Onboard Autonomy on the Intelligent Payload EXperiment (IPEX) Cubesat Mission: A pathfinder for the proposed HyspIRI Mission Intelligent Payload Module

Steve Chien*, Joshua Doubleday*, David R. Thompson*, Kiri L. Wagstaff*
John Bellardo**, Craig Francis**, Eric Baumgarten**, Austin Williams**,
Edmund Yee**, Daniel Fluitt, **Eric Stanton**, Jordi Piug-Suari**

* Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109-8099, USA
e-mail: firstname.lastname@jpl.nasa.gov

** California Polytechnic State University, San Luis Obispo, CA,
Email:{bellardo,jpuigsua}@calpoly.edu

Abstract

The Intelligent Payload Experiment (IPEX) is a cubesat that successfully launched in December 2013 and is currently flight validating autonomous operations for onboard instrument processing and product generation for the Intelligent Payload Module (IPM) of the Hyperspectral Infra-red Imager (HyspIRI) mission concept.

We first describe the ground and flight operations concept for HyspIRI IPM operations. We then describe the ground and flight operations concept for the IPEX mission and IPEX operations are validating the proposed HyspIRI IPM operations. We then detail the current status of the IPEX mission and results from the mission thus far.

1 Introduction

Future space missions will produce immense amounts of data. A single image from the HiRise camera on the Mars Reconnaissance Orbiter (MRO) spacecraft is 16.4 Gigabits (uncompressed). The HyspIRI mission concept under study [HyspIRI] proposes to have two instruments - the HyspIRI thermal infrared imager (TIR) projected to produce 1.2 million pixels per second with 8 spectral bands at 4 and 7.5-12 microns per pixel and the HyspIRI visible shortwave infrared (VSWIR) projected to produce 300 thousand pixels per second with 220 spectral bands per pixel in the 0.4-2.5 micron range. Keeping up with these data rates would require efficient algorithms, streamlined data flows and careful systems engineering.

The HyspIRI mission concept baselines using Direct Broadcast technology [GSFC] to rapidly deliver this data to application users on the ground. However, in order to leverage the existing DB network, this downlink path is limited to approximately 10 million bits per second. The Intelligent Payload Module (IPM) concept for the proposed HyspIRI mission is an onboard processing system intended to intelligently decide which data to downlink when, in order to maximize the utility of the DB system.

The HypsIRI IPM concept would involve both ground and flight automation (See Figure 0). On the ground, users would use Google Earth ™ to specify geographical and seasonal areas of interest. These requests would be automatically combined with predicted overflights to develop a schedule for onboard product generation and downlink [Chien et al. 2009]. Additionally onboard the spacecraft, the instrument data would be analyzed to search for specific event or feature signatures such as a forest fire, volcanic eruption, or algal bloom. These detected signatures could generate alerts or products that would be merged on a priority basis to drive spacecraft operations.

2 IPEX Cubesat Overview

IPEX is a 1 unit (1U) cubesat (Figure 1) [Chien et al. 2012] to flight validate technologies for onboard instrument processing and autonomous operations for NASA's Earth Science Technologies Office (ESTO).

As a 1U cubesat, IPEX is approximately 10cm x 10cm x 10cm. To support the IPEX primary flight software, IPEX carries a 400MHz Atmel ARM9 CPU (no hardware floating point) with 128MB RAM, 512MB flash memory, a 16 GB Micro SD card, and utilizes the Linux operating system. All six sides of the IPEX spacecraft have solar panels for electrical power generation providing 1-1.5W power generation when not

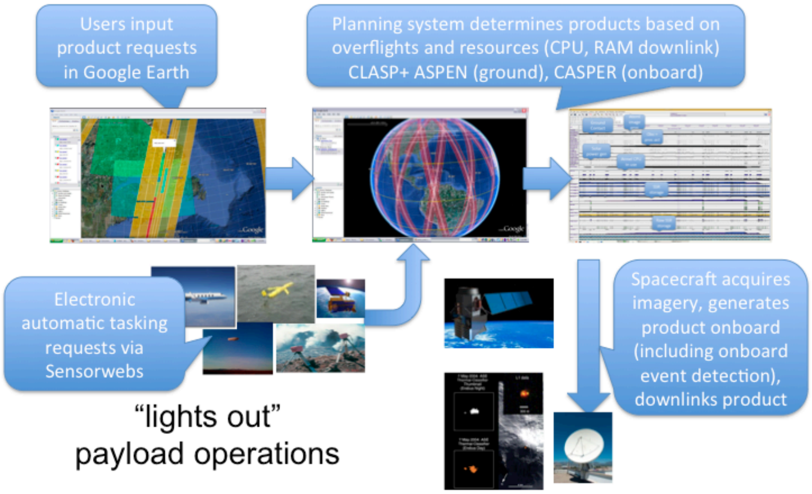

Figure 0: HyspIRI IPM Operations Concept

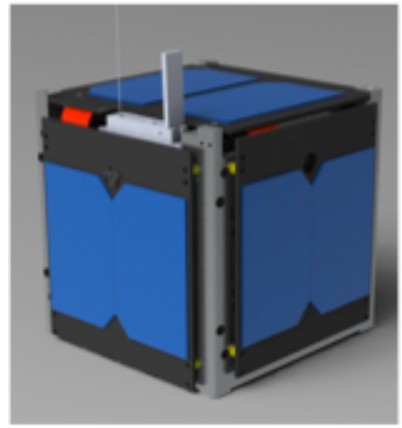

Figure 1: IPEX Model

in eclipse. The IPEX spacecraft uses passive magnetic attitude control to stabilize the CubeSat in low earth orbit. The spacecraft carries several batteries to enable operations in eclipse and continuous processing modes. IPEX carries five Omnivision OV3642 cameras, each producing images at 2048 x 1536 pixel resolution, 3 megapixels in size, with a finest instantaneous field of view of 0.024 degrees. With the IPEX orbit these cameras enable approximately 200m/pixel imagery of the Earth's surface at nadir.

IPEX also carries a Gumstix Earth Storm computer-on-module [Gumstix 2013] which includes an 800 MHz ARM processor, 512MB RAM, 512 MB NAND flash, utilizing the Linux operating system. The Gumstix utilizes less than 1W power and is on the majority of the time.

3 IPEX Ground and Flight Operations

IPEX demonstrates automated ground and flight operations of onboard autonomous processing of instrument data. In order to achieve this end, a range of capabilities and software are required.

3.1 IPEX Ground Mission Planning

The ground mission planning software for IPEX uses the CLASP [Knight & Hu 2009, Rabideau et al. 2010, Knight et al. 2012] planning system to determine the processing and downlink requests based on the projected overflight of the spacecraft.

These requests are then handled in a priority-based fashion by the ASPEN [Rabideau et al. 1999] system to generate a baseline schedule for several days operations as a forward looking baseline schedule. ASPEN must manage the ground contact schedule, eclipse schedule, observation activities, and onboard image processing activities. The onboard image processing activities involve a range of constraints including CPU usage, RAM usage, and downlink product size. The primary activities of image-acquisition and image-processing can also require significant data storage resources based on when the image is acquired versus when the Gumstix is powered on (thermal & power constrained) to process the image.

3.2 IPEX Onboard Planning

Onboard the spacecraft, the CASPER [Chien et al. 2000] planner manages spacecraft resources. CASPER models all of the same resources and constraints as ASPEN and modifies IPEX operations in response to deviations from the ground predicted plan such as: using more or less power than expected, activities taking longer or shorter than expected, or image products being larger or smaller than expected. CASPER also responds to onboard analysis of instrument data such as detection of features or events in imagery. Onboard processing is used to detect data of little value (e.g. images of dark space) early in processing activity. This analysis saves processing time, data-storage, and energy that would have been spent processing these less interesting images. In response, CASPER can schedule follow-on acquisitions from event or feature detection, or previously unscheduled lower priority data acquisition goals.

The CASPER model for IPEX represents a number of software processing workflows and a number of operations constraints.

The basic processing flow of the IPEX spacecraft is as follows.

1. Acquire imagery with a camera (ideally of a ground specified target area)

2. Process the image with a preliminary assessment which scores the image as likely of the Earth

3. Process the image on the Atmel processor with a range of selected image processing algorithms

4. Process the image on the Gumstix processor with a range of selected onboard algorithms.

5. Compare the generated products to determine if the products vary.

Additionally, at each earth contact, the spacecraft performs a number of actions.

1. Downlink engineering telemetry since the last ground contact

2. Downlink statistics on onboard processing (images acquired, images processed, runtimes, comparison results).

3. Downlink a small subset of the images and/or products for ground validation

The CASPER model for IPEX contains a number of resources including: the communications system, power, battery state of charge (energy), several data stores (Atmel SD flash, Gumstix flash, Gumstix SD flash), Atmel and Gumstix CPU resources, and camera resources.

The CASPER IPEX model also contains a number of activities including power generation (via solar panels), acquiring images, processing images using various algorithms, conversion of image formats, ground contacts, cleaning up file systems, solar view, eclipse, and activities pertaining to downlink.

CASPER onboard generally schedules ground-requested imaging, and onboard generated imaging requests and associated image processing along with each set of images acquired. CASPER onboard receives imaging time windows within which IPEX is allowed to image and process. This is to account for the constraint that when the IPEX payload board is powered (e.g. camera or gumstix usage) noise from this card reduces the ability of IPEX to receive uplinked signals.

With real-time telemetry updates from the system, CASPER can also take corrective actions when there are observed deviations and conflicts from the in-memory projections. For example, if the battery state-of-charge (SOC) reaches a critically low level, CASPER will periodically self issue windows of no-gumstix-processing, extending some time until after the SOC reaches a high-threshold. Disk utilization may also not follow the modelled usage, for example due to underperforming downlink bandwidth. As a configurable behaviour, CASPER will generate a cleaning activity to purge old data that has an unlikely chance of ever being downlinked, to make room for new data and progress towards mission milestones and goals.

3.3 IPEX Base Flight Software

The base flight software on IPEX is based on extensions and adaptation of the Linux operating system. The well-known System V *init* process is used directly to start, and restart if necessary, the principal components of the flight software: system manager for health monitoring, watchdog, beacon for real-time distribution of telemetry, datalogger for logging and archiving of telemetry and a sequence execution processes for real-time, time-based, and event-based commanding of the spacecraft.

4 IPEX Onboard Instrument Processing

IPEX validates a wide range of onboard instrument processing algorithms. The vast majority are variations of pixel mathematics, e.g. normalized difference ratios, band ratios, and similar products. For example, many flooding (surface water extent) classifications are based on band ratios [Brakenridge et al. 2005, Ip et al. 2006, Carroll et al. 2009]. Snow and ice products also use simple band processing formulae [MODIS]. Thermal anomaly detection algorithms such as for volcano [Wright et al. 2003, 2004, Davies et al. 2006] and active fire mapping [Justice et al. 2002] also involve computationally efficient slope analysis of spectral signals. Finally, a wide range of vegetation indicators also involve difference ratios or similar computations [Perry and Roberts 2008].

IPEX is also flying more computationally complex image processing technologies. These include: Support Vector Machine Learning Techniques [Cortes and Vapnik 1995, Doggett et al. 2006], spectral unmixing techniques [Bornstein et al. 2011], and TextureCam [Thompson et al. 2012] Random Decision forest classification techniques.

4.1 TextureCam Onboard Image Analysis

IPEX incorporates scene analysis based on the TextureCam image processing suite [Bekker 2014, Wagstaff 2013, Foil 2014]. It generates a pixelwise map of four pixel categories: (1) Clear surface, which could be land or ocean; (2) the planetary limb, or haze; (3) clouds; and (4) outer space. The classification reveals the image fraction subtended by the planetary disk, and the fraction of that surface which is cloud-free terrain or ocean. It assists with downlink prioritization for IPEX, and also demonstrates the analysis technique as a precursor for future use onboard rovers and surface spacecraft [Francis 2014]. Figure 2 shows an example of a typical input image and the classification result

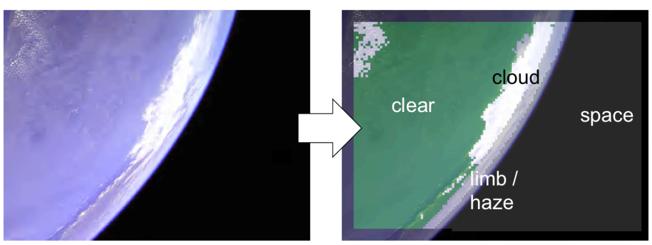

Figure 2: TextureCam image analysis uses a random forest model to classify image pixels. This simple scene required 29 seconds to classify onboard.

The analysis is a machine learning approach based on a random forest classifier [Shotton 2008], an ensemble of simple "decision tree" models fit to subsets of hand labeled training data. Each decision tree is a branching sequence of simple threshold tests. At runtime, control flow begins at the top node and descends down the tree until reaching a terminal leaf node where a classification probability is assigned. The intermediate nodes are binary tests on the results of specific pixel arithmetic operations such as ratios, differences and sums. Each decision tree independently estimates the probability that a pixel belongs to each of the four classes. These probabilities are aggregated across 16 trees to produce a final classification decision.

Decision trees have flown previously on the EO-1 spacecraft [Chien 2005]. The IPEX decision forest advances this strategy in two main ways. First, it runs multiple trees in parallel as described above. This provides statistical regularization [Breiman 2007] without the need for explicit pruning. Second, it analyzes spatial neighborhoods to incorporate local morphology and texture. Each tree node's binary operation applies to two specific pixels at defined locations relative to the pixel being classified. These pixels lie in a local radius-20 neighborhood, and all possible combinations produce hundreds of thousands of potential features. We train each decision tree on a random subset of these features using an expected information gain criterion. The training process grows each tree from the root node outward, selecting the feature from its set which can be thresholded to produce the best expected information gain to the population of pixels reaching that node. For further explanation of the random forest applied to computer vision and planetary science, we refer the reader to previous work [Shotton 2008, Foil 2013, Wagstaff 2013, Bekker 2014].

To reduce runtime, the software classifies every 10th pixel along vertical and horizontal directions, filling in the remainder with nearest-neighbor interpolation. The classifier takes less than a minute to fully classify an image on the IPEX Overo processor. In addition to classifying scene content, the IPEX software uses a connected components analysis to determine the center of each major contiguous class region. The centers are used as the locations of thumbnail subimages downlinked with the telemetry as a parsimonious description of image content. After using connected components to identify contiguous image areas, a distance transform finds the centerpoint of the thumbnail that is farthest away from any neighbor region. This approach is similar to the method used for target selection in previous autonomous science onboard the Mars Exploration Rovers [Estlin 2012].

The IPEX random forest was trained prior to launch using just four hand-labeled images from a high altitude balloon flight. Figures 2 and 3 show results from the onboard classification. Figure 2 is a simple scene that requires approximately 29 seconds. Figure 3 is a more complex scene that requires 48 seconds to complete. Accuracy for these images is better than 95%, but performance varies by image; the classifier was not intended for use over dark terrain or on images with significant sun glare artifacts, so it fails (as expected) in these cases. However, our initial tests suggest it performs well for favorable imaging conditions. To our knowledge this is the first time a machine learning system has been trained on a sub-orbital flight and then successfully used on orbit. For future work we will consider revising the decision forest by retraining with a wide range of orbital scenes.

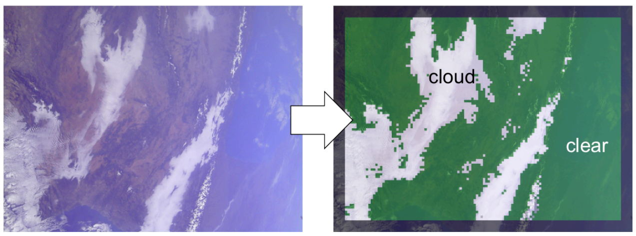

Figure 3: TextureCam result for a more complex scene. The onboard classification executed in 49 seconds.

4.2 Salience Onboard Image Analysis

IPEX also employs an unsupervised method for identifying images with potentially interesting content. It serves as a complement to supervised methods such as TextureCam that can highlight areas that correspond to classes defined *a priori*. Instead, this algorithm proceeds without knowledge of potential pixel classes and highlights areas within an image that are statistically salient (i.e., areas that stand out from their surroundings). Koch and Ullman [1985] first proposed the use of a visual salience map to model human visual attention. Itti and Koch [2000] combined multiple color, intensity, and orientation features into a global salience map, and Elazary and Itti [2008] showed that this salience correlates well with human judgments of what is most interesting in a scene.

We used a pixel-based measure of visual salience that incorporates local context, as proposed by Wagstaff et al. [2008]. The salience score S for a pixel p is defined as the contrast-weighted difference between that pixel's intensity and the distribution of intensities found in a spatial window around the pixel:

$$S = \frac{1}{M}\sum_{i} |p - i| P_w(i)$$

where i ranges over the possible intensity values and $P_w(i)$ is the probability of observing intensity i in the surrounding window. The sum is normalized by M, the maximum possible salience value for any pixel in the specified window, so all scores range from 0 to 1.

Onboard the spacecraft, we compute salience scores across each image that is collected. We apply the algorithm to a downsampled version of the image using a 32 x 32 pixel window to identify the five most salient regions within the image, stepping the window by half of its size across the image. Thumbnails of these regions, along with their salience scores, are saved out for downlink and examination on the ground. If the thumbnails are sufficiently interesting, we can request that the entire full-resolution image be sent down. If these most salient regions are uninteresting (low salience scores and/or no features of interest present), the image can be skipped and the bandwidth allocated to other data.

Figure 4 shows an image collected onboard IPEX and the salience algorithm's corresponding output. The broad cloud-free region is the Taklamakan Desert in China. The first two sub-regions selected by the algorithm exhibit interesting cloud structure. Regions 3, 4, and 5 pick out interesting ground features, including three lakes in Tibet (Lungmu Co, Gozha Co, and Bangdad
Co).

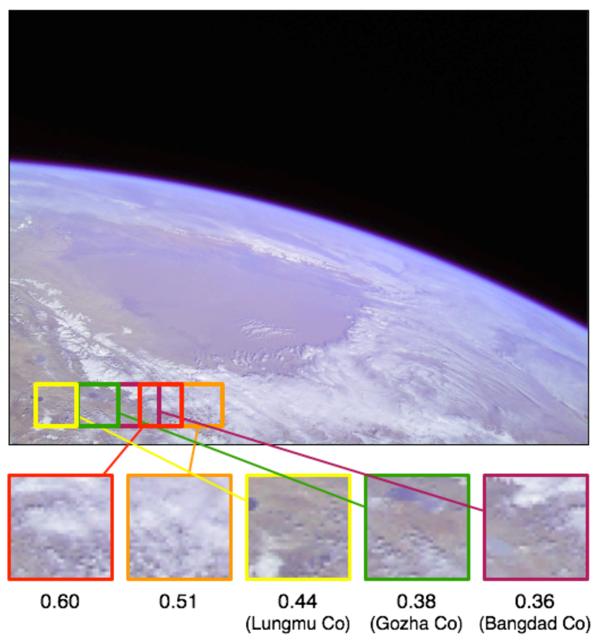

Figure 4. Example IPEX image and the top five most salient regions selected in an unsupervised fashion, along with their salience scores. These regions include interesting cloud structure, and several lakes.

To our knowledge, this is the first time image salience analysis has been performed onboard a spacecraft. Since the technique is based solely on image pixel statistics, it requires no prior training and executes deterministically. It provides an objective assessment of the visual features contained in the images, and the output scores can be used to aid in ranking items for downlink when data collection exceeds available bandwidth. High salience scores could also be used as an indicator that follow-up imaging at higher resolution is merited. More generally, the region selection process provides a valuable focus-of-attention guide to zoom in on features of interest.

5 IPEX Mission Results

The IPEX mission success criteria consisted of toeo parts: (1) demonstrate autonomous onboard product generation and (2) demonstrate autonomous operations of payload operations. As this paper goes to press (April 2014), the IPEX mission has achieved its full mission success technology validation criteria. With respect to onboard product generation, we have autonomously generated and validated over 30,000 image products (detailed statistics shown below).

Product Type	# of Products Generated
Band Ratio Images and Histograms	17108
Ground Loaded images validated	12300
TextureCam Thumbnails	3290
Salience Thumbnails	5920

With respect to autonomous operations, as this paper goes to press, IPEX has acquired and processed over 450 images, operated over 40 days of autonomous operations, acquiring over 93 autonomous response scenes and scores of idle/filler imaging requests fulfilled.

6 Related Work, Future Work, and Conclusions

The Remote Agent controlled the Deep Space One spacecraft for approximately two days in 1999 [Muscettola et al. 1998]. The Autonomous Sciencecraft on the Earth Observing One (ASE) spacecraft has pioneered onboard instrument data analysis [Chien et al. 2005]. In particular ASE highlighted onboard product generation for volcanology [Davies et al. 2006], flooding [Ip et al. 2006], and cryosphere [Doggett et al. 2006] disciplines. However, ASE did not have to deal with high data rate streams that challenge IPEX and the proposed HyspIRI mission and HyspIRI Intelligent Payload Module.

Onboard the Mars Exploration Rovers, the WATCH software enables automatic processing of imagery to track dust devils and cloud features [Castano et al. 2008]. Also onboard the MER rovers the AEGIS software enables onboard retargeting for targets of geological interest [Estlin et al. 2012].

We have described the IPEX mission to flight validate autonomous operations and onboard instrument processing. The IPEX mission demonstrates low cost, autonomous ground and flight mission operations enabling end users to specify image processing and product requests.

Acknowledgements

Portions of this work were performed by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

The IPEX project was funded by NASA's Earth Science Technology Office.

Contributors to the TextureCam codebase include Alphan Altinok, Dmitriy Bekker, Brina Bue, Daniel Howarth, Greydon Foil. The TextureCam project is supported by the NASA Astrobiology Science and Technology Instrument Development program (NNH10ZDA001N-ASTID).

The image salience algorithm development was supported by the AMMOS technology program.

References

D. L. Bekker. D. R. Thompson, W. J. Abbey, N. A. Cabrol, R. Francis, K. S. Manatt, K. F. Ortega, K. L. Wagstaff (2014). A Field Demonstration of an Instrument Performing Automatic Classification of Geologic Surfaces. Astrobiology (in press).B. Bornstein, R. Castano, S. Chien, D. Thompson, B. Bue, "Spectral segmentation and endmember detection onboard spacecraft," IEEE WHISPERS Workshop, Lisbon, Portugal, June 2011.

G.R. Brakenridge and E. Anderson, MODIS-based flood detection, mapping, and measurement: the potential for operational hydrological applications. In:

Transboundary Floods, Proc. of NATO Advanced Research Workshop, Baile Felix – Oradea, Romania, May 4-8, 2005.

Breiman, L. (2001) Random forests. *Machine Learning* 45:5–32, doi:10.1023/A:1010933404324.

M. Carroll, Townshend, J., DiMiceli, C., Noojipady, P., Sohlberg, R. (2009), A New Global Raster Water Mask at 250 Meter Resolution. International Journal of Digital Earth, December 2009.

A. Castano, A. Fukunaga, J. Biesiadecki, L. Neakrase, P. Whelley, R. Greeley, M. Lemmon, R. Castano, S. Chien, "Automatic detection of dust devils and clouds at Mars," *Machine Vision and Applications*, October 2008, vol. 19, No. 5-6, pp. 467-482.

C. Cortes and V. Vapnik, "Support-Vector Networks", Machine Learning, 20, 1995.

S. Chien, R. Knight, A. Stechert, R. Sherwood, and G. Rabideau, "Using Iterative Repair to Improve the Responsiveness of Planning and Scheduling," *Proceedings of the Fifth International Conference on Artificial Intelligence Planning and Scheduling*, Breckenridge, CO, April 2000.

S. Chien, R. Sherwood, D. Tran, B. Cichy, G. Rabideau, R. Castano, A. Davies, D. Mandl, S. Frye, B. Trout, S. Shulman, D. Boyer, "Using Autonomy Flight Software to Improve Science Return on Earth Observing One, *Journal of Aerospace Computing, Information, & Communication*, April 2005, AIAA.

S. Chien, D. Silverman, A. Davies, D. Mandl, Onboard Science Processing Concepts for the HyspIRI Mission, *IEEE Intelligent Systems*, November/December 2009 (vol. 24 no. 6), pp. 12-19.

S. Chien, J. Doubleday, K. Ortega, T. Flatley, G. Crum, A. Geist, M. Lin, A. Williams, J. Bellardo, J. Puig-Suari, E. Stanton, E. Yee, "Onboard Processing and Autonomous Operations on the IPEX Cubesat Mission," Intl Symposium on Space Artificial Intelligence, Robotics, and Automation for Space, Turin, Italy, Sep 2012.

A.G. Davies S. Chien, V. Baker, T. Doggett, J. Dohm, R. Greeley, F. Ip, R. Castaño, B. Cichy, G. Rabideau, D. Tran and R. Sherwood, "Monitoring Active Volcanism with the Autonomous Sciencecraft Experiment (ASE) on EO-1," Remote Sensing of Environment, vol. 101, no. 4, 2006, pp. 427–446.

T. Doggett,, R. Greeley, A. G. Davies, S. Chien, B. Cichy, R. Castano, K. Williams, V. Baker, J. Dohm and F. Ip (2006) Autonomous On-Board Detection of Cryospheric Change. *Remote Sensing of Environment*, Vol. 101, Issue 4, pp. 447-462.

L. Elazary, and L. Itti (2008). Interesting objects are visually salient. *Journal of Vision,* 8(3), 1–15.

T. Estlin, B. Bornstein, D. Gaines, R. C. Anderson, D. Thompson, M. Burl, R. Castaño, M. Judd, AEGIS Automated Targeting for the MER Opportunity Rover, ACM Transactions on Intelligent Systems and Technology, V. 3, No. 3, 2012.

T. Flatley, SpaceCube Processor, <http://hyspiri.jpl.nasa.gov/downloads/public/2.9%20Flatley%20-%20OnBoard%20Processing%20Technology.pdf

G. Foil, D.R. Thompson, W. Abbey, and D. Wettergreen, Probabilistic surface classification for rover instrument targeting. *IEEE/RSJ International Conference on Intelligent Robotic Systems* 775-782, doi:10.1109/IROS.2013.6696439.

R. Francis, K. McIsaac, D. R. Thompson, and G. R. Osinski, (2014) Multiclass linear discriminant analysis as a tool for image segmentation in geology. *IEEE International Conference on Robotics and Automation*, submitted.

Goddard Space Flight Center, http://directreadout.sci.gsfc.nasa.gov/

Gumstix Earth Storm, https://www.gumstix.com/store/product_info.php?products_id=264

HyspIRI Mission, http://hyspiri.jpl.nasa.gov

L. Itti and C. Koch. (2000). A saliency-based search mechanism for overt and covert shifts of visual attention. *Vision Research,* 40, 1489–1506.

F. Ip, J. Dohm, V. Baker, T. Doggett, A. Davies, R. Castano, S. Chien, B. Cichy, R. Greeley, R. Sherwood, D. Tran, G. Rabideau, Flood detection and monitoring with the Autonomous Sciencecraft Experiment onboard EO-1, Remote Sensing of Environment, Volume 101, Issue 4, 30 April 2006, Pages 463-481.

C. O. Justice, L. Giglio, S. Korontzi, J. Owens, J. Morisette, D.P. Roy, J. Descloitres, S. Alleaume, F. Petitcolin, Y. Kaufman, Y. 2002. The MODIS fire products. Remote Sensing of Environment 83:244-262.

C Koch and S. Ullman, (1985). Shifts in selective visual attention: towards the underlying neural circuitry. *Human Neurobiology,* 4, 219–227.

R. Knight, D. Mclaren, S. Hu, "Planning Coverage Campaigns for Mission Design and Analysis: CLASP for the proposed DESDyni Mission," Proceedings International Symposium on Artificial Intelligence, Robotics, and Automation for Space, Turin, Italy, September 2012.

R. Knight and S. Hu. "Compressed Large-scale Activity Scheduling and Planning (CLASP)," Proceedings of the Sixth International Workshop in Planning and Scheduling for Space (IWPSS-2009), July 19th-21st, Pasadena, California, 2009.

MODIS Snow Ice Web site, http://modis-snow-ice.gsfc.nasa.gov/

N. Muscettola, P. Nayak, B. Pell, B. Williams, "**Remote agent: To boldly go where no AI system has gone before**, Artificial Intelligence, 1998 – Elsevier.

E. Perry and D. Roberts, "Sensitivity of Narrow-Band and Broad-Band Indices for Assessing Nitrogen Availability and Water Stress in an Annual Crop, Agronomy Journal, vol. 100 Issue 4, 2008, pp. 1211-1219.

G. Rabideau, S. Chien, D. Mclaren, R. Knight, S. Anwar, G. Mehall, "A Tool for Scheduling THEMIS Observations, International Symposium on Space Artificial Intelligence, Robotics, and Automation for Space, Sapporo, Japan, August 2010.

G. Rabideau, R. Knight, S. Chien, A. Fukunaga, A. Govindjee, "Iterative Repair Planning for Spacecraft Operations in the ASPEN System," *International Symposium on Artificial Intelligence Robotics and Automation in Space*, Noordwijk, The Netherlands, June 1999.

J. Shotton, M. Johnson, and R. Cipolla (2008) Semantic texton forests for image categorization and segmentation. *IEEE Conference on Computer Vision and Pattern Recognition (2008)* 259(4 Pt 2):1–8.

D. R. Thompson, A. Allwood, D. Bekker, N. A. Cabrol, T. Estlin, T. Fuchs, K. L. Wagstaff, "TextureCam: Autonomous image analysis for astrobiology survey," Lunar and Planetary Sciences Conference, Houston, TX, (2012).

D.R. Thompson, B. Bornstein, S. A. Chien, S. Schaffer, D. Tran, B. Bue, R. Castano, D. Gleeson, and A. Noell (2013). Autonomous spectral discovery and mapping onboard the EO-1 spacecraft. Geoscience and Remote Sensing, IEEE Transactions on, 51(6), 3567-3579.

K.L. Wagstaff, J. Panetta, A. Ansar, R. Greeley, M. Pendleton Hoffer, M., Bunte, and N. Schörghofer, (2012). Dynamic landmarking for surface feature identification and change detection. *ACM Transactions on Intelligent Systems and Technology*, 3(3), article number 49.

K. L. Wagstaff, D.R. Thompson, W. Abbey, A. Allwood, D.L. Bekker, N.A. Cabrol, T. Fuchs, and K. Ortega (2013) Smart, texture-sensitive instrument classification for in situ rock and layer analysis. *Geophysical Research Letters* 40(16):4188–4193.

Wright, R., Flynn, L. P., Garbeil, H., Harris, A., & Pilger, E. (2003). Automated volcanic eruption detection using MODIS. Remote Sensing of Environment, 82, 135– 155.

Wright, R., Flynn, L. P., Garbeil, H., Harris, A. J. L., & Pilger, E. (2004). MODVOLC: Near-real-time thermal monitoring of global volcanism. Journal of Volcanology and Geothermal Research, 135, 29–49.