

Science-driven Spacecraft Autonomy

Paul Stolorz, Victoria Gor and Richard Doyle
Information and Computing Technologies Research Section
Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109
818-306-6197
(pauls,gor,rdoyle)@aig.jpl.nasa.gov

Clark Chapman, Randy Gladstone, William Merline and Alan Stern
Southwest Research Institute
1050 Walnut St, Suite 426
Boulder, CO 80302-5143
303-546-9670

Abstract --Autonomy concepts now exert a powerful influence on the development of a wide range of aerospace systems. Their value for automating spacecraft functions such as guidance, navigation and control is well-recognized. However, relatively little effort has been devoted in the past to their use in the context of payload data processing and understanding. In this paper, we focus on the use of autonomy to enable and enhance scientific goals for spaceborne missions. We argue that the general notions of autonomy can be applied directly to a broad range of scientific problems in ways that have never before been considered. We begin by outlining our general philosophy, and by describing the ways in which autonomy can be used to transform the ways in which spaceborne science is conducted. This is followed by a description of two novel systems that we have developed to exploit this philosophy for planetary missions.

1. INTRODUCTION

The development of autonomous spacecraft based on the use of powerful microelectronics and flight computers is a major thrust within the overall NASA goal of launching smaller and lighter missions, as exemplified by initiatives such as JPL's New Millennium Program,

Much of the initial focus for spacecraft autonomy has been on developing new

software and systems concepts to automate engineering functions of the spacecraft: guidance, navigation and control, fault protection, and resource management. However, the ultimate objectives of NASA missions are science objectives. We propose here a new framework for performing science data evaluation and observation planning autonomously onboard spacecraft. The future NASA mission set will feature smaller and more numerous spacecraft in an environment of highly constrained uplink and downlink communications, requiring substantial onboard computation to achieve its goals. The proposed paradigm will enable mission activities to be directed by scientists without the assistance of a ground sequencing team, robust capture and redirection in making discoveries at the target body, accommodation of the realities of limited communication links, and the return of quality science products from missions.

We have identified several autonomy capabilities that are particularly important in the context of spaceborne science:

- Autonomous identification of features and objects of known interest in onboard acquired images and spectra.
- Prioritization and/or edit downlink on the basis of reliable recognition of such features and objects.
- Systematic capture of transient science events through integration of autonomous onboard science data processing with

autonomous onboard capabilities for retargeting admission planning.

- Efficient Pi-driven redirection of mission activities following scientific discoveries at the target body.

A number of intelligent systems technologies can be used to build such an autonomy capability for science. They include data mining technologies, especially pattern recognition, machine learning and knowledge discovery techniques, as well as other capabilities of an autonomous spacecraft, particularly onboard planning. Our own work builds upon previous work on ground-based automated science data processing, notably the SKICAT system for automating the generation of a comprehensive sky object catalog from Mt. Palomar observatory data [2],[4], the JARTool system for automating the detection of volcanos in data returned from the Magellan spacecraft at Venus [1],[6], and the QUAKEFINDER system for automatic detection and measurement of ground movement in remote sensing imagery [7].

To illustrate the potential for science-driven autonomy, we have designed and implemented two prototype software systems designed to tackle two well-identified scientific calculations in a spaceborne setting. The first system, *Satellite-Detector*, is designed to automatically detect and flag natural satellites of asteroids. By performing these identifications onboard a spacecraft in near real-time, *Satellite-Detector* is able to provide inputs to autonomous spacecraft executives to enable mission replanning and retargeting of the spacecraft and/or its detectors if a scientific object of great interest and importance is found. We describe the design and implementation of a prototype ground-based system, and discuss possible scenarios for its use in scientific exploration.

We also describe its application to a specific test case, namely the identification of the moon Dactyl of asteroid Ida from images taken by JPL's Galileo spacecraft in 1993. Using 4 frames of Ida taken by Galileo, our system was able in a blind test to identify Dactyl, with zero false alarms, when the spacecraft was sufficiently far away that

Dactyl consisted of a single image pixel barely registering above background. The success of this prototype in reliably distinguishing the true satellite from background detector and cosmic ray noise, using a totally automated system implementable onboard, demonstrates the power of this approach in an important domain of serendipitous scientific discovery.

The second system, *UV Spectral-Analyzer*, is designed to support the implementation of autonomous spaceborne UV spectrometer experiments. It is based upon automatic iterative identification of spectral lines using estimates of UV-detector point spread functions. By constructing a system that is completely automatic, the analyzer can be deployed onboard spacecraft to improve scientific return. Rapid onboard analysis of low-resolution spectra can give valuable clues to the possible existence of interesting or unexpected spectral lines in scientific targets such as cometary comae and planetary atmospheres. This allows regions of special interest to be targeted autonomously for high-resolution imaging. In the absence of intriguing spectral makeup, the spectrometer can instead be directed to scan across the comet, a course of action which is likely to lead to greater scientific dividends in this case. We describe the structure and implementation of a prototype automatic UV spectral analyser, and explore scenarios for its use in the context of autonomous science missions.

2. AUTONOMOUS SATELLITE DETECTION - BACKGROUND

Until relatively recently, the possible existence of natural satellites orbiting asteroids was a controversial notion in the planetary science community [3],[5],[8],[9]. Regarded generally as a relatively rare event, it has not been a prominent consideration in the design of deep space missions. However, the exciting discovery of the satellite Dactyl orbiting around asteroid Ida by JPL's Galileo spacecraft has spurred a re-evaluation of the prevailing view about the likely abundance of natural satellites in the solar system.

This new development has ushered in a mindset in which systematic searches for natural satellites can be contemplated as a feasible

scientific goal. This perspective suggests several interesting possibilities for scientific discovery driven by spacecraft autonomy. For example, an automated onboard satellite detector offers the potential to detect and flag interesting and unexpected satellites for inspection and retargeting during the course of a mission. We describe in this paper a prototype system that has been designed to perform this task. We plan to extend this initial system to allow the determination of satellite orbits in addition to flagging their existence.

Considerable scientific benefits can be expected to accrue from successful automated detection of satellites onboard autonomous spacecraft. One possibility will be to follow the satellite for sufficient duration to determine its orbit, and from that information to infer the mass of the central asteroid. The size, orbital parameters and composition of a satellite hold significant clues to the origin and age of the asteroid itself. They also have larger implications for understanding the evolution of asteroids as a whole, and consequently understanding their role in the evolution of our solar system. Another possibility is to produce high-resolution images by retargeting the spacecraft or its detectors in order to study important issues such as cratering history, satellite shape and surface geology. In addition spectral data can be obtained in order to study satellite composition.

We are concerned in this paper with the fundamental task of identifying candidate satellites in situations in which they consist of a very small number of image pixels (perhaps only one) registering barely above background. This is the situation that exists as a spacecraft first approaches a known asteroid target at far-field, and is the most critical time to flag a new satellite if the spacecraft is to be given time to react to a detection. In these circumstances, satellites cannot be detected from individual images as it is impossible to distinguish them from transient noise sources such as cosmic rays. Several images are required in order to recognize a persistent object. In addition, detector defects and background stars must be removed as sources of noise.

The most general form of the satellite detection problem is somewhat more complex than this stationary case. It is necessary in general to account for the fact that both the spacecraft and the potential satellite may be in motion over the time period spanned by the relevant image series. Our prototype, which we describe below, has considered the simplified situation in which both the spacecraft and the satellite are stationary. We will briefly discuss the more general case at the end of the paper.

3. AUTONOMOUS SATELLITE DETECTION - STATIONARY CASE

The basic idea of the approach is to automatically search the "1 Hill radius" around a target body for possible natural satellites. The Hill radius is approximately 100 times the radius of the target body. It is the distance at which the gravitational field of the body itself is balanced by the field of the sun.

The *Satellite-Detector* system is built around an algorithm which detects a satellite by searching for persistence targets across several frames in order to separate it from transient cosmic ray and detector noise. In general, such a system must account for stellar background as well by comparing the observed field with an onboard star catalogue, but in our initial tests there was no stellar interference. The system first ingests several images containing the central asteroid body, and then implements the following steps.

- The initial (and most computationally expensive) task consists of co-registering each image with respect to the asteroid target.
- Suitably registered, each image is then binarized using a threshold determined by histogramming of individual images.
- In a perfect world persistence could now be established by simple pixel-wise addition of each binarized image and the application of a suitably stringent threshold. However, in practice it is necessary to account for less-than-perfect co-registration effects. This is accomplished by a spatial clustering algorithm.


- For the initial dataset of Ida/Dactyl, a cluster of nearest neighbor pixels in each direction was enough to extract the satellite from the dataset.

The prototype was successfully tested on the Ida-Dactyl images taken by Galileo, for which background stars, detector defects, and cosmic ray hits are the main source of noise. At the farthest range the satellite Dactyl is one pixel large and only a few intensity levels higher than background, and lower in intensity than many of the cosmic ray hits. The detection is performed long before the encounter in order to encompass the entire Hill sphere, to avoid rapid changes in geometry due to spacecraft motion, and to allow time for non-intrusive processing and replanning. The process is capable of incorporating short image


sequences where images are taken in different frequency bands, although we have not yet tested this aspect.

The software was tested on all the available images of Ida and Dactyl. Its performance was perfect. No manual parameter selection was necessary (all the parameters were selected autonomously by built-in procedures). The farthest sequence available was collected at a distance of 171,318 km from the asteroid center, 3 hours 50 minutes before the encounter. The results are shown in Figure 1 below. Shown are the IR-8890 image and the Green image, together with the final detections. It is anticipated that data at even farther ranges can be successfully processed, while it might be necessary to use subpixel registration for especially far ranges.


(a)



(b)



(c)



4. AUTOMATED ANALYSIS AND DECISION MAKING FOR ULTRAVIOLET SPECTRA

Another important use of automated onboard science processing is in the analysis and decision making for the onboard analysis of planetary ultraviolet spectra. Ultraviolet spectroscopy is one of the most commonly used and powerful techniques for the exploration of planetary and cometary atmospheres.

UV spectrographs have been carried to every planet visited by spacecraft, and are planned components of several future Discovery missions. Cassini orbiter, the Rosetta Orbiter and the Pluto mission. UV spectra derived from these instruments are useful for obtaining the compositions and first-order physical properties (e.g., temperature, pressure) of upper planetary atmospheres and cometary comae. UV spectra are also useful for categorizing the excitation mechanisms of planetary airglows, electrolights, aurora, and other phenomena, and for certain types of surface photometry.

In this paper we demonstrate the automated first-order interpretation of UV spectra for compositional information. The described experiment is in the context of the New Millennium Deep-Space 1 mission, a comet and asteroid flyby mission. This experiment involves the autonomous recognition of certain features of scientific interest at cometary and asteroidal targets. An autonomous capability to examine UV spectra provides valuable savings and improvements in mission execution.

From a scientific standpoint, the development of automated interpretation of UV spectra allows the spacecraft to decide for itself which to pursue when there are many potential applications of real value. As one example, this capability detects when the spectrum of a given region has sufficient signal-to-noise (so the aperture can be moved on to the next locale). As another example, this capability allows near-real-time determination of whether faint features are present in the spectrum (so that the instrument executive

can decide whether it's profitable to remain on this location (or spectral bandpass), or whether it would be wiser (i.e., of higher value) to explore other bandpasses or spatial locations.

From an operations standpoint, the automated interpretation of UV spectra to determine their quality (i.e., S/N at various wavelengths) and content (i.e., the identification of the emissions being detected) enables the intelligent prioritization of datasets for transmission to the ground (or for varying degrees of compression prior to transmission).

We view the approach suggested here as potentially opening up an entirely new realm of inquisitive spaceborne scientific experiments. The approach provides obvious benefits in terms of enabling data editing and prioritization for transmission of important data to earth. Of greater potential, this approach also stresses a novel ingredient, namely sophisticated onboard science processing with the specific goal of enabling automatic decisions by the software to choose between distinct science experiment actions.

5. ONBOARD UV SPECTRAL ANALYSIS

There are two fundamental recognition tasks that must be performed here: the first one is automatic detection of statistically significant peaks in a one-dimensional function, and the second one is comparison of the location and heights of detected peaks with pre-loaded libraries of spectral signatures. The resulting automatic onboard decisions concerning target composition can be exploited in several ways. For example, they can be made immediately available to science PI's using limited telemetry resources to allow experiment decisions, since they represent extremely compact encodings of the detector spectra. They can also be used to trigger retargeting actions of the spacecraft automatically, selecting for example between two or three pre-determined modes of detector/spacecraft deployment.

6, CONCLUSIONS

One obvious improvement for the future development of autonomous satellite detection is the incorporation of spacecraft and satellite motion into the detection and tracking tasks. These problems can be addressed by a general Bayesian probabilistic formalism that allows modelling of satellite trajectories under uncertainty. A positive target identification in any given 2-dimensional image constrains the full 3-dimensional position of the target to a cigar-like volume corresponding to a Gaussian prior distribution for its position in the depth dimension. The intersection of these cigars can then be used to constrain the trajectories of interest,

A number of other scientific goals might be tackled by the science autonomy approach. One possibility is the onboard analysis of asteroidal and planetary craters by automatic means, enabling regions with particularly interesting features to be imaged at high resolution in the case of orbiter and flyby missions, and to be visited in the case of lander missions. Measurement of temporal processes also becomes feasible for a number of situations. These include for example possible sand-dune motion on the surface of Mars, analysis of its ice-cap motions, and studies of Martian storm conditions.

Our approach to science-driven autonomy is specifically designed as the first step in the development of a general robust intelligent onboard processing capability. It includes all the characteristics desirable in novel and expensive experiments such as simplicity, robustness, fast processing speed and reliability. The principles underlying the implementation of adaptive science processing tools are all generalizable across various missions. We believe that they can dramatically increase the power and scientific value of autonomy concepts in the context of space missions, enhancing their cost-effectiveness by supporting vastly reduced communications bandwidth to earth, while at the same time increasing the value of scientific information returned.

REFERENCES

- [1] M.C. Burl, U.M. Fayyad, P. Perona, P. Smyth and M.P. Burl, "Automating the hunt for volcanoes on Venus" in *Proc. Computer Vision and Pattern Recognition Conference, CVPR-94*, 1994.
- [3] D.L. Durda, "The formation of asteroidal satellites in catastrophic collisions." *Icarus*, 120, in press, 1996
- [4] U. Fayyad, N. Weir, and S.G. Djorgovski, "SKICAT: a machine learning system for automated cataloging of large scale sky surveys." *Proc. of Tenth Int. Conf. on Machine Learning*, Morgan Kaufman, 1993.
- [5] T. Gehrels, J.D. Drummond and N.A. Levenson, "The absence of satellites of asteroids." *Icarus* 70, 257-263, 1987.
- [6] J.W. Head et al., "Venus volcanic centers and their environmental settings: recent data from Magellan," *EOS* 72, p. 175, American Geophysical Union Spring meeting abstracts, 1991.
- [7] P. Stolorz and C.T. Dean, "Quakefinder - A Scalable Datamining System for Detecting Earthquakes from Space", *Proc. 2nd International Conf. on Knowledge Discovery and Datamining*, AAAI Press, pp 208-213, 1996.
- [8] T.C. Van Flandern, E.F. Tedesco and R.P. Binzel, Satellites of asteroids. In "Asteroids", ed. T. Gehrels (Tucson: Univ. of Arizona Press), pp. 443-465, 1979.
- [9] S. Weidenschilling, P. Paolicchi and V. Zappala, Do asteroids have satellites? In "Asteroids II", eds. R. P. Binzel, T. Gehrels, and M. S. Matthews (Tucson: Univ. of Arizona Press), pp. 643-658, 1989.

ACKNOWLEDGMENTS

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.