cases the dual modality SLAM approach has the best performance. Table 3 shows us that on trajectory 5 the performance of the *visual* and *dual* SLAM modalities are statistically equivalent. On trajectories 1, 3, and 6 the *thermal* and *dual* SLAM modalities have 82%, 56%, and 84% probabilities of being equivalent, respectively. Further, the effect size of the difference in performance (Table 2) in these cases is not large (< 0.8), and in the case of trajectories 1 and 3, the effect size is very small (< 0.2).

We can conclude with confidence from this analysis that from a pose estimation perspective, that using the *dual* modality SLAM does not cause a degradation in performance. Further, on the trajectories where either the *visual* or *thermal* modalities performs substantially better than the other (all but trajectory 10 and arguably 4), using the *dual* modality would successfully mitigate the risks from choosing only one sensing modality for pose estimation.

Trajectory	Vis vs T-IR		Vis	vs Dual	T-IR vs Dual		
ITAJECTOLY	I(%)	$\overline{I(\%)}$ Cohen's d		Cohen's d	I(%)	Cohen's d	
1	47.5	1.83	46.4	1.76	-2.2	0.09	
2	23.0	0.84	29.7	1.54	8.6	0.26	
3	61.3	2.16	58.2	2.41	-8.0	0.10	
4	-10.5	0.27	27.5	0.70	34.4	1.34	
5	-19.3	1.19	-2.4	0.15	14.2	1.25	
6	64.3	2.52	55.0	2.07	-25.0	0.61	
7	27.5	1.52	33.4	1.59	8.1	0.39	
8	-8.9	0.34	3.9	0.15	11.7	0.55	
9	35.3	1.37	58.2	2.48	35.4	0.80	
10	-5.5	0.37	2.4	0.18	7.5	0.52	
Average	21.5	_	31.2	_	8.5	_	

Table 2: Average relative percent improvement I for (Vis vs T-IR) visible and thermal-infrared, (Vis vs Dual) visible and dual modality, and (T-IR vs Dual) thermal-infrared vs dual modality. The overall average in error rate improvement is given, but the overall effect size is not reported.

Trajectory	Vis vs T-IR			Vis vs Dual			T-IR vs Dual		
Inajectory	Mode	HDI	% in ROPE	Mode	HDI	% in ROPE	Mode	HDI	% in ROPE
1	5.786	[4.707; 6.498]	0	5.656	[4.421; 6.926]	0	-0.182	[-0.823; 0.502]	82
2	3.276	[1.979; 4.711]	0	4.219	[3.164; 5.220]	0	0.600	[-0.591; 1.852]	40
3	4.892	[4.136; 5.928]	0	4.720	[3.953; 5.489]	0	-0.423	$[-1.276 \ 0.524]$	56
4	-0.501	[-1.131; 0.117]	50	1.305	[0.568; 2.036]	2	1.758	[1.349; 2.208]	0
5	-0.700	[-0.938; -0.455]	5	-0.094	[-0.309; 0.122]	100	0.624	[0.425; 0.813]	11
6	2.580	[2.221; 2.943]	0	2.154	[1.748; 2.580]	0	-0.386	[-0.612;-0.167]	84
7	2.006	[1.471; 2.501]	0	2.389	$[1.929 \ 2.824]$	0	0.393	[0.027; 0.762]	71
8	-1.104	[-2.432; 0.136]	15	0.443	[-0.646; 1.579]	48	1.530	[0.664; 2.565]	1
9	1.358	[0.948; 1.806]	0	2.225	[1.932; 2.532]	0	0.872	[0.416; 1.333]	5
10	-0.220	[-0428; 0.001]	99	0.092	[-0.061; 0.245]	100	0.299	[0.101; 0.499]	98

 $\frac{1}{5}$

Table 3: Paired Bayesian t-tests over the difference in rate of position error accumulation. Plots supporting these data are found in Appendix A.

Below, we present the error in the estimated vehicle position over time. Fig. 4 to Fig. 13 show the performance for trajectorys 1 to 10, respectively. For all plots, the error bars correspond to 1 95% confidence interval, computed for 30 trials. As can be seen, depending on the terrain the vehicle flew over, the estimated error can change dramatically.



Figure 4: Thermal-IR data has as good as or better error than the visible light camera. The fused posed estimation performs as good as or better than the Thermal pose estimation, until after 140 seconds, where the increasing error from the visible camera begins to dominate.



Figure 5: On trajectory 2, the dual sensing modality SLAM algorithm performs at least as good as or better than the thermal SLAM, which is in turn as good as or better than SLAM using the visible camera.



Figure 6: In trajectory 3 we see that the error in the visible light SLAM becomes very large. In this case we see that the dual modality rejects the visible light camera info and relies exclusively on the thermal camera data. This is evidenced by the statistically indistinguishable performance between the thermal-only SLAM and the dual modality SLAM.



Figure 7: Initially the pose estimation error of all three modalities is statistically indistinguishable until 140 seconds. From 140 seconds on, we see that the dual modality outperforms both the thermal and visible SLAM.



Figure 8: Here the thermal-IR SLAM begins to perform worse than either the visible SLAM or the dual modality SLAM at about 160 seconds. The dual modality SLAM is robust to this error. However, while the differences between Thermal-IR and the other SLAM approaches is statistically significantly different, the error is not substantially different, compared to the differences in performance on other trajectories.



Figure 9: Like trajectory 3, on trajectory 6 we see that the dual modality rejects a sudden sharp uptick in error of the visible light SLAM. Prior to 160 seconds the performance of all three approaches is statistically indistinguishable.



Figure 10: From 110 seconds to 130 seconds, the thermal-IR SLAM has a slight advantage over the visible light and dual modality SLAM. The error in the Thermal SLAM begins to increase, along with the visible camera, but we see that the dual modality SLAM is able to slow that increase in pose estimation error.



Figure 11: Along trajectory 8 the estimated pose error increases after 80 seconds for all SLAM approaches. The thermal SLAM ultimately performs the worst, followed by the visible light camera, but the dual modality SLAM manages to minimize the error.



Figure 12: The error in the visible camera SLAM starts to increase linearly from around 100 seconds. The thermal camera's error slows down after 100 seconds, which is mirrored by the dual modality. Ultimately the dual modality SLAM achieves the best performance, with statistical significance.



Figure 13: Over the course of the trajectory the performance of the SLAM approaches are statistically indistinguishable.

6 Conclusions

We have demonstrated two principal points in this paper. Firstly, a multisensor approach based on existing fundamentals (e.g. Brunner et al. (2011)) can be extended to behave in a sensible way on icy surfaces. Secondly, we have shown that the thermal-IR component, in particular, can significantly improve localization on simulated Europa data.

In the trajectories that were executed in experimentation, we have demonstrated that adding thermal imagery can drastically improve performance relative to unimodal visual SLAM in many cases. There were also some cases where the visible light camera performed better than unimodal thermal-infrared. Thus, we advocate for the inclusion of both modalities to maximize robustness and safety on operations over icy surfaces.

Our results also show that there is a trajectory dependence on the *utility* of the thermal-infrared imagery. This makes sense, as the effects of glare depend on the relative angles between the viewer and the illumination source, for example. Trajectory-dependent performance is an argument in favour of multi-modal fusion, otherwise, one would have to plan descent trajectories

which avoid glare, possibly reducing the areas which can be safely reached by a lander. Use of our approach can increase the number of feasible landing locations, thus increasing the ability to conduct science on icy moons.

6.1 Future Work

The results presented above demonstrate that using thermal imagery can help with camera-based localization on icy surfaces. However, the preceding work was conducted in a simulated environment, and the generalizability of the performance is dependent on the fideltiy of the simulation. To validate our results, this work should be continued in analog environments to further demonstrate the utility of thermal and visual-thermal navigation.

The SIFT feature detector is a reliable feature detector, however it can be computationally expensive. Given constrained space operations, being able to use simpler feature detectors would make the application more plausible in flight-forward missions.

However, before adopting alternative feature detectors, we would need to answer at least two questions. First, does the proposed feature detector negatively effect the algorithm's performance? If a new feature detector does not find stable landmarks, or it easily gets confused between landmarks, then it would not be suitable in this algorithm.

These claims are, of course, limited by the experimental setup. The simulated environment is only an approximation of the real world environment, so experimentation with actual icy surfaces is warranted. Moreover, unique edge cases and failure conditions should be further analyzed to validate robustness over a wider range of conditions and scenarios.

6.1.1 Improvements to the Landmark Selection Algorithm

In this work we used a spatial diversity score, D, (Section 3.2) to determine which landmarks should be added to the set of tracked landmarks. However, there are other features which could be used to bias that selection process, a few of which are discussed below. In future work we should experiment with these features to determine which, if any, improve the localization performance.

At least two additional features should be added to the weighting function. We include the expected duration in the field of view (FOV), T_{FOV} , and the local contrast of the landmark point, C. These features could be combined with the spatial diversity score using a simple weighting function (Eq. (11)). The values of the weighting terms (α , β , and γ) would need to be determined empirically.

$$W(l) = \alpha D(l) + \beta T_{FOV}(l) + \gamma C(l)$$
(11)

The reason for choosing the expected duration of the landmarks in the field of view is to maximize the number of frames that a landmark appears in. The longer a landmark is observed, the more accurate the estimation of its 3D position, making points with potentially longer durations of visibility more valuable.

The local contrast of the landmark is to make the landmarks more stable with respect to the particular feature detector we are using in this algorithm. SIFT landmarks with greater local contrast are more robust fluctuations in sensor readings, increasing the likelihood of the feature being observed for more frames.

However, the choice of local contrast is tailored to the SIFT algorithm. Should another feature detector be chosen, feature would need to be redesigned.

In Fig. 14 we plot the three proposed weighting terms for all the pixels in the frame, given the previous landmark tracking points. As above, we computed the score for every pixel in the frame. In practice, however, these values would only be computed for the landmarks identified by the feature detectors. As can be seen in Fig. 14a, this would only be for a small fraction of locations in a scene.



(a) Original frame

(b) Spatial Diversity (D)



(c) Duration in FOV (T_{FOV})

(d) Local contrast (C)

Figure 14: The components which make up the feature sampling weight. For illustrative purposes in the above images we have computed the values at every pixel. In the algorithm we only compute these values at candidate landmark sites.

Acknowledgements

We acknowledge support for this work from NASA's Planetary Science Division through grant SC-COLDTech16-0065. We acknowledge the support of NASA's International Internship (I^2) Program and the Foundation for Science and Technology (FCT Portugal). We thank Rick Elphic and Orkan Omurhan for assistance in sensor and planetary models.

References

- Lorenz Roth, Joachim Saur, Kurt D Retherford, Darrell F Strobel, Paul D Feldman, Melissa A McGrath, and Francis Nimmo. Transient water vapor at europas south pole. *Science*, 343(6167):171–174, 2014. 1
- WB Sparks, KP Hand, MA McGrath, E Bergeron, M Cracraft, and

SE Deustua. Probing for evidence of plumes on europa with hst/stis. *The* Astrophysical Journal, 829(2):121, 2016. 1

- CC Porco, P Helfenstein, PC Thomas, AP Ingersoll, J Wisdom, R West, G Neukum, T Denk, R Wagner, Th Roatsch, et al. Cassini observes the active south pole of enceladus. *science*, 311(5766):1393–1401, 2006. 1
- JR Spencer, JC Pearl, M Segura, FM Flasar, A Mamoutkine, P Romani, BJ Buratti, AR Hendrix, LJ Spilker, and RMC Lopes. Cassini encounters enceladus: Background and the discovery of a south polar hot spot. *science*, 311(5766):1401–1405, 2006. 1
- Candice J Hansen, L Esposito, AIF Stewart, J Colwell, A Hendrix, W Pryor, D Shemansky, and R West. Enceladus' water vapor plume. *Science*, 311 (5766):1422–1425, 2006. 1
- Amanda R Hendrix, Terry A Hurford, Laura M Barge, Michael T Bland, Jeff S Bowman, William Brinckerhoff, Bonnie J Buratti, Morgan L Cable, Julie Castillo-Rogez, Geoffrey C Collins, et al. The nasa roadmap to ocean worlds. *Astrobiology*, 19(1):1–27, 2019. 1
- Cynthia B Phillips and Robert T Pappalardo. Europa clipper mission concept: exploring jupiter's ocean moon. *Eos, Transactions American Geophysical Union*, 95(20):165–167, 2014. 1
- Jennifer Eigenbrode, Robert E Gold, Christopher P McKay, Terry Hurford, and Alfonso Davila. Searching for life in an ocean world: The enceladus life signatures and habitability (elsah) mission concept. In 42nd COSPAR Scientific Assembly, volume 42, 2018. 1
- KP Hand, AE Murray, JB Garvin, WB Brinckerhoff, BC Christner, KS Edgett, BL Ehlmann, CR German, AG Hayes, TM Hoehler, et al. Europa lander study 2016 report, europa lander mission. NASA Jet Propuls. Lab., La Cañada Flintridge, CA, USA, Tech. Rep. JPL D-97667, 2017. 1
- Robert Cook and Kenneth Torrance. A reflectance model for computer graphics. *Transactions on Graphics*, 1(1):7–24, 1982. 1, 3.3
- Thomas B Schön and Jacob Roll. Ego-motion and indirect road geometry estimation using night vision. In 2009 IEEE Intelligent Vehicles Symposium, pages 30–35. IEEE, 2009. 2
- Emil Nilsson, Christian Lundquist, Thomas B Schön, David Forslund, and Jacob Roll. Vehicle motion estimation using an infrared camera. *IFAC Proceedings Volumes*, 44(1):12952–12957, 2011. 2

- Christopher Brunner, Thierry Peynot, and Teresa Vidal-Calleja. Combining multiple sensor modalities for a localisation robust to smoke. In 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 2489–2496. IEEE, 2011. 2, 3, 3.3, 6
- Stephen Vidas and Sridha Sridharan. Hand-held monocular slam in thermalinfrared. In 2012 12th International Conference on Control Automation Robotics & Vision (ICARCV), pages 859–864. IEEE, 2012. 2
- Erika Emilsson and Joakim Rydell. Chameleon on firethermal infrared indoor positioning. In 2014 IEEE/ION Position, Location and Navigation Symposium-PLANS 2014, pages 637–644. IEEE, 2014. 2
- Christos Papachristos, Frank Mascarich, and Kostas Alexis. Thermal-inertial localization for autonomous navigation of aerial robots through obscurants. In 2018 International Conference on Unmanned Aircraft Systems (ICUAS), pages 394–399. IEEE, 2018. 2
- S. Kohlbrecher, O. von Stryk, J. Meyer, and U. Klingauf. A flexible and scalable slam system with full 3d motion estimation. In 2011 IEEE International Symposium on Safety, Security, and Rescue Robotics, pages 155–160, Nov 2011. doi: 10.1109/SSRR.2011.6106777. 2
- J. Kerekes, A. Goodenough, S. Brown, J. Zhang, B. Csatho, A. Schenk, S. Nagarajan, and R. Wheelwright. First principles modeling for lidar sensing of complex ice surfaces. In 2012 IEEE International Geoscience and Remote Sensing Symposium, pages 3241–3244, July 2012. doi: 10.1109/IGARSS.2012.6350733. 2
- Jeffrey S. Deems, Thomas H. Painter, and David C. Finnegan. Lidar measurement of snow depth: a review. *Journal of Glaciology*, 59(215):467479, 2013. doi: 10.3189/2013JoG12J154.
- Jong-Hyuk Kim and Salah Sukkarieh. Airborne simultaneous localisation and map building. In 2003 IEEE International Conference on Robotics and Automation (Cat. No. 03CH37422), volume 1, pages 406–411. IEEE, 2003. 3.1
- Tim Bailey. Constrained initialisation for bearing-only slam. In *IEEE Interna*tional Conference on Robotics and Automation, volume 2, pages 1966–1971. IEEE; 1999, 2003. 3.1
- Mitch Bryson and Salah Sukkarieh. Bearing-only slam for an airborne vehicle. In *Australasian Conference on Robotics and Automation*, volume 4, page 18. Citeseer, 2005. 3.1

- Thomas Lemaire, Simon Lacroix, and Joan Sola. A practical 3d bearing-only slam algorithm. In 2005 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 2449–2454. IEEE, 2005. 3.1
- Ngai Ming Kwok and Gamini Dissanayake. An efficient multiple hypothesis filter for bearing-only slam. In 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)(IEEE Cat. No. 04CH37566), volume 1, pages 736–741. IEEE, 2004. 3.1
- Joan Sola, André Monin, Michel Devy, and Thomas Lemaire. Undelayed initialization in bearing only slam. In 2005 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 2499–2504. IEEE, 2005. 3.1
- Joan Sola. Towards visual localization, mapping and moving objects tracking by a mobile robot: a geometric and probabilistic approach. PhD thesis, 2007. 3.1
- JM Martínez Montiel, Javier Civera, and Andrew J Davison. Unified inverse depth parametrization for monocular slam. Robotics: Science and Systems, 2006. 3.1
- Javier Civera, Andrew J Davison, and JM Martinez Montiel. Inverse depth parametrization for monocular slam. *IEEE transactions on robotics*, 24(5): 932–945, 2008. 3.1
- Daniele Marzorati, Matteo Matteucci, Davide Migliore, and Domenico G Sorrenti. Monocular slam with inverse scaling parametrization. In *BMVC*, pages 1–10, 2008. 3.1
- Joan Sola. Consistency of the monocular ekf-slam algorithm for three different landmark parametrizations. In 2010 IEEE International Conference on Robotics and Automation, pages 3513–3518. IEEE, 2010. 3.1
- Joan Sola, Teresa Vidal-Calleja, Javier Civera, and José María Martínez Montiel. Impact of landmark parametrization on monocular ekf-slam with points and lines. International journal of computer vision, 97(3):339–368, 2012. 3.1
- Joan Sola, André Monin, Michel Devy, and Teresa Vidal-Calleja. Fusing monocular information in multicamera slam. *IEEE transactions on robotics*, 24(5):958–968, 2008. 3.1
- Christopher G Harris, Mike Stephens, et al. A combined corner and edge detector. In Alvey vision conference, volume 15, pages 10–5244. Citeseer, 1988. 3.2

- Gabriel Takacs, Vijay Chandrasekhar, Sam Tsai, David Chen, Radek Grzeszczuk, and Bernd Girod. Unified real-time tracking and recognition with rotation-invariant fast features. In 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, pages 934–941. IEEE, 2010. 3.2
- David G Lowe. Distinctive image features from scale-invariant keypoints. International journal of computer vision, 60(2):91–110, 2004. 3.2
- Claude Elwood Shannon. A mathematical theory of communication. *Bell* system technical journal, 27(3):379–423, 1948. 3.3
- Mark Allan, Uland Wong, Terry Welsh, P. Michael Furong, Scott McMichael, Arno Rogg, Brian Gerkey, Ian Chen, Steven Peters, Morgan Quigley, Matthew Deans, and Terrence Fong. Planetary rover simulation for lunar exploration missions. In 2019 IEEE Aerospace Conference. IEEE, 2019. 4.1

Appendix A

Supporting Bayesian Analysis Plots



Figure A1: On trajectory 1 the *dual* and *thermal* modalities are equivalent with 83% probability. Both *dual* and *thermal* are improvements on the *visible* modality.

A.1 Trajectory 1





Figure A2: On trajectory 2 both the *thermal* and *dual* modalities are improvements on the *visible* modality. The *dual* modality is an improvement on the *thermal* modality with 60% probability.



A.3 Trajectory 3

Figure A3: On trajectory 3 the *dual* and *thermal* modalities have equivalent performance, with 56% probability. The *visible* modality is worse than either the *thermal* or *dual* modalities with 100% probability.





Figure A4: On trajectory 4 the *visible* and *thermal* modalities are equivalent with 50% probability. The *dual* modality improves on both the *visible* and thermal modalities with 98% and 100% probability, respectively.





Figure A5: On trajectory 5 the *visible* modality is better than the *thermal* modality with 95% probability. The *visible* and *dual* modalities have statistically equivalent behavior. The *dual* modality is better than the *thermal* modality with 89% probability.



Figure A6: On trajectory 6 the *dual* and *thermal* modalities are statistically equivalent with 84% probability. The *thermal* and *dual* are both statistically significantly better than the *visual* modality.

Trajectory 6

A.6



Figure A7: For trajectory 7 the *thermal* and *dual* modalities significantly improves on the *visible* modality. The *thermal* and *dual* modalities are statistically equivalent with 71% probability.

A.7 Trajectory 7





Figure A8: For trajectory 8 the *visible* modality is better than *thermal* with 85% probability. The *dual* modality is better than *visible* modality with 52% probability and the *thermal* modality with 99% probability.



Figure A9: On trajectory 9, all three sensing modalities perform statistically significantly different from one another. The *thermal* and *dual* modalities are better than the *visible* modality. We can say that the *dual* modality is better than the *thermal* modality with 95% probability.

Trajectory 9

A.9



Figure A10: On trajectory 10 none of the sensing modalities are statistically significantly different from one another.

A.10 Trajectory 10

REPORT DOCUMENTATION PAGE					Form Approved OMB No. 0704–0188			
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Department of Defense, Washington Headquarters Services, Directorate for Information and Reports (0704-0188), 1215 Jetferson Davis Highway, Suite 1204, Arlington, VA 2202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.								
1. REPORT DAT	E (DD-MM-YYYY)	2. REPO	RT TYPE ical Memorandum			3. DATES COVERED (From - To)		
4. TITLE AND S	UBTITLE	Ittelli			5a. CON	ITRACT NUMBER		
Fusion of vi	isible and the	ermal-infrared	l imagery for SLAM	for landing				
on icy moons					5b. GRANT NUMBER SC-COLDTech16-0065			
					5c. PROGRAM ELEMENT NUMBER			
6. AUTHOR(S)					5d. PROJECT NUMBER			
					5e. TASK NUMBER			
5f.						ORK UNIT NUMBER		
7. PERFORMING	GORGANIZATION	NAME(S) AND A	DDRESS(ES)		1	8. PERFORMING ORGANIZATION		
NASA Ames Research Center						REPORT NUMBER		
9. SPONSORING	G/MONITORING A	GENCY NAME(S)	AND ADDRESS(ES)			10. SPONSOR/MONITOR'S ACRONYM(S)		
National A Washington	eronautics an	d Space Adm	inistration			NASA		
washington	I, DC 20040-0	5001				11. SPONSOR/MONITOR'S REPORT		
						NASA/TM-2019-220241		
12. DISTRIBUTION	ON/AVAILABILIT Y l-Unlimited	STATEMENT						
Subject Ca	tegory							
Availability: NASA STI Program (757) 864-9658								
An electronic version can be found at http://ntrs.nasa.gov.								
14. ABSTRACT This paper Due to the make tradit visible light address pro surface sime methods. In to a visible average, con 15. SUBJECT T	addresses the possibility of ional vision-lo- cameras wit blems arising ulation which n this framew camera. Mon mpared to us ERMS	e problem of specular refl- based navigat h a thermal-i g from the ap n models the vork, we show reover, fusing ing a visible 1	localization for land ection as well as hig ion systems relying nfrared camera usin pearance of icy moo appearance and the r that thermal featur both sensing modal light camera alone.	ing on the su h bulk albed on visible im g inverse-dep ns. Results v thermal prop res improve l lities increase	urface of lo, icy su nagery u oth para were obt perties s localizat es the in	f icy moons, like Europa or Enceladus. urfaces present new challenges that unreliable. We propose augmenting ameterized monocular EKF-SLAM to tained from a novel procedural Europa simultaneously from physically-based tion by 23% on average when compared nprovement in localization to 31% on		
thermal SLAM, icy moons, gjgjgjgjgj								
16. SECURITY CLASSIFICATION OF: 17. LIMITAT		17. LIMITATION OF	18. NUMBER	19a. NAME OF RESPONSIBLE PERSON				
a. REPORT	b. ABSTRACT	c. THIS PAGE	ADOINAUI	PAGES .	S11 Information Desk (help@sti.nasa.gov)			
U	U	U	UU		19b. TEL (757) 8	EPHONE NUMBER <i>(Include area code)</i> 864-9658		
			•			Standard Form 298 (Rev. 8/98) Prescribed by ANSI Std. Z39.18		