Icarus 282 (2017) 136–151

Contents lists available at ScienceDirect

Icarus

journal homepage: www.elsevier.com/locate/icarus

Regularization of Mars Reconnaissance Orbiter CRISM along-track oversampled hyperspectral imaging observations of Mars

C.D. Kreisch[a,b,*], J.A. O'Sullivan[c], R.E. Arvidson[a], D.V. Politte[c], L. He[c], N.T. Stein[a,d], J. Finkel[a], E.A. Guinness[a], M.J. Wolff[e], M.G.A. Lapôtre[d]

[a] Department of Earth and Planetary Sciences, Washington University in Saint Louis, Saint Louis, MO 63130, USA
[b] Max Planck Institute for Astrophysics, 85748 Garching bei München, Germany
[c] Department of Electrical and Systems Engineering, Washington University in Saint Louis, Saint Louis, MO 63130, USA
[d] Division of Geological and Planetary Sciences, California Institute of Technology, Pasadena, CA 91125, USA
[e] Space Science Institute, Boulder, CO 80301, USA

ARTICLE INFO

Article history:
Available online 23 September 2016

Keywords:
Mars
Mineralogy
Image processing
Data reduction techniques
Spectroscopy

ABSTRACT

Mars Reconnaissance Orbiter Compact Reconnaissance Imaging Spectrometer for Mars (CRISM) hyperspectral image data have been acquired in an along-track oversampled (ATO) mode with the intent of processing the data to better than the nominal ~18 m/pixel ground resolution. We have implemented an iterative maximum log-likelihood method (MLM) that utilizes the instrument spectral and spatial transfer functions and includes a penalty function to regularize the data. Products are produced both in sensor space and as projected hyperspectral image cubes at 12 m/pixel. Preprocessing steps include retrieval of surface single scattering albedos (SSA) using the Hapke Function and DISORT-based radiative modeling of atmospheric gases and aerosols. Resultant SSA cubes are despiked to remove extrema and tested to ensure that the remaining data are Poisson-distributed, an underlying assumption for the MLM algorithm implementation. Two examples of processed ATO data sets are presented. ATO0002EC79 covers the route taken by the Curiosity rover during its initial ascent of Mount Sharp in Gale Crater. SSA data are used to model mineral abundances and grain sizes predicted to be present in the Namib barchan sand dune sampled and analyzed by Curiosity. CRISM based results compare favorably to in situ results derived from Curiosity's measurement campaign. ATO0002DDF9 covers Marathon Valley on the Cape Tribulation rim segment of Endeavour Crater. SSA spectra indicate the presence of a minor component of Fe^{3+} and Mg^{2+} smectites on the valley floor and walls. Localization to 12 m/pixel provided the detailed spatial information needed for the Opportunity rover to traverse to and characterize those outcrops that have the deepest absorptions. The combination of orbital and rover-based data show that the smectite-bearing outcrops in Marathon Valley are impact breccias that are basaltic in composition and that have been isochemically altered in a low water to rock environment.

1. Introduction

Visible to infrared spectroscopic observations of Mars from telescopic, orbital, and landed observatories have been critically important for establishing the predominantly basaltic composition of the surface (e.g., McCord et al., 1982; Bandfield et al., 2000; Bell et al. 2004a, b; Bibring et al., 2005; Mustard et al., 2005; Poulet et al., 2009) and discovering and mapping evidence for mineral phases formed in the presence of water (e.g., Christensen and Ruff, 2004; Gendrin et al., 2005; Bibring et al. 2006; Bishop et al., 2008; Poulet et al., 2008; Murchie et al., 2009; Ehlmann and Edwards, 2014). The Compact Reconnaissance Imaging Spectrometer for Mars (CRISM) on the Mars Reconnaissance Orbiter (MRO) began operations in 2006 (Murchie et al., 2007). CRISM is a push-broom hyperspectral imaging spectrometer operating from 0.392 to 3.91 μm with a 6.5 nm spectral band spacing and a smallest pixel size projected onto the surface of 18 to 20 m. Images are acquired using an optical sensor unit that can be gimbaled in the along-track direction, thereby tracking a target to produce contiguous pixels from line to line as MRO moves forward in its

* Corresponding author.
E-mail address: ckreisch@wustl.edu (C.D. Kreisch).

http://dx.doi.org/10.1016/j.icarus.2016.09.033

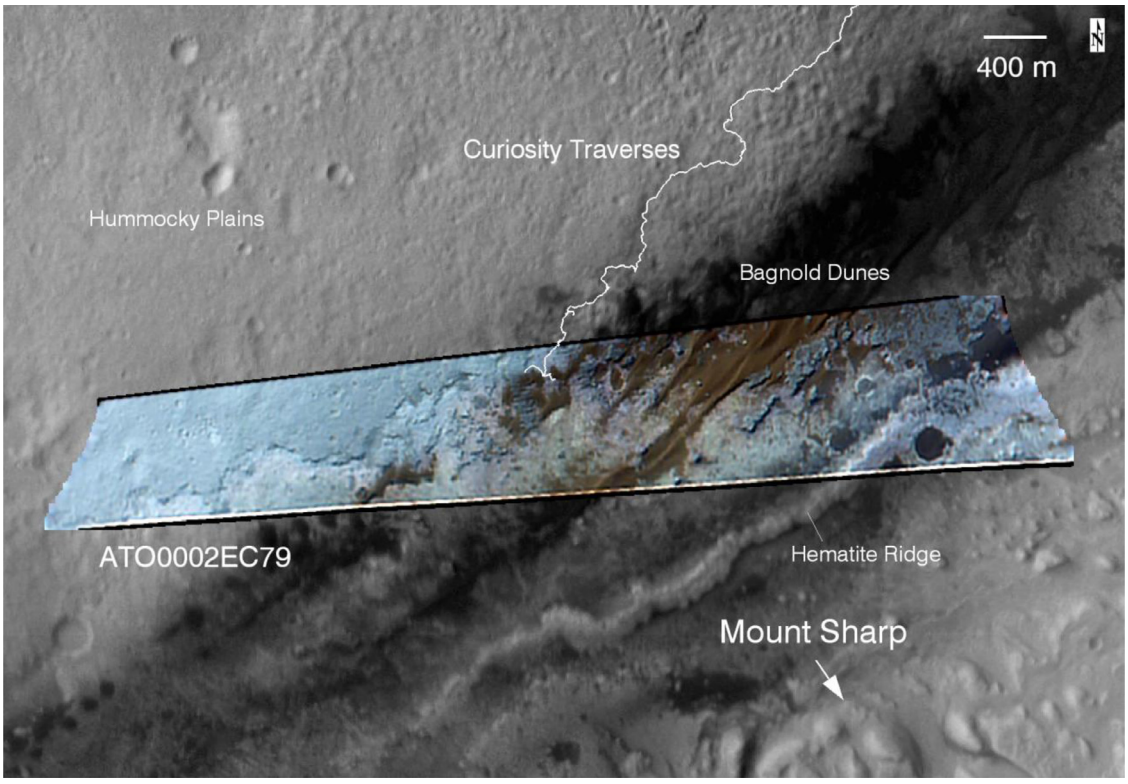

Fig. 1. False color map-projected single scattering albedo (SSA) image of CRISM L data for ATO0002EC79 is shown covering a portion of Gale Crater. The image is overlain on a HiRISE-based mosaic. The CRISM image has been regularized to 12 m/pixel. Traverses for the Curiosity rover are also plotted through sol 1281 and include paths into the barchan dunes in the Bagnold Dune field. Curiosity is on the lower slopes of Mount Sharp which rises ~5 km to the southeast of the area shown in the figure. The Namib Dune is where the rover was commanded to conduct a detailed remote sensing and measurement campaign to characterize the bedforms, composition, and mineralogy of the dune sands. The hematite-dominated ridge (Fraeman et al., 2013) is a prime target for Curiosity, to be followed by continued ascent of Mount Sharp. RGB assigned to CRISM bands located at 2.5288, 1.5059, and 1.0793 μm wavelengths. L data are shown using these wavelengths for subsequent figures. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

orbital path. Since 2010 a new gimbaled motion has been used that employs significant pixel overlap in the along-track direction. This paper presents an algorithm for regularizing these along-track oversampled (ATO) observations using an iterative maximum log-likelihood method (MLM) to retrieve the best estimate of scene information in the presence of noise, with an image pixel size projected onto the surface of 12 m/pixel. The paper also discusses preprocessing the spectral radiance image cubes to surface single scattering albedos and presents as examples reduction and analysis of ATO data over the Curiosity and Opportunity rover locations for direct comparisons to rover-derived data (Figs. 1–4).

2. Processing methodology

2.1. Retrieval of single scattering albedos

The primary derived CRISM hyperspectral image cube archived at NASA's Planetary Data System (PDS) has units of spectral radiance ($W\,m^{-2}\,sr^{-1}\,\mu m^{-1}$) that are based on preflightcalibrations that are updated using an on-board integrating sphere and lamps (Murchie et al., 2007; see also the CRISM Data Products Software Interface Specification document (DPSIS), available online at http://pds-geosciences.wustl.edu/missions/mro/crism.htm). A second PDS product has units of I/F (spectral radiance divided by the solar spectral radiance at the time of observation) and has been highly processed to suppress detector noise (see online the DPSIS). For our purposes we use spectral radiance cubes because the algorithm described in this paper is based on retrieving the optimum estimate of surface single scattering albedos in the presence of noise that has a Poisson statistical distribution. Because the I/F data in the PDS have been highly processed we cannot use these files because the inherent noise patterns have been distorted by the processing. The spectral radiance data are first converted to I/F units by dividing by the solar spectral radiance at Mars at the time of the observation, resampled through the CRISM spectral band passes. Discrete Ordinates Radiative Transfer (DISORT)-based processing (Stamnes et al., 1988) and the Hapke function for surface scattering (Hapke, 2012) are then used to retrieve surface single scattering albedos (SSAs), i.e., the ratio of scattering efficiency to scattering plus absorption efficiencies of a single particle. The procedure was described in Arvidson et al. (2014) but we will provide additional relevant details in this section of the paper. The overall processing flow is outlined in Fig. 5.

The DISORT code models radiative transfer streams associated with dust and ice aerosols, CO_2, H_2O, and CO gases, and a surface boundary layer simulated in our case using the Hapke function (Fig. 5). Parameters used in the Hapke function are presented in Table 1 and were derived from consideration of scattering from typical surfaces based on Pancam data from Spirit and Opportunity (Johnson et al., 2006a,b). The free parameter in the function is the SSA value for each CRISM band and spatial location. Use of DISORT is iterative, beginning with dust and ice aerosols based on inferences from Opportunity and Curiosity observations for CRISM scenes acquired over the rover sites, together with data from other MRO instruments (Wolff et al.,2009). The CO_2 pressure for the scene center is based on a DISORT-computed retrieval using the depths of the 2 μm CO_2 absorption band triplet (Smith et al., 2009). Input CO and H_2O abundances are based on retrieval from CRISM data (Smith et al., 2009). A range of CO_2 pressures is then calculated using DISORT, centered about the estimated pressure for the

Table 1
Parameters used in Hapke Photometric Function. Detailed explanation of parameters found in Hapke (2012).

Single Scattering Albedo	Two Term Henyey Greenstein Single Scattering Albedo		Opposition Effect		Macroscopic Roughness Parameter
	Asymmetry Factor	Forward Fraction	h width	B_0 height	
Derived from data	0.26	0.30	1.0	0.06	15°

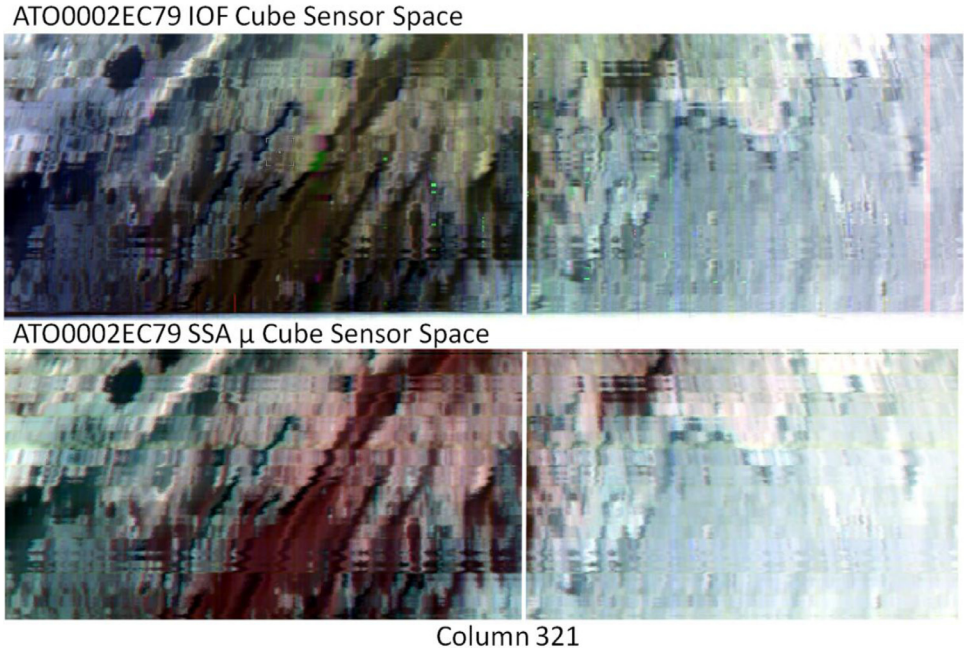

Fig. 2. Sensor space false color composites of I/F L data for ATO0002EC79 generated from spectral radiance data without any further processing (top) and as SSA data after removing extrema and regularization using the maximum log-likelihood algorithm (bottom). The MLM SSA data show less noise elements, especially column-dependent noise. Vertical white lines show the column for which row to row pixel spacing is plotted in Fig. 8. The presence of column-dependent artifacts in the unprocessed data made it difficult to generate a composite with color patterns identical to the processed data. North is roughly in the downward direction because CRISM acquires data during its ascending orbits. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

scene center, and varied as a function of elevation using Mars Orbiter Laser Altimeter (MOLA) data and a nominal pressure scale height. The output is a table of I/F values as a function of SSA values, CO_2 pressures, and incidence, emergence, and phase angle ranges (from Derived Data Records (DDRs), see also the online DPSIS) for the scene of interest. The output is computed at 1 nm spacing and convolved with the CRISM band passes using center wavelengths and a Gaussian assumption for the band pass shapes. Center wavelengths shift with the temperature of the window covering the instrument. These offsets are included in the processing and are based on shifts of the CO_2 bands in the data relative to predictions from DISORT models. Initial SSA cubes are examined for residual 2 μm CO_2 bands. The central pressure is adjusted and the model is run iteratively until residual CO_2 bands are minimized across the scene.

2.2. Statistical tests and extrema removal

CRISM data products have multiple sources of noise, including Poisson detector counting statistics for the S and L detectors, and thermal noise for the L detectors. In particular, for recent scenes thermal noise is relatively large because the detectors are operating ~15 K hotter than at the beginning of the mission due to cryogenic cooler degradation (Murchie et al., 2007). There is also observation and column dependent noise associated with one (or more) detector at a given wavelength (or set of wavelengths) due to calibration errors that leads to vertical striping in the sensor space data (Figs. 2 and 4). Additional stochastic noise is evident as spikes in the spectral domain (Fig. 6). These spikes are not characteristic of absorption features and can be singular or span a few spectral bands due to a temporary coupling of detectors. The spikes are often associated with crossing a sharp surface brightness boundary, where a given detector or set of detectors is biased from the previous acquisition.

The SSA data were found to follow a Poisson distribution if the extrema (i.e., spikes) were first removed. A detailed analysis is presented in Appendix A. Spikes were removed using a rendition of the bit error filter designed by Eliason and McEwen (1990) in which the local standard deviation in the spectral domain is calculated for a given filter width and band center. If the center value exceeds a predetermined number of standard deviations above the local standard deviation, together with an absolute tolerance value, it is replaced by the median value within the local filter domain. Iterative runs with tighter tolerances successfully removed the extreme spectral values, as shown in Fig. 6. This technique is similar to what was employed by Carter et al. (2013) for CRISM data.

3. Maximum log-likelihood method

3.1. Overview

In the maximum log-likelihood method (MLM) with Poisson distributed data, the mean $\mu(\lambda, x, y)$ represents the blurred version of the actual values $c(\lambda', x', y')$ for the CRISM SSA data $d(\lambda, x, y)$ (e.g., Snyder et al., 1992):

$$d(\lambda, x, y) \sim Poisson[\mu(\lambda, x, y)]. \quad (1)$$

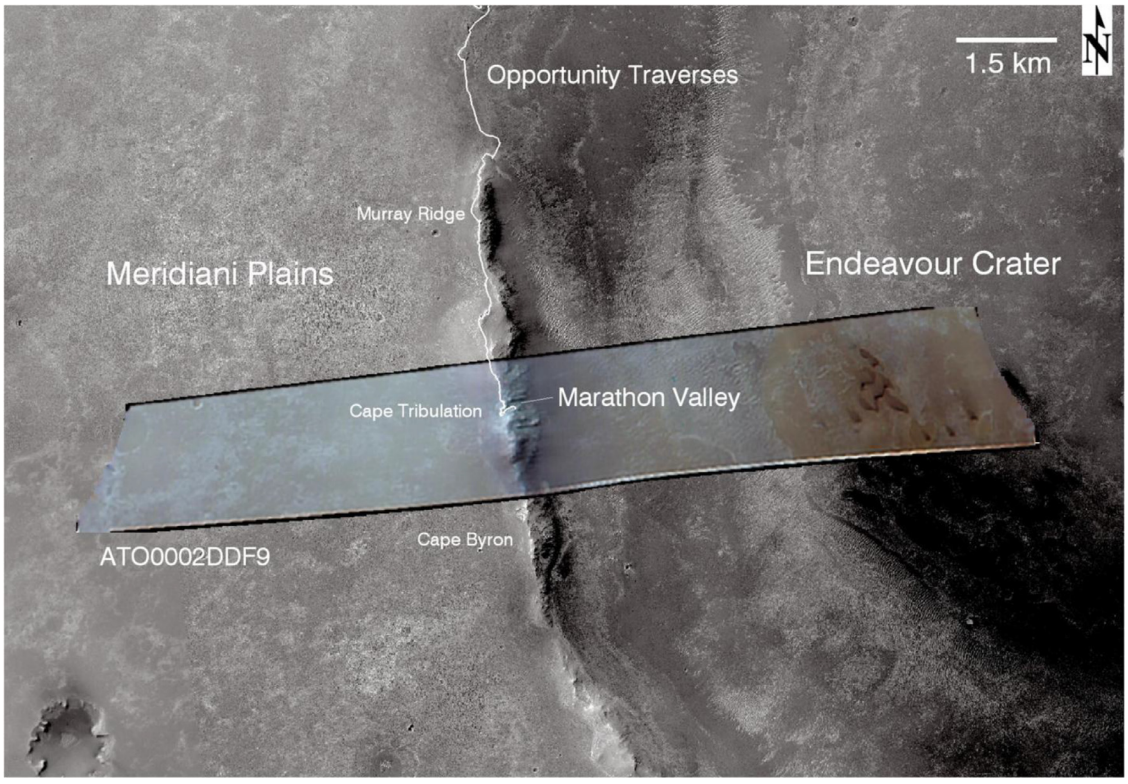

Fig. 3. False color map-projected single scattering albedo (SSA) image of CRISM L data for ATO0002DDF9 covering a portion of Endeavour Crater is shown overlain on a HiRISE-based mosaic. The CRISM image has been regularized to 12 m/pixel. Traverses for the Opportunity rover are also plotted through sol 4320 and include paths within Marathon Valley on the Cape Tribulation rim segment of the 22 km diameter Endeavour Crater. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

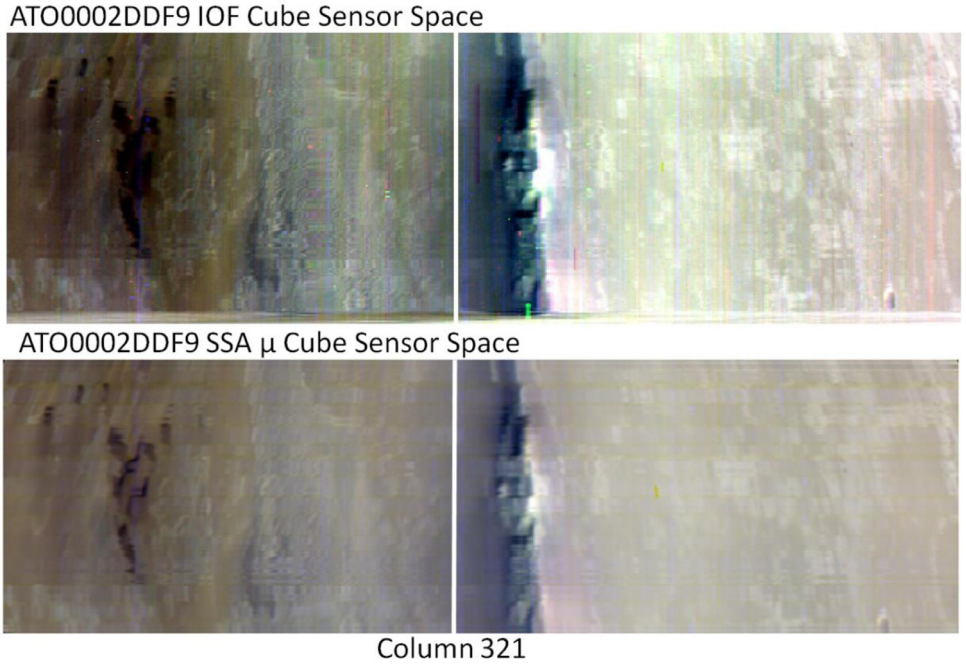

Fig. 4. Sensor space false color composites of I/F L data for ATO0002DDF9 generated from spectral radiance data without any further processing (top) and as SSA data after removing extrema and regularization using the log maximum likelihood algorithm (bottom). The MLM SSA data show less noise elements, especially column-dependent noise. Vertical white lines show the column for which row to row pixel spacing is plotted in Fig. 8. North is roughly in the downward direction because CRISM acquires data during its ascending orbits. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.).

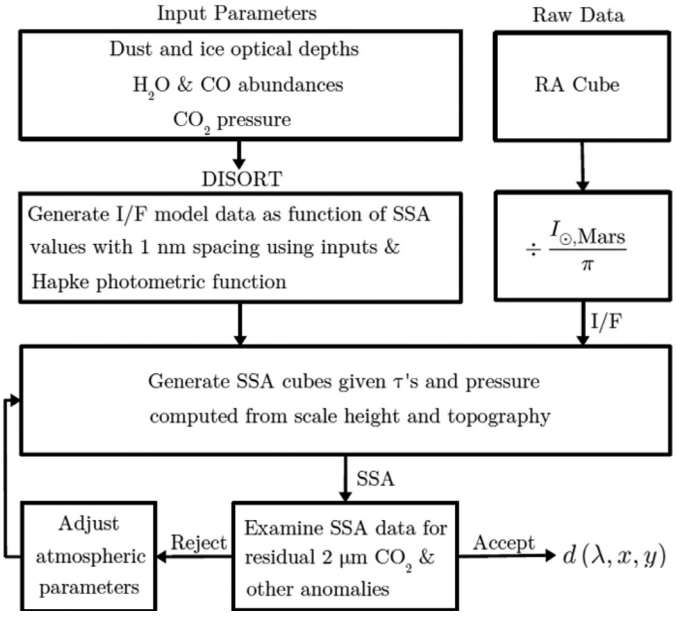

Fig. 5. Flow chart illustrating the approach used to retrieve SSA estimates from CRISM spectral radiance (RA) files. The DISORT radiative code is used to model atmospheric gases and aerosols along with the Hapke photometric function as a surface boundary condition. After a series of iterations to minimize residual CO_2 gas absorptions the output is an SSA data cube ($d(\lambda, x, y)$, where λ is wavelength and x, y represent spatial dimensions of the hyperspectral image cube).

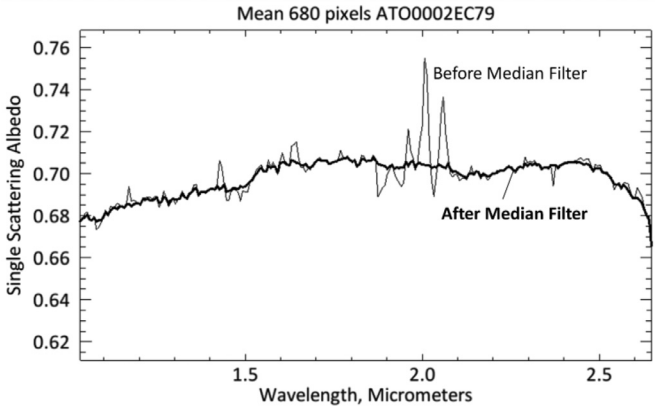

Fig. 6. SSA spectra are shown for a mean of 680 pixels from the sensor space image cube ATO0002EC79 covering a spectrally homogeneous area on the hummocky plains. Area is shown in Appendix Fig. A1. Data are shown with and without removal of extreme values using a median filter as described in the text.

The log-likelihood function for Poisson distributed data is given by

$$l(d|\mu) = \sum_{\lambda\in\Lambda}\sum_{x\in\mathfrak{X}}\sum_{y\in\mathfrak{Y}}\{d(\lambda,x,y)\ln[\mu(\lambda,x,y)] - \mu(\lambda,x,y)\} \quad (2)$$

where

$$\mu(\lambda,x,y) = \sum_{\lambda'\in\Lambda'}\sum_{x'\in\mathfrak{X}'}\sum_{y'\in\mathfrak{Y}'} h\left(\lambda,x,y|\lambda',x',y'\right)c\left(\lambda',x',y'\right) \quad (3)$$

with

$$h\left(\lambda,x,y|\lambda',x',y'\right) = h_{spec}\left(\lambda|\lambda',y\right)h_{spat}\left(x,y|x',y'\right) \quad (4)$$

where Λ, $\mathfrak{X}$, and $\mathfrak{Y}$ are sets of wavelengths, rows, and columns in sensor space, respectively. Primed coordinates are projected on the Martian surface and the h variables represent the spectral and spatial transfer functions for the CRISM instrument, which vary with wavelength and column for the 2D detector array. To reiterate, $\mu(\lambda, x, y)$ is the best estimate for the observed data taking into account the instrument spatial and spectral transfer functions, $c(\lambda', x', y')$ is the best estimate of the pure data, and $d(\lambda, x, y)$ is the measured CRISM data.

I-divergence is an information-theoretic measure of discrepancy between two nonnegative vectors or functions (Csiszár, 1991). The I-divergence between $\mu(\lambda, x, y)$ and $d(\lambda, x, y)$ is given by

$$I(d|\mu) = \sum_{\lambda\in\Lambda}\sum_{x\in\mathfrak{X}}\sum_{y\in\mathfrak{Y}}\left\{d(\lambda,x,y)\ln\left[\frac{d(\lambda,x,y)}{\mu(\lambda,x,y)}\right] - d(\lambda,x,y) + \mu(\lambda,x,y)\right\} \quad (5)$$

where $I(d|\mu) \geq 0$, with equality if and only if

$$d(\lambda,x,y) = \mu(\lambda,x,y)\ \forall\lambda,x,y \in \Lambda,\mathfrak{X},\mathfrak{Y}, \text{ respectively.} \quad (6)$$

By inspection, minimizing the I-divergence is equivalent to maximizing the log-likelihood for the above functions.

For the projected SSA image cube $c(\lambda', x', y')$ to be a minimizer of the I-divergence the following relationship must hold:

$$\sum_{\lambda\in\Lambda}\sum_{x\in\mathfrak{X}}\sum_{y\in\mathfrak{Y}}\frac{h\left(\lambda,x,y|\lambda',x',y'\right)d(\lambda,x,y)}{\mu(\lambda,x,y)} \begin{cases} = H_0\left(\lambda',x',y'\right) \text{ for } c\left(\lambda',x',y'\right) > 0 \\ \leq H_0\left(\lambda',x',y'\right) \text{ for } c\left(\lambda',x',y'\right) = 0 \end{cases} \quad (7)$$

where

$$H_0\left(\lambda',x',y'\right) = \sum_{\lambda\in\Lambda}\sum_{x\in\mathfrak{X}}\sum_{y\in\mathfrak{Y}} h\left(\lambda,x,y|\lambda',x',y'\right). \quad (8)$$

The above statements are known as the Karush-Kuhn-Tucker (KKT) conditions (Kuhn and Tucker, 1951; Karush, 1939).

Consider iterative updates for estimates of $c(\lambda', x', y')$ defined by

$$c^{(k+1)}\left(\lambda',x',y'\right) = \frac{c^{(k)}\left(\lambda',x',y'\right)}{H_0(\lambda',x',y')}\sum_{\lambda\in\Lambda}\sum_{x\in\mathfrak{X}}\sum_{y\in\mathfrak{Y}} h\left(\lambda,x,y|\lambda',x',y'\right) \times\frac{d(\lambda,x,y)}{\mu^{(k)}(\lambda,x,y)} \quad (9)$$

where

$$\mu^{(k)}(\lambda,x,y) = \sum_{\lambda''\in\Lambda''}\sum_{x''\in\mathfrak{X}''}\sum_{y''\in\mathfrak{Y}''} h\left(\lambda,x,y|\lambda'',x'',y''\right)c^{(k)}\left(\lambda'',x'',y''\right) \quad (10)$$

and where $\{\Lambda', \mathfrak{X}', \mathfrak{Y}'\} = \{\Lambda'', \mathfrak{X}'', \mathfrak{Y}''\}$ (the sets are identical). The sequence of images $\{c^{(k)}(\lambda', x', y')\}$, $k=0, 1, 2\ldots$ monotonically converges to the I-divergence minimum that satisfies the KKT conditions (Eq. (7)). This sequence is an expectation-maximization (EM) algorithm for Poisson data (Dempster et al., 1977; Shepp and Vardi, 1982; Snyder and Miller, 1991). It is also referred to as the Lucy-Richardson algorithm (Richardson, 1972; Lucy, 1974).

The convex nature of the data-fit term $l(d|\mu)$ is important to note. If the only element in the null space of Eq. (4) is $c(\lambda', x', y')=0$, then the method guarantees the sequence $\{c^{(k)}(\lambda', x', y')\}$ converges to a unique I-divergence minimizer. Snyder et al. (1992) and others rigorously prove this property.

The presence of noise in CRISM data and the ill-posed nature of the forward operator create instabilities in estimates for $c^{(k)}(\lambda', x', y')$. To overcome this problem we include a convex penalty function in the Huber class in the spatial and spectral domains to guar-

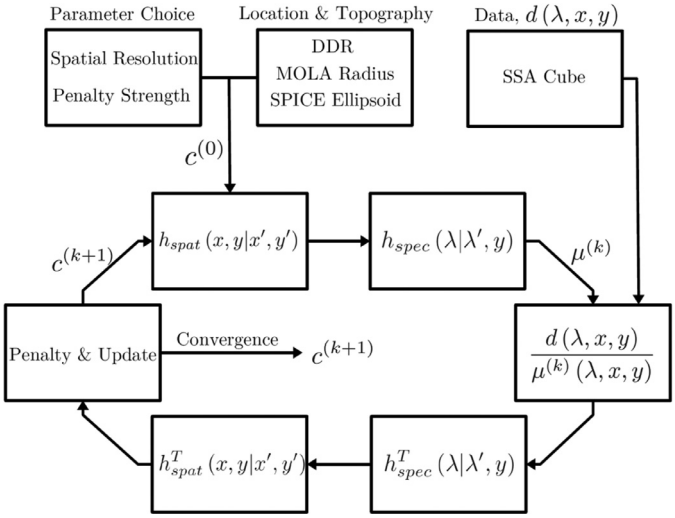

Fig. 7. Flowchart illustrating the iterative maximum log-likelihood with penalty function used to retrieve the best estimate of SSA values in the presence of Poisson noise. Inputs include the despiked version of the SSA cube, the CRISM DDR that contains latitude and longitude values for each sensor space pixel, the location of MRO as a function of time during which the data were acquired, and topographic information from MOLA data. The h variables are the spatial and spectral transfer functions, μ is the estimate of the actual SSA values, blurred by the transfer functions, and c is the projected version.

antee a global solution (e.g., Green, 1990; De Pierro, 1994; Bouman and Sauer, 1993):

$$\Phi(c) = \sum_{\lambda'_i, x'_i, y'_i} \sum_j \left(\beta_{spat} \delta^2_{spat} \ln \left\{ \cosh \left[\frac{c(\lambda'_i, x'_i, y'_i) - c(\lambda'_i, x'_j, y'_j)}{\delta_{spat}} \right] \right\} + \beta_{spec} \delta^2_{spec} \ln \left\{ \cosh \left[\frac{c(\lambda'_i, x'_i, y'_i) - c(\lambda'_j, x'_i, y'_i)}{\delta_{spec}} \right] \right\} \right) \quad (11)$$

where β is the weighting, δ is the threshold for transition between quadratic and linear regimes, $c(\lambda'_i, x'_i, y'_i)$ is the pixel of interest, and $c(\lambda'_j, x'_j, y'_j)$ is a neighboring pixel. Fixed values for β and δ were chosen based on testing multiple scenes with varying locations, lighting and viewing conditions, and spanning several years of observations on Mars. The spectral and spatial values for β are 0.04 and 0.01, respectively, and for δ are 0.9 and 4.0, respectively, and achieve a balance between smoothing the spectral noise and maintaining unique spectral features. Thus the overall approach is minimizing a penalized I-divergence $\min_c I(d|\mu) + \Phi(c)$, which is a convex function of the hyperspectral data cube c, guaranteeing that the iterative algorithm converges. For small value differences, the Huber penalty reduces to Tikhonov regularization (see e.g. Fraeman et al., 2012 for early work on using Tikhonov regularization with CRISM), while for large differences it reduces to total variation regularization.

3.2. Forward projection

To implement the MLM approach we begin with an initial map projected version $c(\lambda', x', y')$ of the SSA data populated with a constant value, i.e., an initial guess, and then forward project to sensor space to generate an initial blurred version of the scene $\mu(\lambda, x, y)$ (Fig. 7). The output spatial resolution for $c(\lambda', x', y')$ is chosen based on the typical along track pixel separation as shown in Fig. 8. A pixel size of 12 m is reasonable for the two example ATOs discussed in this paper, given that the median along-track spacing of the pixels ranges from ~11 to 13 m. We use a Mars equirectangular projection because both data sets were acquired near the equator.

The spectral and spatial transfer functions can be modeled as either single or multiple Gaussians based on preflight calibrations. We have found that the single Gaussian is a good approximation to the transfer functions as shown below for the spectral function:

$$h_{spec} = \frac{2\sqrt{\ln(2)}}{\mathrm{FWHM}_{spec}\sqrt{\pi}} e^{-\ln(16)\left(\lambda'_{c,i} - \lambda_0\right)^2 / \mathrm{FWHM}^2_{spec}} \quad (12)$$

where the FWHM_{spec} is given for each detector in nanometers, $\lambda'_{c,i}$ is the wavelength for pixels $p_{c,i}$ in $c^{(k)}(\lambda', x', y')$, and λ_0 is the center wavelength for the pixel in $d(\lambda, x, y)$.

The spatial transfer function is also modeled as a Gaussian with the spectral FWHM converted to the spatial domain. For the case that the target pixel is directly beneath the spacecraft on a flat surface (nadir viewing geometry), the weight for a neighboring pixel is

$$h_{spat} = S(\lambda') A_{c,i} \frac{\ln(16)}{\pi \mathrm{FWHM}^2_{spat}} \times e^{-\ln(16) R^2_{\sigma} \{[\mathrm{lat}(x'_{c,i}) - \mathrm{lat}(x_0)]^2 + [\mathrm{lon}(y'_{c,i}) - \mathrm{lon}(y_0)]^2\} / \mathrm{FWHM}^2_{spat}} \quad (13)$$

where

$$\begin{aligned} \mathrm{FWHM}_{spat} &= \mathrm{FWHM}_{spec} \frac{\text{spatial sampling}}{\text{spectral sampling}} \\ &= \mathrm{FWHM}_{spec} \frac{2\xi \tan(\mathrm{IFOV}/2)}{6.55 \text{ nm/channel}} \\ &\approx \mathrm{FWHM}_{spec} \frac{\xi \, \mathrm{IFOV}}{6.55 \text{ nm/channel}} \end{aligned} \quad (14)$$

and where $x'_{c,i}$ and $y'_{c,i}$ correspond to coordinates for pixels $p_{c,i}$ in $c^{(k)}(\lambda', x', y')$ that neighbor the kernel center p_c, x_0 and y_0 are the coordinates for the pixel of interest p_d in $d(\lambda, x, y)$, $A_{c,i}$ is the effective area of pixel $p_{c,i}$, $S(\lambda')$ is the normalization factor, R_{σ} is the radius of mars, ξ is the instrument's altitude above the Martian surface, IFOV = 61.5 μrad, 6.55 nm gives the average spacing between CRISM bands, and FWHM_{spec} is the FWHM of the single Gaussian approximation to the spectral transfer function. The normalization is done band by band. The single Gaussian spectral FWHM is provided as a PDS data product, but the spatial FWHM is not and so must be estimated by converting the spectral FWHM to spatial units, which gives a reasonable estimate. The forward projection then maps $c(\lambda', x', y')$ to the blurred sensor space version of the data, $\mu(\lambda, x, y)$.

In general the target pixel is not located at a nadir view. To incorporate how the spatial transfer functions sample the surface for off-nadir viewing and surface topography, we map the Gaussian kernel onto an image plane centered at the target pixel and perpendicular to the line of sight from CRISM. The position and orientation of this plane are derived from several data sets. SPICE data give the spacecraft position as a function of time. The DDR provided with CRISM archives supplies elevation (above the areoid) and areocentric latitude and longitude for the target pixel p_d, which we combine with the baseline areoid from MOLA Mission Experiment Gridded Data Records (MEGDRs), to calculate the vector position of the target pixel. The DDR also provides slope magnitude and azimuth for each CRISM pixel, from which we derive the unit normal vector of the target pixel, allowing projection of each pixel area onto the image plane and proportionally weighting its contribution. Elevations and unit normal vectors for neighboring pixels $p_{c,i}$ are derived via interpolation. In the off-nadir transfer function, the exponential argument is the squared distance between target and neighbor pixels along the image plane, rather than along a spherical surface.

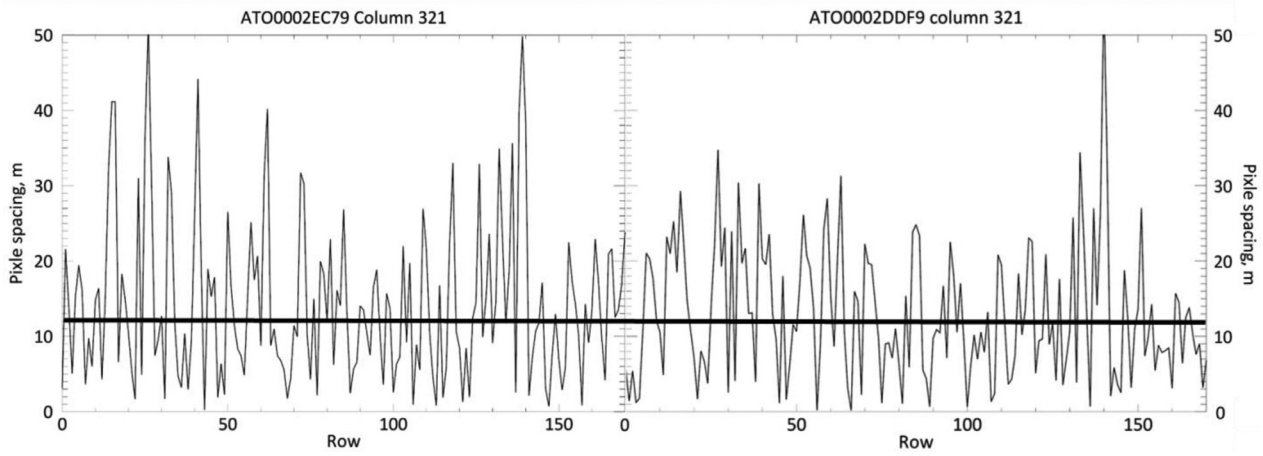

Fig. 8. Plot of row to row pixel spacing projected onto the surface for column 321 for ATO0002EC79 (Left) and ATO0002DDF9 (Right). The large variation is a consequence of the gimbal actuator jitter as it rotates CRISM's optical system to follow an overlapping path along the ground track of the spacecraft. Dark horizontal lines are shown at 12 m spacing and approximately represent the median value which was then used to regularize the *c* cube (projected SSA data) at 12 m/pixel.

3.3. Backward projection and iterations

The back projection in our MLM procedure maps the blurred version of the SSA data in sensor space back to the map-projected version of the scene. The $d(\lambda, x, y)$ SSA data cube is divided by the $\mu^{(k)}(\lambda, x, y)$ cube and back projected through the transfer functions to generate a new estimate of $c^{(k+1)}(\lambda', x', y')$ as shown in Fig. 7, with appropriate application of the penalty function described by Eq. (11). Deblurring is an inverse problem, and it is imperative that the transfer function convolution orders are consistent throughout the algorithm. The back projection begins with the spectral followed by the spatial transfer functions. The new estimate of the scene SSA cube is then used in an iterative fashion through the same procedure to produce a series of estimates of the *I*-divergence (Eq. (5)), decreasing to the best estimate of the SSA projected cube in the presence of Poisson noise, as shown by Snyder et al. (1992). We compute:

$$\begin{aligned}\underset{c}{\operatorname{argmin}} \Big(& -H_0(\lambda', x', y')c(\lambda', x', y')_{EM}^{(k+1)} \ln\left[c(\lambda', x', y')\right] \\ & + H_0(\lambda', x', y')c(\lambda', x', y') + \beta \sum_{\tilde{x}\in\mathcal{N}(\lambda', x', y')} \delta^2 \\ & \times \ln\left\{\cosh\left[\frac{2c(\lambda', x', y') - c^{(k)}(\lambda', x', y') - c^{(k)}(\tilde{x})}{\delta}\right]\right\}\Big)\end{aligned} \tag{15}$$

where we first compute the update without the penalty from Eq. (9) to yield $c_{EM}^{(k+1)}$. The first two terms are the non-trivial *I*-divergence terms when minimizing with respect to $c(\lambda', x', y')$, whereas the third term incorporates the penalty function into the minimization. This third term is the upper bound on the penalty which is decoupled across all spatial-spectral coordinates, allowing for a parallel implementation to find the solution (De Pierro, 1994). The notation in Eq. (15) in the summation for the penalty $\tilde{x} \in \mathcal{N}(\lambda', x', y')$ is over the set of spatial-spectral coordinates in the neighborhood of the point (λ', x', y'). This neighborhood set includes a larger domain in the spectral dimension than in the spatial.

Newton updates to find the stationary points of the argument in Eq. (15) are used with a trust region. The final results are the best estimates of the scene SSA values blurred by the spectral and spatial transfer functions in sensor space (μ cube) and the map projected version with the transfer functions removed (*c* cube) (Figs. 1–4). The power and robustness of the MLM algorithm is best seen with a direct comparison to previous processing techniques and in combination with other measurements, both in situ and orbital.

4. ATO0002EC79 covering curiosity traverses in Gale Crater

Curiosity landed on the hummocky plains north of Aeolis Mons (informally known as Mount Sharp) in August 2012 and crossed onto the lower slopes of Mount Sharp in 2015, characterizing the fluvial-deltaic-lacustrine basal deposits with its complement of remote sensing, contact science, and analytical instrumentation (Grotzinger et al., 2015). Portions of the northern outcrops of Mount Sharp are covered with wind-blown sands informally named the Bagnold Dunes. In 2015–2016 the Namib Dune within the Bagnold Dune Field was the site of an extensive campaign focused on understanding the grain size distributions and mineralocurrences and abundances (Ehlmann et al., 2016). The rationale was to understand modern wind-blown sand dynamics to help better understand the environments of deposition for the sandstones encountered during Curiosity's traverses. The Namib Dune measurement campaign also provided an unprecedented opportunity to use ATO0002EC79 (Table 2, Figs. 1, 2, 9, 10) regularized SSA spectra, spatially localized to 12 m/pixel, to model the grain sizes and mineral abundances using Hapke theory for direct comparison to Curiosity's measurements (Lapôtre et al., 2016).

To illustrate improvements using our methodology we also projected SSA data at 12 m/pixel using a popular map projection code (from the Environment for Visualizing Images (ENVI) Version 5.3) that does not utilize spectral or spatial transfer functions or retrieval of the best estimates of the scene in the presence of Poisson noise. The resultant product and our MLM rendition are both shown as false color images in Fig. 9. The MLM product sharpens morphological details and removes noise, most noticeably in the form of vertical stripes. For example, the texture of the washboard terrain that is located south and southwest of Curiosity's location on Sol 1298 becomes evident with MLM processing (Fig. 9). Comparing the Namib Dune and surrounding washboard terrain in the MLM product with the corresponding HiRISE image (McEwen et al., 2007), with its spatial resolution of 25 cm/pixel, also illustrates the high level of textural detail in MLM processed CRISM images (Fig. 10) as opposed to the normally processed data. In addition, the MLM product better localizes spectral information through use of the spatial transfer function to localize the SSA values to specific 12 m/pixel locations.

The utility of increased localization of spectral signatures using the MLM methodology is illustrated using Curiosity's measure-

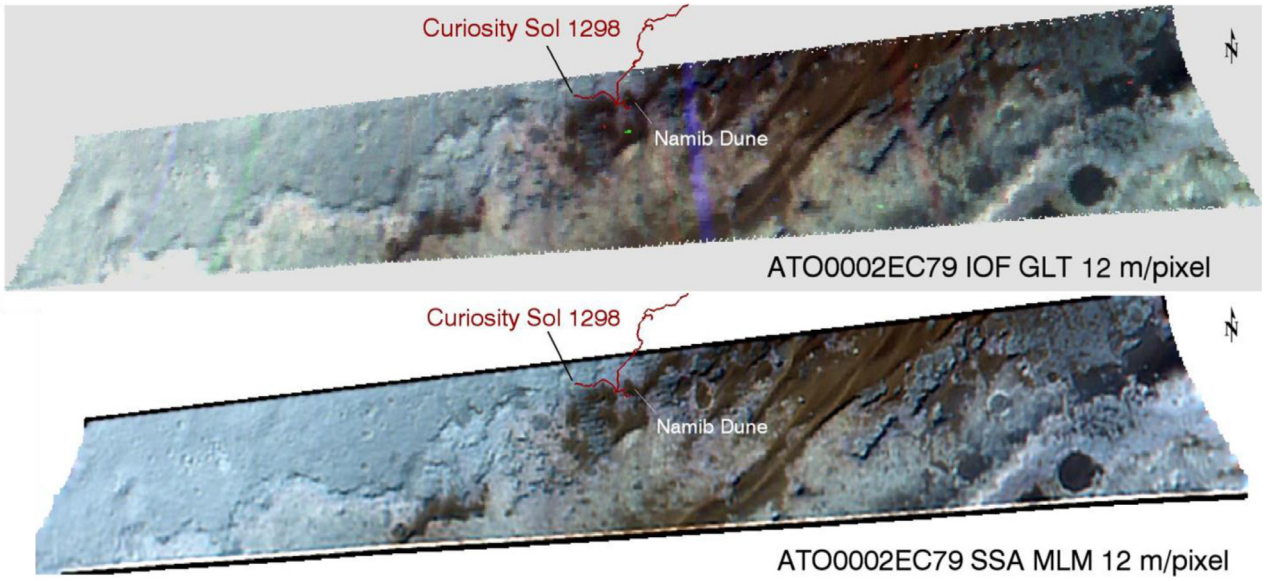

Fig. 9. (Top) False color L SSA data ATO0002EC79 projected at 12 m/pixel using a standard projection tool without removing spikes, use of the log maximum likelihood approach for retrieving the best estimate of SSA values, or removing the blurring associated with spatial and spectral transfer functions. (Bottom) Equivalent product after extrema removal and use of the maximum log-likelihood approach, projected at 12 m/pixel. Note the increase of sharpness for the latter product. Both figures are overlain with Curiosity's traverses. The Namib Dune, site of an extensive Curiosity measurement campaign, is labeled on both figures. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

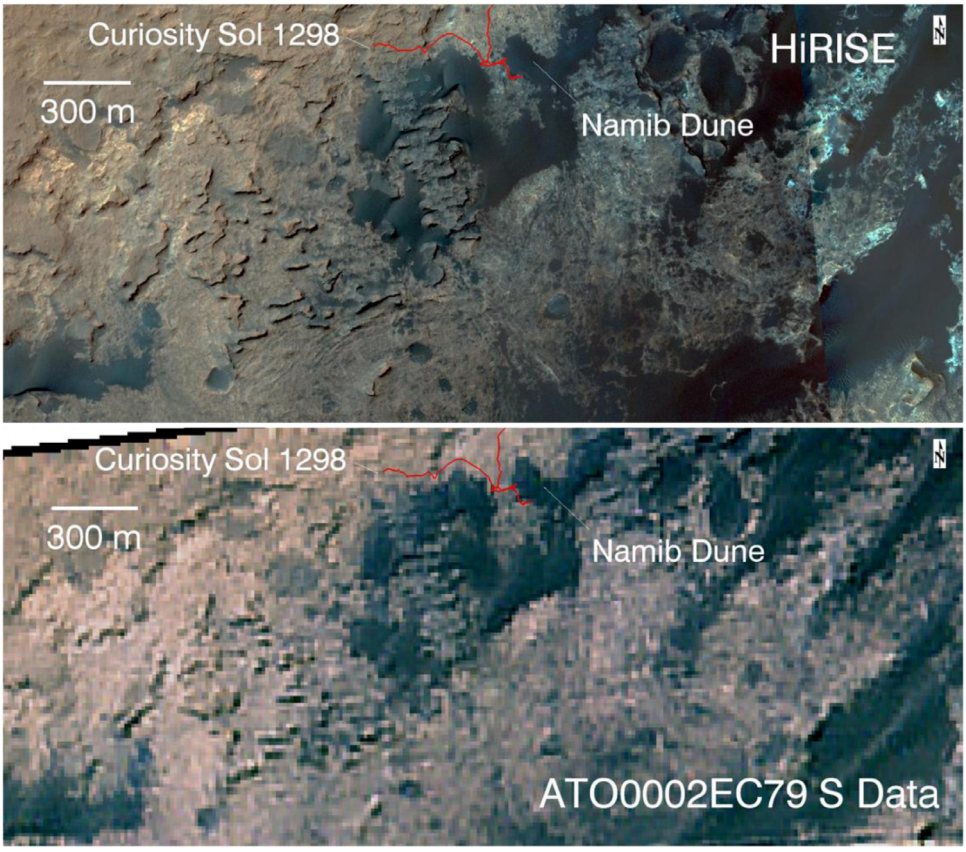

Fig. 10. HiRISE color mosaic is shown at the top, with Curiosity's traverses and the location of the Namib Dune labeled. Bottom frame shows ATO0002EC79 S data regularized to 12 m/pixel with similar annotation. Both scenes have been enhanced to approximately the same false color views. HiRISE RGB colors are derived from the red channel (0.57 to 0.83 μm) and blue-green channel (<0.58 μm), and blue has been estimated from the red and blue-green data. For the CRISM color image RGB corresponds to bands located at 0.7097, 0.5989, and 0.5337 μm. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 2
Summary of Observational Parameters for CRISM Data. L_s is the solar longitude or season with $L_s = 0°$ corresponding to the southern hemisphere fall equinox. Sol is the Mars day after Curiosity (Gale Crater) and Opportunity (Endeavour Crater) rover landings. The next four rows correspond to atmospheric optical depths of dust and ice aerosols, CO_2 surface pressure for the central portion of the scenes, and the amount of water vapor that best models the I/F values in DISORT. H_2O ppt stands for the thickness of water that would be precipitated onto the surface if converted from vapor to liquid phases. Detector temperature is provided next. **δλ** is the time-dependent wavelength shift for CRISM observations. The range of elevation, incidence, emergence, and phase angles are then provided for the scene, relative to the areoid.

Parameter	ATO0002EC79 Gale Crater	ATO0002DDF9 Endeavour Crater
Acquisition Date	March 26, 2014	February 9, 2014
L_s	103°	88°
Center Latitude	4.69° S	2.32° S
Center Longitude	320.95° E	354.66° E
Curiosity/ Opportunity Sol	583	3573
Curiosity/ Opportunity τ (at 0.90 μm)	0.45	0.54
τ_{dust} (at 0.90 μm)	0.40	0.45
τ_{ice} (at 0.32 μm)	0.05	0.09
Pressure	8.10 mb	7.20 mb
H_2O ppt	9.3 μm	9.0 μm
T Detector	−148 °C	−148 °C
$\delta\lambda$	−0.5297 nm	−0.4322 nm
Incidence Angle	57°	55°
Emergence Angle Range	2–32°	4–32°
Phase Angle Range	59–82°	58–80°

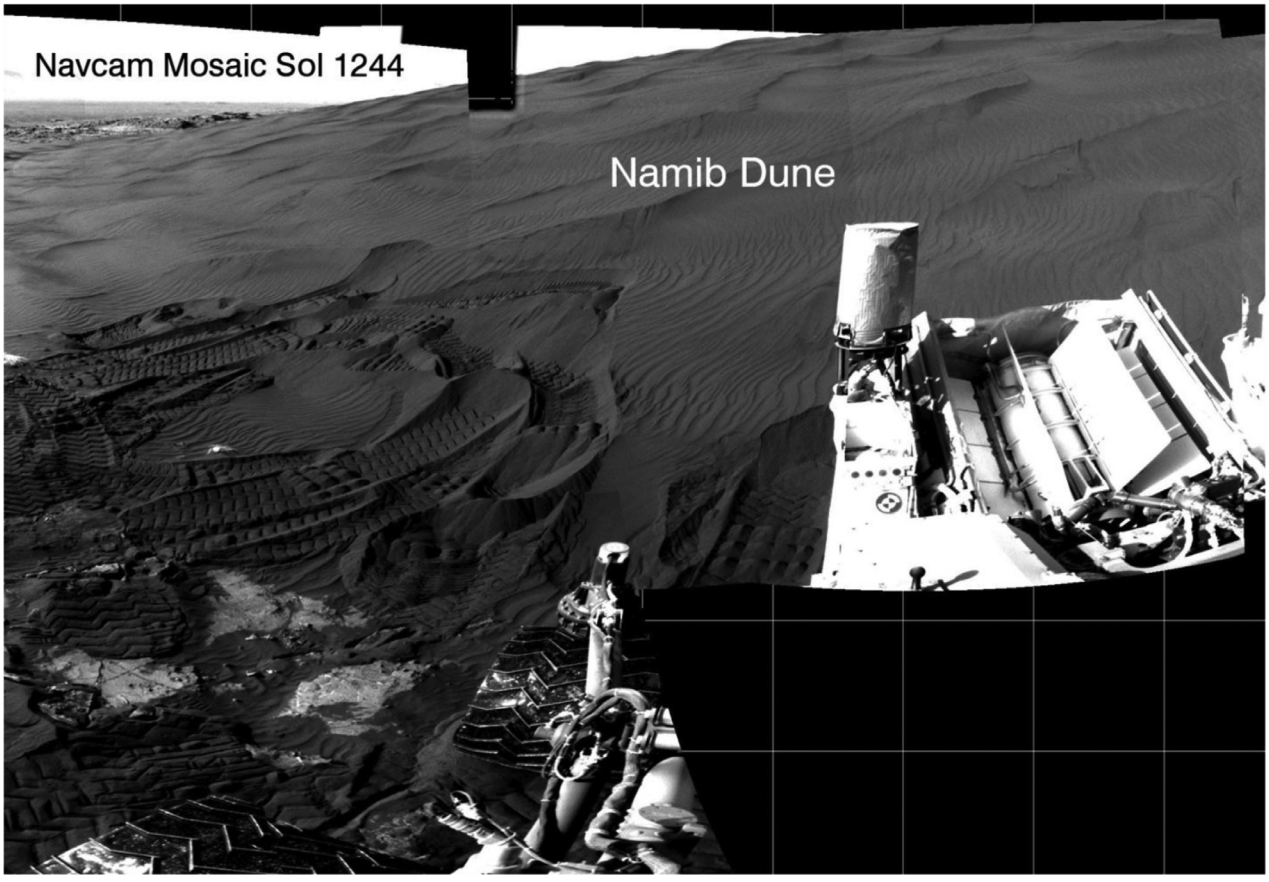

Fig. 11. Portion of a Navcam mosaic of Namib Dune acquired in the location where Curiosity did an extensive remote sensing, contact science, scooping and delivery to CheMin, the on-board X-ray diffraction spectrometer. Tracks show motions associated with positioning for scooping sand with the robotic arm.

ment campaign on the western side of Namib Dune, a feature well-resolved in the MLM-processed data. The measurement campaign included scooping and sieving dune sands, and retrieving mineralogy using the Chemistry and Mineralogy (CheMin) instrument's transmission X-ray diffraction (XRD) capabilities (Achilles et al., 2016) (Fig. 11). For modeling the dune sands using ATO0002EC79, an 89 pixel region of interest was mapped over the Namib Dune and SSA spectra were retrieved for both S and L data for this region (Fig. 12). Individual SSA spectra for minerals of a given grain size combine linearly based on Hapke theory (Hapke, 2012), and thus the problem becomes one of a robust retrieval of a unique set of minerals and grain sizes to replicate the regularized SSA

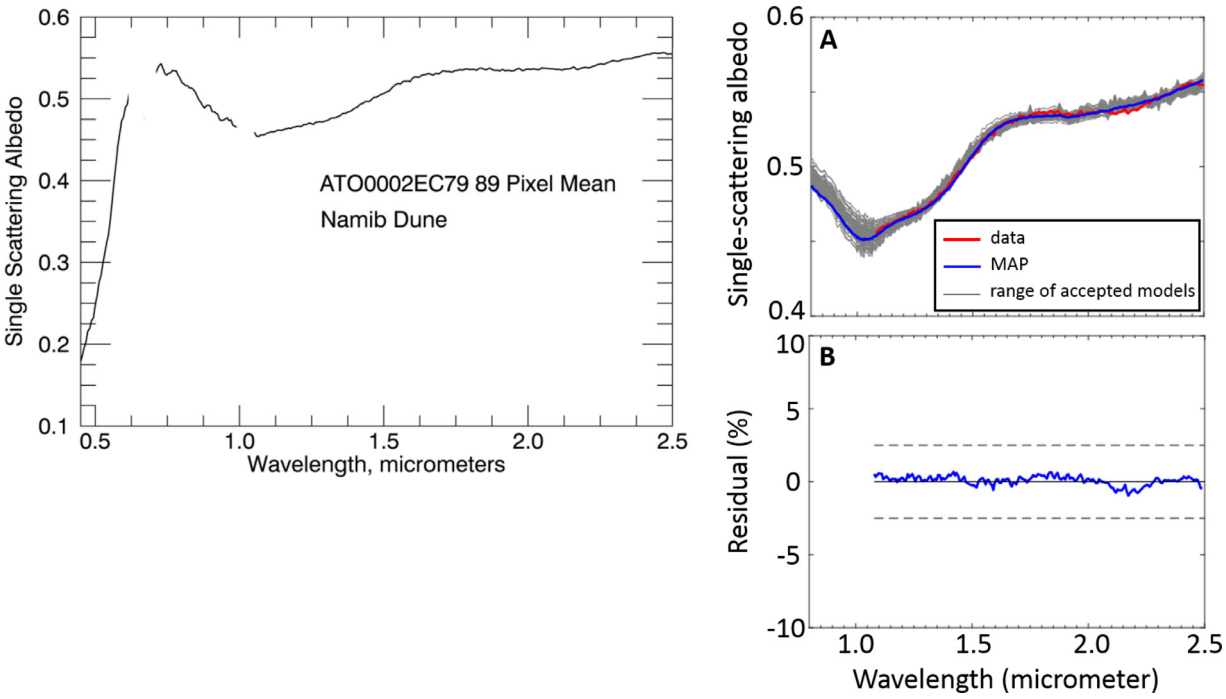

Fig. 12. (Above) SSA spectrum of Namib Dune where Curiosity conducted its measurements and shown for a mean of 89 pixels from the MLM processed image cube ATO0002EC79. Blanked regions correspond to the blocking filter for the S data, where there are known artifacts, and poorly defined values between the S and L detectors where low S/N leads to poor SSA retrievals. (A) SSA spectrum of Namib Dune (L-data, red), MAP (blue), and subset of accepted models (gray). (B) Residual (data/model) of the MAP. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 3
Namib Dune Mineral Abundances and Grain Sizes. Mineral abundances and grain sizes describing the MAP as inverted from the regularized SSA of sands at the Namib Dune. Corresponding 95% confidence intervals for the mean model are given in parentheses.

Mineral	Abundance (wt %)	Grain size (μm)
Olivine	24.4 (1.3--19.5)	773 (171--787)
High Ca Pyroxene	7.8 (0.3–16.7)	102 (75–782)
Low Ca Pyroxene	5.7 (0.2–37.4)	262 (74–782)
Plagioclase	16.3 (9.7–50.0)	146 (82–741)
Magnetite	3.9 (0.7–29.2)	133 (26–194)
Basaltic Glass	41.9 (12.2–52.9)	193 (158–663)

data (Lapôtre et al., 2016). A Markov-Chain Monte Carlo (MCMC) algorithm (Minson et al., 2013) is implemented to invert for probability densities of mineral abundances and corresponding grain sizes from the SSA spectrum, using optical constants of olivine, two pyroxenes (low and high Ca), a plagioclase, magnetite, and basaltic glass as inputs, and Hapke's theory (Hapke, 2012) for forward modeling and evaluation of model likelihood. This technique allows mapping the likelihood of given mineral assemblages and calculating associated confidence intervals. The inversion was performed from the L data only to avoid the strong effects of ferric oxides in dust at shorter wavelengths. Table 3 summarizes weight abundance and grain size parameters describing the maximum a posteriori probability (MAP) estimate, i.e. the most probable model for the 12-dimensional posterior probability density function (PDF) resulting from the MCMC inversion. 95% confidence intervals are presented that correspond to the mean model, the mean of all 500,000 accepted models (Lapôtre et al., 2016). Whereas the inversion is mostly insensitive to grain size, inverted mineral abundances favorably compare to those measured with CheMin in the less than 150 μm-sand fraction, with a mean error of 7.5 wt% and a maximum error of 14.9 wt% for olivine (Achilles et al., 2016; Lapôtre et al., 2016). The normally processed data (i.e., "volcano scan" processed data) could not be fit in a satisfying way by any mixture of the same mineral endmembers. SSA values were estimated from the volcano-scan corrected data by normalizing by the cosine of the incidence angle. Inversions using ATO SSA data without MLM regularization produce similarly good mineral assemblages. However, the regularized data enables the use of a stricter covariance criterion in the MCMC inversion, and the corresponding MAP residual has an RMS about four times lower than that of the MAP for the non-regularized ATO SSA data.

5. ATO0002DDF9 covering opportunity traverses in Endeavour Crater

Opportunity has been exploring the rim of the 22 km wide, Noachian age Endeavour Crater since 2010 with a focus on characterizing evidence for aqueous alteration of the Shoemaker formation impact breccias and the underlying finely-layered Matijevic formation (Squyres et al., 2012; Arvidson et al., 2014, 2016). CRISM ATO0002DDF9 was acquired to complement a suite of four ATO observations over the Cape Tribulation rim segment, where Wray et al. (2009) and Noe Dobrea et al. (2012) used non-ATO CRISM observations (i.e., 18 m/pixel) to detect evidence for dehydrated Fe^{3+} and Mg^{2+} smectites. Overlapping and regularized ATO observations projected at 12 m/pixel allowed detection and mapping within the Marathon Valley portion of Cape Tribulation of 2.3 and 2.4 μm M-OH combination absorptions uniquely associated with Fe^{3+} and Mg^{2+} smectites (Fox et al., 2016). ATO0002DDF9 is one of the observations used by Fox et al. (2016). It has the highest degree of overlap and will be used here to illustrate the advantages associated with using our methodology. Fig. 13 shows false color images for data projected at 12 m/pixel using standard map projection code and our MLM methodology. The morphology

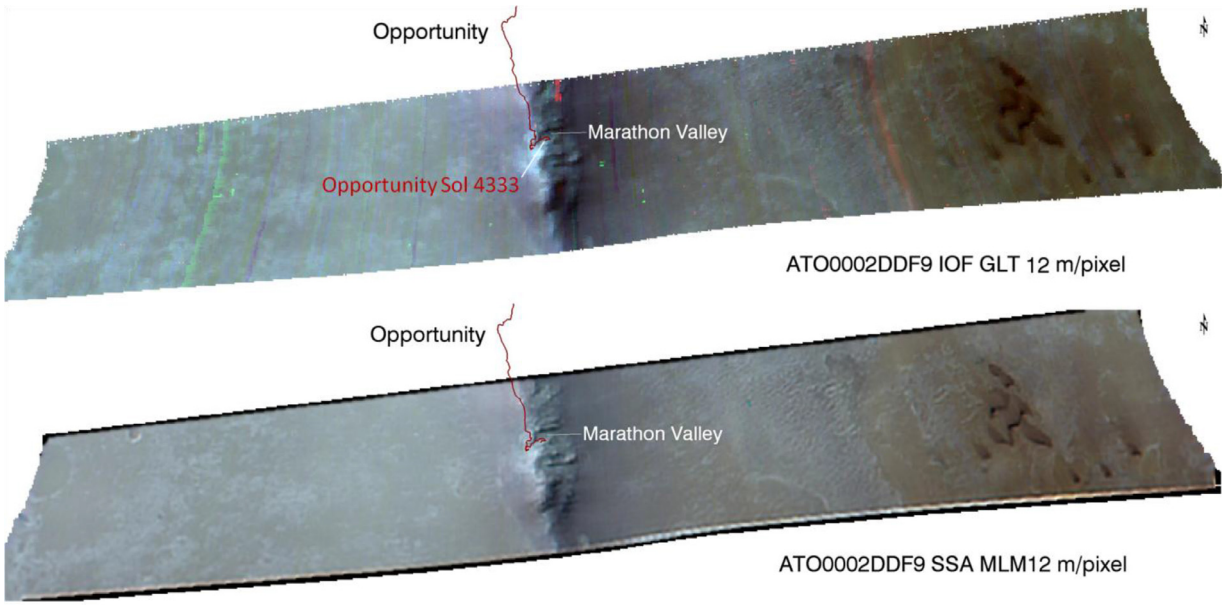

Fig. 13. (Top) False color L SSA data ATO0002DDF9 projected at 12 m/pixel using a standard projection tool without removing spikes, use of the log maximum likelihood approach for retrieving the best estimate of SSA values, or removing the blurring associated with spatial and spectral transfer functions. (Bottom) Equivalent product after extrema removal and use of the maximum log-likelihood approach, projected at 12 m/pixel. Note the increase of sharpness for the latter product. Both figures are overlain with Opportunity's traverses. Marathon Valley, a site of extensive exploration by Opportunity, is labeled in both figures. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

of Cape Tribulation and Marathon Valley is much sharper in the MLM as opposed to the other product. Further, use of standard processing methods limit smectite detections to a broad swath across Marathon Valley, with weak and noisy spectral retrievals. The MLM-processed version allows much better spatial localizations of SSA spectra, making it possible to discriminate at the outcrop level between regions of the valley floor that carry the smectite signature and regions that do not.

The smectite detections derived from the SSA spectra for scene ATO0002DDF9 are located in specific areas within Marathon Valley, Spirit of Saint Louis crater, and its surrounding apron (Fig. 14). The absorption depths are quite shallow relative to laboratory-derived spectra of pure phyllosilicates when both data sets are shown as continuum removed spectral radiance factors for direct comparisons. Opportunity was directed to characterize these smectite-bearing localities and found that they are associated with relatively planar, polygonally-fractured impact breccia outcrops (Fig. 15). The breccias, as characterized by Opportunity, have a basaltic composition, with a slight enrichment in Mg and S relative to other breccias measured on Endeavour's rim (Mittlefehldt et al., 2016; Fox et al., 2016). Rocks within fracture zones between the polygons tend to be redder and are enriched in Al and Si relative to surrounding rocks, suggestive of leaching by fracture-enhanced fluid flow (Mittlefehldt et al., 2016; Fox et al., 2016). The large areal extent of the polygonal outcrops relative to the fractures strongly argues that the outcrops carry the smectite signature and thus that the aqueous alteration was isochemical, with low water to rock ratios. The exception would be more extensive leaching along the fractures, consistent with enhanced fluid flow along these features. Similar results were found for the Matijevic formation on the Cape York Endeavour rim segment, with evidence from CRISM ATO data for Fe^{3+} smectites associated with planar outcrops that Opportunity found to be of basaltic composition (Arvidson et al., 2014). Evidence for extensive leaching was confined to the Esperance fracture zone, again implying alteration in a low water to rock environment, with the exception of enhanced fluid flow along fractures (Arvidson et al., 2014).

6. Conclusions

Mars Reconnaissance Orbiter Compact Reconnaissance Imaging Spectrometer for Mars (CRISM) hyperspectral image data have been acquired in an along-track oversampled (ATO) mode with intent of processing the data to better than the nominal ~18 m/pixel ground resolution. An iterative maximum log-likelihood method (MLM) was developed and implemented that utilizes the instrument spectral and spatial transfer functions and includes a penalty function to regularize the CRISM data, both in sensor space and as projected hyperspectral image cubes at 12 m/pixel. Preprocessing steps include retrieval of surface single scattering albedos (SSA) using the Hapke Function and DISORT-based radiative modeling of atmospheric gases and aerosols. Resultant SSA cubes are despiked to remove extrema and tested to ensure that the remaining data are Poisson-distributed, an underlying assumption for the MLM algorithm implementation.

Two examples of regularized ATO data sets are presented. ATO0002EC79 covers the route taken by the Curiosity rover during its ascent of Mount Sharp in Gale Crater. SSA data are used to model mineral abundances and grain sizes predicted to be present in the Namib barchan sand dune sampled and analyzed by Curiosity. CRISM-based results with high fidelity localization and use of spatial transfer functions to remove spectral effects from surrounding areas compare favorably to results derived from Curiosity's dune sand measurement campaign. ATO0002DDF9 covers Spirit of Saint Louis Crater and Marathon Valley on the Cape Tribulation rim segment of Endeavour Crater. High fidelity localization of SSA spectra demonstrate that Fe^{3+} and Mg^{2+} smectite signatures are located on specific planar impact breccia outcrops. Conversion of SSA spectra for these regions to spectral radiance factors allows direct comparisons to absorption depths for laboratory spectra and demonstrates that the outcrops have only a minor smectite component. Opportunity's exploration and characterization of the smectite-bearing outcrops show that these impact breccias are basaltic in composition and have been isochemically altered, most likely in a low water-to-rock environment.

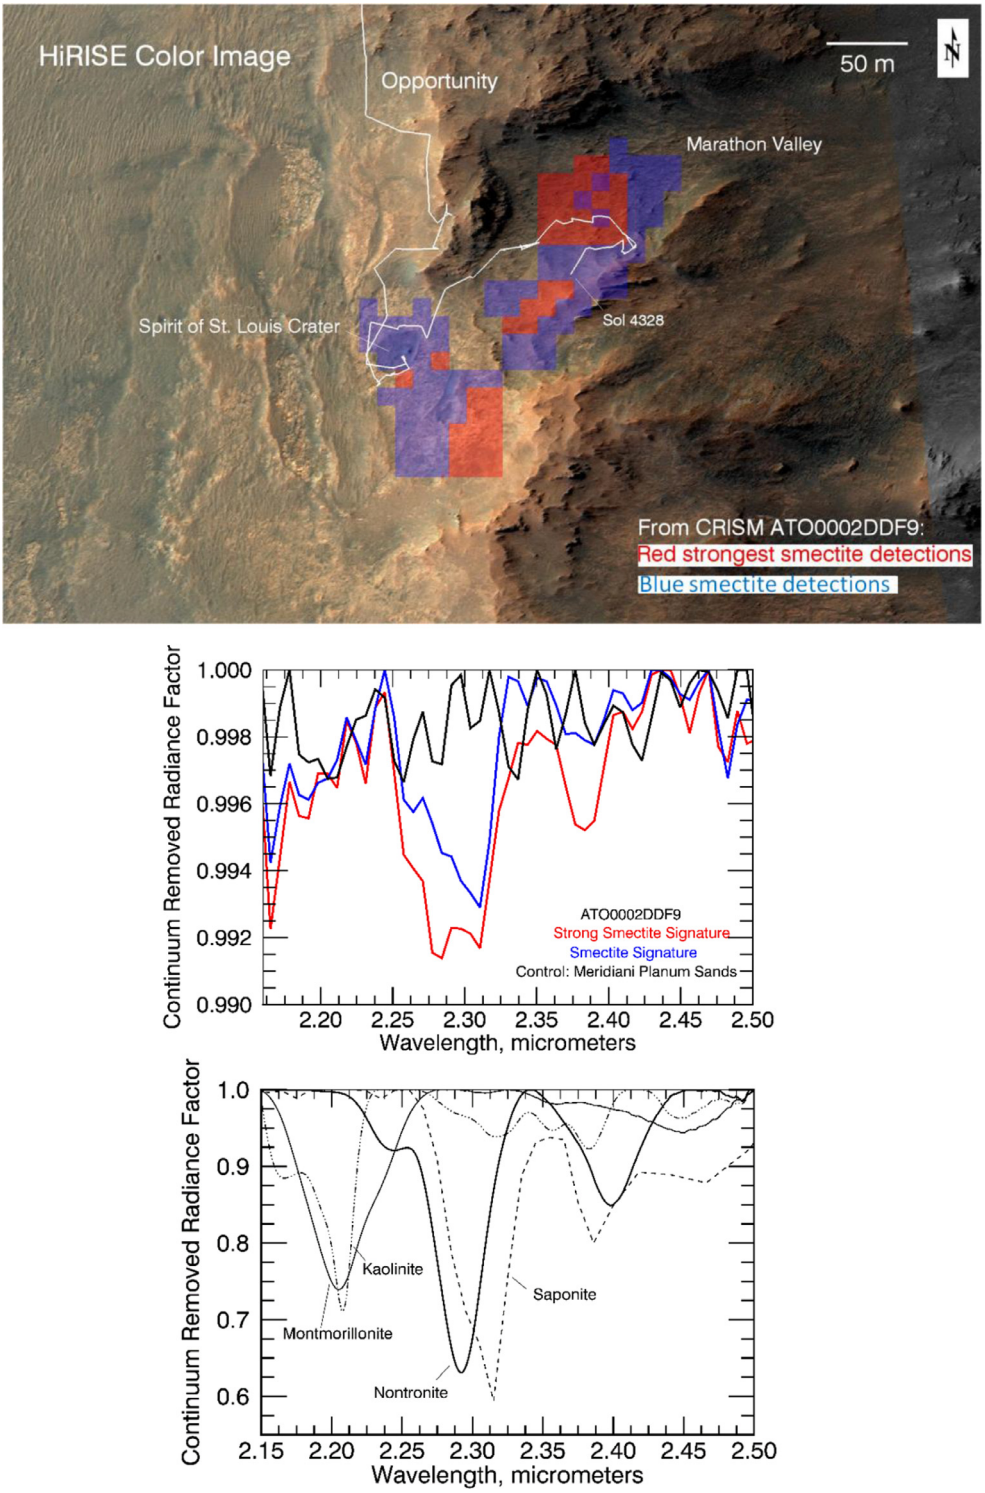

Fig. 14. (Above) Portion of a HiRISE color image covering Marathon Valley that shows mapped concentrations of Fe^{3+} and Mg^{2+} smectites derived from the maximum log-likelihood SSA cube projected at 12 m/pixel. Red areas correspond to locations in which both ~2.3 and 2.4 μm absorptions were twice the noise depth, whereas blue areas correspond to regions with absorptions at least 1.5 times the noise levels. (Upper right) Continuum-removed spectra from the SSA cube showing 2.3 and 2.4 μm absorptions due to metal-OH combination bands and indicative of Fe^{3+} and Mg^{2+} smectites are shown as means for the areas delineated in the HiRISE image. The Hapke function was used to first compute the radiance factor values for direct comparison to lighting and viewing conditions that pertained to collecting the laboratory data shown in the lower right plot. (Lower right) Laboratory-based spectra of pure phyllosilicate mineral powders. Both sets of spectra are shown in continuum removed form for radiance factors. For the CRISM data the radiance factors were estimated from SSA spectra and the Hapke function to correspond to lab spectra, which were acquired in a radiance factor lighting and viewing geometry of 0° emergence and 30° incidence angles. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

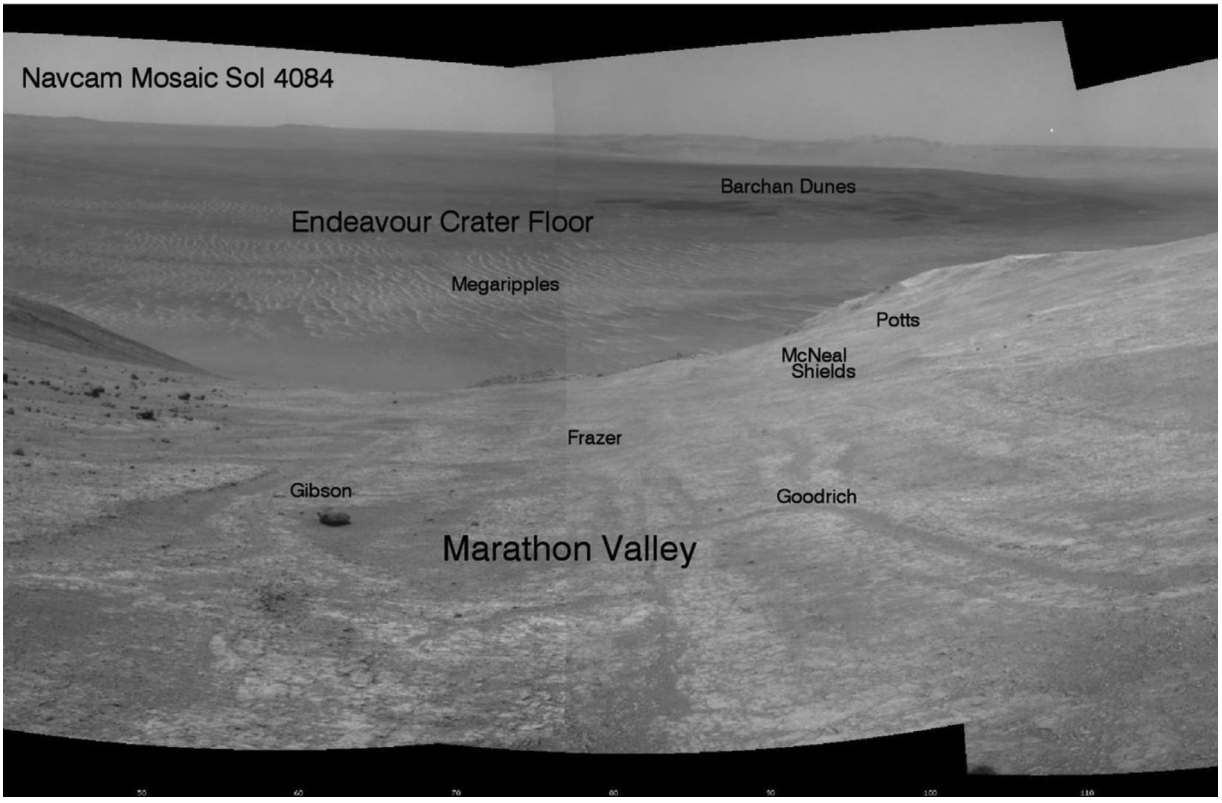

Fig. 15. Portion of an Opportunity Navcam mosaic looking to the east into Marathon Valley and the floor of Endeavour Crater. Several targets are shown for which Opportunity acquired detailed imaging and compositional data. Note the relatively planar and polygonal nature of the outcrops.

Acknowledgements

This work was supported in part from NASA through Contract 110609 from The Johns Hopkins University/Applied Physics Laboratory for participation in the CRISM Science Team, Cornell University Subaward 77301–10534 for participation in the Opportunity rover mission, and Contract 110609 from The California Institute Technology/Jet Propulsion Laboratory for participation in the Curiosity rover mission. We thank K. Li for optimization of the MLM code and A. A. Fraeman for early work on regularization. Many thanks to Valerie Fox, Kathryn Powell, and Susan Slavney for comments on the manuscript. Data analysis made use of ENVI (Exelis Visual Information Solutions, Boulder, Colorado). Computations were performed in part using the facilities of the Washington University Center for High Performance Computing, which were partially provided through grant NCRR 1S10RR022984-01A1.

Appendix A. Poisson hypothesis test

The MLM algorithm is based on an assumption that the data $d(\lambda, x, y)$ can be modeled as the product of Poisson distributed random variables $m(\lambda, x, y)$ and a scale factor $\alpha(\lambda, x, y)$

$$d(\lambda, x, y) = \alpha(\lambda, x, y) m(\lambda, x, y)$$

$$m(\lambda, x, y) \sim Poisson\left(\frac{\mu(\lambda, x, y)}{\alpha(\lambda, x, y)}\right) \quad \text{(A1)}$$

where the scale factor $\alpha(\lambda, x, y) = \alpha$ is assumed to be constant and independent of spectral band and spatial location, and $\mu(\lambda, x, y)$ is the blurred version of the scene, as before. To test this assumption we retrieved data from ATO0002EC79 for a spatially homogeneous area on the hummocky plains (Fig. A1) and found for despiked data that $\alpha = (6.47 \pm 2.73) \times 10^{-5}$, so α is essentially invariant (Fig. A2). It is important to choose a spatially homogeneous area to perform the Poisson test because (1) samples in a homogeneous area have the same mean and (2) a sufficient number of samples with the same mean is needed to estimate α. Fig. A3 shows a direct comparison between the expected Poisson cumulative and probability distribution functions and the corresponding histograms for scaled data $d(\lambda, x, y)/\alpha$ retrieved from the homogeneous area of Fig. A1 at wavelength 1.99 µm. Because data at all bands (with the exception of beginning and end bands, affected by boundary effects) show similar results, only one wavelength is presented as an example. Although there are some differences between the shape of the Poisson distribution and the histogram near its peak, the cumulative distributions match closely, providing qualitative motivation that the region is Poisson distributed.

Given the homogeneous region appears Poisson distributed, we scale the entire regularized sensor space SSA scene $\mu(\lambda, x, y)$, which is the mean of the data $d(\lambda, x, y)$, by the scale factor α. Chi-square tests are applied to robustly evaluate whether the data follow a Poisson or Gaussian distribution. The scene contains a large number of random variables with different means, due to differences in morphology and mineralogy. A random variable could be, for example, a given texture in the scene. Pixelsrepresenting the same mineralogical unit are samples of this random variable. These samples have the same distribution, and all random variables are assumed to be independent. Testing the distribution of a large number of samples from a single random variable has been studied extensively through, for example, the chi-square hypothesis test, the Kolmogorov-Smirnov Hypothesis test (Massey Jr., 1951), and the D' Agostino-Pearson test for Normality (D' Agostino and Pearson, 1973). However, CRISM image cubes are characterized by many random variables with a limited number of samples, or even one sample in most cases, for each random variable. This situation has not been well studied.

Fig. A1. Homogeneous area marked in pink for ATO0002EC79 sensor space SSA data, shown with the same RGB wavelengths as used in the main test. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.).

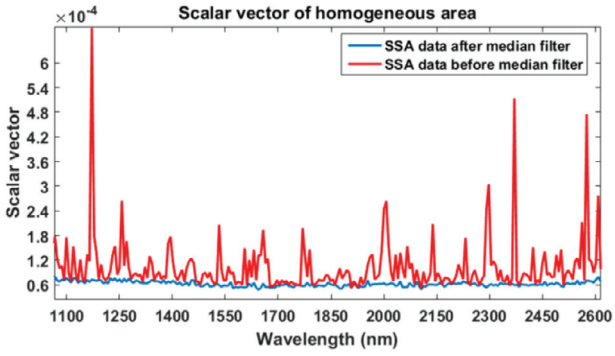

Fig. A2. Scalar vector of homogeneous area before and after median filter for data from the homogeneous area shown in Fig. A1. Note that the vector is essentially constant for all wavelengths after a median filter has been applied. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.).

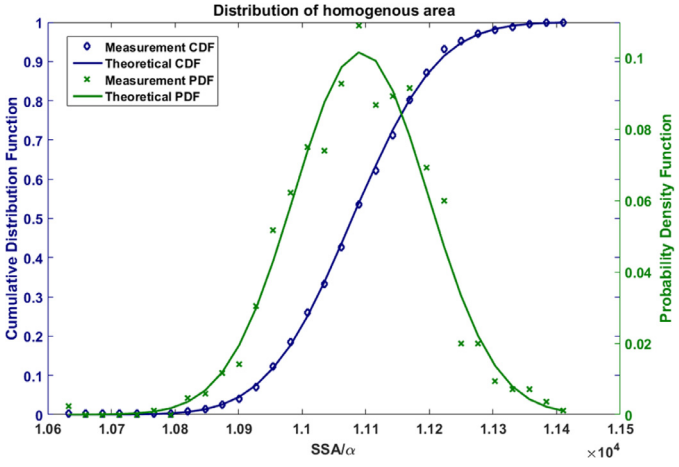

Fig. A3. Distribution of samples in the homogeneous area at wavelength 1.99 µm after scaling by the expected scale vector in Fig. A2. Cumulative Distribution Function (CDF) is shown in blue while Probability Density Function (PDF) is shown in green. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.).

The histogram of G-statistics can be used in goodness-of-fit testing. The null hypothesis that all random variables are Poisson corresponds to

$$H_{null} : m_i \sim Poisson(\lambda_i)\ i = 1, 2, 3, \cdots, N \tag{A2}$$

where m_i is the Poisson distributed random variable with mean λ_i and N is the total number of random variables. Hence, the probability mass function of m_i is $q_{\lambda_i}(m_i) = \lambda_i^{m_i} e^{-\lambda_i}/(m_i!)$. Harremoës and Tusnády (2012) show that for a Poisson distribution, the test G-statistic

$$\begin{aligned} G^2(m_i) &= -2 \ln q_{\lambda_i}(m_i) \\ &= 2(\lambda_i + \ln m_i! - m_i \ln \lambda_i) \\ &\overset{m_i \to \infty}{\to} 2\left(m_i \ln \frac{m_i}{\lambda_i} - m_i + \lambda_i\right) \\ &= 2I(m_i|\lambda_i) \end{aligned} \tag{A3}$$

is approximated by a χ^2-distribution with one degree of freedom. The limit $m_i \to \infty$ is valid because m_i is proportional to photon counting. Numerically, α is $O(10^{-4})$ implying m_i is $O(10^4)$, since SSA values are $O(1)$, which is sufficiently large. A second possible null hypothesis is that the data are Gaussian distributed random variables with means ρ_i and equal variances σ^2

$$H_{null} : m_i \sim N\left(\rho_i, \sigma^2\right)\ i = 1, 2, 3, \ldots, N. \tag{A4}$$

Here each random variable has a different mean but the same variance. It is important that the variances are equal because each random variable has a limited number of samples, making the variance difficult to accurately compute. An additional physical motivation is that the magnitude of the additive Gaussian noise throughout the scene should be similar. This is known as the basic assumption for additive white Gaussian noise (AWGN). The corresponding test G-statistic is

$$G^2(m_i) = \frac{(m_i - \rho_i)^2}{\sigma^2} \tag{A5}$$

which has a χ^2-distribution with one degree of freedom.

If G^2 is a χ^2 random variable, the probability that G^2 will exceed the observed value g^2 is

$$\Pr\left[G^2 > g^2\right] = 1 - F_{\chi^2}\left(g^2\right) \tag{A6}$$

where $F_{\chi^2}(\cdot)$ is the cumulative distribution of a χ^2 random variable with 1 degree of freedom. This probability is reported as the P-value of each measurement:

$$pv(m_i) = 1 - F_{\chi^2}\left[G^2(m_i)\right]. \tag{A7}$$

Given a number η between 0 and 1, if the P-value is less than η, the test is rejected, and if the P-value is larger than η, the test is accepted. As $N \to \infty$, η indicates the probability that the test is rejected if the null hypothesis is true (Devore, 2001):

$$\eta = \Pr\left[H_{null} \text{ is rejected} \middle| H_{null} \text{ is true}\right]. \tag{A8}$$

This probability is often reported as the type I error, or significance.

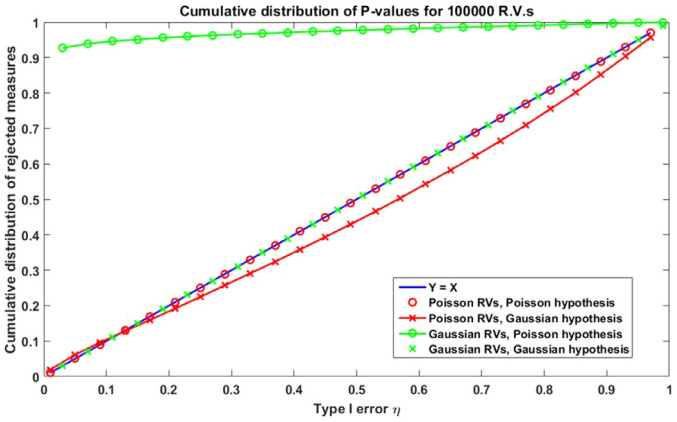

Fig. A4. Cumulative distribution of P-values for various cases (simulation). The case for Poisson random variables (RVs) is in red while Gaussian is in green. The hypothesis is likely to be true if the data fall close to the diagonal line (blue). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.).

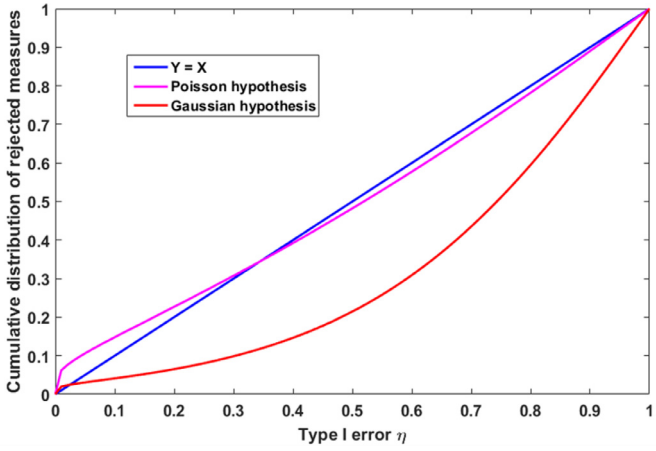

Fig. A5. Cumulative distribution of P-values for Poisson and Gaussian tests. The Poisson hypothesis (magenta) is more accurate than the Gaussian hypothesis with the same variance (red) for the whole scene. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.).

To demonstrate this concept, Fig. A4 shows the results of a simulation of 100,000 random variables with mean values ranging over three orders of magnitude. When generating Poisson random variables, 100,000 means are first uniform-randomly generated and then 100,000 Poisson samples are generated from the corresponding mean (one sample for one mean). Means are generated in the same way for Gaussian random variables, while variance is chosen asthe expectation of 100,000 means. A hypothesis that accurately describes the random variables should yield cumulative distributions that approximate a diagonal line through the plot. Tests in which the hypothesized distribution is false yield curves that are not along the diagonal. Assuming a Poisson distribution when the variables are actually Gaussian distributed is worse than assuming the variables are Gaussian distributed when they are actually Poisson distributed. One possible explanation is that when the mean becomes sufficiently large, the Poisson distribution can be approximated by a Gaussian distribution, but this is not valid in the opposite sense.

In Fig. A5 we apply this method to the scaled, median filtered ATO0002EC79 IR SSA data cube $d(\lambda, x, y)/\alpha$ for each of the two hypotheses: 1) the data are Poisson distributed; 2) the data are Gaussian distributed with the same variance for all pixels. We use SSA estimates derived from our MLM method for the means assuming a Poisson distribution. For the Gaussian hypothesis an altered MLM algorithm that minimizes square error is implemented, with themedian filtered SSA data cube again as input. The Poisson null hypothesis yields a cumulative distribution much closer to the diagonal than the Gaussian null hypothesis with the same variance for all pixels. The remarkable linearity of the Poisson cumulative distribution indicates the filtered CRISM data distribution is well approximated by a Poisson distribution.

References

Achilles, C.N., et al., 2016. Mineralogy of eolian sands at Gale Crater. Lunar Planet. Sci. XLVII, 2532 (abstract).

Arvidson, R.E., et al., 2014. Ancient aqueous environments at Endeavour Crater, Mars. Science 343 (6169), 1248097. doi:10.1126/science.1248097.

Arvidson, R.E., et al., 2016. High concentrations of manganese and sulfur in deposits on murray ridge, Endeavour Crater, Mars. Am. Mineral. 101 (6), 1389–1405. doi:10.2138/am-2016-5599.

Bandfield, J.L., Hamilton, V.E., Christensen, P.R., 2000. A global view of martian surface compositions from MGS-TES. Science 287, 1626–1630.

Bell III, J.F., et al., 2004a. Pancam multispectral imaging results from the spirit rover at gusev Crater. Science 305, 800–806.

Bell III, J.F., et al., 2004b. Pancam multispectral imaging results from the opportunity rover at meridiani planum. Science 306, 1703–1709.

Bibring, J.-P., et al., 2005. Mars surface diversity as revealed by the OMEGA/Mars express observations. Science 307, 1576–1581.

Bibring, J.-P., et al., 2006. Global mineralogical and aqueous Mars history derived from OMEGA/Mars express data. Science 312, 400–404. doi:10.1126/science.1122659.

Bishop, J.L., et al., 2008. Phyllosilicate diversity and past aqueous activity revealed at mawrth Vallis, Mars. Science 321, 830–833. doi:10.1126/science.

Bouman, C., Sauer, K., 1993. A generalized gaussian image model for edge-preserving map estimation. IEEE Trans. Image Process. 2, 296–310.

Carter, J., Poulet, F., Murchie, S., Bibring, J.P., 2013. Automated processing of planetary hyperspectral datasets for the extraction of weak mineral signatures and applications to CRISM observations of hydrated silicates on Mars. Planet. Space Sci. 76, 53–67.

Christensen, P.R., Ruff, S.W., 2004. Formation of the hematite-bearing unit in meridiani Planum: evidence for deposition in standing water. J. Geophys. Res. 109, E8003. doi:10.1029/2003JE002233.

Csiszár, I., 1991. Why least squares and maximum entropy? An axiomatic approach to inference for linear inverse problems. Ann. Statist. 19, 2032–2066. doi:10.1214/aos/1176348385.

D'Agostino, R., Pearson, E.S., 1973. Tests for departure from normality. Empirical results for the distributions of B2 and $\sqrt{}$ B1. Biometrika 60 (3), 613–622.

Dempster, A.P., Laird, N.M., Rubin, D.B., 1977. Maximum likelihood from incomplete data via the EM algorithm. J. R. Statist. Soc. Series B (methodological) 1–38.

De Pierro, A.R., 1994. A modified expectation maximization algorithm for penalized likelihood estimation in emission tomography. IEEE Trans. Med. Imag. 14, 132–137.

Devore, M.D, 2001. Recognition Performance from Synthetic Aperture Radar Imagery Subject to System Resource Constraints Thesis. Washington University, St. Louis, Missouri.

Ehlmann, B.L., Edwards, C.S., 2014. Mineralogy of the martian surface. Annu. Rev. Earth. Planet Sci. 42, 291–315.

Ehlmann, B.L., et al., 2016. Chemistry and mineralogy in-situ at the bagnold sand dunes: evidence for aeolian sorting and size-dependence in sand composition. Lunar Planet. Sci. XLVII, 1536 (abstract).

Eliason, E.M., McEwen, A.S., 1990. Adaptive box filters for removal of random noise from digital images. Photogramm. Eng. Remote Sens. 56, 453–458.

Fox, V.K., et al., 2016. Smectite deposits in marathon Valley, Endeavour Crater, Mars, identified using CRISM hyperspectral data. Geophysical Res. Lett. 43, 4885–4892. doi:10.1002/2016GL069108.

Fraeman, A.A., et al., 2012. Increasing the spatial resolution of oversampled CRISM images at Gale Crater. Lunar Planet. Sci. XLIII, 2123 (abstract).

Fraeman, A.A., et al., 2013. A hematite-bearing layer in Gale Crater, Mars: mapping and implications for past aqueous conditions. Geology 41, 1103–1106.

Gendrin, A., et al., 2005. Sulfates in martian layered terrains: the OMEGA/Mars express view. Science 307, 1587–1591.

Green, P.J., 1990. Bayesian reconstructions from emission tomography data using a modified em algorithm. IEEE Trans. Med. Imag. 9, 84–93.

Grotzinger, J.P., et al., 2015. Deposition, exhumation, and paleoclimate of an ancient lake deposit, Gale Crater, Mars. Science 350. doi:10.1126/science.aac7575.

Hapke, B., 2012. Theory of Reflectance and Emittance Spectroscopy, 2nd ed. Cambridge University Press Cambridge Books Online.

Harremoës, P., Tusnády, G., 2012. Information divergence is more χ 2-distributed than the χ 2-statistics. In: 2012 IEEE International Symposium on Information Theory Proceedings (ISIT), pp. 533–537. doi:10.1109/ISIT.2012.6284247.

Johnson, J.R., et al., 2006a. Spectrophotometric properties of materials observed by pancam on the mars exploration Rovers: 1. Spirit. J. Geophys.-Res. 111. doi:10.1029/2005JE002494, E2, E02S14.

Johnson, J.R., et al., 2006b. Spectrophotometric properties of materials observed by pancam on the mars exploration Rovers: 2. Opportunity. J. Geophys. Res. 111, E12S16. doi:10.1029/2006JE002762.
Karush, W., 1939. Minima of Functions of Several Variables with Inequalities as Side Constraints. M.Sc. Dissertation, Dept. of Mathematics, Univ. of Chicago, Chicago, Illinois.
Kuhn, H.W., Tucker, A.W., 1951. In: Nonlinear Programming. Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability, 1950. University of California Press, Berkeley and Los Angeles, pp. 481–492.
Lapôtre, M.G.A., et al., 2016. A quantitative assessment of aeolian fractionation at the bagnold dunes of Gale Crater, Mars, from orbit to ground, lunar planet. Science. XLVII, 1513 (abstract).
Lucy, L.B., 1974. An iterative technique for the rectification of observed distributions. Astron. J. 79 (6), 745–754. doi:10.1086/111605.
Massey Jr., F.J., 1951. The kolmogorov-smirnov test for goodness of fit. J. Am. Statist. Assoc. 46 (253), 68–78.
McCord, T.B., Clark, R.N., Singer, R.B., 1982. Mars: Near-infrared spectral reflectance of surface regions and compositional implications. J. Geophys. Res. 87 (B4), 3021–3032. doi:10.1029/JB087iB04p03021.
McEwen, A.S., et al., 2007. Mars reconnaissance orbiter's high resolution imaging science experiment (HiRISE). J. Geophys. Res. (Planets) 112, 5.
Minson, S.E., Simons, M., Beck, J.L., 2013. Bayesian inversion for finite fault earthquake source models I—Theory and algorithm. Geophys. J. Int. 194 (3), 1701–1726.
Mittlefehldt, D., et al., 2016. Alumina+silica±germanium alteration in smectite-bearing marathon valley, Endeavour Crater rim, Mars, Lunar planet. Science XLVII (2086) (abstract).
Murchie, S., et al., 2007. Compact reconnaissance imaging spectrometer for Mars (CRISM) on Mars reconnaissance orbiter (MRO). J. Geophys. Res. (Planets) 112, 5.
Murchie, S.L., et al., 2009. A synthesis of martian aqueous mineralogy after one mars year of observations from the Mars reconnaissance orbiter. J. Geophys. Res. 114, E00D06. doi:10.1029/2009JE003342.
Mustard, J.F., et al., 2005. Olivine and pyroxene diversity in the crust of Mars. Science 307, 1594–1597.
Noe Dobrea, E.Z., Wray, J.J., Calef, F.J., Parker, T.J., Murchie, S.L., 2012. Hydrated minerals on Endeavour Crater's rim and interior, and surrounding plains: new insights from CRISM data. Geophys. Res. Lett. 39, L23201. doi:10.1029/2012GL053180.
Poulet, F., et al., 2008. Mineralogy of terra meridiani and western arabia terra from OMEGA/MEx and implications for their formation. Icarus 165, 277–300.
Poulet, F., et al., 2009. Quantitative compositional analysis of martian mafic regions using the MEx/OMEGA reflectance data. 2. Petrological implications. Icarus 201, 84–101. doi:10.1016/j.icarus.2008.12.042.
Richardson, W.H., 1972. Bayesian-based iterative method of image restoration. JOSA 62, 55–59. doi:10.1364/JOSA.62.000055.
Shepp, L.A., Vardi, Y., 1982. Maximum likelihood reconstruction for emission tomography. IEEE Trans. Med. Imag. 1 (2), 113–122.
Smith, M.D., Wolff, M.J., Clancy, R.T., Murchie, S.L., 2009. Compact reconnaissance imaging spectrometer observations of water vapor and carbon monoxide. J. Geophys. Res. 114, E00D03. doi:10.1029/2008JE003288.
Snyder, D.L., Miller, M.I., 1991. Random Point Processes in Time and Space, 2nd Ed. Springer-Verlag, New York.
Snyder, D.L., Schulz, T.J., O'Sullivan, J.A., 1992. Deblurring subject to nonnegativity constraints. IEEE Trans. Signal Process. 40, 1143–1150.
Squyres, S.W., et al., 2012. Ancient impact and aqueous processes at Endeavour Crater, Mars. Science 336, 570–576.
Stamnes, K., Tsay, S.-C., Wiscombe, W., Jayaweera, K., 1988. Numerically stable algorithm for discrete-ordinate-method radiative transfer in multiple scattering and emitting layered media. Appl. Opt. 27, 2502–2509.
Wolff, M.J., et al., 2009. Wavelength dependence of dust aerosol single scattering albedo as observed by the compact reconnaissance imaging spectrometer. J. Geophys. Res. 114, E00D04. doi:10.1029/2009JE003350.
Wray, J.J., et al., 2009. Phyllosilicates and sulfates at Endeavour Crater, meridiani planum, Mars. Geophys Res Lett 36. doi:10.1029/2009GL040734.