18th International Conference on Space Operations, Montreal, Canada, 26 - 30 May 2025.

SpaceOps-2025, ID # 129

An Approach For Operating High Data Volume, Targeted Remote-Sensing Missions At Mars

Nicolas Thomas[a]*, Matthew R. Read[a,b], Miguel Almeida[a]

[a] *Space Research and Planetology Division, Physikalisches Inst., University of Bern, Sidlerstrasse 5, CH-3012 Bern, SWITZERLAND,, nicolas.thomas@unibe.ch*
[b] *The Open University, Walton Hall, Milton Keynes, MK7 6AA, UNITED KINGDOM.*
* Corresponding Author

Abstract

The European Space Agency (ESA) is in the process of defining a set of new missions for remote-sensing at Mars in the 2030s and beyond. The first missions in this plan, Lightship and Spotlight, as presented at the 10th International Mars Conference on Mars (July 2024), offer the opportunity to design a holistic operations approach for all payloads onboard to maximize scientific return. The operations approach taken for the Colour and Stereo Surface Imaging System (CaSSIS) onboard the ExoMars Trace Gas Orbiter (TGO) appears to be a robust and effective means of acquiring high quality observations of the surface taking into account the properties of the orbit (including illumination conditions), times excluded for data acquisition (e.g. for reaction wheel off-loads), requests for specific high priority targets (both scientific and for lander operations), and targets selected by algorithms based on previously acquired datasets. The CaSSIS operations system performs its tasks by separating the different possible acquisition possibilities into corresponding pointing possibilities which are then analysed separately. It is expected that Spotlight 1 will carry between 2 and 4 payloads (two imagers plus a further, as yet unselected, remote sensing instruments) that can be treated in a similar way. A full end-to-end framework can be constructed following ESA's standard planning principles (using long-term, medium-term, and short-term plans). The operational approach should be considered in the early stages of mission/payload design. Using a CaSSIS-like framework, could allow optimization of the mission/payload requirements at an early stage without having a significant effect on the measurement approaches taken by the instrument teams. It would also allow a much more integrated approach to data acquisition placing joint observations at the heart of the missions. Finally, operations have tended to be labour-intensive in ESA science and robotic exploration programmes. An approach following that used for CaSSIS can make a significant reduction in labour costs while improving the quality of the data acquired.

.

Keywords: Mars, remote-sensing, low orbit, target selection

Acronyms/Abbreviations

ACS	Atmospheric Chemistry Suite
CaSSIS	Colour and Stereo Surface Imaging System
CaST	CaSSIS Sugestion Targeting
CaTLIST	CaTL Imaging Suggestion Tool
CTF	CaSSIS Targetting File
DVAR	Data Volume Acquisition Rate
ESA	European Space Agency
ESAC	European Space Astronomy Center
GDPR	General Data Protection Regulation
HiRISE	High Resolution Imaging Science Experiment
ITL	Instrument Timeline
MRO	Mars Reconnaissance Orbiter
NOMAD	Nadir and Occultation for Mars Discovery
PFS	Planetary Fourier Transform Spectrometer
ROI	Region of Interest
TGO	Trace Gas Orbiter

1. Introduction

The European Space Agency (ESA) is in the process of defining a set of new missions to fly to Mars in the 2030s and beyond. One of the primary goals of these missions is to provide a communications infrastructure that can be used to support both very high resolution remote-sensing for robotic and human landing site selection and a high performance relay for subsequent landed missions. The Mars Communication and Navigation Infrastructure (MARCONI) has been widely publicised[1,2]. The concept follows on from the highly successful, ExoMars Trace Gas Orbiter (TGO), which carried a 2.2 m X band high-gain antenna with a two-axis pointing mechanism and 65 W radio frequency travelling-wave tube amplifier to communicate with Earth[3]. This communications system supported (and at the time of writing continues to support) the landing and operations of NASA's Perseverance rover in Jezero Crater[4].

An important aspect of this concept is that scientific payload can be placed on the communications orbiters to provide state-of-the-art science. For example, TGO carried 4 main instruments. The ACS and NOMAD were primarily infrared spectrometers with various apertures designed to provide high spectral resolution measurements of the Martian atmosphere. This was stimulated by a (weak) detection of methane in the Martian atmosphere by the Planetary Fourier Transform Spectrometer (PFS) on ESA's Mars Express spacecraft and by subsequent follow-up ground-based observations. A neutron spectrometer, FREND, was included in the payload to provide low resolution maps of the amounts of hydrogen-bearing molecules (such as H_2O and OH bearing minerals) in the sub-surface down to a depth of approximately 1 metre. Finally, the CaSSIS instrument was selected to provide moderately high resolution imaging (4.5 m/px) of the surface in four visible and NIR colour bands and in stereo.

The present approach for the first element of the original MARCONI concept is called Lightship[5], which will not only carry payload itself but will also carry a sub-spacecraft, called Spotlight, that also supports payload including two high resolution imaging systems. Although this mission is still in a study phase, it is already apparent that the data volumes available to the payload will be substantially higher than for TGO. We recall that, at the time of writing, CaSSIS has provided over 46.8% of the data transmitted from the spacecraft in the past 7 years which translates into over 29 Tbits (>50,000 images and >5% of the surface in full colour). Managing the data acquisition at such high data volumes is a challenge given the flexibility that is demanded from interplanetary spacecraft by the community. Here we address this problem based upon the experience with TGO and CaSSIS. In the next section, we look at the fundamental problems associated with data acquisition at planetary targets. Although we focus on Mars, similar issues will soon become evident for missions to the Moon, large asteroids, and the Galilean satellites. In the following sections, we sketch out possible approaches to maximize the scientific return from orbiter missions. Then we summarize.

2. Basic issues with data acquisition

Spacecraft entering low orbit about planetary targets can be designed and operated in many different ways depending upon the scientific objectives. Mission goals could include, for example, the complete mapping of a planetary target at visible wavelengths. One can easily envisage this being achieved by continuous imaging over a wide swath from a sun-synchronous orbit. The operational task in this case is quite straightforward. However, difficulties begin to arise when the spatial resolution requirement becomes tighter and when multi-spectral coverage is required. Taking Mars as an example, a 5.7 terapixel mosaic of the surface of Mars rendered at 5.0 m/px has been produced[6,7] using data from the Context Camera (CTX) on Mars Reconnaissance Orbiter (MRO) The mosaic covers 99.5% of Mars from 88°S to 88°N. Imaging of Mars is now pushing to much higher resolutions (e.g. observations with HiRISE onboard MRO at 0.25 m/px)[8] and using multiple spectral bands following the CaSSIS high resolution imaging observations of minerals near current and potential landing sites[9,10]. A global mapping requirement is already challenging at 4 m/px without long duration missions (MRO is approaching 20 years in Mars orbit) At resolutions suitable for landing site certification (<<1 m/px) with mineralogical context (e.g in 6-8 colours), a global mapping requirement becomes subject to significant compromise. In addition, TGO has shown that scientifically non-sun-synchronous orbits can provide new information because of the possibility of performing observations at different local times. This can be useful for determining surface phase functions, showing evidence of atmospheric variability, and eliminating effects of shadows in reconstructed maps and mosaics. This is complimented by studies of seasonal variability (e.g. seasonal ice cap evolution[11], geyser activity[12], etc.) These scientific aims result in science requests for repeat imaging. By the end of 2024, the total surface imaging acquired in colour with CaSSIS was 6.769% but the overlap was such that only 4.994% of the surface had actually been covered in colour. In other words, more than 1/3rd of the data acquisition was repeat coverage. It is also apparent that as the initial surveys of solar system objects have been completed, we have become aware that some regions on planetary targets are more interesting than others. This (rather trivial) statement implies that we should focus our high resolution observing strategies on selected regions of interest (ROIs) rather than trying to attempt global mapping at marginally higher resolutions.

Further constraints on operation arise from spacecraft activities such as reaction wheel off-loading as well as communication passes with landed assets. In addition, other instruments on the spacecraft may have other scientific goals requiring spacecraft operations that are no longer compatible with those of high resolution surface imaging (e.g. limb scans and/or occultation experiments).

It is apparent from the work performed on HiRISE and CaSSIS that future, high data volume, targeted imaging will require a fully automated workflow. While this is immediately of importance for the upcoming ESA Mars missions, missions to other planetary objects, where first surveys have been completed and now more detailed observations are required (e.g. Ceres, Ganymede, Europa) facilitated by higher data rates, would also benefit from such an approach. In the following sections outline an approach and the associated requirements.

3. ROI selection

The challenge of ROI selection has already been faced by the HiRISE and CaSSIS science teams. The HiRISE team developed a target database which eventually became the public tool, HiWISH[13]. As a result of the cooperation between the universities of Arizona and Bern, a similar science team ROI selection tool (CaST) was implemented for CaSSIS. While these tools have proven to be very effective and strongly supported by the respective science teams, long mission durations and high data rates have resulted in a paucity of targets in the database and particularly in regions of slightly lower scientific interest which leads to an absence of selectable targets on ~20% of orbits.

Good selection of highly relevant targets is time-consuming. Furthermore, there is significant flexibility available in the tools requiring the requestor to choose whether this might a recurring target (e.g. for studying variability at high latitude), a target best observed in a specific season, or a target requiring multiple colours and/or stereo. Giving the target a title and a description has also been found to be necessary to provide some sort of explanation of the scientific goal. It also forces the suggester to think rather than aimlessly clicking on a region. Allowing public access to support target selection is valuable as a public relations action but is not a solution because all public targets need to be validated by a scientist to ensure that it is both scientifically justifiable and that it has been correctly entered into the system. The complexity of compliance with General Data Protection Rights (GDPR) is an additional burden.

The above issues have resulted in efforts to import databases and to use lower resolution global maps of thermal inertia, surface albedo/reflectance and surface composition to select ROIs with programmable algorithms[14]. The key to this approach is to develop a clear set of priorities for each source of ROIs. In the case of CaSSIS, this has been put in place during the mission and consequently has some inconsistencies. Planning this prioritization in advance can be of benefit to the target selection process. A table illustrating the priority concept for a Mars camera is shown in Table 1.

Table 1 Prioritization of ROIs in the target database according to their source.

Priority	Description	Source	Examples	~percentage of total targets in database
10	Reserved for Principal Investigator; special targets to be acquired immediately	Human individual request	ROI needed to be acquired at a specific date/time to complete a high profile paper.	0.01%
9	Future and currently active landing sites	Human targeting by scientists with knowledge of landing site needs	ROI helping decide traverse of, for example, Perseverance in Jezero Crater	0.1%
8	Targets of high scientific interest that may produce a significant increase in our knowledge of the object	Human targeting of individual ROIs	Studies of Nili Fossae or another area of great significance	4%
7	Targets of moderate scientific interest that may produce an increase in our knowledge of the object	Human targeting of individual ROIs	Studies of areas with chloride deposits	8%

6	Targets of lower scientific interest that may support our knowledge of the object	Human targeting of individual ROIs	Mapping of banded terrain in Hellas Basin	15%
5	Targets related to surveys of particular phenomena	Database of sites exhibiting a particular phenomenon	Concentric crater fill	33%
4	Targets related to automated target identification algorithms	Calculations using algorithms working on databases or maps to identify ROIs	Extraction of sites showing bright nighttime IR and low blue albedo to find rocky terrains	~40%

4. “Flying the plan”?

The term “flying the plan” was coined by the CaSSIS team to differentiate the orbit control approach used by TGO when compared to that used by the Mars Reconnaissance Orbiter (MRO). In the case of MRO, the orbit is occasionally corrected to maintain the Sun synchronous property selected for MRO but the actual ground-track is not known until 3-4 weeks before execution. Consequently, the target acquisition planning occurs quite late. In the case of TGO, the orbit is effectively frozen several months in advance as part of the long-term plan (LTP) and orbit correction manoeuvres (OCMs) are made every Saturday (if necessary) at a weekly planning boundary to put the spacecraft back on this pre-determined trajectory. We refer to this as “flying the plan”.

The two approaches have advantages and disadvantages (Table 2). The key difference is in the timing of activities given that, in both cases, there is a need to minimize the work of the spacecraft and instrument operations teams to reduce cost and there is also a need to fix the spacecraft rolls (to access off-nadir targets) in good time so that the commanding of these critical manoeuvres can be planned and validated.

Table 2 Advantages and disadvantages of "flying the plan"

Flying the plan	Orbital drift
Observations can be planned early. There is no pressure to define the trajectory late in the process.	The trajectory must be solved close to the time of the event.
Orbit exclusions can be identified early in the process.	Orbits that need to be excluded from the observational plan are found relatively late in the process.
Only a correction manoeuvre is required and this has no influence on the targeting team who have already prepared their plan.	The trajectory must be provided to the targeting teams fairly late on a regular cycle. Only then can the targeting teams develop their plan.
This approach may not be robust against varying atmospheric drag at low altitudes.	Robustness against varying atmospheric drag at low altitudes is good.
Late changes to the plan (particularly the spacecraft roll) are discouraged to avoid duplicating work.	Flexibility in the spacecraft roll is maintained until relatively late in the process allowing late target selection.

The comparison of TGO and MRO operations has illustrated these differences well. The planning process for TGO occurs over a much longer time period than MRO but lacks a degree of flexibility. For MRO, planning is late and must be performed to strict deadlines but is flexible and can be adjusted to updated priorities factoring in orbit exclusions (for communications with landers, reaction wheel off-loading, special pointings, etc.) relatively late.

The effort put into developing and checking a final plan is large and it has become evident that both strategies can benefit from automating the target selection process.

5. Plan construction and scientific oversight

Work within the CaSSIS team in recent years has led us to establish a series of tools that combine to provide automated systems in which data can be commanded, executed, downlinked, calibrated, and archived with almost no human intervention. The CaSSIS case is described in full by Almeida et al. (this issue). As part of this scheme, one tool accesses databases with prioritized targets, the instrument specific constraints, and spacecraft and other instrument constraints. The tool then places targets on the instrument timeline (ITL) in the best way possible to make a complete target file. This tool (called CaTLIST) is presented in Read et al. (this issue). The TGO/CaSSIS approach takes the spacecraft/instrument constraints and chooses and fits the plan around them so that it is the science that gets the priority according to science team defined priorities. If one considers the situation with MRO, a key difference is seen. MRO uses a tool (called TOS) which accepts target plans from the main instruments onboard and discards observations until it has a final plan for the spacecraft observations that meets the constraints.

A further benefit of the CaSSIS procedure is that scientific oversight becomes more straightforward. The CaSSIS and HiRISE approach has been to appoint a scientist for a fixed period (with CaSSIS this is one week, for HiRISE it is two) who effectively takes responsibility for the imaging plan in that period. In the case of HiRISE, the plan requires more manual work by the scientist to construct the plan in the first place. The adaptation of this approach made by the CaSSIS team is that the plan is constructed by computer and then it is checked and optimised by the scientist. This significantly reduces the workload and is particularly practical as the mission becomes more mature with reductions in funding/support (and to some extent, enthusiasm). Obviously, the higher the quality of the computer-generated plan, the less work there is for the scientist.

6. Multiple instruments

It is fairly straightforward to extend this approach to multiple remote-sensing instruments on a single spacecraft if the individual instrument constraints and targeting can be defined in a similar way. The main question then revolves around the prioritization of the instruments with respect to each other. With the current TGO system, this is done by engineers at the European Space Astronomy Centre (ESAC) who define data volume constraints and exclusion zones to each instrument when off-pointings (e.g. occultation observations) are planned. The approach for MRO was to use a scheduler software that picked an observation for each instrument from a prioritized stack on a round-robin basis. If the top of the stack was an observation that no longer met the necessary constraints or collided with a previously scheduled observation from another instrument, then it was thrown out and the next observation in the stack was processed. It has been recognized that this is not the optimum approach (a large number of observations that have been worked on and consumed human time are thrown out – it is not efficient in this sense) but a higher level of automation has not yet been attempted/demonstrated.

On the other hand, the current Lightship/Spotlight concept is ideal for tackling this problem because of the relatively small number of instruments per spacecraft (Spotlight currently foresees two images with up to two additional payloads only if resources allow) – the problem should not be that complex to solve and can be more easily tested because a CaTLIST type tool makes replanning faster and deterministic. The complexity arises because each instrument has its own data acquisition approach – something that should not be compromised or interfered with simply to make operational planning easier. Indeed, most instruments have more than one acquisition "mode" and each of these (with their associated requirements) must be modelled.

Instruments with high data volume acquisition rates (DVAR) will usually operate for shorter periods on the orbit while a lower DVAR implies longer along-track acquisition. Onboard data management is also a necessity while in addition, the pointing capability of the spacecraft and the speed with which it can move between pointings must be accurately modelled. The latter is one of the challenges of operating TGO because the spacecraft is relatively slow and a lot of time is consumed slewing from one pointing to another.

The resulting constraints and modelling requirements result in a table addressing all the aspects needed for the plan. This in itself is nothing new[15] – codes have been developed in the past by both ESA and NASA for science mission planning (e.g. MAPPS[16]). However, these codes have required substantial manual support to produce plans compatible with the constraints. The planning flow is illustrated in Figure 1 which is similar in content to that given in Vallat et al.[15].

Table 3 Constraints to modelling of the spacecraft and instrument resources that must be modelled.

Constraints	Comments
Data management	Must be modelled both internally in each instrument and externally on the spacecraft and matched to agreed data volume partitioning at system level.
Power consumption	In the general case, power consumptions must be modelled and matched to agreed spacecraft constraints.
Pointing profiles	Slew magnitude and slewing rates must be accurately modelled between all instrument pointings.
Spacecraft maintenance activities	Exclusions and constraints resulting from spacecraft activities must be integrated into the planning tool.

One important aspect should also not be underestimated. In the planning of CaSSIS observations, the mutual goodwill and willingness to compromise shown by the instrument teams of ACS, NOMAD, and CaSSIS has been of substantial importance. Without this willingness to agree to small modifications to the detailed planning, the scientific output would have been considerably reduced. This has also resulted in the teams having a better understanding of what will be necessary for future missions.

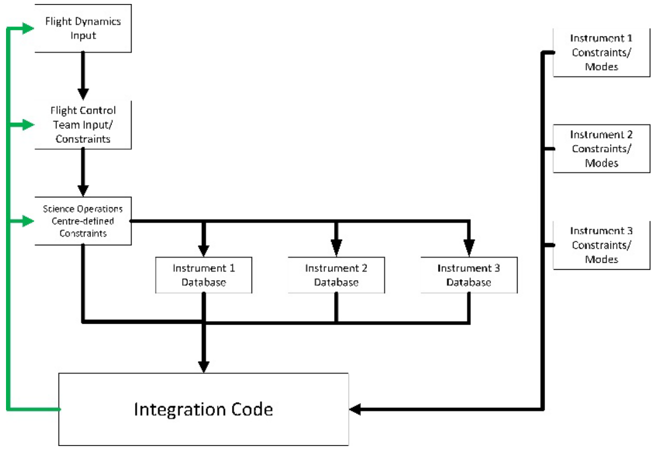

Figure 1 Schematic diagram of information flow to and from a code that integrates targeting databases with spacecraft constraints for a future ESA mission such as Spotlight.

7. Future ESA Mars missions

The European Space Agency's MARCONI concept will provide a large increase in data volume but the drive towards high spatial resolution implies that imaging systems and imaging spectrometers will still be pointed instruments. They will not perform continuous mapping with a constant data acquisition and data flow. Furthermore, targeting of specific areas such as landing sites will remain a very high priority as will seasonal and diurnal coverage. Consequently, prioritization of selected targets will still be necessary for at least the next decade.

Another point of note is that the MARCONI concept envisages multiple launches to produce a communications network at Mars with each relay spacecraft taking with it a sub-spacecraft (such as Spotlight). While this increases the data relay possibilities further, it also increases the number of instruments operating at Mars and does so with sub-spacecraft on different orbits which are optimised for the scientific return. With more "traditional" approaches to operations, we would expect a substantial increase in operations personnel and associated costs. Methods to reduce this burden by automating labour-intensive activities will be necessary otherwise the operations costs alone will choke the financing of future missions. This is particularly true given the lifetimes of the spacecraft currently being touted. If Spotlight really does stay in orbit collecting data for 12 Earth years and following missions have similar ambitions, then potentially there will be 12 spacecraft in operation at Mars if ESA can launch once every opportunity. While this is probably financially unrealistic, half that frequency would still be 3 times the current situation where ESA "only" has Mars Express and TGO in orbit. Consequently, a concerted effort to reduce labour-intensive operation through automation is almost a necessity.

8. Other objects/missions

The exploration of our planetary system is entering a new phase in which we have established the basic properties of objects in our Solar System and we have obtained reasonably complete coverage of objects of interest at low spatial and spectral resolution. Our communication systems are expected to improve substantially in the next years with higher bandwidth and the use of laser communication terminals. New missions will therefore have higher spatial resolution driving higher data volumes from most future targets. This implies that optimising approaches to targeting and data acquisition will become increasingly important to maximize the scientific return and also potentially to reduce operations costs. For objects such as main belt large asteroids, Mercury, the Moon, and the Jovian and Saturnian moon systems, these approaches could also be applied assuming low orbit can be achieved.

9. Conclusions

Automation of targeted planning of selected sites on the Martian surface is rapidly becoming a necessity as data volumes and spatial resolutions increase. The workflows for such approaches can be established but require scientists to have the discipline to provide target databases with priorities. These databases can comprise targets of different priority and, with the support of target identification software, can contain sufficient possible targets to ensure that all instruments onboard a spacecraft can be serviced correctly. The development of such automation has been illustrated with the CaTLIST tool for the CaSSIS instrument on TGO. Upcoming missions such as Lightship/Spotlight, which will have a relatively low number of remote sensing instruments, would be ideal testing grounds for expanding this approach to multi-instrument systems. The computing power is now such that such approaches could supercede the current more manual workflows that have been used heretofore. As other planetary targets become better understood through low resolution global mapping (e.g. Ceres), such a tool would be beneficial for follow-up missions to these targets which might use high resolution, low total spatial coverage to address high interest science.

Acknowledgements

This work was produced during the development and operation of the CaSSIS project. CaSSIS is a project of the University of Bern funded through an ESA PRODEX grant awarded by the Swiss Space Office. The contributions of the Lunar and Planetary Laboratory of the University of Arizona and the U.S. Geological Survey to CaSSIS database and pipeline development are gratefully acknowledged.

References

[1] Williams, A., ESA plots spacecraft passenger mission for LightShip tug, 12 March 2025, https://www.electronicsweekly.com/news/esa-plots-spacecraft-passenger-mission-for-lightship-tug-2025-03/ (accessed 25.03.25)

[2] ESA, Towards low-cost missions to Mars, 21 January 2025 https://www.esa.int/Enabling_Support/Preparing_for_the_Future/Discovery_and_Preparation/Towards_low-cost_missions_to_Mars (accessed 25.03.25).

[3] ESA, ExoMars Trace Gas Orbiter (TGO), 12 March 2020, https://exploration.esa.int/web/mars/-/46475-trace-gas-orbiter, (accessed 25.03.25).

[4] Spaceref, ESA Mars orbiters support NASA Perseverance landing, 11 February 2021, https://spacenews.com/esa-mars-orbiters-support-nasa-perseverance-landing/ , (accessed 25.03.25).

[5] ESA, Seven things you didn't know about Europe's tug to Mars, 17 September 2024, https://blogs.esa.int/to-mars-and-back/2024/09/17/seven-things-you-didnt-know-about-europes-tug-to-mars/ (accessed 26.03.25).

[6] Murray Lab., https://murray-lab.caltech.edu/CTX/ (accessed 26.03.25)

[7] Malin, M.C., et al. (2007) Context Camera investigation on board the Mars Reconnaissance Orbiter, J. Geophys. Res., 112, E05S04. doi:10.1029/2006JE002808.

[8] McEwen, A.S. et al. (2007) Mars Reconnaissance Orbiter's High Resolution Imaging Science Experiment (HiRISE), J. Geophys. Res., 11

[9] Thomas, N., and 60 colleagues, (2017), The Colour and Stereo Surface Imaging System (CaSSIS) for the ExoMars Trace Gas Orbiter, Space Science Reviews, 212, 1897, doi:10.1007/s11214-017-0421-1.

[10] Fawdon, P., Grindrod, P., Orgel, C., Sefton-Nash, E., Adeli, S., Balme, M., … Volat, M. (2021). The geography of Oxia Planum. Journal of Maps, 17(2), 621–637. https://doi.org/10.1080/17445647.2021.1982035

[11] Langevin, Y., J.-P. Bibring, F. Montmessin, F. Forget, M. Vincendon, S. Dout{\'e}, F. Poulet, and B. Gondet, (2007), Observations of the south seasonal cap of Mars during recession in 2004-2006 by the OMEGA visible/near-infrared imaging spectrometer on board Mars Express, Journal of Geophysical Research (Planets), 112, E08S12, doi:10.1029/2006JE002841.

[12] Hansen, C.J., N. Thomas, G. Portyankina, A. McEwen, T. Becker, S. Byrne, K. Herkenhoff, H. Kieffer, and M. Mellon, (2010), HiRISE observations of gas sublimation-driven activity in Mars' southern polar regions: I. Erosion of the surface, Icarus, 205, 283, doi:10.1016/j.icarus.2009.07.021.

[13] McEwen, A.~S., E. Eliason, V.~C. Gulick, Y. Spinoza, R.~A. Beyer, and HiRISE Team, (2010), HiRISE: The People's Camera, AGU Fall Meeting Abstracts, 2010, ED23A-0712.

[14] Read et al. (2025) this issue.

[15] Vallat, C., N. Altobelli, B. Geiger, B. Grieger, M. Kueppers, C. Muñoz Crego, R. Moissl, M.G.G.T. Taylor, C. Alexander, B. Buratti, M. Choukroun, and the RSGS Liaison scientists Group, (2017), The science planning process on the Rosetta mission, Acta Astronautica, 133, 244, doi:10.1016/j.actaastro.2017.01.018.
[16] Pérez-Ayúcar, M., M. Ashman, M. Almeida, M.~C. Sitjá, J.~J. García Beteta, R. Hoofs, M. Kueppers, J.M. Yaseli, D. Merritt, F. Nespoli, and E.S. Suarez, (2018), The Rosetta science operations and planning implementation, Acta Astronautica, 152, 163, doi:10.1016/j.actaastro.2018.07.049.