AGU PUBLICATIONS

Earth and Space Science

RESEARCH ARTICLE
10.1002/2016EA000219

Key Points:
- The MSL Mastcam instruments have undergone a robust and detailed preflight and in-flight calibration process
- The calibration of the Mastcam instruments results in quantitative estimates of radiance and radiance factor in calibrated images
- The robust and detailed calibration of the Mastcam instruments enables quantitative multispectral and stereoscopic analyses on Mars

Correspondence to:
J. F. Bell III,
jim.bell@asu.edu,

Citation:
Bell, J. F., III et al. (2017), The Mars Science Laboratory *Curiosity* rover Mastcam instruments: Preflight and in-flight calibration, validation, and data archiving, *Earth and Space Science*, *4*, 396–452, doi:10.1002/2016EA000219.

Received 23 SEP 2016
Accepted 8 JUN 2017
Accepted article online 28 JUN 2017
Published online 30 JUL 2017

The Mars Science Laboratory *Curiosity* rover Mastcam instruments: Preflight and in-flight calibration, validation, and data archiving

J. F. Bell III[1], A. Godber[1], S. McNair[2], M. A. Caplinger[2], J. N. Maki[3], M. T. Lemmon[4], J. Van Beek[2], M. C. Malin[2], D. Wellington[1], K. M. Kinch[5], M. B. Madsen[5], C. Hardgrove[1], M. A. Ravine[2], E. Jensen[2], D. Harker[2], R. B. Anderson[6], K. E. Herkenhoff[6], R. V. Morris[7], E. Cisneros[1], and R. G. Deen[3]

[1]School of Earth and Space Exploration, Arizona State University, Tempe, Arizona, USA, [2]Malin Space Science Systems, Inc., San Diego, California, USA, [3]Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California, USA, [4]Texas A&M University, College Station, Texas, USA, [5]Niels Bohr Institute, University of Copenhagen, Copenhagen, Denmark, [6]USGS Astrogeology Science Center, Flagstaff, Arizona, USA, [7]NASA/Johnson Space Center, Houston, Texas, USA

Abstract The NASA *Curiosity* rover Mast Camera (Mastcam) system is a pair of fixed-focal length, multispectral, color CCD imagers mounted ~2 m above the surface on the rover's remote sensing mast, along with associated electronics and an onboard calibration target. The left Mastcam (M-34) has a 34 mm focal length, an instantaneous field of view (IFOV) of 0.22 mrad, and a FOV of 20° × 15° over the full 1648 × 1200 pixel span of its Kodak KAI-2020 CCD. The right Mastcam (M-100) has a 100 mm focal length, an IFOV of 0.074 mrad, and a FOV of 6.8° × 5.1° using the same detector. The cameras are separated by 24.2 cm on the mast, allowing stereo images to be obtained at the resolution of the M-34 camera. Each camera has an eight-position filter wheel, enabling it to take Bayer pattern red, green, and blue (RGB) "true color" images, multispectral images in nine additional bands spanning ~400–1100 nm, and images of the Sun in two colors through neutral density-coated filters. An associated Digital Electronics Assembly provides command and data interfaces to the rover, 8 Gb of image storage per camera, 11 bit to 8 bit companding, JPEG compression, and acquisition of high-definition video. Here we describe the preflight and in-flight calibration of Mastcam images, the ways that they are being archived in the NASA Planetary Data System, and the ways that calibration refinements are being developed as the investigation progresses on Mars. We also provide some examples of data sets and analyses that help to validate the accuracy and precision of the calibration.

Plain Language Summary We describe the calibration and archiving of the images being obtained from the Mastcam multispectral, stereoscopic imaging system on board the NASA Curiosity Mars rover. Calibration is critical to detailed scientific analysis of instrumental data, and in this paper we not only describe the details of the calibration process and the steps in our resulting data calibration pipeline but also present some examples of the kinds of scientific analyses and discoveries that this calibration has enabled.

1. Introduction

The Mast Camera (Mastcam) instrument on the NASA Mars Science Laboratory (MSL) rover *Curiosity* consists of a pair of focusable digital CCD cameras (detectors, optics, and filter wheels) that can acquire multispectral (400–1000 nm), stereoscopic images of the Martian surface and atmosphere at two specific fixed focal lengths. An externally mounted calibration target enables the relative reflectance calibration of the images, and two electronics boards in the rover body enable data processing and transmission of images to the rover's central computer. The cameras are mounted atop a 2 m tall mast that enables them to be rotated 360° in azimuth and ±90° in elevation.

The primary objective of the Mastcam investigation is to characterize and determine details of the history and processes recorded in geologic material at the MSL landing site within Gale crater, particularly as they pertain to habitability. Gale is an ~154 km diameter ancient impact crater centered near 5.4°S, 137.8°E, along the topographic dichotomy between the heavily cratered southern highlands and the lower, younger northern plains. A major goal of the mission is to explore evidence for past habitability within Aeolis Mons, the ~5 km tall central mound of layered sedimentary rocks (informally called "Mount Sharp") within Gale.

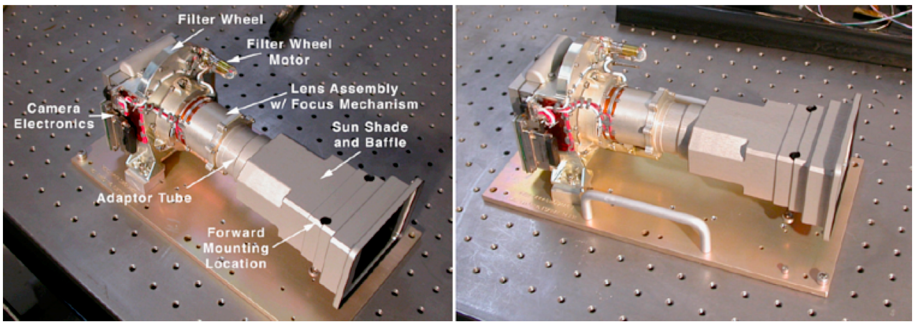

Figure 1. The Mastcam (left) M-34 camera head and (right) M-100 camera head. For scale, spacing of the optical table's bolt hole pattern is 1" (2.54 cm).

Mastcam instrument characteristics and observational strategies have been designed to address the following six specific objectives: (1) observe landscape physiography and processes; (2) examine the properties of rocks; (3) study the properties of fines; (4) view frost, ice, and related processes; (5) document atmospheric and meteorologic events and processes; and (6) support/facilitate rover operations, analytical laboratory sampling, contact instrument science, and other MSL science. *Curiosity* was launched in November 2011 and landed in Gale on 6 August 2012. As of this writing, the rover has traversed more than 15 km from the landing site to the base of Mount Sharp over ~1500 Martian solar days (sols), and Mastcam images have played an important role in actively enabling the MSL team to make exciting discoveries about the habitability of ancient Mars [e.g., *Williams et al*., 2013; *Grotzinger et al*., 2014].

Here we describe the series of preflight component-level, stand-alone camera-level, and integrated rover-level tests and calibration activities and analyses that were performed with the Mastcams in order to enable raw Mars images to be geometrically and radiometrically calibrated after being downlinked to Earth. We also describe standard procedures and special tests/observations that have been performed with the cameras on Mars in order to validate the preflight calibrations, to monitor for potential changes in the calibrations over time, and to enable additional calibration of the data to relative reflectance, for more direct comparisons to laboratory reflectance spectra of rocks and minerals. More details about the Mastcam instrument, and the Mastcam science investigation, in general, can be found in *Malin et al*. [2010, 2017] and *Bell et al*. [2012, 2013], and more details about the general goals of the *Curiosity* mission, plus the goals of other payload instruments also carried by the rover, can be found in *Grotzinger et al*. [2012].

2. Brief Instrument Description

2.1. Cameras

Mastcam (Figure 1) [*Malin et al*., 2017] consists of two focusable color cameras mounted on the rover's Remote Sensing Mast (RSM). The two cameras have different focal lengths and different sets of narrowband science filters (Tables 1 and 2). The cameras use the Kodak KAI-2020 CCD (7.4 μm square pixel pitch, with a

Table 1. MSL/Mastcam Electronic and Operational Characteristics

Characteristic	M-34 (Left)	M-100 (Right)
Field of view (FOV)	20° × 15°	6.8° × 5.1°
Baseline stereo separation	24.5 cm	
Spatial scale	450 μm/pixel at 2 m, 22 cm/pixel at 1 km	150 μm/pixel at 2 m, 7.4 cm/pixel at 1 km
Angular instantaneous FOV	0.22 mrad/pixel	0.074 mrad/pixel
Focal length	34 mm	100 mm
f/number	8	10
Focus range	~0.5 m to infinity	~1.6 m to infinity
Number of spectral filters	Seven Plus RGB Bayer Pattern	Seven Plus RGB Bayer Pattern
CCD detector gain	16.0 e^-/DN	15.8 e^-/DN
CCD detector read noise	18.0 e^-	15.8 e^-
CCD detector full well	26150 e^-	23288 e^-
Linearity (to 90% full well)	$r^2 > 0.999$	$r^2 > 0.999$

Table 2. MSL/Mastcam Filter Wavelengths, Band Passes, and Other Parameters[a]

M-34 (Left)					M-100 (Right)				
Filter	λ_{eff} ± HWHM (nm)	F_{ref} (DN)[b]	F_{Sun} (W/m^2/nm)[c]	Precision (%)[d]	Filter	λ_{eff} ± HWHM (nm)	F_{ref} (DN)[b]	F_{Sun} (W/m^2/nm)[c]	Precision (%)[d]
L0[e]	590 ± 88	--	--	--	R0[e]	575 ± 90	--	--	--
L0R	640 ± 44	9343	0.8620	1.2	R0R	638 ± 44	5980	0.8653	1.9
L0G[f]	554 ± 38	10089	0.9771	0.3	R0G[f]	551 ± 39	6457	0.9802	1.5
L0B	495 ± 37	9802	1.0145	5.7	R0B	493 ± 38	6273	1.0175	2.5
L1	527 ± 7	1796	1.0064	4.3	R1	527 ± 7	1149	1.0064	3.7
L2	445 ± 10	2016	0.9381	51.0	R2	447 ± 10	1290	0.9886	24.5
L3	751 ± 10	1045	0.6654	0.3	R3	805 ± 10	454	0.5913	3.4
L4	676 ± 10	1635	0.7924	0.1	R4	908 ± 11	171	0.4630	0.4
L5	867 ± 10	364	0.5101	0.3	R5	937 ± 11	103	0.4402	0.5
L6	1012± 21	104	0.3770	1.0	R6	1013 ± 21	67	0.3769	1.0
L7	880 ± 10, ND5[g]	--	--	--	R7	440 ± 10, ND5[g]	--	--	--

[a]Shaded cells indicate filters where the narrowband response partially or completely blocks one or more of the Bayer RGB filter responses (see text for details).
[b]Reference DN level of a perfectly diffuse white sunlit surface at Mars perihelion distance (1.38 AU), with no atmospheric attenuation, at zero incidence angle, at an exposure time of 10 ms, and observed through this filter. See section 5.2.7.
[c]Solar spectral irradiance at the top of the Martian atmosphere at perihelion (1.38 AU), in W/m^2/nm, convolved to each Mastcam filter band pass. Solar spectral irradiance data at 1 AU from *Colina et al.* [1996].
[d]Conservative estimate of the relative (filter to filter) precision for each filter. See section 3.2.6.
[e]Broadband near-IR cutoff filter used for Bayer filter RGB imaging. Data provided for reference only, as this cutoff filter can only be used in combination with the Bayer RGB filters.
[f]There are two green filters per 2 × 2 Bayer unit cell (Figure 1), with essentially identical characteristics.
[g]ND5 means 10^{-5} neutral density coating for solar imaging.

microlens on each pixel) to acquire images of up to 1648 × 1200 pixels (including dark columns) and are capable of relatively high frame rate acquisitions ("video" up to about four frames per second). The cameras acquire color via Bayer pattern filters on the CCD (Figure 2) and also have selectable science filters that image through the Bayer pattern filters (Figure 3). Characteristics of the Mastcam optics useful in the calibration and analysis of data products are described in Table 1.

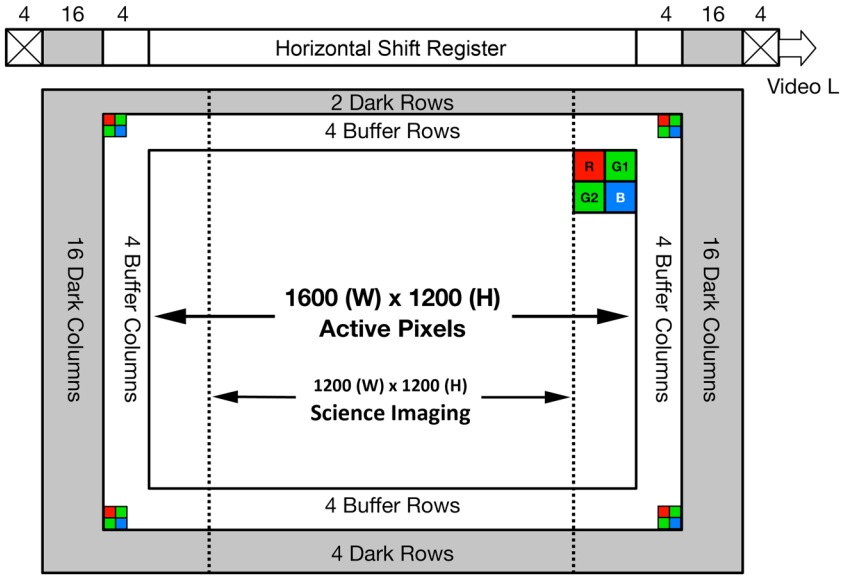

Figure 2. Simplified schematic of the Kodak KAI-2020 interline transfer CCD used in the MSL Mastcams, along with a graphical representation of each of the 2 × 2 Bayer RGB filter unit cells bonded directly onto the CCD's active pixels. Pixels are not depicted to scale, in order to display all of the relevant RGB unit cell, dark shielded, buffer, science imaging, and horizontal shift register regions described in the text. Pixel clocking is up then to the right in this representation of the image coordinate space, and so the first pixel clocked out is a Green 1. The figure is based on original Kodak vendor data available online at http://www.stargazing.net/david/QSI/KAI-2020LongSpec.pdf.

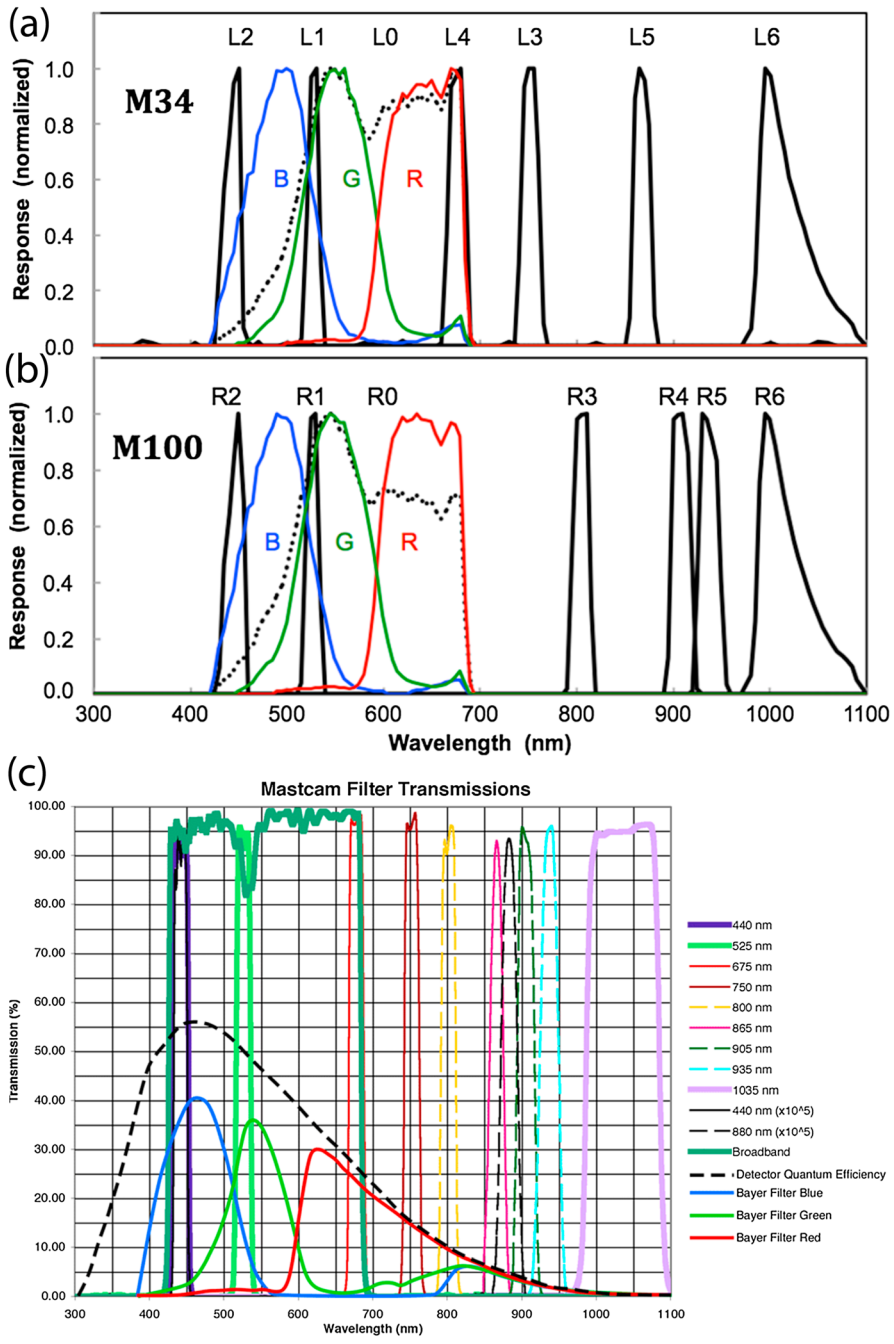

Figure 3. Normalized MSL/Mastcam system-level spectral response profiles for (a) the left eye M-34 camera and (b) the right eye M-100 camera. When the filter wheel is set to Filter 0 (broadband IR cutoff; dashed profiles shown here), the resulting Bayer filter RGB response profiles are as shown here in red, green, and blue. (c) When the filter wheel is set to a nonzero number, however, the response profiles of the Bayer filters are as shown here, based on component-level detector QE (%) and filter transmission data provided by the vendor.

Each Mastcam camera images through an eight-position filter wheel, actuated by an Aeroflex (now Cobham) stepper motor. One of the positions in each filter wheel (Table 2) is a broadband infrared cutoff filter for use with the Bayer broadband RGB color capability of the CCD, at filter position "0." Bayer RGB imaging across the entire sensor was included in order to provide a simple way to include color information as an additional component of assessing the general geologic context of the scene. Twelve of the 16 other filter positions provide multispectral imaging capability at nine unique additional wavelengths spanning ~400 to

Table 3. Types of Raw and Processed Images That Can Be Generated by the Mastcams

Type	Product	Form	Encoding
A	Image	Image	Raster 16 bit, uncompressed
B			Raster 8 bit, uncompressed but companded
C			Losslessly compressed (Huffman)
D			JPEG grayscale (luminance only)
E			JPEG 4:2:2 YCrCb chrominance subsampling
F			JPEG 4:4:4 YCrCb chrominance subsampling
G		Thumbnail	Raster 8 bit, uncompressed but companded
H			JPEG grayscale (luminance only)
I			JPEG 4:4:4 YCrCb chrominance subsampling
J	Video	Image	Raster 8 bit, uncompressed but companded
K			Losslessly compressed (Huffman)
L			JPEG grayscale (luminance only)
M			JPEG 4:2:2 YCrCb chrominance subsampling
N			JPEG 4:4:4 YCrCb chrominance subsampling
O		Thumbnail	Raster 8 bit, uncompressed but companded
P			JPEG grayscale (luminance only)
Q			JPEG 4:4:4 YCrCb chrominance subsampling
R	ZStack	Focus Merge Image	JPEG 4:4:4 YCrCb chrominance subsampling
S		Range map Image	JPEG grayscale (luminance only)
T		Focus Merge Thumbnail	JPEG 4:4:4 YCrCb chrominance subsampling
U		Range Map Thumbnail	JPEG grayscale (luminance only)

1100 nm and including three filters shared by each camera (Table 2). Most multispectral images are acquired using the central 1200 × 1200 pixel "science imaging" area of the CCD (Figure 2). Additionally, two filters (one on each camera) with neutral density coatings provide direct solar imaging capability in two colors. The spectral bandwidths described in Table 2 are based on the system-level monochromator measurements described in section 3.2.4 below.

The left "eye" of the Mastcam uses an eight-element, focusable, 34 mm effective focal length (EFL) lens that provides images with an angular instantaneous field of view (IFOV) of 218 μrad/pixel. The right eye uses a nine-element focusable 100 mm EFL lens that provides an IFOV of 74 μrad/pixel. Both cameras have an outer sapphire window that provides protection from dust. Using internal lens triplets that can be actuated over a small range of motion (~ ±1 mm, also using Aeroflex stepper motors), the M-34 and M-100 cameras can acquire in-focus images from a range of ~0.4 and ~1.6 m from the sapphire window to infinity, respectively. Additional details about the Mastcam optics design, fabrication, and performance can be found in *Ghaemi* [2009].

Each camera also relies on a Digital Electronics Assembly (DEA) which contains DC-DC converters, DRAM memory for temporary image storage, and flash memory (8 GB/camera) for nonvolatile storage, communications interfaces, and a Field Programmable Gate Array (FPGA) with an embedded processor that runs the flight software which implements Mastcam commands. The DEA provides power and data interfaces to the camera's CCD image sensor and its support circuitry and includes motor drivers for the optomechanical system (focus and filter wheel motion).

The Mastcam DEAs provide a high level of flexibility for operational acquisition and return to Earth of raw, compressed, or small "thumbnail" images. Compression can be performed within the DEA to process data using both lossless predictive difference and lossy discrete cosine transform (JPEG) compression, the latter using either 4:2:2 or 4:4:4 chroma subsampling or grayscale sampling of luminance only [e.g., *Malin et al*., 2013]. The DEA's large flash memories permit each camera to acquire and store many images and internal flight software and hardware permit images to be processed both into and out of the flash memory. Although generalized downsampling, sometimes also called subsampling or pixel summing, cannot be performed (e.g., the Mastcams cannot acquire summed images, which would be lower resolution but smaller in size), small thumbnails that are 1/8 the linear dimensions and 1/64 the data volume can be generated from the initial, full-sized images. The various types of "raw" and processed images that can be generated by the Mastcams are described in Table 3. For additional details on the processing of full size and thumbnail JPEG images as well as factors concerning their operational use, see *Malin et al*. [2013, 2017].

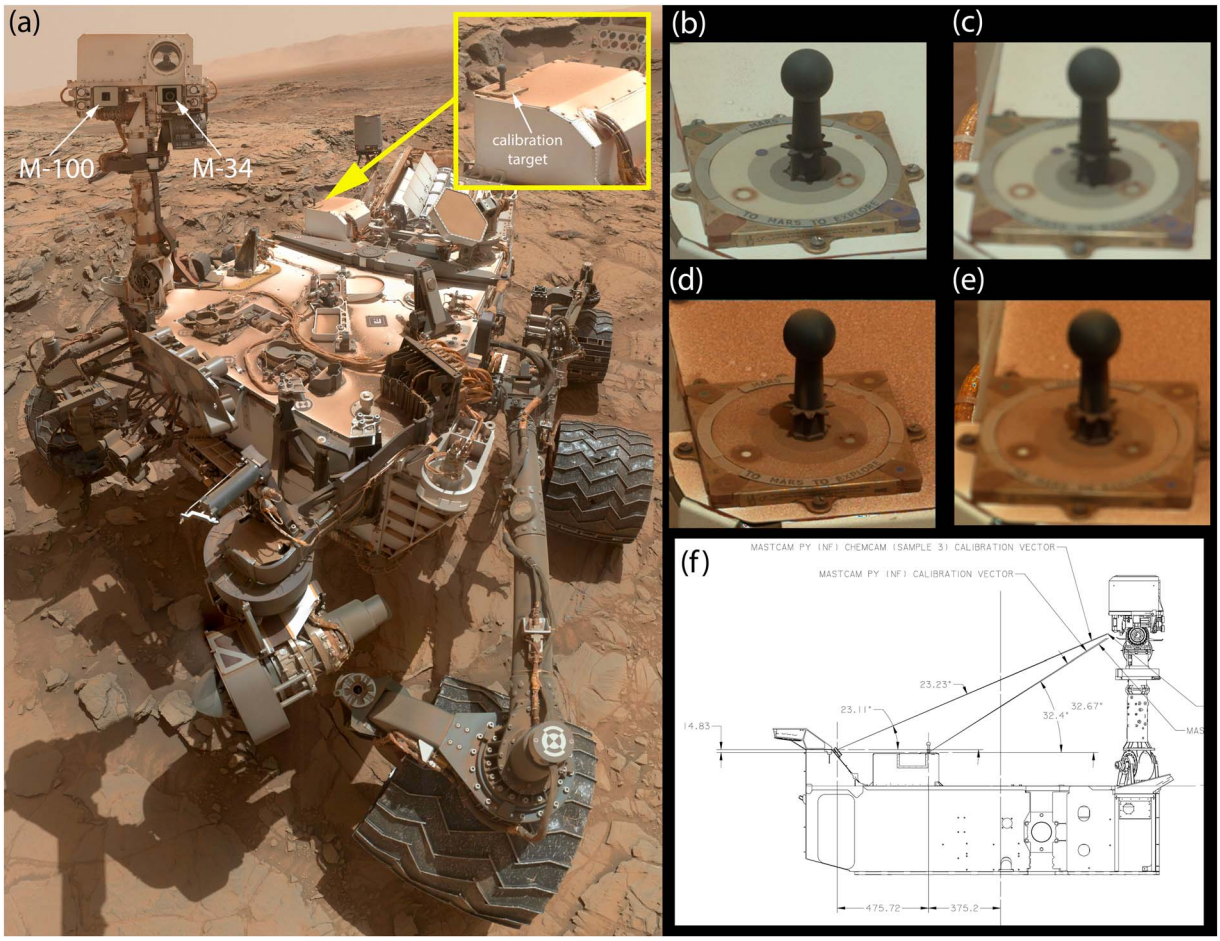

Figure 4. (a) Locations of the M-34 and M-100 Mastcam cameras as well as the Mastcam calibration target (inset) within the sol 1126 MAHLI "selfie" taken at the "Big Sky" drilling site within Gale crater. (b and c) The M-34 and M-100 views, respectively, of the calibration target on sol 69, when they were generally free of air fall dust. Dust deflected by the sweep magnets embedded below the calibration surface (see section 2.2) has been captured into rings in the corners and on the white and gray patches of the target. (d and e) The M-34 and M-100 views, respectively, of the calibration target on sol 807, when they had become appreciably dusty (like much of the rest of the rover deck in Figure 4a). (f) The elevation angle from which the calibration target is viewed relative to the Mastcams (≈32.4° up from horizontal). For reference, the ChemCam calibration target [*Wiens et al*., 2012] (visible in the upper right of the inset in Figure 4a, behind the Mastcam calibration target) is at an angle of ≈23.1° up from horizontal, when viewed by the Mastcams.

2.2. Mastcam Calibration Target

Mastcam includes an 8 × 8 cm wide, 6 cm tall, external color calibration target, mounted on the top of the Rover Pyro Firing Assembly (RPFA) box on the rover's deck, designed to enable imaging of a reference set of color and gray scale materials of known spectral reflectances and photometric properties (Figure 4). The target provides a way to validate the instrument's radiometric calibration and a way to enable tactical-timescale calibration of the radiance images to relative reflectance or Lambert albedo. The shadow cast by a central post (gnomon) provides a way to assess the direct versus diffuse components of solar and sky irradiance incident on the target.

The target flown on MSL is a modified flight spare of the Pancam calibration target flown on the Mars Exploration Rovers (MER) *Spirit* and *Opportunity* [*Bell et al*., 2003, 2006]. One science-related modification to the Mastcam cal target was to embed six "sweep magnets" [*Madsen et al*., 2003; *Bertelsen et al*., 2004] approximately 1 mm under the surface of the color chips and two of the gray scale rings, to enable small parts of the target to remain relatively dust free. A second set of modifications included changes in the target's function as an education and public outreach sundial [e.g., *Bell and Sullivan*, 2004]. Specifically, the dial's motto, fabrication date, and side-panel message were all modified for MSL compared to MER (Figure 4). Other E/PO aspects of the target's design remain the same as for MER. Details of the reflectance and photometric properties of the target materials can be found in *Bell et al*. [2003, 2006].

2.3. Integration With the Rover

The Mastcam cameras are mounted on the RSM at a height of 1.97 m above the Martian surface. The left and right cameras are separated by 24.2 cm, and they are positioned symmetrically relative to an azimuth actuation axis that is located on the front right corner of the rover, 559 mm starboard of the vehicle's centerline. Elevation actuation of both cameras occurs along an axis that is located 64.6 mm below the camera boresights or 1.91 m above the surface.

When the cameras are pointed at the Mastcam calibration target (rover frame azimuth of 189.2° and 176.2° for the M-34 and M-100 cameras, respectively, with 0° azimuth corresponding to straight in front of the rover; rover frame elevation of −32.4° downward from horizontal), the base of the gnomon is ~1.2 m from the front sapphire windows of the Mastcam optics. This placement is too close to be in good focus for the M-100 camera, but the larger number of pixels on the target's calibration surfaces adequately compensates for the blurred nature of those images and does not compromise the use of the target as a relative calibration source for M-100 images.

3. Preflight Camera Testing and Calibration: Methods, Data Sets, and Results

3.1. Introduction and Philosophy

The primary goals of Mastcam preflight testing and calibration were to develop a detailed understanding of the performance of the cameras under a range of environmental conditions relevant to Mars, to validate preassembly predictions of instrument performance so that models could be constructed to interpolate or extrapolate expected performance on Mars to conditions where preflight testing was not possible, and to acquire data sufficient to enable the conversion of measured Data Number (DN) values on Mars to an accurate estimate of radiance on sensor, in physical units such as $W/m^2/nm/sr$.

To that end, a series of tests were run at the instrument level (stand-alone assembled 34 mm and 100 mm flight camera systems driven by Ground Support Equipment (GSE) designed to simulate their respective DEAs), as well as some tests at the system level (cameras mounted to the rover in their flight configuration), to characterize the cameras' radiometric and geometric properties. These tests included characterization of the electronic and noise properties of each CCD detector, the spectral throughput of the system, the level of pixel-to-pixel responsivity ("flatfield") variations, the absolute radiometric responsivity, focus performance, geometric performance, and image quality, and stray and scattered light assessment. Where appropriate, data were obtained through all of the wideband and narrowband multispectral filters in both cameras. Here we describe the detailed nature of those preflight tests and the results derived from them. Those results, along with additional results from tests conducted in flight (section 4), are then used as part of the Mastcam data reduction and calibration pipeline discussed in section 5.

Raw Mastcam Experiment Data Records (EDRs) and radiance-calibrated Reduced Data Records (RDRs) have been archived to the NASA Planetary Data System (PDS) since the beginning of the MSL mission in August 2012 (see http://pds-imaging.jpl.nasa.gov/data/msl and section 5.5). Those currently archived RDRs, created using the data sets and methods described here and in the PDS archive itself [*Caplinger*, 2013], should be considered to be "Version 1" Mastcam RDR data products, produced primarily for tactical decision-making and quick-look science analyses. Refinements to the Version 1 Mastcam calibration, many of which are described here, can in many cases improve the fidelity of the radiance calibration and/or allow the data to be further calibrated to greater accuracy to radiance factor or to estimated Lambert albedo. We are working to archive in the PDS all of these further-refined "Version 2" Mastcam RDR data products. However, the information, data, and algorithms presented here are intended to allow individual Mastcam data users to perform these and other higher-fidelity calibrations themselves starting from the original PDS-archived EDRs, should that be necessary for their particular scientific goals.

3.2. Preflight Instrument-Level Testing

3.2.1. CCD Characterization

We characterized the gain (electrons, e^-, per DN), read noise (e^-), and full well capacity (e^-) of the Mastcam signal chain by imaging a diffuse illuminated integrating sphere target and by subsequent analysis of photon transfer curves using the technique described by *Janesick et al.* [1987]. Like many CCD imaging systems, the Mastcams use correlated double sampling to measure the small difference in voltage between a

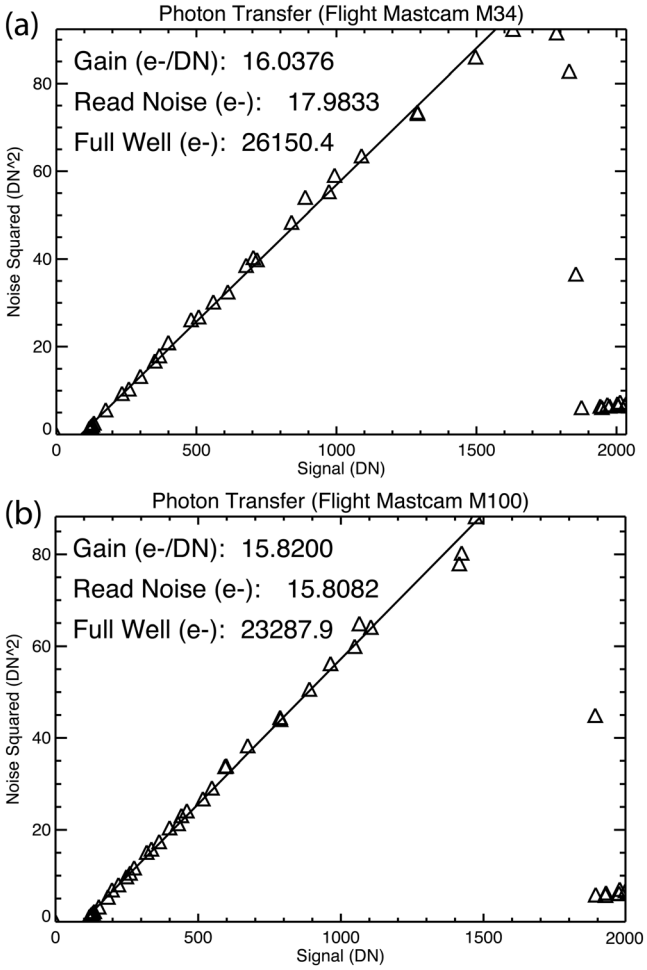

Figure 5. MSL/Mastcam (a) 34 mm and (b) 100 mm preflight CCD linearity, gain, and full well photon transfer curve data sets, with resulting CCD parameters derived using the methods of *Janesick et al.* [1987].

reference ("reset") and scene (video) signal. Reset and video signals are read separately as 12 bit data values which are then subtracted. The resulting difference is encoded as an 11 bit (DN = 0–2047) output signal. Uncompressed 11 bit image data were acquired at room temperature in a cleanroom laboratory setting on 30 July 2009. We obtained pairs of images at exposure durations of 0, 5, 10, 20, and 50 ms at sphere lamp currents of 0, 6, 7, and 8 A. The photon transfer curve is shown in Figure 5, and the results for gain, read noise, and full well are listed in Table 1. Figure 6 demonstrates the linear response of the Mastcam detectors, over the range from zero signal up to ~90% full well, based on images of a diffusely illuminated constant radiance integrating sphere target acquired at room temperature in a cleanroom laboratory setting on 17 August 2009. Operationally, manual and/or automatic exposure times (the latter using the autoexposure algorithm described in *Maki et al.* [2003], which is the same algorithm as used on the Mars Hand Lens Imager (MAHLI) camera [see *Edgett et al.*, 2012, section 7.3.2]) are selected to avoid signal levels above the known linear range. In postprocessing of Version 2 calibrated RDRs, raw EDR signal levels above ~1800 DN (240 DN when companded to 8 bit data; see section 5.3.2) are flagged as having a nonlinear response.

3.2.2. Bias and Dark Current Characterization

We characterized the rate of accumulation of dark current (charge generated in the CCD from thermal background effects) as a function of temperature during preflight testing. The Mastcam camera heads were designed to operate on Mars over an allowable flight temperature (AFT) range of −55°C to +50°C and the Mastcam DEAs over an AFT range of −40°C to +50°C. Figure 7 shows the actual range of flight temperatures experienced by the M-34 Mastcam during the first ~1000 sols of *Curiosity*'s mission on Mars. Maximum CCD temperatures experienced in flight are around +12°C and maximum DEA temperatures are around +40°C.

Thermal vacuum tests were conducted during preflight testing in July 2009 to characterize the bias signal (from the DEAs) and dark current signal (from the CCDs) over a range of approximately −60°C to +60°C. Figure 8 shows the preflight data that were collected and analyzed to create the following models for the background bias + dark current levels for each of the cameras:

$$M-34:\ \mathrm{DN_{back}} = \left[t_{\mathrm{exp}}\cdot 2.9e^{0.08T}\right] + 121.5 \quad (1)$$

and

$$M-100:\ \mathrm{DN_{back}} = \left[t_{\mathrm{exp}}\cdot 2.5e^{0.08T}\right] + 122.0 \quad (2)$$

where T is the reported or inferred focal plane array (CCD) temperature for each camera, in °C, and t_{exp} is the exposure time of the image, in seconds.

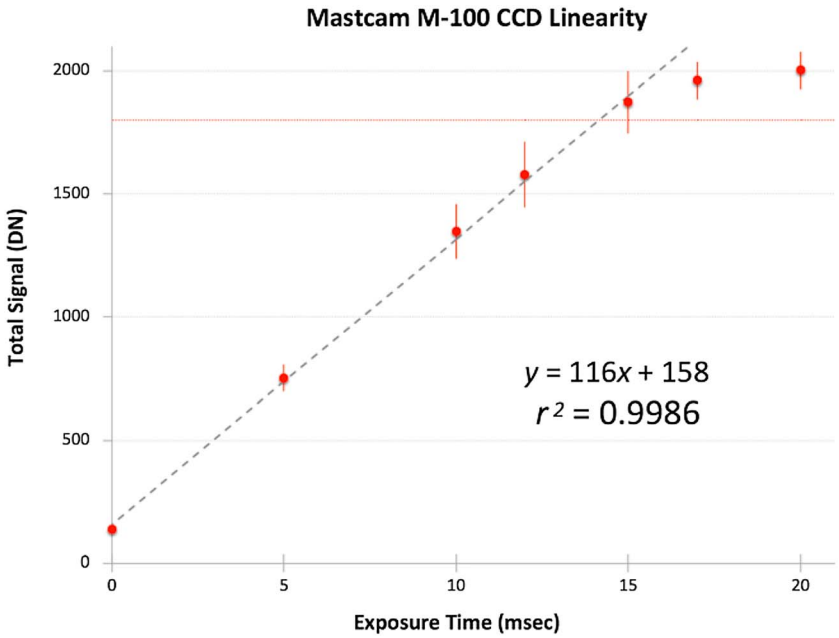

Figure 6. Example MSL/Mastcam CCD linearity test data results from integrating sphere data taken through the Bayer red filter, demonstrating linear performance of the M-100 detector from zero to ~90% full well. Performance is similar for the M-34 CCD. The r^2 value of the linear fit is for the first four data points only. In postprocessing, signal levels above 1800 DN (red horizontal line) are flagged as nonlinear.

The bias level alone, estimated from zero-second exposure images, exhibits only a slowly varying DC offset of about 120 ± 5 DN over the range of temperatures relevant to Mastcam operations (Figure 8). Inside the camera electronics, the value of this background signal is a commandable parameter and is subtracted from all pixels in the image during the 11 bit to 8 bit companding process. The commanded background

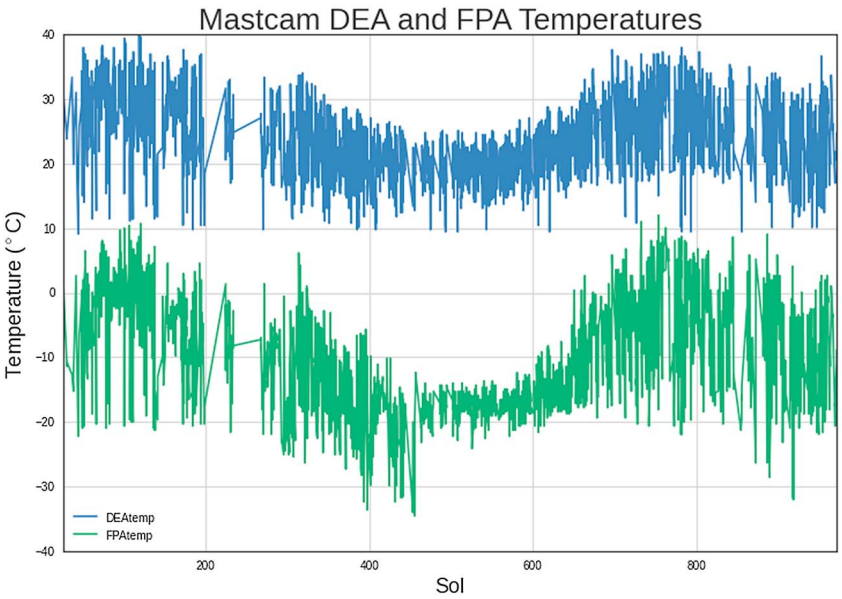

Figure 7. Actual range of flight temperatures experienced by the M-34 Mastcam Focal Plane Array (FPA; green lower data) and Digital Electronics Assembly (DEA; blue upper data) during the first ~1000 sols of *Curiosity*'s mission on Mars. Temperature telemetry for the M-100 Mastcam FPA and DEA is not downlinked as frequently, and so the behavior of the M-100 optics heater temperature sensor (HTR-1 in the PDS archived data labels) or sometimes the M-34 camera's FPA temperature is often used as a proxy for the temperature history of the M-100 (see section 4.2.1).

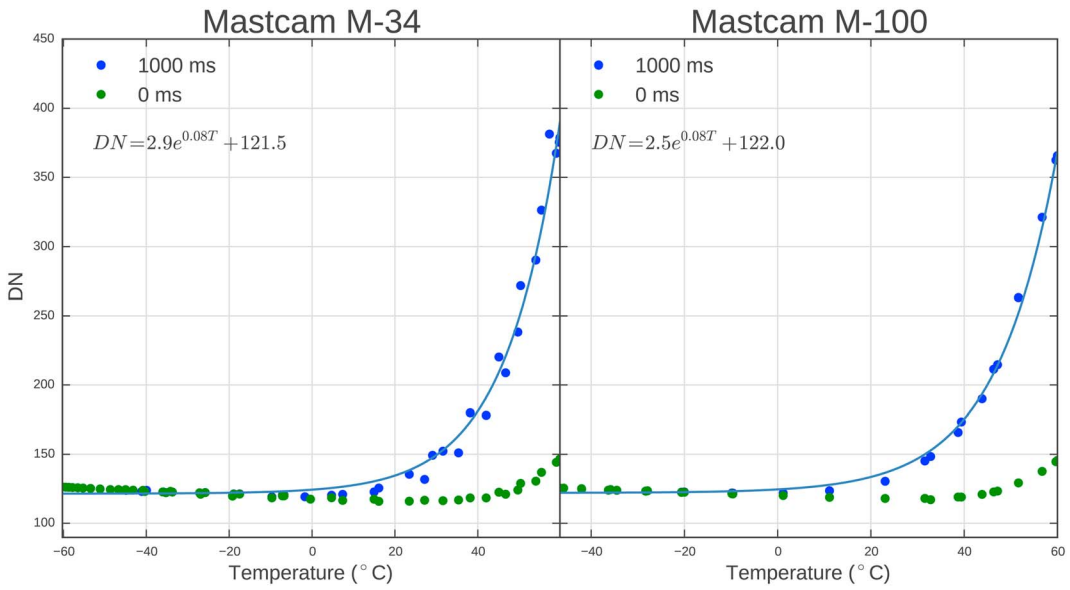

Figure 8. Zero-exposure bias signal (green data points) and 1000 ms dark current (blue data points) data for the (left) M-34 and (right) M-100 Mastcam CCDs, as a function of Focal Plane Array (FPA) temperature. Fits to the dark current data show the functional form of the Mastcam dark current model, in DN/s.

level is reported in the DARK_LEVEL_CORRECTION processing parameter keyword in the PDS archived image label for each Mastcam image and is usually set to 117 DN during flight operations on Mars.

Additionally, Figure 8 shows that dark current is not significant at the tested exposure time of 1000 ms below about 20°C. Since Mastcam FPA temperatures in flight have not exceeded +12°C, and exposure times rarely exceed 1000 ms (most exceptions being longer exposures specifically designed to characterize dark current in flight or for long and very cold nighttime exposures; see section 5.4.3), dark current accumulation can almost always effectively be ignored in processing and calibration of Mastcam images.

3.2.3. Electronic Shutter Smear

The Mastcam CCD sensor is a progressive scan, interline transfer device [e.g., *Howell*, 2000; *Truesense Imaging, Inc.*, 2012]. Charge from the photosensitive area of each pixel is transferred into an adjacent vertical (aligned along CCD columns) light-shielded shift register. Those charges are then clocked down, one CCD line (row) at a time, into a horizontal shift register, from which each pixel is then clocked out of the device horizontally for subsequent digitization. While the shift of charge from the photosensitive pixels into the vertical shift registers is essentially instantaneous (μs) and simultaneous for all pixels, clocking of the measured charges down into the horizontal register and out of the device takes significantly more time.

For the Mastcam CCDs, the system clock rate is typically set to 20 MHz, and the readout time for a full-frame 1648 × 1200 pixel image is about 420 ms. Mastcam does not have a mechanical shutter, and thus, during the readout a small amount of incident light leakage into the vertical shift registers can occur. However, the sensitivity of the vertical registers is between 5600 and 10,000 times (75–80 dB) lower than the sensitivity of each pixel's photosites [*Eastman Kodak*, 2009; *ON Semiconductor, Inc.*, 2014; 2015], and thus, the leaked signal is very small. Regardless, this light leakage still introduces an excess signal into the data, called electronic shutter smear because the excess charge is "smeared" down the array as it is being clocked out, producing dim vertical bands extending below bright features in the scene. The magnitude of this smear is proportional to the scene intensity in each pixel, and thus, the smear image mimics the scene image, but at a much fainter level and with a superimposed linear ramp of charge added, since photosensitive pixels farthest from the horizontal shift register (which is at the "bottom" of the array) accumulate charge during the readout longer than those at the top [e.g., *Bell et al.*, 2003]. A complication in cameras like the Mastcams is that the sensitivity of each pixel in each line varies because the Bayer filters impose an additional wavelength-dependent pattern onto each image (e.g., Figure 2; see also discussion in *Edgett et al.* [2015]).

The amount of smear signal in a given image is proportional to the ratio of the frame readout time to the exposure time. For example, for typical Mastcam exposure times of 10 ms and for scenes without signal levels over the sensor's full well, the worst-case electronic shutter smear level for pixels farthest from the horizontal register (which take the full 420 ms to read out) and with vertical shift registers 75 dB less sensitive than the pixel's photosites would be [(420 ms/10 ms)/5600] = 0.75% of the scene signal or at most about 15 DN for a signal near full well. Given that most Mastcam images are commanded to have average signal levels only near half to three fourths full well, electronic shutter smear is not expected to be a significant effect. The smear percentage increases, however, as the exposure time becomes shorter or if the scene contains very bright features (such as the Sun glinting off rover structure). In such instances it might be desirable to remove the smear signature. For a static scene, this can easily be done by computing a running sum of the signal level across each column and subtracting the appropriate fraction of it from each pixel.

It should be noted that shutter smear in the Mastcam interline transfer CCDs is different from the shutter smear observed in frame transfer CCD imagers like the MER/Pancam system. As described in *Bell et al*. [2003], frame transfer in the Pancams is a much simpler situation (and a much larger effect). In-flight testing and other aspects of electronic shutter smear assessment and potential removal for the Mastcams are discussed in sections 4.2.2 and 5.3.3 below.

3.2.4. System Spectral Throughput

Each Mastcam camera is equipped with an eight-position filter wheel, positioned close to each camera's CCD (Figure 1 and Table 2). Filter 0 (wideband near-IR cutoff filter) on each camera is used for direct Bayer RGB imaging, and Filter 7 is equipped with a neutral density coating that enables direct solar imaging in one of two colors (near-IR 880 nm for the M-34 and blue 440 nm for the M-100). During preflight calibration on 17–19 August 2009, the system-level throughput (CCD + Bayer filter + optics + filter wheel filter) was characterized for each filter wheel and Bayer filter combination by acquiring images of the output slit of a monochromator operated under ambient room temperature conditions at Malin Space Science Systems (MSSS) in San Diego. For both the M-100 and M-34 cameras, calibration runs consisted of monochromator sweeps from 300 to 1100 nm at 10 nm steps and subsequent 100 nm sweeps at 5 nm steps centered around each filter's effective center wavelength. For Filters L0 and R0, the 5 nm sweeps covered 400 nm of wavelength space from 400 to 800 nm. Monochromator test data were not acquired for Filters L7 and R7 because it was not possible to create a bright enough monochromatic source to detect through the neutral density five coatings on those solar filters. Instead, estimates for the system-level passbands of those filters were derived from component-level and/or vendor-supplied test data. Each output image contains a 1648×96 pixel array subsampled from near the center of the full array, and the image of the monochromator slit was consistently centered in the array for each calibration run.

Postprocessing of the data was conducted to characterize each filter's band pass, estimating the effective center wavelength and filter width from the peak-normalized average of the monochromator slit image signal as a function of wavelength. Effective center wavelength is defined as the wavelength where the total integral under the band-pass curve (defined as the area of the curve above the level of 1% of the peak maximum value) gets to half its total value. Filter width is defined as half the width of the band-pass curve at half of the normalized peak maximum value (half width at half maximum, HWHM). Data were extracted from sets of individual Red (R), Green1 (G1), Green2 (G2), and Blue (B) Bayer pixels (Figure 2) in each monochromator slit image. For Filters L0 and R0, the R, G1, G2, and B curves were analyzed separately, treating them effectively as separate filters (although parameters for the G1 and G2 pixels were statistically the same, and so the average parameters for the two green pixels are plotted and reported here). For Filters L5–L7 and R4–R6 (all of which have effective band centers $>$800 nm), the Bayer filter pattern is effectively transparent (Figure 3), and so all pixels were included in the calculation statistics. For filters where the narrowband filter partially or completely blocks one or more of the Bayer responses (shaded cells in Table 2), only one specific Bayer pixel was used for the statistics: R for Filters L3, L4, and R3; the average of G1 and G2 for Filters L1 and R1; and B for Filters L2, R2, and R7. A similar process is employed with Mastcam images in Filters L1–L4 and R1, R2, and R7 in lossy compressed images downlinked from Mars: kernels are applied within the flight software that completely discard pixels whose Bayer peak is outside the narrowband filter band pass, using only the same specific Bayer pixels listed above (see section 5.2.1). Derived normalized band-pass profiles for the Mastcam filters using these methods are plotted in Figure 3, and band-pass characteristics are listed in Table 2.

Monochromator scan images acquired through the narrowband blue Filters L2 and R2 (band centers near 445 nm) suffered from low signal-to-noise ratio (SNR) compared to the other filters because of the relatively weak signal level of the calibrated light source at that wavelength and thus the long integration times (and thus higher background levels) required. The acquired data for the narrowband blue filters also revealed a much higher level of relatively constant 500–1100 nm rejection band "leakage" for these filters—up to 21% and 9% of the in-band response levels for L2 and R2, respectively (compared to rejection band responses of 1–4% for the other filters). The low SNR of the L2 and R2 data, however, yield high uncertainties on those out-of-bank leakage levels. The end result, as discussed in section 3.2.6 below, is a much higher level of radiometric uncertainty for the L2 and R2 filters compared to the other filters.

We did not perform any corrections for the spectral radiance of the monochromator's illumination source, nor did we calculate an "expected correction" for the shape of the Sun's radiance spectrum, which is the illumination source on Mars. Given the relatively narrow passbands of most of the Mastcam filters, these spectral corrections could be expected to introduce only small (several percent) shifts in the effective band centers estimated here. Given the fact that most of the mineralogic spectral features reasonably expected to be encountered with Mastcam at Gale crater are broad, solid-state absorptions [e.g., *Bell et al*., 2003; *Rice et al*., 2010], such small potential shifts in estimated effective center wavelengths or bandwidths are not expected to have any significant effect on scientific analyses.

3.2.5. Pixel-to-Pixel Responsivity (Flatfield) Characterization

Variations in pixel-to-pixel responsivity in imaging systems can be characterized by acquiring images of so-called "flat" (uniformly illuminated) targets. Observed nonuniformities in the resulting images result from variations induced by optics, filters, dust, and/or intrinsic variations in the radiometric responses of individual pixels. Normalized flatfield images quantitatively characterize the combined magnitude of these nonuniformities and can be used to remove most of their effects on imaging data.

During preflight camera testing at MSSS on 31 July and 6 and 8 August 2009, both Mastcams imaged the uniformly illuminated exit port of a Spectralon-coated integrating sphere under ambient room temperature and pressure conditions. Images were acquired through all of the Mastcam filters except the L7 and R7 solar filters (which could not be adequately illuminated). Dark images were also obtained with the sphere exit port blocked, in order to characterize background signal levels in the sphere image data. To avoid any potential nonuniformity effects associated with the edges of the sphere's exit port, both cameras acquired 1648 × 512 pixel images at three elevation positions that enabled the bottom, middle, and top third of the field of view to be centered in the sphere's exit port. The images were later combined into single full-field 1600 × 1200 pixel flatfield images. Example flatfield image for Filter R0 in both cameras is shown in Figure 9.

Postprocessing of the flatfield calibration data included visual and quantitative inspection to assess data quality, subtraction of the background dark images, combination of the images from the top, middle, and bottom of the fields of view, de-Bayering the data to work with separate R, G1, G2, and B channel data and calculating image statistics from the central region of each image to normalize each final flatfield image to a mean value of 1.0. Maximum flatfield variations detected in the images were approximately ±10%, dominated by variations near the edges of the field where the circular Mastcam filters begin to vignette the CCD's rectangular field of view (Figure 9). Typical variations near the center of the fields of view, and especially within the 1200 × 1200 pixel region defined as the nominal science imaging area of each Mastcam, are less than ±1%. This is why multispectral imaging sequences, which require more robust quantitative calibration than geologically focused RGB imaging, are typically restricted to the science imaging area of the CCD.

Initial flatfield files currently being used for the rapid tactical calibration of images and initial PDS release versions of the data are archived with the PDS (see section 5.2.5). Further refined and processed versions of the normalized Mastcam flatfield images for both cameras and all filters except the solar ND filters (filter #7 in each camera) are being developed from in-flight calibration measurements (see sections 4.2.3 and 5.3.4).

3.2.6. Radiometric Responsivity

To produce spectral or true-color data products, we need to characterize how the camera, filter, and optics encode the scene radiance into the digital output for each possible camera and filter combination. The combination of these properties can be represented by a single coefficient, often known as the radiometric calibration coefficient (Table 4). Initial estimates of the radiometric calibration coefficients

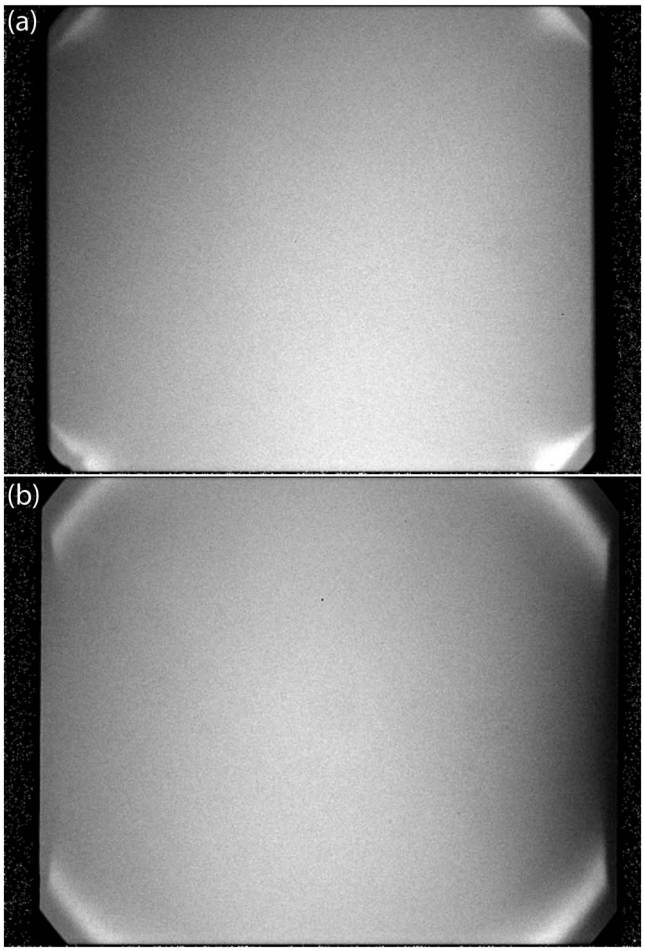

Figure 9. Example 1648 × 1200 pixel (a) M-34 and (b) M-100 Mastcam Filter 0 flatfield images. These images correspond to the de-Bayered G1 pixels from each camera. Nonuniformities noted in visual inspection include occasional dead pixels, blurring at left and right edges, brightening near corners, and vignetting in corners. Pixel (0, 0) is at the top left of each image. The problematic corner regions are not included in standard 1200 × 1200 pixel science imaging data acquired by Mastcam (see text). These normalized flatfield images are contrast enhanced so that black is ≤0.85 and white is ≥1.05.

[see *Caplinger*, 2013], which are being used currently in the tactical Version 1 calibration of images and initial PDS release versions of the data, were based on a component-level model of the CCD quantum efficiency (QE), filter band-pass response, and optics transmissivity, based on piece-part testing and vendor data. These initial estimates agree very well with the more detailed estimates based on Version 2 calibrated images obtained with the actual flight cameras and the data processing methods described here.

During the summer of 2009, we imaged the exit port of the same integrating sphere described in section 3.2.5, after acquiring a National Institute of Standards and Technology (NIST)-certified radiometric calibration of the sphere's radiance as a function of wavelength and input current. Using the sphere's 8 A setting, we acquired multiple images of the radiance from the sphere in both Mastcams, under ambient room temperature and pressure conditions. Each image was a 1648 × 512 pixel subframe taken from the center of each camera's field of view at various exposure settings.

Postprocessing included de-Bayering the input images into separate R, G1, G2, and B images; calculating the mean value from a relatively uniform region within the image of the sphere's exit port; subtracting the mean of a dark, unexposed region of each image in order to remove any background bias + dark current signal; dividing the resulting background-corrected DN values by the exposure time to arrive at an estimate of DN/s; and then dividing the DN/s by the estimated radiance of the integrating sphere in each filter, derived by convolving the NIST-certified sphere radiance with the system spectral response functions calculated in section 3.2.4 (Figure 10). The standard deviation of the mean DN values extracted from the sphere images was propagated through the calculations and combined with a conservative estimate of ±10% uncertainty on the sphere's calibrated radiance, based on NIST calibration of the monitoring diode. The final radiometric coefficients, in units of [(W/m^2/nm/sr)/(DN/s)], along with their estimated uncertainties, are listed in Table 4. In almost all cases, the ±10% uncertainty in the estimate of the sphere radiance dominates the radiometric uncertainty. However, for the narrowband blue (L2 and R2) filters, additional uncertainties in the characterization of their spectral throughput, discussed below, lead to much higher absolute calibration uncertainties than for the other Mastcam filters (Table 4).

An estimate of the relative filter-to-filter uncertainty (a measure of the precision of normalized "spectra" ultimately derived from Mastcam multispectral measurements) can be made by dividing out the assumed absolute radiance spectrum of the calibration lamp in Figure 10 and assessing the levels of the resulting filter-to-filter deviations. These relative filter precision estimates are also listed in Table 2. The close

Table 4. Refined MSL Mastcam Multiplicative Radiometric Calibration Coefficients

Filter Number	Effective Center Wavelength (nm)	Red	Green 1	Green 2	Blue
		[W/m^2/nm/sr)/(DN/s)] ± 1σ Uncertainty			
		M-34 (Left Mastcam)			
L0	590	3.56e−07 ± 3.6e−08	3.39e−07 ± 3.4e−08	3.39e−07 ± 3.4e−08	4.47e−07 ± 4.5e−08
L1	527	5.62e−05 ± 5.7e−06	2.38e−06 ± 2.4e−07	2.38e−06 ± 2.4e−07	8.69e−06 ± 8.7e−07
L2	445	1.60e−04 ± 1.1e−04	4.78e−05 ± 3.2e−05	4.54e−05 ± 3.0e−05	2.81e−06 ± 1.8e−06
L3	751	2.61e−06 ± 2.6e−07	8.96e−06 ± 9.0e−07	8.94e−06 ± 9.0e−07	2.86e−04 ± 3.0e−05
L4	676	2.04e−06 ± 2.0e−07	4.05e−05 ± 4.1e−06	4.28e−05 ± 4.3e−06	1.35e−04 ± 1.4e−05
L5	867	6.29e−06 ± 6.3e−07	6.51e−06 ± 6.5e−07	6.50e−06 ± 6.5e−07	6.51e−06 ± 6.5e−07
L6	1012	1.29e−05 ± 1.3e−06	1.28e−05 ± 1.3e−06	1.28e−05 ± 1.3e−06	1.29e−05 ± 1.3e−06
		M-100 (Right Mastcam)			
R0	575	6.36e−07 ± 6.4e−08	6.08e−07 ± 6.1e−08	6.07e−07 ± 6.1e−08	7.98e−07 ± 8.0e−08
R1	527	9.06e−05 ± 9.1e−06	3.83e−06 ± 3.8e−07	3.83e−06 ± 3.8e−07	1.37e−05 ± 1.4e−06
R2	447	2.23e−04 ± 8.7e−05	6.32e−05 ± 2.4e−05	6.23e−05 ± 2.4e−05	4.09e−06 ± 1.6e−06
R3	805	6.96e−06 ± 7.0e−07	9.99e−06 ± 1.0e−06	9.94e−06 ± 9.9e−07	1.14e−05 ± 1.1e−06
R4	908	1.36e−05 ± 1.4e−06	1.38e−05 ± 1.4e−06	1.38e−05 ± 1.4e−06	1.39e−05 ± 1.4e−06
R5	937	1.81e−05 ± 1.8e−06	1.80e−05 ± 1.8e−06	1.80e−05 ± 1.8e−06	1.82e−05 ± 1.8e−06
R6	1013	2.15e−05 ± 2.1e−06	2.14e−05 ± 2.1e−06	2.14e−05 ± 2.1e−06	2.14e−05 ± 2.1e−06

similarity of typical Mastcam Mars surface relative reflectance spectra derived from this calibration (e.g., section 5.4.4) to spectra acquired from previous Mars surface missions over the same wavelengths indicates that these relative precision estimates are likely to be overly conservative (especially for the narrowband blue Filters L2 and R2).

Because of schedule constraints, it was not possible to acquire data to estimate the Mastcam responsivity coefficients at more flight-like temperatures and pressures. However, we did perform limited responsivity tests of a sample of KAI-2020 CCD sensors from the same lot as the flight sensors, over temperatures down to −105°C. We did not notice any significant variations that would be consistent with a large change in QE over that temperature range, which is much wider than the range experienced by the CCDs on Mars (Figure 7). In addition, experiences with other Mars surface CCD imaging systems (e.g., MER/Pancam)

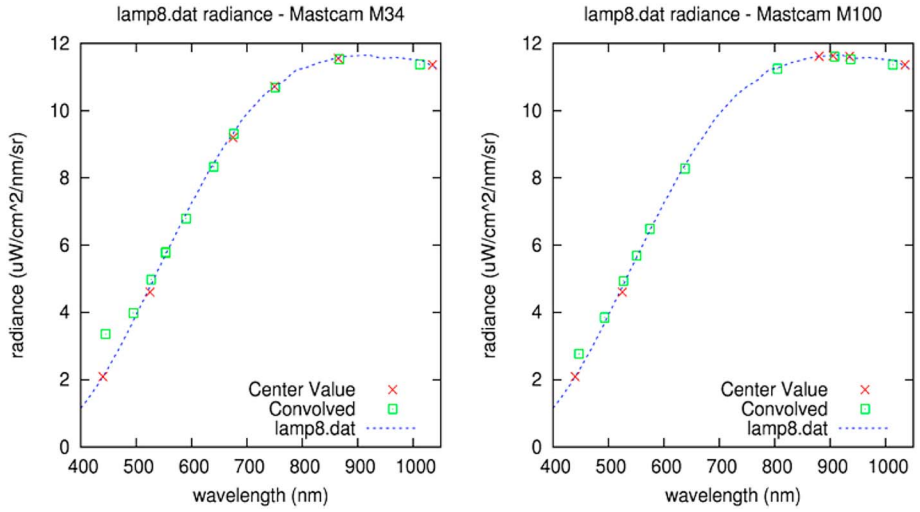

Figure 10. Estimated integrating sphere input radiances used for deriving (left) M-34 and (right) M-100 radiance calibration coefficients for each of the filters in each camera (Table 4). "Center value" estimates are used in the current tactical-timeline radiance calibration of Mastcam images. Refined estimates based on a full convolution of the Mastcam spectral response profiles derived in section 3.2.4 are shown as "Convolved" (weighted) data points. "lamp8.dat" refers to the NIST-calibrated spectral radiance of the integrating sphere. Significant deviations in the fit of the weighted L2 and R2 (445 and 447 nm) data points to the input lamp spectrum result in higher relative uncertainties in the radiance calibration coefficients for those filters.

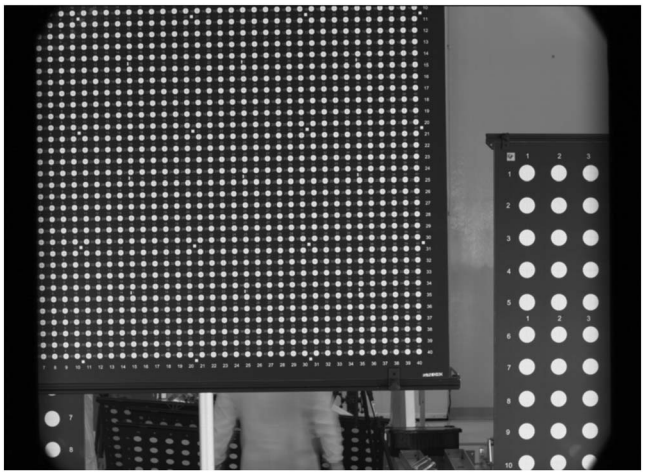

Figure 11. Example preflight Mastcam M-34 image of geometric calibration targets obtained during initial rover-level testing at JPL.

[*Bell et al*., 2003, 2006], as well as in-flight Mastcam radiance factor validation measurements (e.g., sections 5.2.7 and 5.4.4), show that the effects of varying QE over temperature are likely small and within the ~10% overall estimated uncertainty of the Mastcam radiance estimates. The radiometric responsivities of the G1 and G2 filters are the same within the uncertainties and can be averaged in practice to estimate an overall green responsivity. Filters with band passes common to both cameras also exhibit comparable radiometric response coefficients (within ±1% to ±7% or within the estimated ±10% typical absolute radiance uncertainties) after compensating for the cameras' difference in focal ratios $[(f/8)^2/(f/10)^2 = 0.64]$. As a further check on the relative accuracy of these derived coefficients, we also found that the Mastcam coefficients in Table 4 for near-IR wavelengths greater than 800 nm that are comparable to the wavelengths used for the MER/Pancam filters [*Bell et al*., 2003] are essentially the same as those derived for the Pancams at room temperature, after compensating for differences in the focal ratios and CCD quantum efficiencies among these imaging systems.

3.2.7. Preflight Focus Calibration

While the M-34 and M-100 Mastcams are each fixed focal length cameras, they are *focusable*, over a relatively wide range of focus distances (Table 1). Analysis of preflight Mastcam images of a geometric calibration target (e.g., Figure 11) was used to derive an initial calibration of the focus of each camera as a function of distance. Those calibrations (Table 5) [*Caplinger*, 2013] can be used to make an initial estimate of distances to targets of interest (and thus the sizes of those targets, using the known IFOV of each pixel; Table 1) and the FOCUS_POSITION_COUNT keyword in the PDS labels of archived Mastcam images. That initial calibration has been subsequently refined in flight to also include the effects of temperature on the focusing of the cameras, as described in section 4.2.5 below.

3.2.8. Geometric Performance and Resolution

We acquired a small set of images of a standard geometric calibration target (the "SVG Squares and Wedges pattern" target from *Imatest* [2015]; Figure 12) over several focus positions and under ambient laboratory conditions at MSSS using the M-34 and M-100 flight cameras on 17 and 20 August 2009. We calculated the Modulation Transfer Function (MTF; also known as the Spatial Frequency Response or SFR) using the freely available MATLAB code sfrmat3 [*Burns*, 2015] on the images that appeared to be visually in best focus.

Table 5. Preflight Mastcam Focus Motor Count Versus Distance Data

Mastcam 34 mm (Left)		Mastcam 100 mm (Right)	
Distance (m)	Motor Count	Distance (m)	Motor Count
0.502	1802	1.682	1276
0.858	2007	1.827	1610
1.236	2105	2.006	1759
2.050	2224	3.004	2241
2.566	2249	4.196	2549
3.514	2303	4.257	2570
4.15	2310	4.654	2634
5.827	2315	5.051	2648
7.085	2320	6.037	2786
infinity	2404	7.025	2853
		infinity	3301

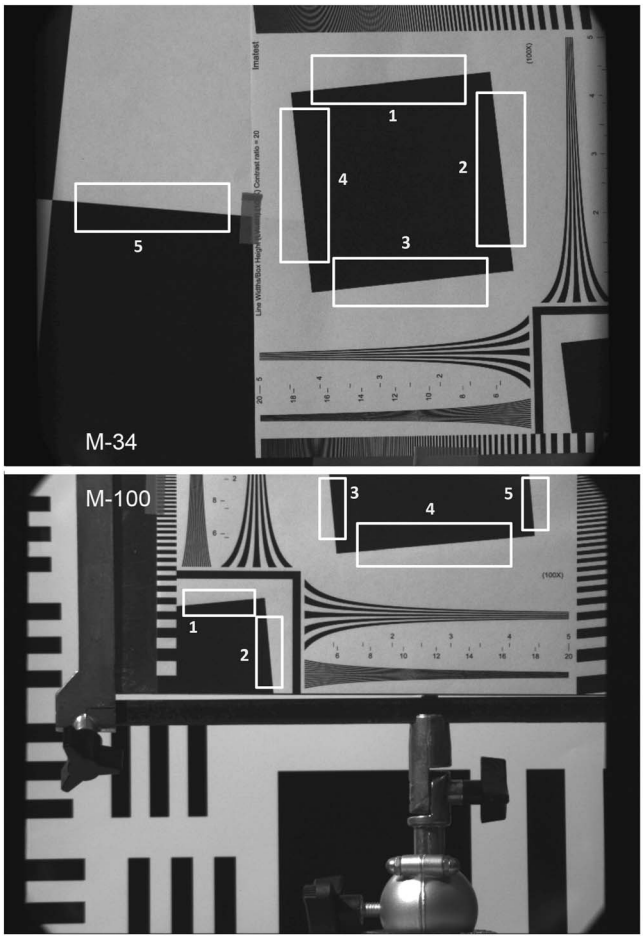

Figure 12. Example flight Mastcam (top) 34 mm and (bottom) 100 mm preflight calibration interpolated Green Bayer pixel images of the SVG Squares and Wedges pattern target [*Imatest*, 2015] used to assess the MTF performance of the flight cameras. Five regions of interest, numbered on each image, were analyzed.

The algorithm used in this software follows the International Organization for Standardization 12233 procedure for measuring the resolution of electronic still cameras using the slanted edge method [*Burns*, 2000]. We defined rectangular regions of interest spanning five suitable slanted bright-dark edges on the test chart [*Imatest*, 2015], primarily near the center half of the field of view. The raw images of the targets were interpolated to remove the Bayer pattern, converting each image into separate images of the red, green (Figure 12), and blue channels. Each single-band image was used as input to the sfrmat3 code to determine the MTF for each band separately. The MTFs for each of the five regions were within about 10% of each other, and so they were averaged to obtain a single, averaged MTF for each band in each camera. The detector pixel scale (7.4 μm per pixel) was used to convert the MTF from units of line pairs per pixel (lp/pix) to line pairs per millimeter (lp/mm).

Our analysis (Figure 13) reveals that the average (over the Bayer R, G, and B bands) MTF at 50% contrast occurs at a spatial frequency of 0.185 lp/pix for the M-34 camera and 0.155 lp/pix for the M-100 camera, corresponding to image scales of 5.4 and 6.4 pixels, respectively, for these cameras. Differences between the R, G, and B Bayer bands are generally small and within the noise of the analysis. For reference, Figure 13 also shows that the values of the MTF at Nyquist for the M-34 and M-100 cameras are ~6% to ~10% and ~4% to ~7%, respectively, assuming that "Nyquist" for a color Bayer imaging system is defined not as the standard 0.5 lp/pix of monochrome imaging systems but as $0.5(\sqrt{2}/2) = 0.35$ lp/pix [e.g., *Lensation GmbH*, 2013]. No aliasing was observed.

Several important caveats apply to this analysis. First, because of strict time limits on the testing of the flight Mastcams, imaging of the target at a large number of fine-scale focus positions was not possible, and thus, the particular images analyzed here, while visibly having the best focus of the available data set, might not be at the absolute best possible focus. Second, resolution is known to vary as a function of radial distance from the optic axis in these cameras, a fact apparent even from visual inspection of the images, but this effect has not been included in this analysis. For example, the corners of the M-100 rectangular field of view are slightly blurry relative to the central regions, since these areas are outside the nominally designed 1200 × 1200 science imaging area of the field of view (Figure 2). However, this effect generally has not impacted the quality of the science returned from Mars, as imaging requiring the highest calibration fidelity (or for mosaicking purposes) includes only the central 1200 × 1200 science imaging area.

The MTF and resolution values derived above are likely to be conservative estimates of the cameras' quantitative imaging performance. Indeed, the MTF and resolution performance may seem poor compared to other imaging systems; this is at least partly a result of the effect of the Bayer pattern color filter array

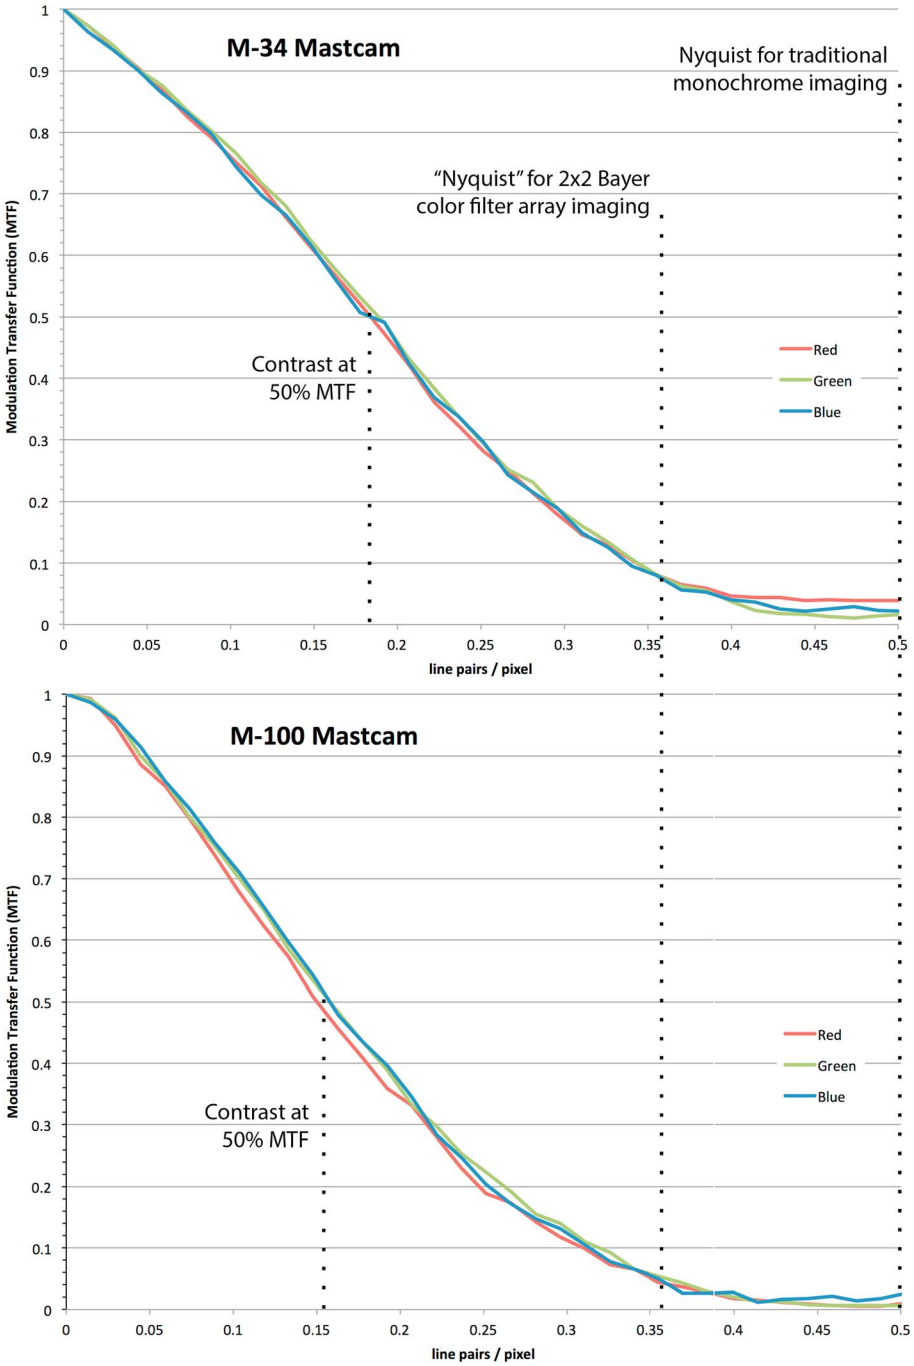

Figure 13. Modulation Transfer Function (MTF) curves for the (top) M-34 and (bottom) M-100 Mastcams, based on analyses of the resolution target images shown in Figure 12. See text for details.

bonded on top of the Mastcam CCDs. Bayer pattern filters like this have been shown to create a significant decrease (a factor of 2 or more) in imaging system MTF compared to monochrome systems [e.g., *Yotam et al.*, 2007].

3.2.9. Stray and Scattered Light Testing

A qualitative assessment of the level of stray (from out of field sources) and scattered (from in-field sources) light susceptibility for both Mastcams was conducted during ambient pressure and temperature testing at MSSS on 31 August 2009 (for the M-100) and 30 September 2009 (for the M-34). Images were collected while a bright fiber optic point source was manually moved just inside and outside of each camera's field of view.

Table 6. Known and Consistently Bad Pixels on the MSL/Mastcam CCDs

M-34 (x, y)[a]	Kind[b]	M-100 (x, y)[a]	Kind[b]
(140, 167)	gray	(300,198)	gray
(448, 242)	hot	(821, 315) to (822, 319)	dead
(692, 611)	gray	(734, 363)	gray
(755, 612)	gray	(422, 588)	hot
(1034, 634)	gray		
(1446, 800)	gray		
(228, 924)	gray		
(792, 980)	hot		
(1354, 1027) to (1355, 1028)	gray		
(380, 1102)	gray		
(560, 1106)	hot		
(1416, 1193)	gray		

[a]Where (0, 0) is the upper left pixel of each image.
[b]See text for details.

While the exact stray and scattered light levels cannot be quantified from these test images, visual inspection of the data did not reveal any obvious levels of stray light from bright sources just outside the fields of view of the cameras or obvious ghosts or glints from scattered light within the camera itself. Stray or scattered light has also not proven to be a significant concern in flight data acquired on Mars, although weak and offset ghost images of the Sun can be detected in solar filter imaging (see section 4.2.7).

3.2.10. Bad Pixels and/or Particulates on the Detectors

CCD image sensors often have individual or groups of pixels that are "dead" (nonresponsive), "gray" (responsive but either slightly more or less than average), or "hot" (saturated) compared to their typical neighbors, regardless of the exposure time or incident light level. Gray pixels that are somewhat less responsive than average are often caused by dust or other small semiopaque particulates adhering to the detector. CCD images, especially long-exposure images, also often include an additional population of hot or gray pixels created via cosmic ray or other high-energy particle interactions during the image exposure. Collectively, these kinds of anomalous pixels are often referred to as bad pixels. The effects of bad pixels sometimes propagate to their surrounding neighbors (±3–5 pixels in all directions) because of, for example, charge bleeding, electronic shutter smear, Bayer pattern interpolation, or 8 × 8 pixel JPEG compression. Dead and hot pixels (and potentially their contaminated neighbors) are typically not correctable via calibration, while gray pixels sometimes are.

The Mastcam detectors exhibit a small fixed number of dead and gray pixels per camera and a variable number of hot pixels that depends on the exposure time of the images and the vagaries of the solar cycle and cosmic ray flux activity. Hot pixels are not corrected in any way as part of the Mastcam calibration pipeline, although they can be easily filtered and replaced for cosmetic purposes using a standard low-pass or "dust and scratches" kind of filter found in many commercial image processing applications. Gray pixels can be corrected by flatfielding, if their responsivity is not substantially higher or lower than average and they have not changed since the last time flatfield calibration images were acquired. Table 6 provides a list of the relatively few currently known and consistently dead, gray, or hot pixels for each of the Mastcam cameras, based on preflight calibration measurements and in-flight performance on Mars. A few additional pixels on both CCDs exhibited hot or gray behavior for a limited period of time and then went back to normal behavior (perhaps having self-annealed with continued operation). Interpretation of morphologic or compositional variations in regions near and even around any of these bad pixels (especially in the same column) should be avoided.

3.2.11. Geologic Samples and Color/Reflectance Standards Imaging

In order to qualitatively assess the Mastcam system's abilities to resolve fine-scale detail and color/multispectral variations within natural geologic materials, we also conducted a series of ambient temperature and pressure multispectral imaging tests of a variety of diverse geologic samples, as well as a set of reflectance, geometric, and colorimetric standards. Our test target (Figure 14) is similar to targets that had been used for the same kind of "natural validation" imaging by the Mars Exploration Rover Pancam instruments during their preflight calibration and testing [e.g., *Morris and Graff*, 2002; *Bell et al*., 2003].

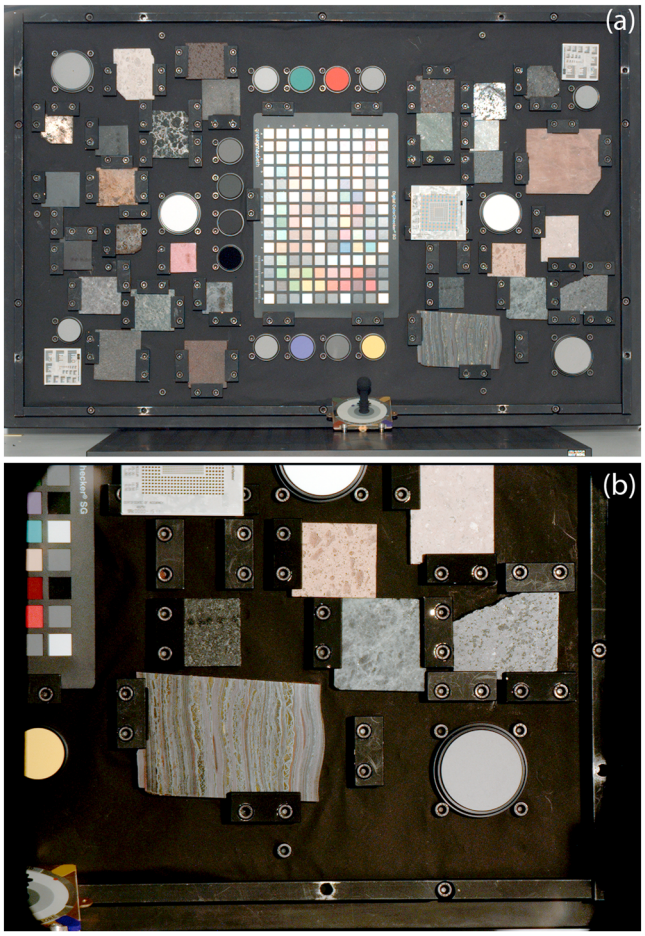

Figure 14. Geometric, color/reflectance, and geologic materials target imaged under ambient pressure and temperature conditions during MSL/Mastcam preflight calibration at Malin Space Science Systems, Inc., during late September 2009. (a) Entire target imaged by the wide angle M-34 flight instrument. (b) Lower right of target imaged by the narrow-angle M-100 flight instrument at 3 times higher resolution. Images have been white balanced using reference reflectance standards in the scene. For scale, the gray circular target at lower right is 5 cm in diameter. An engineering model of the Mastcam calibration target was also imaged in the foreground.

Analysis of images of this target with both the M-34 and M-100 multispectral filters demonstrated that relative reflectance (relative to standards on the target) could be derived to better than 10% accuracy even at room temperature and that spectral variations in natural geologic samples could be detected and spatially mapped using our narrow multispectral passband filters, even though the signal had also passed through the overlying RGB Bayer filter pattern (e.g., Figure 15).

3.3. Preflight Geometric/Camera Model Testing and Validation

The Mastcam lenses exhibit almost no distortion [*Ghaemi*, 2009]. One can observe the undistorted nature of Mastcam images by examining pictures of dot targets like that shown in Figure 11. These kinds of images show visually that the dot spacing and shape do not change visibly anywhere in the field of view, for either camera. However, a more accurate approach to geometric characterization involves determination of the instrument's camera model.

A geometric camera model is a set of equations that transform a 3-D point in space to a 2-D position in an image (pixel location). The model can also be inverted to transform a pixel in the image to a set of points in space that would map to that pixel (the imaging locus). The model contains the camera position and pointing vector, and it also models lens distortion and the interior geometry of the instrument. A camera model facilitates accurate monoscopic and stereoscopic measurement of points in an image or images, precise mosaicking of images, and creation of "linearized" (i.e., geometrically corrected) image products (section 5.2.9) for which lens distortion has been removed. To accommodate multiple kinds of existing end user software for geometric analyses, two kinds of initial, mathematically similar camera models have been developed for Mastcam and are described here.

3.3.1. Four-Vector MSSS-Produced CAHV Camera Model

The first is a four-vector "CAHV" photogrammetric model produced at MSSS that is routinely used by the Mastcam team in the creation of geometrically corrected Mastcam PDS archival data products (see section 5.2.9). The CAHV camera model acronym refers to the vectors which permit transformation from object to image coordinates [e.g., *Yakimovsky and Cunningham*, 1978; *Gennery*, 2001; *Di and Li*, 2004] (see also the inset of Table 8 below). The one-letter CAHV model term definitions are the following: C—camera center vector from the ground coordinate system origin to the camera perspective center; A—camera axis unit vector perpendicular to the image plane; and H, V—horizontal and vertical information vectors.

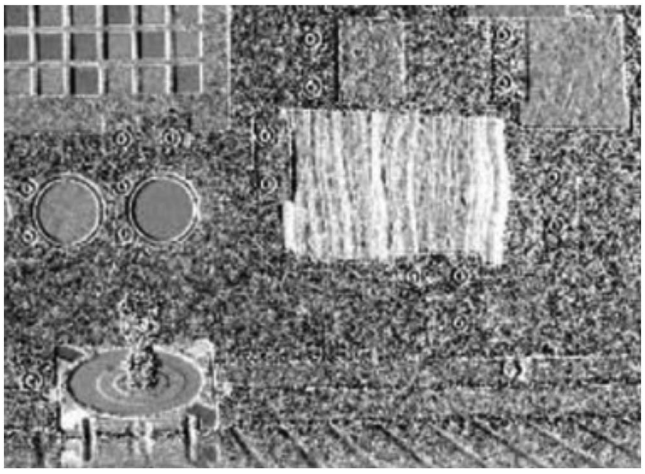

Figure 15. Band depth map at 865 nm (relative to continuum wavelengths at 751 nm and 1012 nm) generated from a portion of the Mastcam M-34 preflight geologic target multispectral imaging campaign (Figure 14). This spectral parameter is sensitive to the presence of crystalline hematite and indeed highlights the hematite-bearing color chip on the lower left of the Mastcam calibration target and the hematite-bearing layers in the Banded Iron Formation geologic sample.

The CAHV model can be losslessly constructed from a pinhole camera model. Specifically, via the CAHV model, if P is a point in the scene, then the corresponding image location x, y (column, row) can be computed as follows:

$$x = ((P–C)\,H)/((P–C)\,A), \quad \text{and} \quad y = ((P–C)\,V)/((P–C)\,A). \tag{3}$$

To fit a photogrammetric camera model to the MSL Mastcams, we acquired a small set of Assembly, Test, and Launch Operations geometric calibration target images at Jet Propulsion Laboratory (JPL) on 6–7 December 2010. For this activity, both Mastcam camera heads were viewing the same collection of JPL-provided dot targets (e.g., Figure 11). The images have known (measured via laser metrology) object space (3-D) coordinates for a set of points captured in the images. The 2-D image space coordinates corresponding to the points were measured as accurately as possible using standard centroiding methods, and we used a nonlinear weighted least squares (NLS) algorithm (also known as bundle adjustment) [e.g., *Brown*, 1958] to determine the set of parameters that minimize the difference between the 3-D points projected into the image using the camera model and the points measured from the image.

We determined and correlated the center pixel location (to subpixel accuracy) for each dot in the available dot target images with the position of that point in the rover reference frame (determined by JPL-Caltech engineers, who surveyed the location of the four corners of each target). Following the bundle adjustment procedures described by *Brown* [1958] and *Mikhail et al*. [2001], we determined the exterior parameters (target positions and orientations), interior parameters (focal length and detector geometry), and lens distortion coefficients. The position for the farthest target was not adjusted so as to fix the scale of the system. We ran a single optimization that used data from all calibration images to model the effect of focus motor count where needed.

In our analysis, the term *focal length* refers to the value of a theoretical, ideal camera without distortion; this differs from the physical, effective focal length determined by lens design. Mastcam focal length (f_m) was modeled as a linear function dependent on each camera's focus motor count. The linear effect of motor count on focal length is referred to as the focus factor (f_f) and has units of millimeters of focal length per focus motor count. The function is

$$f_m = f_{\text{nom}} + f_f(m_{\text{nom}}–2800) \tag{4}$$

where f_m is the modeled focal length at motor count m and f_{nom} is a reference focal length at a nominal focus motor count of m_{nom} = 2800. The value of 2800 is somewhat arbitrary; we selected it because it is between the minimum and maximum focus positions for the M-100 camera. For the M-100 camera, the NLS procedure yielded f_{nom} = 100.446803 and f_f = −2.547113e−03, yielding M-100 focal lengths ranging from 104.329 mm to 99.170 mm for focus motor counts ranging from 1276 to 3301, corresponding to distances of ~1.7 m to infinity, respectively. For the M-34 camera, focal length did not exhibit a statistically significant change with focus motor count in our NLS modeling; instead, a constant focal length of 34.367 mm was derived, as was expected from the theoretical design of the optics.

The origin of the image focal plane coordinate system (x, y) is at the principal point, that is, the location at which the optical axis intersects the focal plane. This is defined in millimeters, with the x direction to the right (increasing columns) and the y direction up (decreasing rows). The boresight is the same point but defined in pixels (i, j). Relative to the upper left Mastcam CCD pixel, defined as pixel (0, 0), the M-34 boresight is located

Table 7. MSL/Mastcam Geometric Camera Model Parameters

	Mastcam-34 (Left)	Mastcam-100 (Right)
Camera Interior Orientation Parameters		
Focal length f_{nom}, mm (see equation (4))	34.367205	100.446803
Focus factor, f_f (see equation (4))	0.0	−2.547113e−03
Boresight row	588.41	608.81
Boresight column	834.63	836.11
a_{11}, pixels per mm (see equations (7) and (8))	135.153	135.154
a_{12}, pixels per mm (see equations (7) and (8))	−0.0382	−0.0386
Lens Radial Distortion Parameters		
Lens distortion k_1	−1.118977e−04	1.513695e−04
Lens distortion k_2	−1.023513e−06	0.0
Distortion center row, mm	−0.113876	0.262451
Distortion center column, mm	0.152029	−0.250667

at column i_0 = 834.62, row j_0 = 588.41 and the M-100 boresight is located at column i_0 = 836.11, row j_0 = 608.81.

Radial distortion is measured in the focal plane relative to a point in the focal plane called the *center of distortion*. Here we assumed that the center of distortion is equal to the principal point. Given that r, the distance in millimeters of a point (x, y) in the focal plane from the principal point (x_0, y_0) for an ideal camera with no distortion is

$$r = \left((x-x_0)^2 + (y-y_0)^2\right)^{0.5}, \tag{5}$$

then the radius for the actual camera with distortion is calculated as

$$r' = r\left(1 + k_1 r^2 + k_2 r^4\right) \tag{6}$$

Values for the distortion coefficients k_1 and k_2 are reported in Table 7. The value of k_2 for the M-100 camera was found to be statistically insignificant and thus set to zero. Higher-order distortion terms were not needed to accurately model either camera's distortion.

The mapping of points from the focal plane coordinate system (x, y), in millimeters, to the image coordinate system (i, j), in pixels, is defined as an affine transformation in which

$$i = i_0 + a_{11}x + a_{12}y \tag{7}$$

$$j = j_0 - a_{12}x - a_{11}y \tag{8}$$

and for which (i_0, j_0) is the boresight pixel. Because the focal length can be confounded with pixel pitch, the value of a_{11} was set to the number of pixels on the array per millimeter and was not adjusted during the analysis. Note that this indicates that the pixels may be slightly nonsquare or that there may be a slight nonorthogonality of the focal plane axes. The affine coefficients (in pixels per millimeter) determined by our analysis are listed in Table 7.

This initially derived Mastcam camera model has some important limitations. For example, (a) the movement of the principal point and the camera center due to focus motor count have not been modeled, (b) the center of distortion might be offset from the principal point, (c) the change in lens distortion parameters due to motor count/working distance has not been modeled, (d) the effect of Martian environment temperatures on camera model parameters has not been modeled, and (e) how well the linear model applies to focal lengths at minimum working distance and at infinity has not been investigated. Uncertainties associated with those limitations could be decreased with additional analysis of preflight or flight calibration data.

3.3.2. JPL-Produced CAHVOR Camera Model

The second Mastcam camera model is a six-vector CAHVOR model produced at JPL-Caltech to support rover operations. The CAHVOR camera model parameters are reported with all of the Mastcam images archived with the PDS *except* the geometrically corrected products, which use the four-vector model described in section 3.3.1 above (see also section 5.2.9). The CAHVOR model is an extension of the CAHV that includes additional vector terms for optical ("O") and radial ("R") distortion [e.g., *Yakimovsky and Cunningham*, 1978;

Table 8. MSL/Mastcam CAHVOR Camera Model Vector Parameters[a]

Parameter	Component 1	Component 2	Component 3
	Mastcam 34 (Left)		
(C)enter	0.767151	0.433709	−1.971648
(A)xis	0.999664	0.025047	0.006727
(H)orizontal	712.373106	4664.465028	33.182389
(V)ertical	570.612488	−14.279011	4648.733195
(O)ptical	0.999627	0.026908	0.004759
(R)adial	−0.000151	−0.139189	−1.250336
	Mastcam 100 (Right)		
(C)enter	0.872866	0.677208	−1.970931
(A)xis	0.999741	−0.022668	0.001764
(H)orizontal	1162.600942	13353.117702	−93.448986
(V)ertical	512.731806	83.454058	13371.226140
(O)ptical	0.999495	−0.029483	0.011824
(R)adial	−0.000106	1.436779	−0.685884

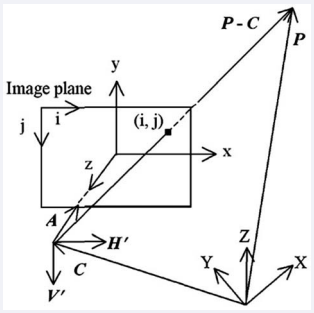

[a]For example, *Yakimovsky and Cunningham* [1978]; inset figure defining CAHV vectors is from *Di and Li* [2004].

Gennery, 2001; *Di and Li*, 2004]. As with the MSL Navcam and Hazcam instruments [*Maki et al*., 2012], a CAHVOR camera model was produced for Mastcam at JPL-Caltech to support *Curiosity* rover operations. The calibration effort followed the procedure described by *Yakimovsky and Cunningham* [1978], as modified by *Gennery* [2001, 2006], and used the same targets that were set up in December 2010 for the MSL Navcam and Hazcam calibration effort [*Maki et al*., 2012] (e.g., Figure 11). The data analyzed were based on the same six Mastcam image pairs acquired on 7 December 2010 that were examined for the four-vector model described above. The derived M-34 and M-100 CAHVOR camera model coefficients are listed in Table 8 and are also listed within the data product label's "GEOMETRIC_CAMERA_MODEL_PARMS" section for each of the PDS archived Mastcam data products (EDR "_XXXX" images as well as RDR "_DRXX" and "_DRCX" products); see section 5.2.7), except for geometrically corrected (linearized) Mastcam archive RDR data products ("DRLX" and "DRCL" products; see Table 13 below), which use the geometric camera model parameters described in Table 7 and section 3.3.1 above (see also section 5.2.9).

Unlike the model in section 3.3.1, the center of distortion in the JPL CAHVOR model is not constrained to the principal point, as evidenced by the difference in the O and A vectors in Table 8, and that the radial distortion (R) includes higher-order terms. However, these are very slight differences that are unlikely to have significant impact when using the models. More significant is that the JPL model is not currently adjusted for focus, even for the M-100. Rather, it is a single compromise model derived for an "average" focus (specifically, focus motor counts of 2315 for the M-34, and 2702 for the M-100, although the models are a compromise at all focus values). Incorporating focus into the JPL CAHVOR model is a task for future model improvements.

3.3.3. Pointing the JPL Camera Model

The CAHVOR camera model combines both intrinsic parameters (camera internal geometry, focal length, lens distortion, etc.) and extrinsic parameters (camera position and pointing vector) into a single model. That means a given CAHVOR model can only be used at one specific pointing. In order to use it at a different pointing, say as the mast is slewed around, the model must be transformed to reflect the new extrinsic parameters. This section discusses the process used by the JPL operations ground software (by MIPL, the Multimission Image Processing Lab), which in turn mimics what is done on board the rover. Thus, the camera models present in the (nonlinearized) PDS labels (see section 5.2.9) can be recreated given the mast azimuth/elevation that is also in the label. Alternatively, adjusted models can be created with refined pointing using methods such as the mosaic bundle adjustment process described in *Deen et al.* [2015].

Details of the algorithm, and an example showing its use, are in Appendix A. In summary, a model of the mast head kinematics is used to derive the camera position (an arbitrary point that moves with the head) and a quaternion expressing mast head rotation. The calibration camera models similarly have a point and a quaternion attached (derived in the same way, representing the extrinsic parameters of the camera pose during calibration). The calibration camera model is then rotated by the inverse of the calibration quaternion and translated to the origin, effectively removing the extrinsic parameters. The model is then rotated and translated to the desired pointing, which results in a final model for a full-frame image. Most

Table 9. MSL Mast Kinematics Parameters

Name	Value	Meaning
Pcal	0.80436 0.55942 −1.90608	Calibration position of camera head, ROVER_NAV frame (meters)
Qcal	0.99999 0.00001 −0.00325 −0.00104	Calibration quaternion of camera head, ROVER_NAV frame (s,V)
az_home	3.167345	Azimuth joint angle when head is forward/level (radians)
el_home	1.588171	Elevation joint angle when head is forward/level (radians)
az_point	0.714803 0.559424 −1.109344	Arbitrary point on azimuth axis, ROVER_MECH frame (meters)
az_axis	−0.001 0.0 1.0	Elevation axis unit vector, ROVER_MECH frame; normalized by the pseudocode
el_point	0.714668 0.561701 −0.785576	Arbitrary point on elevation axis, ROVER_MECH frame (meters)
el_axis	−0.022 −1.0 0.0	Elevation axis unit vector, ROVER_MECH frame; normalized by the pseudocode
az_elaxis	0.0	Azimuth during elevation axis survey, middle of backlash (radians)
epsilon	1.0e−10	Prevents divide by zero
rmech_to_rnav	0.09002 0.0 −1.1205	Add this to a ROVER_MECH frame value to convert it to ROVER_NAV (frame orientation is the same) (meters)

Mastcam images are subframed, in order to return only the science portion of the image. Additionally, thumbnail images are downsampled. These factors are applied to the full-frame camera model in order to get the final model.

It should be noted that the same algorithm is also used to create pointed camera models for the Navcam and Chemcam-RMI instruments. The parameters in Table 9 are identical; only the calibration models themselves differ. Note that a different treatment of MSL camera mast head kinematics is presented in *Peters* [2016].

4. In-Flight/Surface Calibration and Validation: Methods, Data Sets, and Results

4.1. Cruise

Curiosity was launched on 26 November 2011. For instrument checkout purposes, three sets of Mastcam bias/dark current images were obtained during the cruise to Mars, on 13 March, 20 April, and 14 June 2012. The cameras and RSM were stowed against the rover deck within the completely dark confines of the sealed cruise stage spacecraft configuration. For each set of images, two 1 s and two 10 s exposures were acquired from each camera. No significant dark current was detected in the data sets, consistent with the cold temperatures of the CCDs during these imaging sessions (0.8°C, −7.6°C, and −8.9°C, respectively), and the preflight predictions of the dark current behavior (section 3.2.2 and Figure 8). Bias levels were also low, as expected. Differencing of subsequent exposures did not reveal any anomalous noise behavior in the cameras.

4.2. Special In-Flight Calibration and Testing on Mars

4.2.1. In-Flight Bias and Dark Current Testing

On *Curiosity* mission sols 320 and 321 (30 June and 1 July 2013) and again on sols 1350 and 1351 (24 and 25 May 2016), specific sets of Mastcam imaging sequences were commanded to characterize the in-flight dark current performance of the cameras at different temperatures by acquiring a series of images of the ground through the L7 and R7 solar ND filters (Table 10). CCD temperatures for the M-100 camera are not routinely returned in telemetry as part of the MSL telemetry architecture. However, we observed a strong correlation ($r^2 > 0.98$) between the optics heater temperature (HTR1 in the PDS labels) of the M-34 camera when its heater is off and its routinely reported CCD temperature (FPA temperature in the PDS labels). Thus, we used that

Table 10. In-Flight Bias and Dark Current Observational Data Compared to Preflight Model

						DN Observed[b]		DN Modeled[c]	
Sol	Sequence	LTST	t_{exp} (s)	T_{CCD} (°C)[a]	T_{DEA} (°C)	M-34	M-100	M-34	M-100
320	mcam01049	13:37	0.0	−9.5	+20.6	120.0	120.3	121.5	122.0
		13:37	0.1	−9.5	+20.6	--	120.4	--	122.1
		13:37	0.5	−9.5	+21.6	120.7	--	122.2	--
		13:37	1.0	−9.5	+21.6	--	121.7	--	123.1
		13:37	10.0	−9.5	+21.9	132.9	--	135.1	--
320	mcam01050	17:07	0.0	−12.7	+27.8	119.8	120.0	121.5	122.0
		17:08	0.1	−12.7	+27.8	--	120.1	--	122.1
		17:08	0.5	−12.7	+28.1	120.2	--	122.0	--
		17:08	1.0	−12.7	+28.1	--	121.1	--	122.9
		17:08	10.0	−12.7	+28.1	128.1	--	132.0	--
321	mcam01051	10:47	0.0	−26.4	+16.0	121.9	122.6	121.5	122.0
		10:47	0.1	−26.4	+16.0	--	122.6	--	122.0
		10:47	0.5	−26.1	+16.7	122.0	--	121.7	--
		10:47	1.0	−26.1	+16.7	--	122.8	--	122.3
		10:47	10.0	−26.1	+17.0	124.9	--	125.1	--
1350	mcam06485	11:28	0.0	−14.9	+17.0	120.2	122.0	121.5	122.0
		11:28	0.5	−14.6	+17.6	121.4	122.1	122.0	122.4
		11:28	10.0	−14.6	+17.6	129.1	130.5	130.5	129.8
1350	mcam06491	14:36	0.0	−5.7	+25.5	120.0	120.1	121.5	122.0
		14:36	0.5	−5.7	+25.8	120.8	121.9	122.4	122.8
		14:36	10.0	−5.7	+25.8	137.7	138.6	139.9	137.8
1350	mcam06492	16:36	0.0	−8.9	+25.2	120.0	120.3	121.5	122.0
		16:36	0.5	−8.5	+25.2	120.8	122.0	122.2	122.6
		16:36	10.0	−8.5	+25.5	133.8	133.7	136.2	134.7
1351	mcam06495	09:06	0.0	−15.2	+9.8	120.6	122.0	121.5	122.0
		09:06	0.5	−15.6	+10.5	121.8	122.1	121.9	122.4
		09:06	10.0	−15.6	+10.8	128.4	130.8	129.8	129.2

[a]For the M-34 (left) Mastcam CCD. Value for the M-100 (right) Mastcam CCD assumed to be the same.
[b]After adding back the onboard subtracted value of the DARK_LEVEL_CORRECTION label value.
[c]Based on preflight bias + dark current model in section 3.2.2 and Figure 8.

information to estimate a relationship between the routinely reported M-100 HTR1 temperature and its rarely reported FPA temperature. Specifically, we found that an estimate for the M-100 CCD temperature can be made by assuming that CCD temperature in °C = 1.1 × HTR1 + 3.0, for times when HTR1 is off.

A background level of 117 DN (the DARK_LEVEL_CORRECTION value stored in the PDS labels for these data files) was subtracted from all of the images by the DEA prior to downlink. Table 10 shows that the background levels observed on Mars for the M-34 and M-100 cameras are small at the typical operating temperatures of the CCDs and consistent with the predictions of the preflight background model described in section 3.2.2 to within 2.5 DN at the warmest operating temperatures of the CCD and electronics and typically within ±1 DN of the model for typical operating temperatures.

4.2.2. In-Flight Electronic Shutter Smear Characterization

For the small fraction (<10%) of Mastcam images acquired with very short exposure times (typically <6 ms), section 3.2.3 showed that electronic shutter smear might need to be corrected in order to achieve adequate radiometric and/or photometric accuracy, depending on the goals of the observation in question. To enable assessment of the effects of shutter smear in real flight data, as well as possible eventual higher-fidelity modeling and removal of the effects in calibrated images, we acquired two observations on Mars.

The shutter smear test data were acquired on *Curiosity* sols 36 (sequence mcam00155) and 38 (mcam00169) and consisted of M-34 and M-100 images of the Mastcam calibration target and bright rover deck surroundings acquired using very short exposure times (2.3 or 4.3 ms) and immediately followed by identical images acquired at a commanded exposure time of 0.0 ms. An example of the M-100 images from sol 38 is shown in Figure 16. The 4.3 ms image looks normal and does not reveal any obvious evidence of shutter smear (indeed, the analysis in section 3.2.3 showed that the worst-case effect in such an image would be only a 1.7% effect for pixels near the bottom of the image). The zero-second image has an average signal level of about 7% of

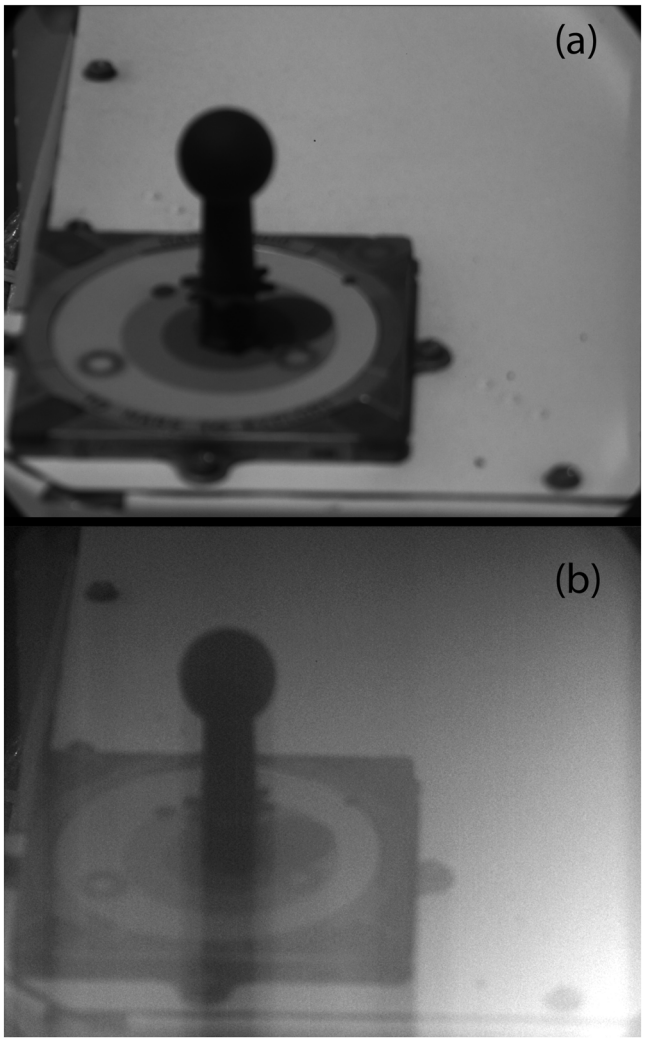

Figure 16. Example of in-flight electronic shutter smear test images acquired on *Curiosity* sol 38 in sequence mcam00169. (a) "Normal" M-100 Filter R0 green Bayer filter image of the Mastcam calibration target using an exposure time of 4.3 ms. (b) Smeared M-100 Filter R0 green Bayer filter image of the same scene acquired immediately after image in Figure 1a at a commanded exposure time of 0.0 ms. The zero-second image shows effects of both shutter smear and a "ghost image" of the original scene. See text for details.

the brightness of the 4.3 ms image and shows several interesting effects. First, the zero exposure image exhibits a noticeable "smear" pattern in the direction toward the bottom of the image (toward the horizontal shift register). But second, there is a faint residual or "ghost" image of the original 4.3 ms scene superimposed upon the smear signal. The former is expected based on the nature of electronic shutter smear in this CCD (section 3.2.3). The latter, however, is a manifestation of a different effect in the Mastcam images, related to the way exposures are implemented in the FPGA logic of the Mastcam DEAs.

Essentially, if a zero-second exposure is commanded for the Mastcams, the resulting exposure time is not actually zero. This is because even if a zero-second exposure is commanded, the CCD still transfers charge from the photosites to the vertical register; it just does so a very short but nonzero amount of time after the photosites have been cleared of charge by the electronic shutter. In the Mastcam DEA FPGA logic, the commanded vertical transfer pulse duration for each pixel is programmed to be about 106 μs wide, from start to finish (see *Truesense Imaging, Inc.* [2012] for electronic shutter timing details). While that time interval between the end of the commanded "zero-second" exposure and the charge transfer into the vertical registers is very short, it is still apparently long enough for observable signal to accumulate in the photosites from conventional photocharge and is thus not smear signal, per se. The fact that the ghost image is a faint mimic of the nonzero exposure image supports this explanation. This explanation can be tested empirically, using the sols 36 and 38 shutter smear test images. Scaling the nonzero exposure time by the ratio of the average DN values in the zero exposure to nonzero exposure images implies an effective actual zero-second exposure time of around 90 μs. This is indeed comparable to the stable duration of the vertical transfer pulse width.

This kind of ghost image, therefore, occurs in all Mastcam images after each exposure is complete but is amplified in the unusual (and unrealistic to implement) circumstance of commanding a zero-second image. Both the ghost signal and the shutter smear are small, however, and not apparent in most normal Mastcam images. The most noticeable effects of shutter smear occur when there is a bright object near or past saturation in a surrounding darker field (such as the calibration target scene in Figure 16). Shutter smear has also been observed in MAHLI "self-portrait" images when looking at white rover structures and in the Mars Descent Imager (MARDI) descent sequence with the bright heat shield against the darker Martian surface, as both of those cameras use the same CCD and timing logic.

4.2.3. Improved Flatfields From In-Flight Sky Measurements

For the purpose of creating a flatfield image, sky images were taken on sols 36–38 (sequences mcam00153, mcam00164, and mcam00173). The images were taken with the Sun low in the west near 1600 Local Martian Solar Time (LMST) and pointed toward the anti-Sun azimuth. During this time, there were no expected or observed discrete clouds and the dust optical depth was about 0.79 [*Lemmon*, 2014]. The concept was to acquire, in each filter, a 2 × 2 mosaic centered on a fifth image. To manage data volume, the mosaic underwent lossy compression and was used to characterize large-scale radiance variations; the central image was losslessly compressed for use as a sky flat. Due to operational time constraints, M-34 and M-100 images were acquired simultaneously. The 2 × 2 mosaic was designed such that the image centers overlapped the corners of the central image.

M-34 images were processed first, taking advantage of the complete overlap of the flatfield image and the mosaic. We first describe the processing for Filters 5 and 6, which have essentially equal response from each component of the Bayer unit cell and can be considered a monochromatic array. All images were read in, an inverse look-up table was used to derive detector counts, mean bias and dark counts were derived from masked portions of the full-frame images and were subtracted from the whole image, cosmic ray strikes and hot pixels were rejected with a selective median filter, a shutter image was calculated and subtracted, and the images (which had been autoexposed) were divided by exposure time and scaled by cosine of the solar zenith angle to match the illumination of the central image. The initial flatfield was taken to be the central image (i.e., the initial assumption was that the sky had uniform brightness). A mapping between detector coordinates and sky coordinates (elevation and azimuth relative to the Sun) was determined from image header and timing information. The following steps were iterated 10 times: the 2 × 2 images were each divided by the flatfield, they were projected into sky coordinates and sampled at the projected geometry of the central image, a polynomial (second order in azimuth and elevation) was fit to the sky mosaic to form a new sky radiance model, the central image was divided by the sky radiance model to produce a new flatfield image, and the result was normalized to the mean of the sky model radiance for the 128 × 128 pixels around the optical axis (adjusted due to Bayer and JPEG constraints to detector lines 528–655 and samples 768–895). For inspection, a final mosaic was created, with the central image on top of the other sky images; the result was then verified to be free of significant artifacts from variations in sky brightness, such as seams or discrete cloud features. The resulting flatfield was stored as an image of the detector (i.e., a 1648 × 1200 monochrome image).

M-34 images from Filters 1–4 were processed similarly, but the analysis had to account for the Bayer sampling. For each, one or more elements in the Bayer unit cell had significant response to light through the filter, and one or more elements had substantially less response. The losslessly compressed central image returned all elements and showed their relative response. The JPEG-compressed 2 × 2 mosaic images had been (on board) sampled only in their responsive cells (green for L1, blue for L2, and red for L3, L4, and R3) and interpolated for the other cells. A mask was created such that interpolated cells were not used in the analysis, but the above process was otherwise followed. For these filters, two flatfield images were created. The first is for lossless images and shows the relative response differences in the Bayer unit cell that come from dividing the losslessly compressed image by the sky model. The second is for JPEG images and samples the flatfield at responsive elements of the unit cell, with bilinear interpolation over the other elements.

M-34 images from Filter 0 contain color information. The processing was similar, except that the three bands (red, green, and blue) were done separately, allowing for different sky models. This also resulted in two distinct flatfield images. The first, for losslessly compressed images, is two-dimensional. Unlike the lossless flatfields discussed above, the mean effect of the Bayer pattern has been divided out due to the normalization by three independent sky models. Thus, the use of this flatfield does not affect the relative signal in red, green, and blue. The flatfield image for JPEG-compressed images is three-dimensional, with each band's flatfield image calculated from the raw image using constant-hue-based interpolation.

M-100 images were originally intended to undergo the same procedure, but accomplishing that would have required a separate grid with ~5° between mosaic images rather than ~15°. For efficient use of time, the images for each camera were taken in parallel. As a result, the spacing of the M-100 images is not ideal for the above procedure. However, for Filters R0 (RGB), R1, R2, and R6, the above process produced detailed sky radiance information for the appropriate time and wavelength from M-34 images. In these cases, the M-100 central image, after the input and preprocessing steps, was divided by the M-34 sky model to

produce a flatfield image. For Filters R3, R4, and R5, the same process was used except that the model was a linear interpolation between M-34 models adjacent in wavelength (751, 867, and 1012 nm for M-34, 805, 908 and 937 nm for M-100). The resulting flatfield was normalized to the mean of the sky model radiance for the 128 × 128 pixels around the optical axis (adjusted due to Bayer and JPEG constraints to detector lines 600–727 and samples 696–823). This resulted in monochromatic flatfield images for R4, R5, and R6; separate lossless and JPEG monochromatic flatfield images for R1, R2, and R3; and a monochrome and three-band flatfield image for R0.

During the acquisition of these sequences, the focus was set to be typical of midfield imaging. Vignetting around the edges of the field (Figure 9) is focus dependent, so masking of the borders of the flatfield is recommended (areas where a responsive Bayer channel's flatfield value is >0.75 are acceptable; those that are $0.25 \leq$ flatfield ≤ 0.75 may be useful for midfield distances; and those that are <0.25 are not expected to be reliable). The flatfield images account for variations in detector response. However, there are some pixels that are insufficiently responsive; these may be identified using the same criteria as for vignetting. The use of a look-up table to compress all images to 8 bits per pixel (prior to JPEG or lossless predictive compression) limits the accuracy of the flatfield to ~1% per pixel for the primary Bayer channels for each filter and ~2% or worse for the other channels.

4.2.4. Validation and Improvement of Radiance Calibration From In-Flight Sky Measurements

Several sky-crossing Mastcam imaging surveys have been obtained at different times during the mission, partly for atmospheric science purposes and partly to provide an independent validation (via modeling) of the Mastcam absolute radiometric calibration (section 3.2.6). For example, on sol 1645, sequences mcam008516 and mcam008517 were designed and run to constrain the absolute radiometric calibration of each Mastcam filter. Specifically, 94 images were obtained through all left eye (M-34) filters (Table 2), as well as right eye (M-100) Filters R3, R4, and R5 (805, 908, and 937 nm). Images including scattering angles of 5°–150° were obtained through M-34 Filters L1, L2, and L4 (527, 445, and 676 nm). The data were acquired at an FPA temperature of approximately –7°C.

The images were calibrated to radiance using the Version 2 pipeline described in this paper, and the radiance calibration coefficients in Table 4, and were then modeled with a retrieval code based on the discrete ordinates radiative transfer code DISORT [*Stamnes et al*., 1988]. For initial parameters, the aerosol model of *Tomasko et al*. [1999] was used for aerosol mean radius, *a*; variance of the size distribution, *b*; slope of internally scattered light, *G*; and angle of minimum internally scattered light, $\theta_{\min}$. These parameters describe the shape of an arbitrary phase function, with the first two parameters controlling forward scattering and the last two parameters controlling side and back scattering. Initial single-scattering albedo as a function of wavelength was updated to values from *Wolff et al*. [2009], obtained from orbit under global dust storm conditions. The L1, L2, and L4 filters (527, 445, and 676 nm) were modeled, allowing an arbitrary adjustment to radiance to achieve a best fit. The model fits were all based on the shape of the radiance profile across the sky, with the absolute calibration normalized out. An implied calibration parameter (radiance per DN/s) was extracted from the fits (Table 11), and errors between the best fit model radiances and as calibrated original radiances were observed to be less than 1%. Values of $a = 1.25$ μm and $b = 0.22$ μm were used for all models, based on results from these three filters. Each of the remaining filters were fit without adjusting the size distribution; left eye modeled radiance fits were generally also subpercent accuracy, while right eye fits had 1–2% errors relative to the original as calibrated radiances. Table 11 compares the laboratory-derived radiance calibration coefficients (Table 4) to the coefficients derived from this sky modeling work. The sky model is considered to have a 10% uncertainty for the purposes of this comparison.

Right eye (M-100) versions of filters in common between the two cameras were not included in the above analysis, because of the preferentially larger larger left eye (M-34) field of view and the desire to acquire sky images spanning a wide range of azimuths and elevations as close together in time as possible. Instead, sky flat images (e.g., section 5.3.4), in which the left and right eyes were commanded as stereo pairs, were used to derive an in-flight left/right eye response ratio. Table 11 shows that ratio for each right eye filter that has a left eye counterpart. The radiometric calibration for the sky model has been derived from the left eye value and the left/right ratio.

We note that despite the difference in temperature between the input data sets (ambient preflight data versus −7°C in-flight data), in general, the laboratory calibration and sky model calibration match closely

Table 11. Validation and Refinement of MSL Mastcam Radiometric Calibration Coefficients Based On Radiance Modeling of In-Flight Mars Sky Observations

Filter	Wavelength (nm)	Preflight Coefficient[a] [W/m^2/nm/sr)/(DN/s)] ± 1σ Uncertainty	Sky Model Coefficient[b] [W/m^2/nm/sr)/(DN/s)] ± 1σ Uncertainty	Left/Right Ratio[b]	Preflight/Sky Model Coefficient Ratio
		M-34 (Left Mastcam)			
L0R	640	3.56e−07 ± 3.6e−08	3.88e−07 ± 3.9e−08	-	0.92
L0G	554	3.39e−07 ± 3.4e−08	3.56e−07 ± 3.6e−08	-	0.95
L0B	495	4.47e−07 ± 4.5e−08	4.43e−07 ± 4.4e−08	-	1.01
L1	527	2.38e−06 ± 2.4e−07	2.38e−06 ± 2.4e−07	-	1.00
L2	445	2.81e−06 ± 1.8e−06	1.85e−06 ± 1.9e−07	-	1.52[c]
L3	751	2.61e−06 ± 2.6e−07	2.62e−06 ± 2.6e−07	-	1.00
L4	676	2.04e−06 ± 2.0e−07	2.05e−06 ± 2.1e−07	-	1.00
L5	867	6.45e−06 ± 6.5e−07	6.47e−06 ± 6.5e−07	-	1.00
L6	1012	1.28e−05 ± 1.3e−06	1.50e−05 ± 1.5e−06	-	0.85
		M-100 (Right Mastcam)			
R0R	638	6.36e−07 ± 6.4e−08	6.32e−07 ± 6.3e−08	0.615 ± 0.018	1.01
R0G	551	6.08e−07 ± 6.1e−08	5.89e−07 ± 5.9e−08	0.607 ± 0.018	1.03
R0B	493	7.98e−07 ± 8.0e−08	7.43e−06 ± 7.4e−08	0.598 ± 0.018	1.07
R1	527	3.83e−06 ± 3.8e−07	3.66e−06 ± 3.7e−07	0.650 ± 0.020	1.05
R2	447	4.09e−06 ± 1.6e−06	3.11e−06 ± 3.1e−07	0.628 ± 0.019	1.32
R3	805	6.96e−06 ± 7.0e−07	6.02e−06 ± 6.0e−07	n/a	1.16
R4	908	1.38e−05 ± 1.4e−06	1.21e−05 ± 1.2e−06	n/a	1.14
R5	937	1.81e−05 ± 1.8e−06	1.66e−05 ± 1.7e−06	n/a	1.09
R6	1013	2.14e−05 ± 2.1e−06	2.24e−05 ± 2.2e−06	0.668 ± 0.020	0.96

[a]From Table 4, using the dominant red, green, or blue channel as indicated in Table 12 or the average of all such channels for the green or infrared filters (see section 3.2.4).
[b]See section 4.2.4 for details.
[c]Shaded cells for Filters L2 and R2 highlight the fact that the sky model-derived radiometric calibration coefficients for these two filters are likely to be much more accurate than the preflight-derived coefficients in Table 4. See section 4.2.4 for details.

(Table 11). There was no bias toward this: the normalization removed any input from the laboratory absolute radiance calibration, and the process was done "blind" to the normalization constant. This modeling work thus provides an independent validation of the required better than ±10% accuracy of most of the Version 2 radiance calibration coefficients listed in Table 4. The exceptions are (a) several near-IR filters (L6, R3, and R4) for which the differences between the preflight and sky model coefficients are consistent with better than ±10% to ±15% accuracy and (b) the narrowband blue filters, L2 and R2 (near 445 nm), which exhibit significantly larger differences in derived radiance coefficients. The larger L2 and R2 uncertainties are consistent with the significantly higher uncertainties in their laboratory radiance calibration coefficients (Table 4), because of their relatively poorer fit to the calibration lamp spectrum (Figure 10) and the lower SNR of the data acquired for their spectral throughput characterization (see section 3.2.4). Therefore, because of the much better agreement between the laboratory preflight and independent in-flight sky model-derived radiometric calibration coefficients for all of the other Mastcam filters, the sky model-derived radiometric calibration coefficients for the L2 and R2 filters in Table 11 are most likely better values to use than those in Table 4 to estimate the absolute radiances of images acquired through those two specific filters.

4.2.5. Focus Testing/Performance

Two sets of tests were conducted to characterize and validate the focus calibration of the Mastcams in flight. First, a series of measurements taken in coordination with the MSL ChemCam instrument were used to refine the preflight calibration of focus motor count versus distance to target (section 3.2.7) over the range of temperatures and other environmental conditions actually experienced on Mars. Then, a series of additional Mastcam images of the same targets over a variety of temperatures typical of flight operations was obtained in order to characterize small deviations in the geometric (camera model) parameters of the cameras over that range of conditions.

4.2.5.1. Focus Versus Motor Count Validation

In order to validate the preflight calibration of Mastcam focus motor count versus distance under Mars conditions, as well as to provide end users a way to estimate distance to imaged targets (and thus their

spatial scale), analysis was conducted on 271 sets of M-100 images and 93 sets of M-34 images for which nearly concurrent distance data were obtained by the ChemCam instrument's laser range finder [*Wiens et al*., 2012]. For the M-100 camera, we found a strong linear fit ($R^2 > 0.993$) between best focus position identified by the camera's onboard autofocus algorithm and 1/distance as determined by ChemCam (the few centimeters difference in distance to the scene between the ChemCam optics and the Mastcam optics (Figure 4) is a small and relatively insignificant fraction of the total distance). While the difference between observed and predicted focus motor counts in this fit is within ±15 focus motor counts (2% of the typical focus count range) for more than half the data analyzed, deviations in the fit of up to ±60 focus motor counts were noted for the rest of the images. Those deviations were found to have a good linear correlation ($R^2 > 0.790$) with the temperature of the M-100 camera (as recorded in the OPTICS_TEMP keyword value in the PDS archive labels). Thus, a temperature-dependent correction was generated for the M-100 camera. For the M-34 camera, the linear fit between best focus and 1/distance was also highly correlated ($R^2 > 0.982$). Deviations with temperature for the M-34 focus calibration were found to span about ±15 focus motor counts for all of the images but were not observed to correlate with the temperature of the camera (consistent with the wider-angle M-34 camera having a substantially lesser sensitivity to focus than the narrow-angle M-100).

The resulting refined in-flight calibrations for the relationship between focus motor count and distance to target were found to be

$$M-34: D = 363.64/(2427.50 - F_{34}) \tag{9}$$

and

$$M-100: \quad D = 3322.3/(3491.9 - 2.58T - F_{100}) \tag{10}$$

where D is the distance to the focused image target in meters, T is temperature of the camera in °C, and F_{34} and F_{100} are the focus motor counts commanded to the M-34 and M-100 cameras, respectively (and stored in the FOCUS_POSITION_COUNT keyword in the archived PDS labels).

These models do an excellent job of matching the "true" distance to the target scene in more than 90% of the cases examined to date. Most of the cases where the model is not giving as good a fit are for the M-100, where the actual temperature of the camera could potentially be more than ±5°C different from the reported optics temperature, especially during active heating (the M-100 temperature sensor is on an external heater; see section 4.2.1). Another possible cause of distance discrepancies is that many of the targets were rocks that do not fill the field of view. Experience using the cameras on Mars shows that the combination of partial field-filing rocks and surrounding materials creates additional scene entropy that can cause the autofocus algorithm (the same as used on the MAHLI camera; see sections 7.2.2 and 7.2.3 in *Edgett et al*. [2012] for details) to "pull" back to lower focus motor counts relative to other kinds of scenes.

Finally, chromatic aberration in the Mastcam optics in the near-IR could also cause some discrepancies in estimated distances to imaged targets. Typically, autofocus is performed on images through the L0 and R0 filters (at visible R, G, and B wavelengths; Table 2) and not at other wavelengths that might be acquired in the same imaging sequence. Specific in-flight tests of autofocus at other wavelengths, however, performed on sols 468, 475, 488, and 493 (sequences mcam01864, mcam01888, mcam01934, and mcam01960) of targets 70, 60, 2.4, and 10 m away, respectively, showed that best focus was achieved 25 focus motor counts lower in the M-34 L6 filter (1012 nm) and 52 motor counts lower in the M-100 R6 filter (1013 nm) than for RGB images acquired in the same sequence. While the slightly defocused nature of the near-IR images in Mastcam multispectral data sets has not proven to be an impediment to scientific analyses (mostly because spectra are typically extracted from groups of tens to hundreds of pixels, rather than single pixels; see section 5.4.4), users interested in quickly estimating distances/spatial scales of targets in multispectral sequences should default to using equations (9) and (10) with the recorded FOCUS_POSITION_COUNT values from the associated RGB filter L0 or R0 images also always acquired in such sequences.

4.2.5.2. Thermal Focus Test

The so-called "thermal focus test" images were obtained in flight on Mars during *Curiosity* sols 917 and 918 (6 and 7 March 2015) to characterize the effects of temperature on the position of the focus mechanism for

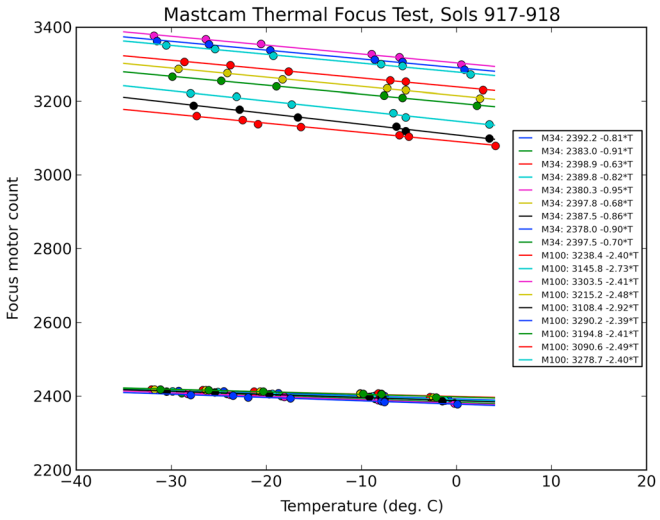

Figure 17. Dependence of focus motor count on Mastcam focal plane array temperature for the M-34 and M-100 Mastcams, based on six in-flight mosaic sequences acquired on Mars on *Curiosity* sols 917 and 918.

focused Mastcam images. The test was motivated by slight differences between the laboratory-calibrated focus motor count values and the in-flight focus motor count values for any given distance. The proposed hypothesis was that the differences in motor positions were correlated to differences in the temperature of the camera head and that thermal contraction of the camera head at the colder temperatures in flight was causing the focus position to shift slightly.

The image observations consisted of a 3 × 3 raster of losslessly compressed images acquired with both the M-34 and M-100 cameras of the same target region (identical RSM pointing angles) at several different times of day, refocused using autofocus for each image in the mosaic. The rationale behind the 3 × 3 raster was that it allowed statistical testing of the results. The different times of day were chosen to capture a range of diurnal temperatures at roughly 10°C intervals between about −30°C and 0°C. A high-entropy scene was selected for the test, located ~15 m (±2 m) due south of the rover so that the shadows were symmetric across the day. The motor count values from each of the autofocus images were used to characterize the change in focus as a function of temperature (Figure 17). The results show that the focus changes by ~0.8 motor count/°C for the M-34 and ~2.5 motor counts/°C for the M-100.

4.2.6. Compression Testing/Performance

Most Mastcam images have been returned from Mars using some level of lossy JPEG compression (for details, see both Appendix E and section 4.4.5.4 in *Malin et al.* [2013]). Several studies have been conducted by the Mastcam operations team using flight Mastcam data to assess the performance of the compressor and its influence on the returned science data. For example, one study determined the average compressed data volume of JPEG-compressed M-34 and M-100 images of hundreds of representative Martian scenes as a function of JPEG quality factor (ranging from 0 to 100) [see *Pennebaker and Mitchell*, 1992]. Scene activity, image dimensions (depth of field), and focus quality were all found to influence compression efficiency to different degrees. The results (Figure 18, which also shows compression study results from the MAHLI and MARDI cameras) show that images from the M-100 camera generally compress better than those from the M-34 camera for images commanded to the same JPEG quality factor. The 0.5 to 1.0 bit per pixel increase in compressibility of M-100 images compared to M-34 is likely due to the much larger depth of field of the M-34 camera. An important operational implication from Figure 18 is that JPEG compression of M-34 or M-100 images using high-quality factors (e.g., ≥90) can reduce downlinked data volume by up to a factor of 2 compared to lossless compression. And as shown in section 5.4.1 below, the difference in image quality between losslessly compressed images and those JPEG compressed to high-quality factors is usually imperceptible to the human eye.

Another compression test examined the influence of JPEG quality factor on the extraction of 12-color Mastcam spectra of various regions of interest. For example, lossless Mastcam multispectral images from the sol 281 drill hole imaging campaign in "Yellowknife Bay" [*Grotzinger et al.*, 2014] were recompressed after downlink using a copy of the onboard JPEG compression algorithm and converted to compressed images at JPEG quality factors of 85 and 45. Mastcam data from bright reddish dusty surfaces, darker reddish sandy surfaces, and gray drill tailings were extracted from the same multipixel regions of interest (ROIs) at all compression levels. The averages and standard deviations of the resulting 12-color Mastcam spectra of the regions are virtually identical in the lossless and quality 85 and 45 data sets. This is perhaps not surprising, as typical ROIs were extracted from large pixel regions compared to the 8 × 8 pixel compression block

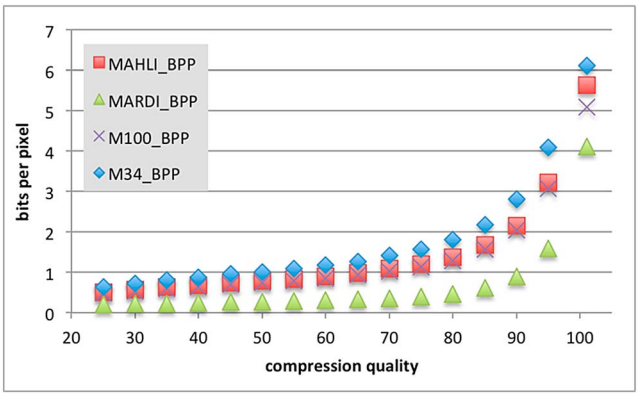

Figure 18. (a) Average compression performance of the Mastcam JPEG compression algorithm [*Malin et al*., 2013] as a function of JPEG quality factor, for hundreds of representative Mars images from the four-color cameras on the *Curiosity* rover. "Bits per pixel" on the *y* axis is relative to the original downlinked 8 bit companded data. Lower *y* axis values mean more compressible images (e.g., 4 bits per pixel is a compression factor of 2:1 and 1 bit per pixel is a compression factor of 8:1). Values at quality factor 101 correspond to losslessly compressed images. While the difference in downlinked bits is large between losslessly compressed images and those that are JPEG compressed to high-quality factors, the difference in image quality is often imperceptible to the human eye (however, see also Figure 29).

of the JPEG algorithm. While large differences in pixel-to-pixel brightnesses can easily be visually detected within small (comparable to the 8 × 8 pixel compression blocks) groups of pixels when images of different JPEG quality factors are examined at high zoom, the JPEG compressor does a very good job of preserving the overall average signal level of the scene when integrated over larger pixel regions. Thus, in practice, care should always be used to extract Mastcam spectra from ROIs that are large compared to the 8 × 8 pixel JPEG compression blocks, to avoid incorrect interpretation of smaller-scale compression artifacts.

4.2.7. In-Flight Scattered Light Testing

Scattered (from in-field sources) light within the Mastcams is generally difficult to quantify during Mars surface operations. However, imaging of the solar disk using the neutral density solar filters provides a special opportunity to characterize scattered light effects, at least for solar imaging data sets. For example, on *Curiosity* sol 33, a series of images of the Sun's disk were taken through both the M-34 and M-100 cameras (sequences mcam00008 and mcam00009). Images of the Sun through the M-100 camera's filter R7 (10^{-5} neutral density filter plus 440 ± 20 nm narrowband filter) show several low-level (≤5% intensity) ghost images of the Sun and a very high level (100% to 300% of the primary solar image's signal level) ghost image associated with a pinhole filter leak at the bottom of the field of view. Fortunately, all of the ghost/leak images are offset substantially (many hundreds of pixels) from the primary image of the Sun, and their signals show no influence on the signal levels of the direct solar image itself. Images through the M-34 camera's Filter L7 (10^{-5} neutral density filter plus 880 ± 10 nm narrowband filter) do not show any similar ghosts or apparent filter leaks, but they do reveal a small "halo" of scattered light surrounding the Sun's disk, at an intensity of ~1% of the Sun's DN level. Because of their relatively low signal levels and/or their wide spatial separation from the direct images of the Sun, however, neither this level of scattered light in the M-34 solar images nor the large ghost images and apparent filter leaks seen in the M-100 solar image data have impeded the ability to derive accurate estimates of the Martian atmospheric opacity using the Mastcams [e.g., *Lemmon*, 2014].

4.3. Mastcam Calibration Target Imaging

As described in section 2.2, Mastcam includes an external calibration target mounted on the top of the RPFA box on the rover deck (Figure 4). Imaging of the target enables a way to quickly calibrate other sequences to a quantitative estimate of radiance factor (*I/F*; see section 5.2.7), a quantity directly comparable to laboratory reflectance spectra of rocks and minerals. Perhaps most importantly, because the target is being illuminated by both direct sunlight and indirect diffuse (and significantly reddened) sky radiance, calibration of the scenery around the rover relative to the calibration target can effectively remove almost all of the spectral reddening effects of the diffuse illumination in the scene, for scenes that are observed at comparable solar incidence angles as the calibration target [e.g., *Bell et al*., 2006]. This enables tactical-timescale quick-look assessment of the reflectance properties of the scene without having to perform detailed (and time-consuming) radiative transfer modeling of the atmospheric component of the observed scene radiance. A comparison of in-flight Mastcam spectra of the calibration target to preflight high-spectral resolution laboratory spectra of the calibration materials by *Wellington et al*. [2017] shows that the multispectral data calibration pipeline described below (section 5.3.6) can reproduce very well the actual spectral properties

of the target materials, increasing our confidence in the pipeline's ability to accurately represent the spectral properties of Martian scenes being imaged.

The Mastcam calibration target is not usually imaged when other sequences are only being acquired through the RGB Bayer filters (filter position 0 in both cameras), because quantitative reflectance assessment of the scene is typically not the goal of RGB-only imaging. Conversely, the calibration target is almost always imaged close in time to multispectral sequences that observe the scene using the narrowband science filters (Table 2) in addition to the RGB Bayer filters. As a general rule, the calibration target is imaged at the same approximate Local True Solar Time (LTST; within about an hour before or after) and with the same filter set as the to-be-calibrated multispectral sequence with which it is associated. If time, power, or data volume are heavily constrained, however, a new calibration target sequence may not have been taken if one already exists from a previous recent sol (within about 5 sols) that was taken at a similar LTST (again, within about 1 h of that time). To save data volume and avoid imaging of unnecessary other parts of the rover deck, downlinked images of the calibration target since *Curiosity* sol 66 (sequence mcam00297) have been subframed to 384 × 320 pixels for the M-34 camera (starting at CCD row 721 and column 625) and 1152 × 944 pixels for the M-100 (starting at CCD row 33 and column 305).

Starting on *Curiosity* sol 3 and extending to sol 1159 (the most recent PDS data release as of this writing), 223 Mastcam imaging sequences of the calibration target were obtained, corresponding to imaging of the target approximately every 5 sols. The total downlinked Mastcam data volume dedicated to calibration target imaging has been less than 0.5% of the total downlinked Mastcam data volume overall. Over time, air fall dust has slowly settled onto the target (Figure 4), decreasing its contrast and necessitating the development of a dust correction model to compensate for the influence of a thin but still semitransparent layer of dust on the color and photometric properties of the calibration standard materials. Section 5.3.6 describes this dust model, as well as the general use of the calibration target in the enhanced Mastcam data calibration pipeline.

5. Data Reduction, Validation, and Archiving

5.1. Introduction and Methodology

Figure 19 provides a simplified overview of the basic steps in the current MSL/Mastcam data reduction and calibration pipeline. Full-frame or subframed (but not thumbnail) images (Table 3) are calibrated to radiance factor (*I/F*; see section 5.2.7) using a variety of preflight and in-flight ancillary data files and images. Section 5.2 describes the initial tactical calibration pipeline being used to generate current PDS-released EDR and RDR archive products from the Mastcam investigation. Section 5.3 describes further planned enhancements to these initial calibrations, in works for future PDS archive releases of the Mastcam data. Section 5.4 describes some examples useful for the validation of the accuracy and precision of the Mastcam calibration. Finally, section 5.5 provides a summary of Mastcam data products and PDS archive products.

5.2. Initial (Tactical) Data Reduction and Calibration Pipeline

An initial version of the calibration pipeline steps in Figure 19 is being used to generate rapid-timescale (approximately daily) versions of calibrated Mastcam RDRs for tactical use by the MSL Science Team, as well as initial Version 1 PDS archive versions of calibrated Mastcam RDRs for use by the broader community. This section describes those initial calibration steps in detail.

5.2.1. Decompression

If the raw EDRs (*.dat files in the PDS archive) were compressed for downlink, then the first step is to decompress the data into the proper spatial domain format. For example, if the data were JPEG color compressed, then the YCrCb coefficients of the 8 × 8 pixel frequency domain JPEG Minimum Coded Units (MCUs) are transformed back into 8 × 8 spatial MCUs and reordered into image arrays for each of the three RGB output color bands, and the output is a 24 bit (three-band sequential channels) color image. For JPEG gray scale compressed data, just the luminance (*Y*) channel is transformed back into a single channel 8 bit gray scale output image. For losslessly compressed data, the Huffman decoding process described in Appendix C of *Malin et al*. [2013] is used to create the decoded images.

An important part of the decompression process for Mastcam images that have been JPEG compressed is performing a Bayer pattern interpolation (also known as "debayering" or in some contexts "demosaicing")

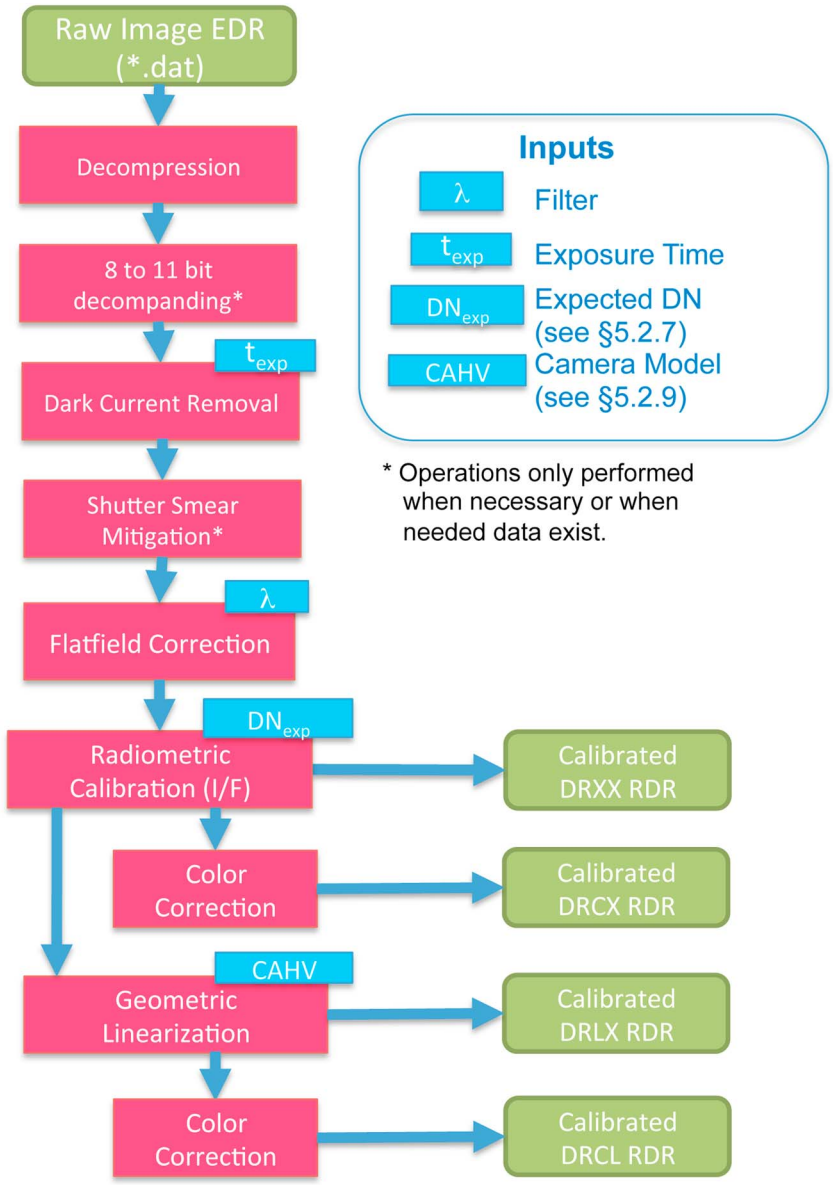

Figure 19. Flow chart of MSL/Mastcam tactical data calibration steps, for initial data sets being archived in the PDS.

to create the separate full-sized RGB channels of the output decompressed JPEG. A raw image from a sensor with a Bayer filter often results in what looks like high-frequency noise in a zoomed out image and while zoomed in may look like a discontinuous checkerboard-like pattern (Figure 20). This appearance is due to the interleaving of the red, green, and blue channels in the Bayer pattern (e.g., Figure 2). For typical viewing, these raw sensor data are passed through an interpolation algorithm that computes the missing Bayer colors for each given pixel of the R, G, and B channels of the image. For instance, where the sensor has a red filter, the interpolation algorithm will compute the data for the same pixel in the corresponding green and blue channels using information from neighboring pixels. All Mastcam images are acquired through the Bayer color filter array pattern, but that pattern can be (optionally) interpolated either within the camera electronics or using ground processing after downlink. For normal RGB imaging through the broadband Filter 0 (Table 2 and Figure 3), the interpolation method employed within the instrument to produce color JPEG products is a hardware implementation of the algorithm developed by *Malvar et al.* [2004]. Section 5.3.1 below compares the default debayering algorithm for Mastcam data to several alternatives.

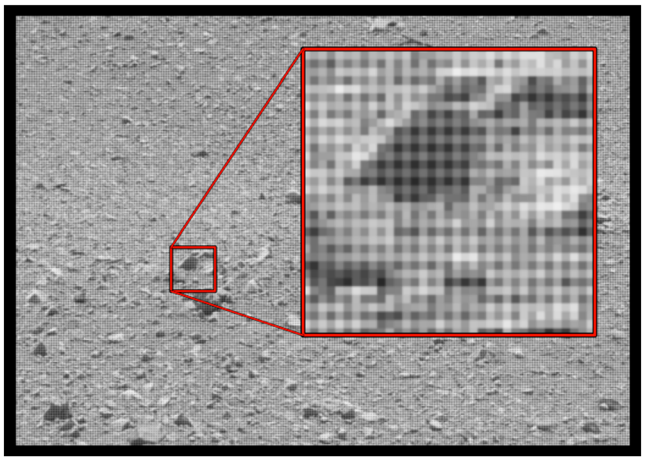

Figure 20. Raw Mastcam image downlinked from Mars, exhibiting apparent high-frequency noise from the Bayer pattern (Figure 2). Inset: Magnified image to show the Bayer pattern more clearly.

Mastcam supports losslessly compressed (first-difference Huffman encoding) image acquisition modes that return raw or near-raw image data that contain the Bayer pattern. Lossless or completely uncompressed non-thumbnail EDR product types A, B, C, J, and K (Table 3) contain the Bayer color filter array pattern and can be interpolated using ground processing. Data compressed with lossy (JPEG) compression modes return images that have been interpolated on board the rover, prior to downlink. Lossy JPEG-compressed images (EDR product types D, E, and F; Table 3) are interpolated based on how well the band pass of the commanded science filter (Filters 1–7, if used) overlaps the red, green, or blue Bayer filter band passes. Specifically, kernels are applied to the science filter data prior to compression either to discard two of the three Bayer colors that are outside the science filter band pass or to enable the use of all three Bayer colors for science filter wavelengths beyond which the Bayer filters are transparent (see section 3.2.4, Figure 3, and Table 2). Table 12 describes which interpolation scheme is used for which camera and filter combination when lossy JPEG compression is used to downlink images from Mars.

Bayer interpolation of uncompressed or losslessly compressed images (EDR product types A, B, and C; Table 3) is accomplished after downlink either real time in analysis tools used by the MSL Science Team or as the first step ("Decompression") in the generation of calibrated RDR and PDS archive products (Figure 19).

5.2.2. Companding and Decompanding

"Companding" is a portmanteau word blend of the words "compressing" and "expanding" and refers to the process of compressing the original 11 bit (0–2047) DN values of each raw Mastcam pixel down to 8 bits (0–255) of dynamic range. The process is sometimes also referred to as resampling [e.g., *Bell et al*., 2006]. Companding is performed because it is usually desirable to scale the data down to a smaller number of bits per pixel so that Poisson (shot) noise is not encoded or downlinked in the telemetry. Because Poisson noise in detectors like CCDs is proportional to the square root of the number of electrons detected, using a square root-based look-up table (LUT) to scale the original 11 bit data down to 8 bits of dynamic range provides a way to decrease the number of bits downlinked without incurring a statistically significant loss of information (i.e., the noise is not quantized). For example, an original 11 bit Mastcam DN value of 1700 corresponds to a signal of ~27,200 e^- (Table 1), which has a Poisson (counting) noise of $\sqrt{27,200}$ or ~165 e^-, or ~10.3 DN. Thus, there is no statistical difference (within the noise of the measurement) between Mastcam DN values of 1695, 1700, or 1705 DN. We therefore "compand" the original DN value of 1700 DN into the 8 bit value of 232. After downlink, the 8 bit value of 232 is "decompanded" back to the 11 bit value of 1698, which is statistically identical to the original measured 11 bit value of 1700 DN. Because of the square root nature

Table 12. Bayer Pattern Interpolation Scheme Used for Lossy JPEG-Compressed Mastcam Data

Camera	Filter 0	Filter 1	Filter 2	Filter 3	Filter 4	Filter 5	Filter 6	Filter 7
M-34 (left)	Malvar[a]	Green[b]	Blue[c]	Red[d]	Red	Identity[e]	Identity	Identity
M-100 (right)	Malvar	Green	Blue	Red	Identity	Identity	Identity	Blue

[a]Malvar means that interpolation using the algorithms of *Malvar et al*. [2004] is performed.
[b]Green means that bilinear interpolation of green Bayer pixels is performed; red and blue pixels are discarded.
[c]Blue means that bilinear interpolation of blue Bayer pixels is performed; red and green pixels are discarded.
[d]Red means that bilinear interpolation of red Bayer pixels is performed; blue and green pixels are discarded.
[e]Identity means that no interpolation is performed; image is returned as a monochrome JPEG that was compressed from raw data with as is Bayer values (because the Bayer filters are transparent at near-IR wavelengths; see Figure 3).

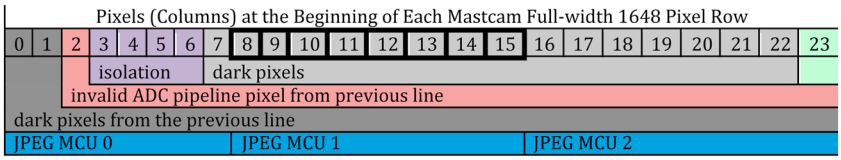

Figure 21. Layout of the first 24 columns of full-width Mastcam images, showing the details of the first 23 masked (dark) pixels in columns 0–22 and the first photoactive pixel (column 23). The JPEG compressor's Minimum Coded Units (MCUs) are also indicated. See text for details, and see also Figure 2.

of Poisson noise, there is an approximately 1-to-1 mapping of 11 bit values to 8 bit values for low DN values and a many-to-1 mapping of 11 bit values to 8 bit values for high DN values. Almost all Mastcam images have been companded on the rover (within the Mastcam DEA) using the square root-based LUT provided here in Appendix B. Details on the 31 other potential Mastcam LUTs are provided in Appendix B of *Malin et al*., 2013. The opposite process, expanding the downlinked 8 bit data back to an estimate of their original 11 bit value, is referred to here as "decompanding" and is essentially part of the decompression process in the pipeline processing of Mastcam images.

5.2.3. Dark Current Modeling and Removal

Under normal conditions, the Mastcam cameras generate only a small amount of dark current. A temperature-dependent model of dark current was generated during thermal vacuum testing, and this model can be applied routinely (section 3.2.2 and Figure 8). A biased offset is also built in to the Mastcam signal chain, to provide room to accommodate variations in electronic behavior of the detector. Estimates of the bias and dark current can be made by using masked (dark) columns on the detector (Figure 2) to extract dark current values or by special bias and dark current full-frame images acquired occasionally in flight (section 4.2.1). There are three ways currently implemented in the Mastcam calibration pipeline to remove the effects of bias and dark current: (a) using masked (dark) pixels, if they were downlinked with the images to be calibrated; (b) using the preflight model of bias and dark current versus FPA temperature; or (c) "manual" bias and dark current correction using analogous measurements from images acquired under similar circumstances.

5.2.3.1. Using Dark Columns

Some rows, columns of pixels along the edges of the detector are masked (Figure 2), specifically corresponding to columns 0–22 and 1631–1647 in full-width 1648 pixel Mastcam images. These pixels have the same bias level and accumulate dark current during an image exposure in the same manner, as photoactive pixels. The detailed layout of the first 24 columns of full-width Mastcam images is provided in Figure 21. Note that JPEG Minimum Coded Units (MCU) for lossy compressed data occur in blocks of 8×8 pixels. For images compressed as JPEGs, the dark pixel in column 7 is compressed with nondark pixels from MCU 0. For this reason, only dark columns 8–15, corresponding to JPEG MCU 1, are used by Mastcam RDR processing. Thumbnail processing also takes into account the compression applied. Specifically, thumbnails effectively average all eight of the JPEG MCU 1 pixels into a single thumbnail pixel. Thus, only the second thumbnail pixel column is used for bias and dark current subtraction, for thumbnails generated from full-frame Mastcam images.

For consistency, all images, even if not JPEG compressed, use the same group of 8 pixels, from dark columns 8–15 in full-frame Mastcam images, or column 1 (second column) in full-frame thumbnail image, for dark

Table 13. Five Kinds of Mastcam Data Products Being Archived in the NASA PDS

Processing Code[a]	D = Decompressed?	R = Radiometrically Corrected?	C = Color Corrected?	L = Linearized?	CODMAC Level[b]
XXXX	no	no	no	no	2
DRXX	yes	yes	no	no	4
DRCX	yes	yes	yes	no	4
DRLX	yes	yes	no	yes	4
DRCL	yes	yes	yes	yes	4

[a]Each letter in the code corresponds to the action in the next four columns. An "X" in the processing code corresponds to "no" for that action. See Appendix D.
[b]See Appendix C.

correction. The average of these 8 pixels (or just the second column, if the image is a thumbnail) along the entire height of the image is used to estimate the combined bias and dark current signal, and this average (recorded in the DARK_LEVEL_CORRECTION processing parameter keyword in the archived PDS labels) is then subtracted from all pixels in the original image. If the image is a full-height (1200 rows) image or a thumbnail, the first two and last two lines (rows) of the data are not included in the average. Subframed images might not have dark columns included if the first column is >15 in the original full-frame image coordinates.

5.2.3.2. Using Preflight Bias and Dark Current Model

Preflight calibration data measured and modeled both bias and dark current charge accumulation as a function of temperature (section 3.2.2 and Figure 8). By dividing by the exposure time of the image to be calibrated, and using the known or inferred temperature of the CCD at the time of image acquisition and the model described in section 3.2.2 and Figure 8, the dark current accumulation in DN/s can be estimated and subtracted. The bias or DC offset component of the signal is assumed to be constant based on preflight data analysis and the typical temperatures of operation of the focal plane on Mars (Figures 7 and 8). The bias level is recorded in the DARK_LEVEL_CORRECTION keyword in the archived Mastcam PDS data label (section 5.5) and has a typical value of 117 DN in Mastcam flight images acquired on Mars. This (commandable) bias level is usually subtracted from the data automatically by the DEA prior to companding, JPEG compression, and downlink of the data.

5.2.3.3. Manual Bias and Dark Current Subtraction

Bias, at least, can always be subtracted using the value of the DARK_LEVEL_CORRECTION keyword in the archived Mastcam PDS data labels. If the dark current level cannot be estimated, however, either because dark pixels or focal plane array temperature data are not available for the specific image in question, then a more manual approach to dark current removal is attempted by the data reduction pipeline. This involves using the dark current estimated either from an image close in time from the other Mastcam camera (if that other camera was also active simultaneously or close in time, and the images from that camera contain the required dark pixel or temperature data) or from an image close in time from the same camera that *does* contain the required dark pixel or temperature data. "Close in time" is rather subjective, but a general metric, based on analysis of data in flight, is that it should be less than about 30 min so that the temperature of the analogous observation is likely to be close to that of the observation that it is designed to calibrate. If dark current cannot be estimated manually, then only bias is subtracted from the image.

5.2.4. Electronic Shutter Removal

The effects of CCD electronic shutter smear (section 3.2.3) are not currently being removed from any of the initial PDS archived Mastcam images. Future enhancements to the pipeline that could accommodate an a posteriori smear correction in the flight data are described in section 5.3.3.

5.2.5. Flatfielding

As described in section 3.2.5, the brightness response recorded by the Mastcam detectors is not uniform as a function of position within the fields of view. In addition to pixel-to-pixel responsivity variations of the CCD, there are several geometric effects that influence the responsivity across the scene. For example, the Mastcam full fields of view are vignetted in their corners by a filter wheel mask (Figure 9; the intention of the design was to utilize just the central 1200 × 1200 pixel science imaging field of the array (Figure 2), although in practice many times the full 1648 × 1200 span of the array is employed when a more rigorous quantitative radiance calibration is not required, in order to cover more of the scene in fewer mast pointing positions), and the Mastcams also display some structure under uniform illumination resulting from small internal reflections (e.g., section 3.2.5 and Figure 9). Laboratory measurements during calibration (section 3.2.5) provide measures of the nonuniformity of the cameras' responses. These calibrations are used to smooth out brightness variations in areas of lower or higher response. Normalized flatfield reference images for every Bayer filter and science filter combination (except for the solar/ND Filters L7 and R7, for which no flatfield calibration data exist) were derived from preflight calibration data (section 3.2.5) and are being used for current PDS-archived calibrated Mastcam images. These flatfield files are also archived in the PDS (for example, at http://pds-imaging.jpl.nasa.gov/data/msl/MSLMST_0001/CALIB). Mathematically, the flatfield files are stored as 1/flat, so correction is multiplicative during processing in the data reduction pipeline. If the image is a sensor subset, then the flatfield reference file is subframed accordingly, and if the image is a thumbnail, then a downsampled flatfield that is 1/64 of the full-resolution flatfield is used instead.

5.2.6. Correction of Dead/Bad Pixels

The initial Mastcam data reduction and calibration pipeline for PDS archive products do not perform any corrections for known anomalous (dead or otherwise bad) pixels on the CCDs (section 3.2.10), as the correction provided by the multiplicative flatfield is deemed adequate, and the number of bad pixels is so small. Future enhancements to the pipeline that will accommodate these cases are described in section 5.3.2.

5.2.7. Initial Radiometric Calibration

Decompanded, bias and dark current subtracted, and flatfielded Mastcam images are converted to an initial estimate of radiance factor or *I/F*, where *I* is equal to the measured scene radiance and πF is equal to the solar irradiance at the top of the Martian atmosphere at the time of the observation, convolved to the particular Mastcam band pass being calibrated. Mathematically, the calibration to *I/F* in the initial PDS-archived "DRXX" format RDRs (Table 12 and Appendix D) is defined as the ratio of the observed calibrated DN level (DN_{obs}) to the expected DN level (DN_{exp}) that would be produced by imaging a perfectly diffuse white surface, illuminated by sunlight, at the heliocentric distance of Mars, with no atmospheric attenuation, at zero incidence angle, and with a reference exposure time of 10 ms. Reference solar DN levels (F_{ref}) at the perihelion distance of Mars (1.38 AU) are listed in Table 2. Note that the narrowband signal levels are computed for the dominant Bayer filter color in cases where the Bayer filters are not uniformly transmitting (Figure 3 and Table 2). For example, for the L1 and R1 527 nm filters, the green Bayer position is used. To compute DN_{exp} for an image with an exposure time of t_{exp} ms and acquired at a solar distance of d_{Sun} AU, we use the following expression:

$$DN_{exp} = F_{ref} \cdot (t_{exp}/10) \cdot (1.38/d_{Sun})^2 \tag{11}$$

The derived *I/F* values archived in the DRXX format Mastcam RDRs are then simply calculated as DN_{obs}/DN_{exp}. Archived Mastcam *I/F* calibrated data are stored in the PDS as 16 bit integers in the DRXX format files (Table 13); conversion to floating point *I/F* values can be performed by multiplying the stored 16 bit values by the constants in the PDS Label keyword "RADIANCE_SCALING_FACTOR" and then adding the constants in the PDS label keyword "RADIANCE_OFFSET."

Users who want to derive an estimate of the absolute spectral radiance of the Version 1 calibrated Mastcam images (I_{obs}, for example, in $W/m^2/nm/sr$) can simply multiply the *I/F* values archived in the PDS calibrated images [$(I/F)_{PDS}$] by the weighted value of the solar spectral irradiance in that band pass at the top of the Martian atmosphere at Mars perihelion (F_{Sun}), scaled for the heliocentric distance d_{Sun} in AU at the time of the observation, and then dividing by π:

$$I_{obs} = (I/F)_{PDS} \cdot F_{Sun} \cdot \left[(1.38/d_{Sun})^2\right]/\pi \tag{12}$$

Reference solar spectral irradiance values (F_{Sun}) at the perihelion distance of Mars for each Mastcam band pass are also listed in Table 2.

5.2.8. Color Correction

To approximate a "white balanced" color view of calibrated RGB Mastcam images, they are linearly scaled by a set of coefficients developed from preflight testing. Specifically, the new approximately white balanced R'G'B' color space is defined relative to the original calibrated RGB color space as

$$\begin{bmatrix} R' \\ G' \\ B' \end{bmatrix} = \begin{bmatrix} 1.20 & 0 & 0 \\ 0 & 1.00 & 0 \\ 0 & 0 & 1.26 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix} \tag{13}$$

These coefficients were derived from preflight Mastcam imaging of a Macbeth color calibration target under terrestrial solar illumination at approximately solar noon. The transformation is performed on the calibrated 11 bit pixel values, which are then companded to 8 bits using the companding look-up table 0 defined in Appendix B, to create the "DRCX" and DRCL format Mastcam RDRs archived in the PDS (Table 13).

5.2.9. Geometric Linearization for Mosaicking and Map Projection

A *linearized* image is one from which the effects of lens distortion have been removed and slight deviations of each pixel from square have been adjusted. The processing involves warping the image pixels to show how the scene would appear if imaged by an ideal camera that has no distortion. The resulting image can be modeled using a simple pinhole camera model or, equivalently, a CAHV or CAHVOR camera model.

The four-vector Mastcam CAHV camera model described above (section 3.3.1) is directly applied to Mastcam images to produce the geometrically linearized RDR products presently being archived in the PDS. The distortion is removed by inverting the nonlinear equation using an iterative procedure such as Newton's method. First, we determine the size of the linearized image by projecting points along the edge (we use the four corners and four midpoints) onto the focal plane. The limits of the projection are used to set the frame size for the new image. A detector is defined having square pixels with the nominal pixel pitch and principal point as the original image. The value for each pixel in the linearized image is calculated by projecting the center of each pixel onto the focal plane, applying the radial distortion model, and transferring the position to a fractional pixel location in the original image. Bicubic interpolation is then used to calculate the actual value. Some of the pixels in the linearized image will project outside of the original image and must be given a missing data value. This value is defined by the MISSING_CONSTANT keyword in the archived PDS data product label.

In the associated PDS image label (.LBL) files for linearized DRLX and DRCL RDR products (Table 13), the camera model parameters used to linearize the image are reported as a four-vector CAHV model in the GEOMETRIC_CAMERA_MODEL_PARMS group within the PDS label. For calibrated but nonlinearized archived Mastcam images (in DRXX and DRCX images; Table 13), the full six-vector CAHVOR model parameters described above (section 3.3.2) are reported in the PDS image label so that users can linearize those images themselves using those parameters, if desired.

5.3. Enhancements to the Initial Mastcam Data Reduction and Calibration Pipeline

Figure 22 describes a modified version of the initial tactical-timescale Mastcam data calibration flow chart that takes advantage of new test and calibration data acquired in flight on Mars, as well as more detailed analyses of preflight calibration data sets.

5.3.1. Improved Bayer Pattern Decompression/Interpolation Methods

Raw or losslessly compressed Mastcam images stored in the DEA are not debayered on board the rover, and thus, they still contain the RGB mosaic pattern generated from the Bayer color filter array built onto the sensor (Figure 20). This means that for data downlinked in raw or lossless form, users have the ability to choose which debayering algorithm to use during calibration and analysis. As described in section 5.2.1 above, the default debayering method to produce interpolated color JPEG products archived in the PDS is the algorithm developed by *Malvar et al*. [2004], known as Improved Linear Interpolation (ILI) or Malvar-He-Cutler demosaicing. We have also experimented with simple bilinear interpolation and Directional Linear Minimum Mean Square-Error Estimation (DLMMSE, also known as Zhang-Wu demosaicing [*Zhang and Wu*, 2005]) algorithms. Examples for a typical Mastcam image are shown in Figure 23.

Each demosaicing algorithm has its own benefits as well as its own set of introduced image artifacts. Bilinear interpolation (Figure 23b) looks the worst due to the reduced sharpness of the image, but it is exceptionally fast and thus could be the default choice in interactive tools where demosaicing needs to be done in real time. The default ILI algorithm (Figure 23c) offers improved sharpness and is a good default choice for non-interactive situations to maintain consistency with the compressed JPEG images that have been demosaiced on board the rover. The ILI algorithm, however, can introduce color-related artifacts in some cases. The DLMMSE algorithm (Figure 23d) was introduced as an alternative to ILI to avoid introducing color-related artifacts, although it can introduce its own pixelated or crosshatch-type artifacts in some cases. When interpreting fine-scale image features (such as stratigraphic layering or sharp albedo boundaries) from an imaging sensor with an attached color filter array, it is critical to realize that any demosaicing algorithm can introduce potential artifacts. Evaluating such features with different algorithms is prudent, then, to produce the highest-quality representation of the original scene.

5.3.2. Correction of Saturated or Hot/Bad Pixels

There are many sources of potentially large localized pixel-to-pixel response variations. For example, some pixels may have saturated in the original image acquisition. We define "saturation" in the enhanced

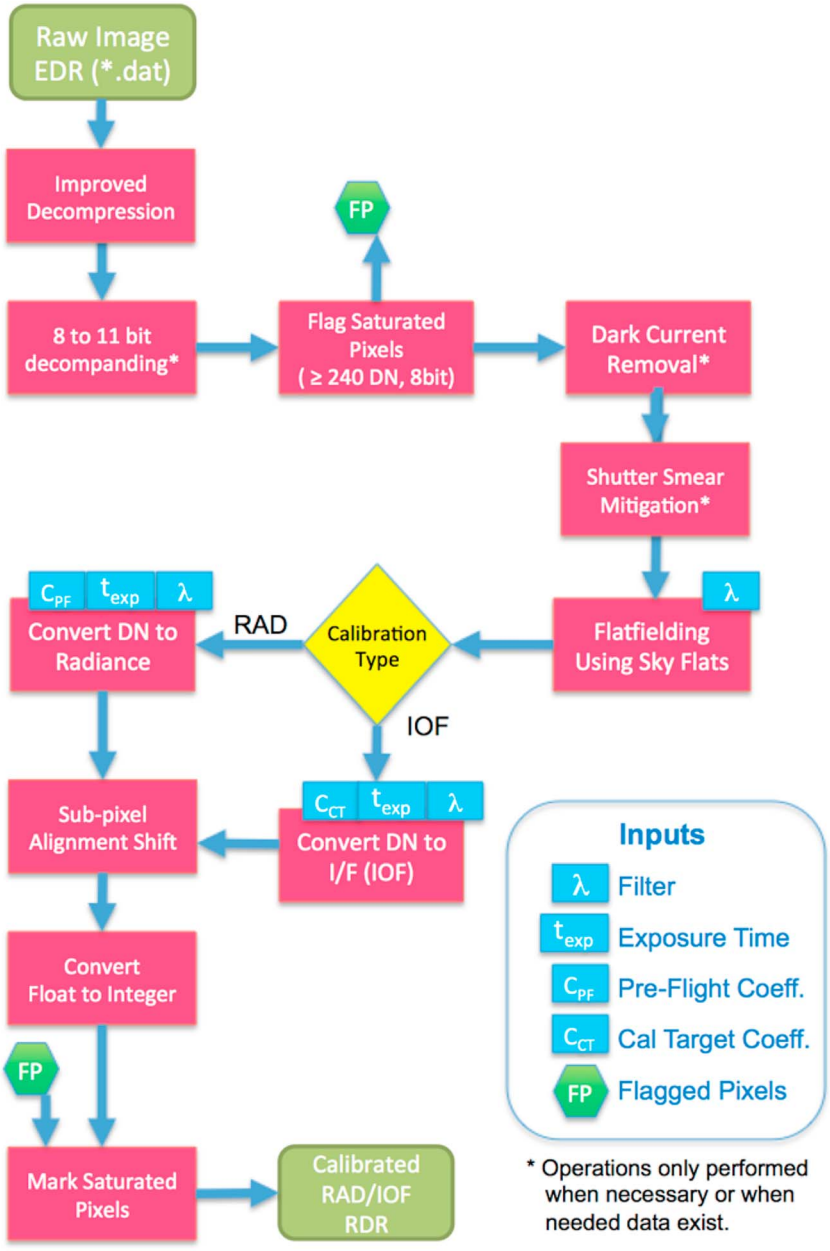

Figure 22. Flow chart describing enhanced MSL/Mastcam data calibration steps.

Mastcam data calibration pipeline as signal levels above the known linear range, which is ~1800 raw DN (240 DN when companded to 8 bit data; see section 3.2.1). When such high pixel values are encountered in raw Mastcam images, we flag them as saturated and replace their value in our pipeline processing with the value of the MISSING_CONSTANT keyword stored in their calibrated file labels. MISSING_CONSTANT is set to a negative value significantly below the minimum value of the valid data in the calibrated data file. Negative pixel values of MISSING_CONSTANT result in easy to identify values to ignore in calibrated data, as negative radiance or *I/F* values otherwise almost never occur within the calibrated Mastcam data set.

Another source of potentially large pixel-to-pixel variations are nonuniformities in the sensitivity of the individual pixel photosites. These were mapped by the detector manufacturer and characterized during calibration (section 3.2.10 and Table 6). These are typically single pixels that are hot (more sensitive and hence brighter) or "cold" (less sensitive and darker than their neighbors). Yet another source of localized

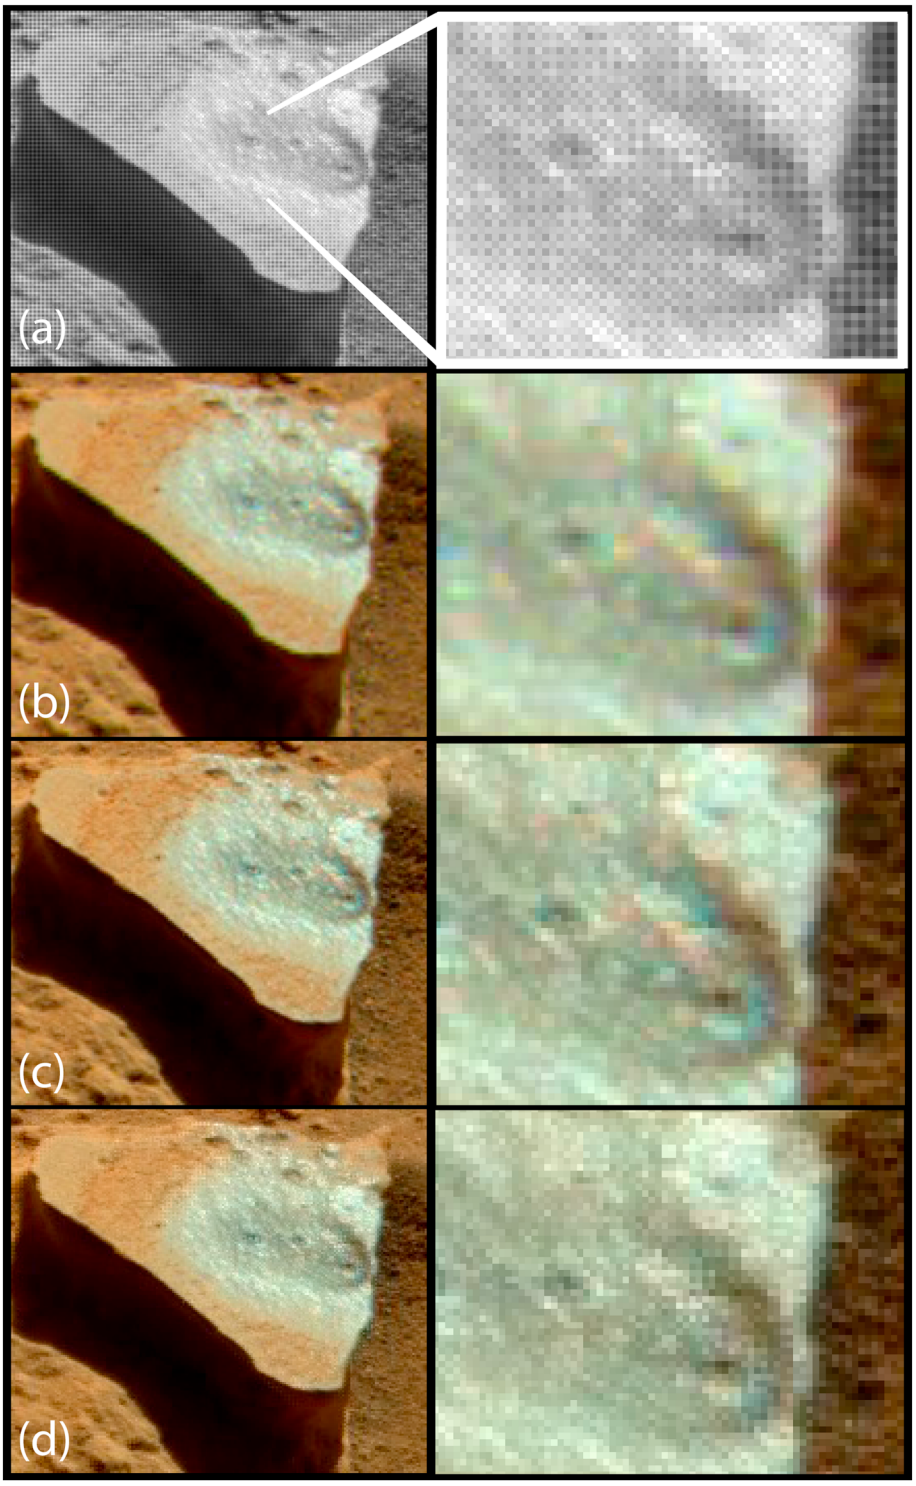

Figure 23. Comparison of (a) raw MSL Mastcam image showing the original Bayer filter pattern, versus application of (b) bilinear interpolation, (c) default Improved Linear Interpolation (ILI) or Malvar-He-Cutler, and (d) Directional Linear Minimum Mean Square-Error Estimation (DLMMSE), or Zhang-Wu, demosaicing algorithms. Portion of Mastcam M-100 image 0620MR0026570000401488C00_DXXX, with zoomed insets on the right.

pixel response variations is contamination on the detector or optics. Such contamination occults light, creating umbral or penumbral shadowing and so-called gray pixel regions, usually a few pixels in size. These were also mapped during calibration (section 3.2.10 and Table 6) and are monitored to make sure they do not migrate with time.

The final potential source of nonuniform pixel response that we consider here is radiation damage, from solar and galactic cosmic rays as well as local sources on the vehicle. Specifically, the MSL spacecraft has two energetic particle sources: the Radioisotope Thermoelectric Generator (a constant source of low flux neutrons) and the Dynamic Albedo of Neutrons experiment that includes an active, pulsed neutron generator. Radiation-induced pixel responses are both brighter and darker than their neighbors. Some of the problem pixels can self-heal, while others become long lived. Dark current images in cruise (section 4.1) and on Mars (section 4.2.1) provide some insight into the nature of such problematic pixels, which can be monitored over time.

While the enhanced calibration pipeline flags and effectively removes saturated pixels from being improperly interpreted in calibrated data, we do not flag or replace hot, cold, or gray pixels. Rather, we monitor them with time and attempt to make data users aware of known problematic pixels in both cameras (Table 6), so as to avoid their improper interpretation. Improvements to the calibration pipeline could involve proactively replacing these pixels with the median of their surroundings, for example.

5.3.3. Improved Electronic Shutter Smear Removal

Currently, the only way to remove the effects of electronic shutter smear from Mastcam images is to acquire and subtract a near-simultaneous zero-second exposure from the nonzero exposure to be corrected. Unfortunately, acquiring such accompanying zero-second exposures is not part of the standard data acquisition sequence for Mastcam imaging (it would double the data volume of an observation because there is no way to do the zero-second subtraction on board), and so this is generally not possible to implement with the flight data from Mars. Fortunately, the effects of shutter smear represent only a small and often insignificant noise source in most Mastcam images (section 3.2.3). If a zero-second exposure was acquired and downlinked immediately before or after a nonzero exposure image of the same scene, then that zero-second image is subtracted from the nonzero image in the enhanced Mastcam data reduction pipeline to completely remove the effect of electronic shutter smear. If a zero-exposure image is not available (as is typically the case), and if the exposure time is less than ~6 ms, then an empirical electronic shutter smear correction algorithm might be needed to estimate and remove the electronic shutter smear from the image. An a posteriori model to remove, or at least to mitigate, the effect of smear from Mastcam images could be devised using a combination of (a) preflight Mastcam smear test images; (b) preflight MAHLI (which uses the same kind of CCD and readout process as the Mastcams [*Edgett et al*., 2012, 2015]) smear calibration observations, which are archived in the NASA/PDS with other archived MAHLI calibration data sets [see *Edgett et al*., 2015]; (c) the in-flight Mastcam shutter smear test images from *Curiosity* sols 36 and 38; and/or (d) a deeper analysis of the specific clocking and physical architecture of the KAI-2020 CCD [*Eastman Kodak Company*, 2009; *Truesense Imaging*, *Inc*., 2012]. For example, a typical deterministic procedure to implement this correction is to linearly subtract a percentage of the scene signal from rows "downstream" in the readout process, to analytically remove the smear component [e.g., *Bell et al*., 2003]. That is, the scene itself can be used to estimate the accumulated effect of shutter smear and to subtract that accumulation from the final image. Future refinements to the Mastcam calibration pipeline will explore implementing such a smear correction model for the small fraction of images that might benefit from this additional correction.

5.3.4. Improved Flatfielding

Occasional daytime observations of the Martian sky in the antisunward direction have been used in flight to acquire additional, higher-fidelity data on the flatfield behavior (and its time variations) of the Mastcam optical system (section 4.2.3). These "sky flat" calibration files are incorporated into the refined Mastcam calibration pipeline (Figure 22), as time-variable updates to the Mastcam flatfield files. For example, a second sky flat data set was acquired on sol 320 (sequence mcam01052, at 14:30 LMST). The images were acquired with the Sun at 64° elevation, centered on the anti-Sun azimuth, and at 30° elevation, during a time when the optical depth was 0.87. The sequence acquired images with the same geometry as the previous sky flat sequences yet reveal small differences compared to the sols 36–38 sky flat images, presumably due to small variations in the pattern and thickness of minor dust

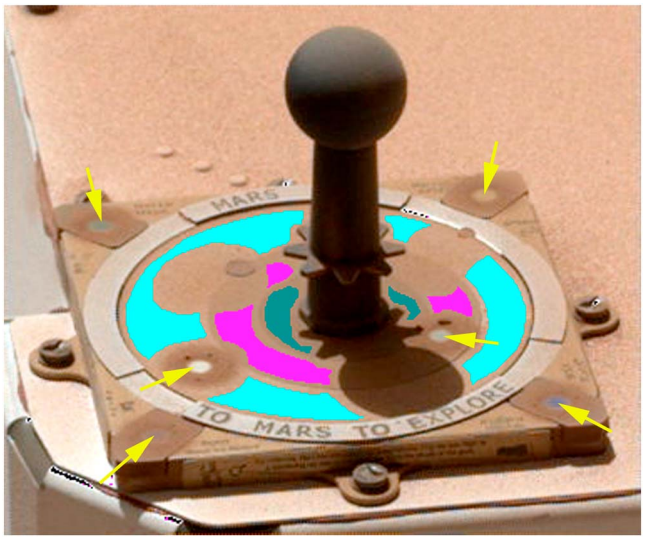

Figure 24. Image of the Mastcam's dusty calibration target acquired around noon local time on sol 514 of the mission (16 January 2014). The image was acquired through the standard Bayer RGB bands of the Mastcam's L0 filter. The effect of the six sweep magnets is apparent with rings of strong dust accumulation around small spots of reduced dust cover (arrows). Drawn on top of the image are selected ROIs from the three gray scale rings away from the magnets and other regions of thicker concentrations of dust on the rings. Turqoise: White ring, Light purple: Gray ring, Teal: Black ring. Image ID: 0514ML0020230030202937D01.

deposition on the front window of the Mastcam optics. Additional sky flat monitoring sequences are being acquired over time (e.g., mcam06606 on sol 1356) and are being used to continually update the flatfield correction files in the refined Mastcam calibration pipeline.

5.3.5. Improved Radiance Calibration

The initially archived Version 1 radiance-calibrated Mastcam RDRs are being calibrated using radiance coefficients based on a component-level model of the CCD QE, filter band-pass response, and optics transmissivity [*Caplinger*, 2013]. However, using the radiance coefficients estimated from the preflight testing described in section 3.2.6 (Table 4), which have been validated or amended based on the in-flight sky modeling work described in section 4.2.4 and Table 11, a better Version 2 estimate of the radiance on sensor can be calculated from decompanded, bias and dark current subtracted, and flatfielded Mastcam images, since those coefficients are based on the actual as-built system-level performance of the cameras. Work is under way to archive these Version 2 high-level Mastcam RDR data products, along with their associated ancillary calibration files and information, in the PDS.

5.3.6. Enhanced Radiance Factor (*I/F*) Calibration Using the Mastcam Calibration Target

Mastcam images can be converted from units of radiance to radiance factor (*I/F*; equal to π times the bidirectional reflectance) by comparison with near-simultaneous images of the Mastcam calibration target (caltarget). The caltarget is mounted on the rover deck behind the remote sensing mast and on the same (right) side of the rover (Figure 4). The enhanced Mastcam *I/F* calibration process is generally similar to that used for the *I/F* calibration of the *Spirit* and *Opportunity* rover Pancam images [*Bell et al*., 2006]. The presence of air fall dust on the caltarget requires a procedure for dust correction, since caltarget reflectances are no longer perfectly known once appreciable amounts of dust have accumulated. Fortunately, the dust cover accumulating on the target has been thin enough during the mission to date to enable the variable reflectance properties of the different substrate calibration materials to still be visible under nominal illumination conditions in Mastcam images.

Patches of seven spectrally distinct materials are mounted on the 8 × 8 cm caltarget base, organized into three gray scale rings (of approximately 20%, 40%, and 60% reflectance) [*Bell et al*., 2003], and four colored corner chips. Beneath the two brightest gray scale rings and underneath each of the four corner chips are six small sweep magnets [*Madsen et al*., 2003; *Bertelsen et al*., 2004]. The magnetic force from each of these magnets produces a small region of highly reduced dust deposition on the caltarget surfaces just above the magnets.

For each caltarget image, ROIs are selected for analysis by a calibration pipeline operator. The less dusty centers of the six magnets are selected, as are three regions of the gray scale rings away from the magnets. The reflectance calibration may be performed either by comparison with the six small, magnet centers (which offer a reduced dust layer at the cost of higher noise due to the low number of pixels) or using the three larger, nonmagnetic regions of the gray scale rings that have more dust deposition but better pixel statistics. Currently, the dust correction and reflectance calibration is performed based on the three "nonmagnetic" regions, but this could change as dust continues to accumulate on the caltarget.

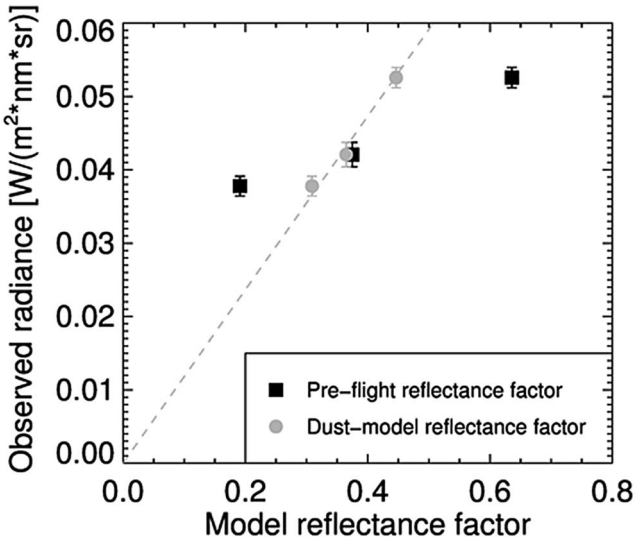

Figure 25. Radiance-reflectance plot for calibration target image acquired through the Mastcam's L3 (751 nm) filter on sol 514 as part of the same sequence as the image in Figure 24. The data points show observed radiances from the black, gray, and white region calibration target ROI's, as shown in Figure 24. The error bars are a simple standard deviation of all pixels within the ROI. Black squares show the three calibrated radiances plotted at the abscissa values of the clean caltarget reflectances that were determined preflight. Those points do not fit a line that would go through zero radiance at zero reflectance, and so clearly the reflectances of the ROIs have changed (the cal target has become dusty). The gray circles are the same three observed radiances but plotted as a function of the reflectances derived from our cal target dust correction model. The gray points fall approximately on a straight line through the origin, telling us that the dust model is indeed working properly. The slope of this line is a measure of the incoming irradiance.

Figure 24 shows an example of ROIs selected on a caltarget image acquired on sol 514 through the Mastcam's L0 (Bayer-RGB) filter. The ROIs shown are for the three nonmagnetic regions of the caltarget rings. Figure 25 shows a plot of the radiances observed in the L3 (751 nm) filter from the three regions as a function of the known (clean) reflectances (black squares) of those three materials [*Bell et al*., 2003]. If the reflectance values were good estimates, the three points would fall on a straight line through the origin, and the slope of this line would measure the incoming irradiance. The presence of dust reduces the contrast so that preflight-measured reflectance values are no longer good estimates of actual caltarget reflectances, and the three points no longer fall on a line through the origin.

The dust correction procedure fits the three points to an analytical two-layer scattering model based on the work of *Hapke* [1993, section 9.D.2]. The model treats single-scattering events in full detail and uses a two-stream formalism [e.g., *Zdunkowski et al*., 2007] to treat multiple-scattering events. The dust model and procedure follows very closely the one described in full detail in *Kinch et al*. [2015], the only difference is that we here do not distinguish between diffuse sky irradiance and solar (directional) irradiance; rather, all irradiance is assumed to come from the direction of the Sun. The utility of this model for analysis of dusty caltarget surfaces was demonstrated by *Johnson et al*. [2006] in laboratory studies. A similar, but simpler, two-stream scattering model is currently used for dust correction on the Mars Exploration Rovers' Pancam [*Kinch et al*., 2007], but work is ongoing to employ this more sophisticated model on that mission as well [*Kinch et al*., 2015]. Preliminary work on employing this dust model for the MSL Mastcam was presented in *Kinch et al*. [2013].

In addition to the unknown incoming irradiance, the dust model as employed here adds one other free parameter to the fit. This is the extinction optical depth of the dust layer on the caltarget τ_{cal}, defined in the standard way so that for radiance J normally incident on the dusty caltarget, the radiance that reaches the underlying caltarget surface without interacting with the dust is $J \cdot \exp(-\tau_{cal})$. In addition, the dust single-scattering albedo must be specified. The gray circles in Figure 25 demonstrate the procedure. These points are placed at the best fit reflectances of dusty caltarget surfaces as determined by the fit to the dust model. The results show that the expected linear relation between observed radiance and model reflectance now holds.

The fit is performed independently for every single caltarget image, and thus, the dust model derives both incoming irradiance and dust optical depth on the caltarget for every image. This history of dust deposition is shown for the L3 (751 nm) filter in Figure 26 and compared with the optical depths of dust in the atmosphere as observed through the Mastcam's 880 nm solar filter (L7). Figure 26 also shows derived incoming solar irradiances. These are shown relative to the known solar irradiance at the top of the atmosphere. The relative stability of these values is an indication that the value for dust single-scattering albedo employed in the model is adequate. If the dust was assumed to be too dark, the

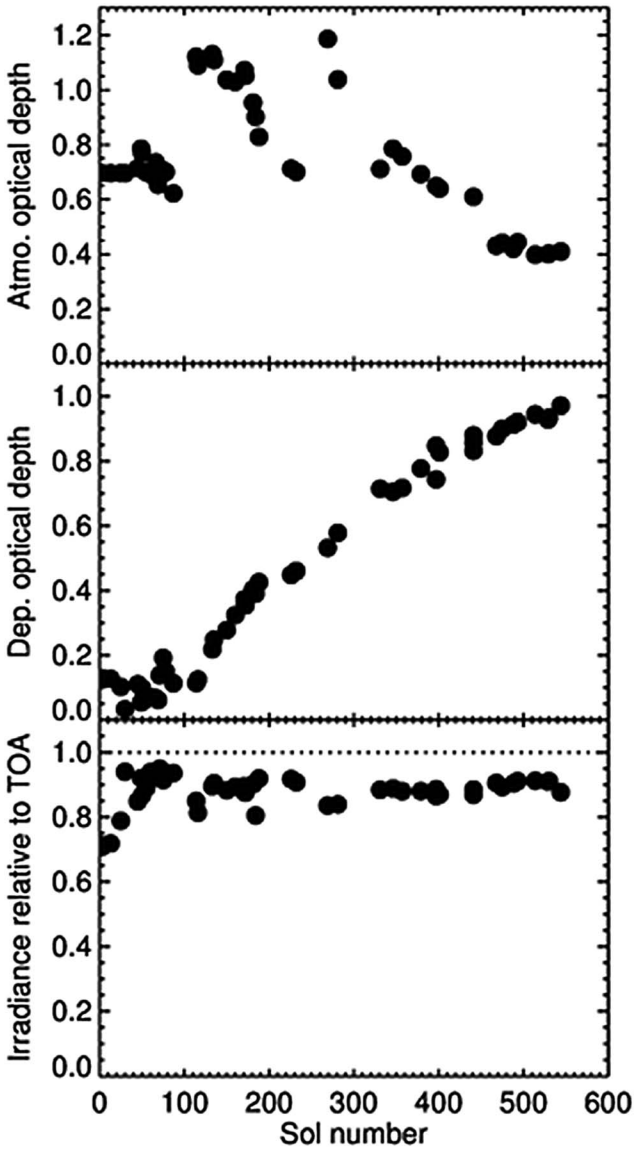

Figure 26. Time evolution of quantities derived by the dust model for all Mastcam calibration target images in the L3 (751 nm) filter and with the Sun at least 45° above the horizon. The first 600 sols of the mission are shown. The middle plot shows the derived extinction optical depth, τ_{cal}, of dust deposited on the caltarget. There is some noise in the early part of the mission, but after that, the plot shows a smooth increase in dust cover, with a slightly lower rate toward the later period. For comparison the top plot shows atmospheric optical depths as derived from observations of the Sun by the Mastcams L7 (880 nm) solar filter. The atmospheric dust load was lower at the later period consistent with the observed lower deposition rate. The bottom plot shows model-derived solar irradiance relative to the known top of atmosphere (TOA) value. Apart from some noise early on, the derived values are stable around 0.9. The noise in the beginning of the mission is probably because of the influence of dust deposited on the caltarget during landing. This material is likely to have had different color properties than the air fall dust.

derived incoming irradiances would drift to higher values as the caltarget gets dustier and vice versa for dust assumed to be too bright.

5.4. Examples: Accuracy and Precision of the Calibration

5.4.1. Example Morphologic Observations: Effects of Focus and Compression

On *Curiosity* sol 17, a 1 column × 34 row Mastcam mosaic sequence (mcam00050) was acquired with very large image overlap (75%) along a near-vertical transect from the surface in front of the rover out to the slopes of Mount Sharp on the horizon (Figure 27). These images were used to characterize the M-100 depth of field in flight; however, they also demonstrate the ability of Mastcam to acquire a set of images at different pointings that can then be combined (focus-merge product) via a "best focus" algorithm to produce a completely in-focus image. The example images demonstrate how the depth of field increases with distance to the scene target (the farthest image shows most of Mount Sharp and the far-field in focus, while the closer images are most in focus near the middle of the scene and are less in focus at the closest and farthest distances.

On sol 193, an eight-image M-100 "z-stack" sequence (mcam01026) was commanded on a laterally extensive outcrop near the region known as "Gillespie." A z-stack is a set of images acquired at a fixed pointing but with changing focus stepping through the scene's depth of field. Flight software within the Mastcam DEA processes the images into a resulting focus merge that preserves the best focused parts of each original input image (see section 7.6 in *Edgett et al*. [2012] for details). The full set of eight images at different focus positions was not downlinked, but the resulting on board-calculated focus-merge product was (Figure 28). Only the final merged z-stack image needs to be downlinked initially, and if individual z-stack images are requested by the science team, they can be downlinked at a later date. This reduces downlinked data, with a minimal impact to operational time. While only used rarely for Mastcam imaging so far in *Curiosity's* mission, z-stacking has been used frequently for microscale imaging by the MAHLI.

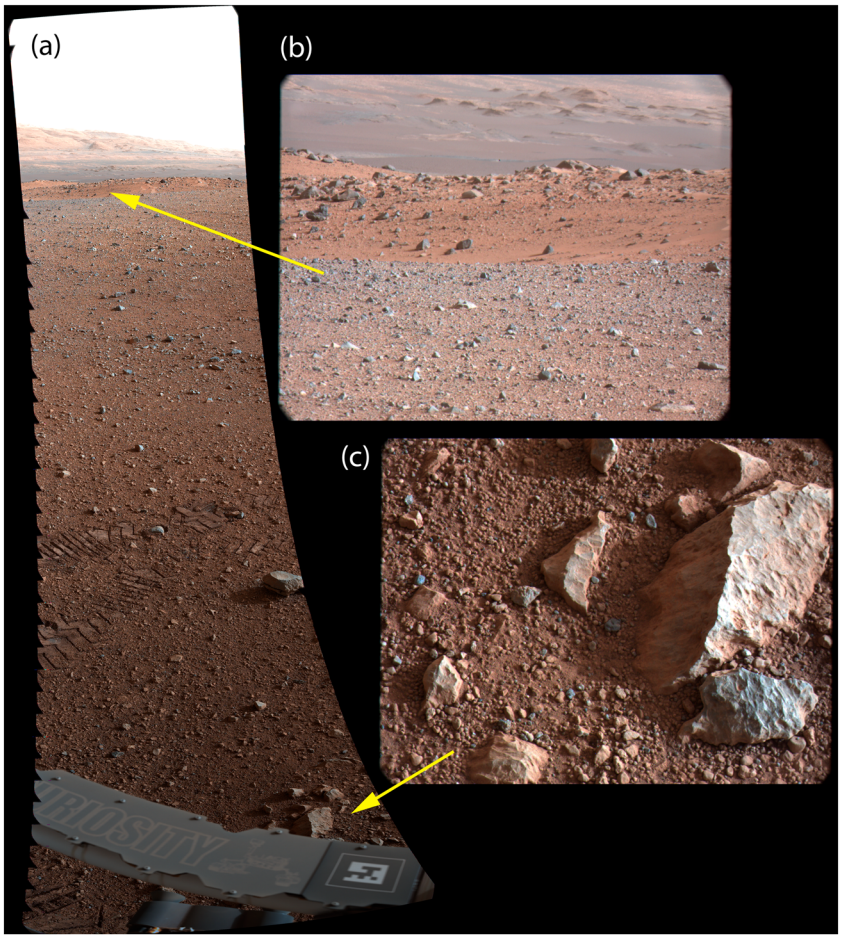

Figure 27. Images from the special 1 column × 34 row focus test mosaic acquired with the Mastcams on *Curiosity* sol 17. (a) Mosaic of 34 M-34 images acquired in the test, each acquired with 75% overlap relative to the previous one. (b) Example of one of the M-100 images from the mosaic with nearer, middle, and far ridges in the same scene. (c) Example of another M-100 image from the mosaic with near-field targets at a variety of distances.

Figure 29 shows a specific qualitative example, from M-34 imaging on *Curiosity* sol 1155, of the quantitative performance of the Mastcam JPEG compressor described in section 4.2.6 and Figure 18. Images at relatively high JPEG quality factors are generally indistinguishable from losslessly compressed images for most scenes imaged by Mastcam. However, some scenes with high-frequency features like layers or lamination show JPEG compression artifacts even at relatively high JPEG quality factors. The MSL science team and the Mastcam operations team actively identify images with JPEG artifacts that could negatively impact the robust scientific analysis of the data. Such images are often flagged for potential lossless retransmission, if and when downlink resources allow, before the original onboard image is scheduled for deletion.

In an even more extreme example, on sol 943 several sets of M-34 sky images were acquired in sequence mcam04146 to search for clouds. This sky imaging was originally downlinked at compression 35 but was brought back losslessly to discern any thin cloud layers. As seen in Figure 30, the original (JPEG 35) images show the deleterious effects of heavy compression on scenes with little variability. This effect is completely removed in the losslessly compressed version.

5.4.2. Example Stereo Products

Even though the two Mastcams have a factor of 3 difference in focal length, it is still straightforward to acquire near-simultaneous stereo pair images. Due to the stereo baseline and toe-in of the cameras, the area of the M-34 field of view also covered by the narrower M-100 field of view is restricted to a horizontal band in the center 1/3 of the M-34 sensor. The simultaneous M-100 field's placement

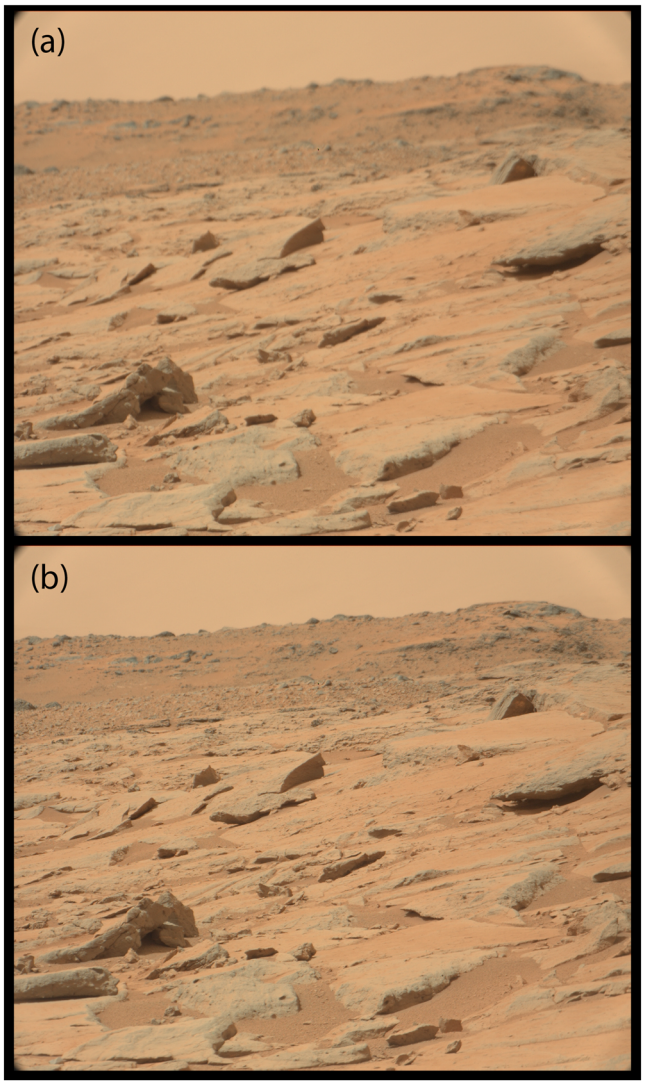

Figure 28. Example results from a special eight-image z-stack M-100 image test conducted on *Curiosity* sol 193 with sequence mcam01026. (a) One of the eight images acquired of the scene, with focus set for the foreground. (b) The resulting focus-merged z-stack image calculated automatically in Mastcam flight software. Note the dramatic removal of depth-of-field effects, especially for the far ridge.

within that horizontal band is dependent on the distance to the target, with closer M-100 targets appearing on the right side of the M-34 field of view and distant M-100 targets appearing in the left side (Figure 31). To simplify sequencing when acquiring stereo mosaics, M-34 images can be subframed vertically to exclude the top 1/3 and bottom 1/3 of the field of view, as well as subframed horizontally to exclude columns outside of the area covered by M-100 targets at distances from 1.9 m to infinity. The process is informally known as "shrink-wrap stereo," and the resulting M-34 frame size is 1152 pixels across by 432 pixels high, roughly centered in the sensor. This reduces the data volume of the M-34 to ~32% of a normal full frame.

Figure 31 also shows some additional dramatic examples of both near-field and far-field Mastcam stereo anaglyphs. The near-field view is a color stereo red-blue anaglyph rendering of a single sol 39 Mastcam left-right stereo pair of the Bradbury Plains conglomerate target named "Hottah," from sequences mcam00177 and mcam00178. The midfield and far-field view is from a sol 938 mosaic using Mastcam sequence mcam04119 that covered a midfield ridge called "Salsberry Peak" and the far-field lower slopes of Mount Sharp.

5.4.3. Sky/Astronomical Imaging Examples

The Mastcams have been used to acquire a variety of daytime and nighttime sky images that have been designed to achieve specific atmospheric science, meteorological, and astronomical objectives [*Malin et al.*, 2017]. The most common sky imaging observations have been relatively frequent direct solar images using the Mastcam L7 and R7 solar neutral density-coated filters [*Lemmon*, 2014], timed to characterize the history of atmospheric dust opacity in Gale crater as compared to the long-term record of opacity derived from similar kinds of direct solar imaging from other rover and lander sites [*Lemmon et al.*, 2015]. Additional sky imaging observations have included time-lapse observations of sunsets to characterize haze layers in the lower atmosphere above Gale crater (Figure 32a) [*National Aeronautics and Space Administration (NASA)*, 2015], time-lapse and video mode Mastcam observations of transits (annular solar eclipses) of Phobos across the Sun's disk (Figure 32c) [*NASA*, 2013a], and nighttime time-lapse imaging of an occultation of Deimos by Phobos (Figure 32b) [*NASA*, 2013b; *Lemmon et al.*, 2013; *Lemmon*, 2015], all designed to help refine the (slowly evolving) orbital parameters of both moons; twilight and nighttime imaging of the Earth and the Moon, Ceres, Vesta, and Deimos designed to search for evidence of nighttime clouds or hazes in the atmosphere;

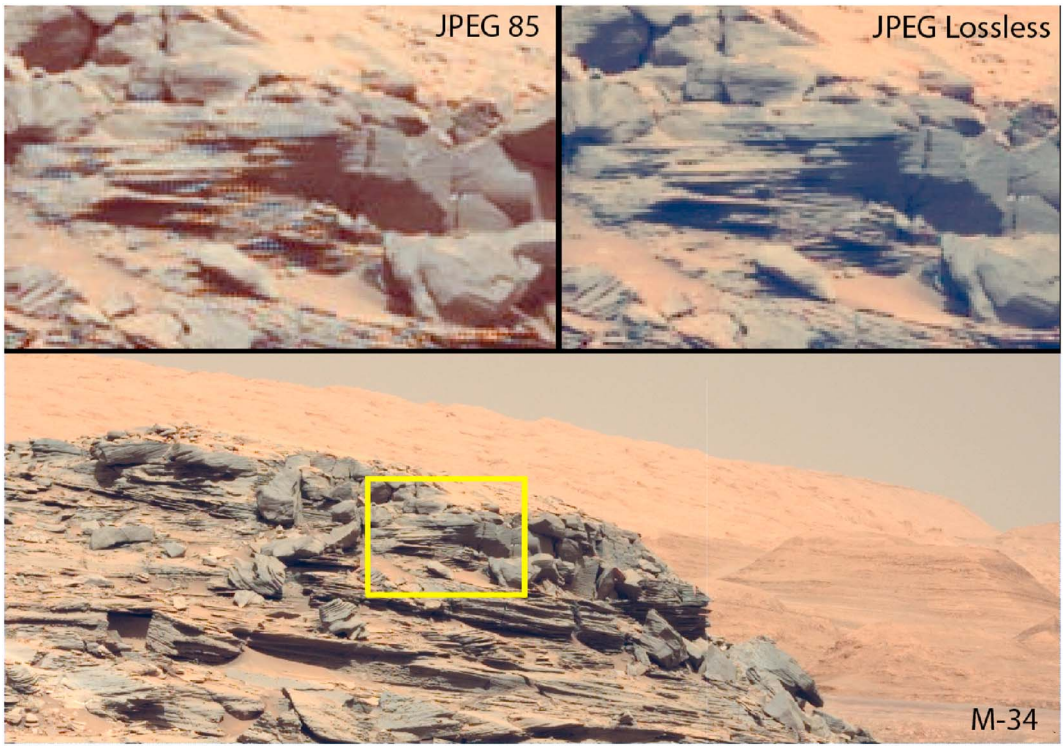

Figure 29. Mastcam M-34 image of finely layered outcrop rocks acquired on *Curiosity* sol 1155 and sequence mcam05219. The upper left inset shows a zoomed-in view of some of the layers in the originally downlinked JPEG quality factor 85 image. JPEG artifacts introduce blockiness and color banding that prevents a detailed assessment of the orientations and spacings of the layers. At upper right is an example of the same scene after redownlinking the onboard image losslessly. While the "cost" of the image was about 2.6 times more downlinked bits than the JPEG 85 version (e.g., Figure 18), the lossless image does not suffer from JPEG artifacts that interfere with fine-scale geologic interpretation.

daytime imaging of sunspots and a rare transit of Mercury as seen from Mars [*NASA*, 2014], partially designed to help monitor sunspot activity in support of other solar monitoring spacecraft; and nighttime observations of the close encounter of comet C/2013 A1 (Siding Spring) with Mars in October 2014 [*Lemmon et al.*, 2014].

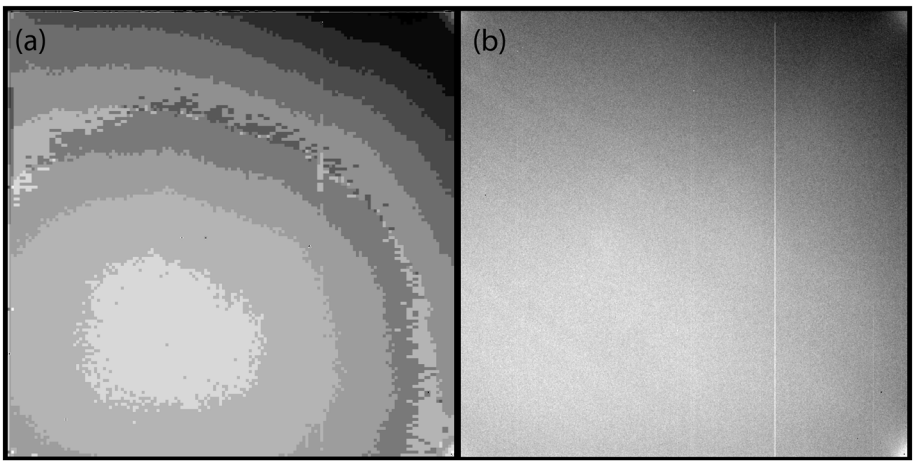

Figure 30. Example comparing very heavy Mastcam JPEG compression to lossless compression, for a relatively uniform scene, from sky images taken on *Curiosity* sol 943 in Mastcam sequence macam04146. (a) Highly compressed image of the sky downlinked at JPEG quality factor 35. (b) Losslessly compressed view of the same scene, downlinked later at a cost of ~15.5 times as many bits as Figure 30a. The white vertical line is the result of the readout of a cosmic ray strike on the detector within some of the topmost rows of the image.

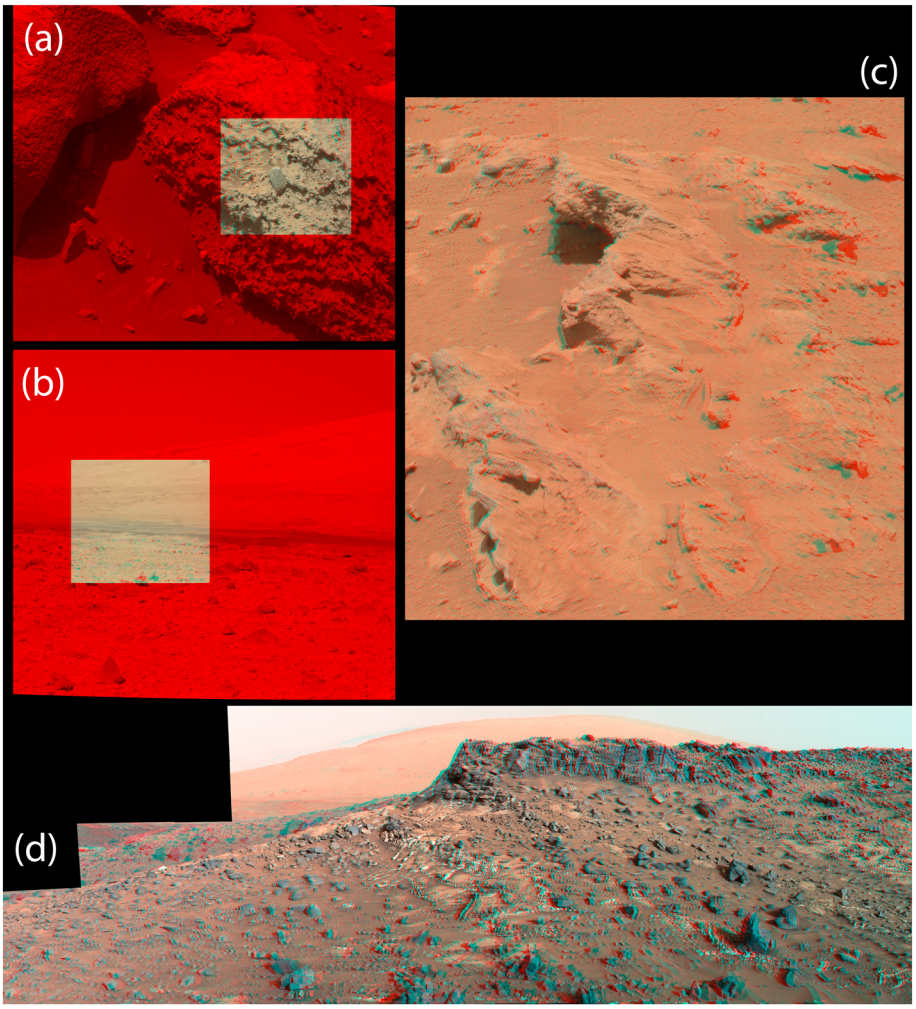

Figure 31. Example near-field and far-field stereo image results from Mastcam imaging. (a) Simultaneous stereo acquisition of M-34 (red) and M-100 (color anaglyph) from sequence mcam06894, on sol 1408, of the Tumba target at a distance of ~2.1 m. Note the M-100 placement in the center right of M-34 frame. (b) Same for simultaneous stereo acquisition of M-34 and M-100 from sequence mcam01985 on sol 505, including midfield targets at a distance of ~20 m to the base of Mount Sharp at a distance of several kilometers. Note the M-100 placement in the center left of M-34 frame. (c) Color red-blue stereo anaglyph of the conglomerate target Hottah acquired on sol 39 using Mastcam sequences mcam00177 (M-100 data) and mcam00178 (M-34 data). (d) Color stereo mosaic of the vein-rich ridge known as Salsberry Peak, with a second ridge and Mount Sharp in the background, acquired on sol 938 using Mastcam sequence mcam04119. Use standard red-blue stereo glasses to view the three-dimensional nature of these scenes.

5.4.4. Quantitative Radiance Factor Multispectral Observations

Mastcam has acquired several hundred multispectral sequences at the time of this writing, making use of the narrowband science filters on each camera to characterize the visible to near-infrared reflectance properties of rocks, soils, and other materials encountered by the rover. Due to the spectral influence of reddish dust, multispectral observations frequently target relatively cleaner surfaces, which may be undisturbed materials with less dusty rock faces but and include surfaces brushed, drilled, or otherwise disturbed by the rover. Example reflectance spectra from sol 762 (the "Confidence Hills" drill site) which have been calibrated to *I/F* by the methods described in section 5.3.6 are shown in Figure 33. Spectra from the drill tailings and the surface brushed by the Dust Removal Tool (DRT) show a strong absorption feature near the 527 nm filter compared with nearby dust and soil spectra, consistent with an enrichment in hematite. A variety of different spectral shapes and absorption features have been observed in reflectance spectra from other multispectral observations acquired along the traverse [e.g., *Wellington et al*., 2017].

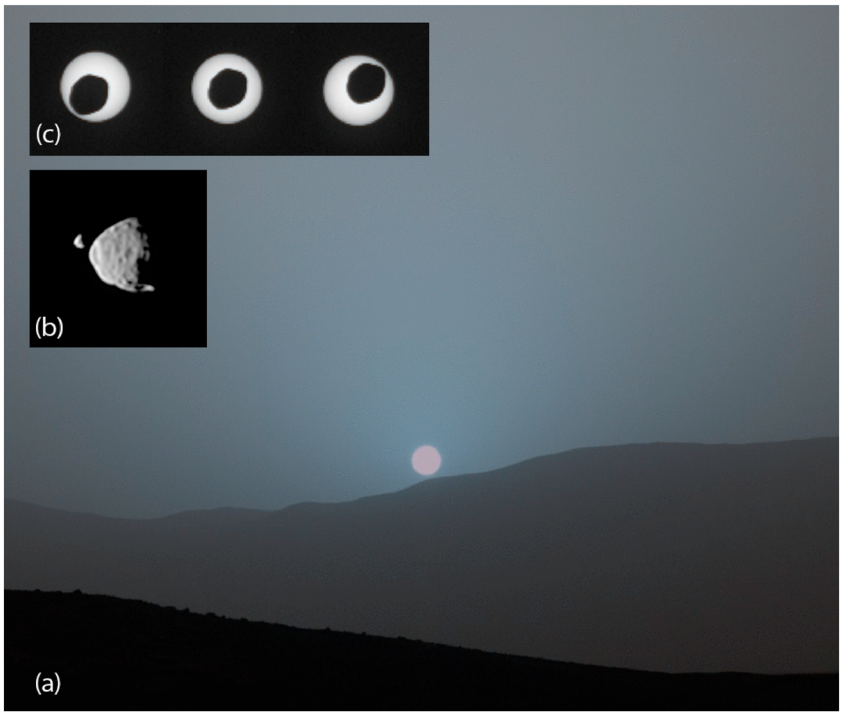

Figure 32. Examples of MSL/Mastcam daytime, twilight, and nighttime imaging for atmospheric science and astronomical observation goals. (a) White-balanced M-34 Filter 0 RGB composite image of a Martian sunset, acquired at twilight on sol 956 (15 April 2015) [*NASA*, 2015]. (b) M-100 Mastcam Filter 0 gray scale image of an occultation of Deimos by Phobos, observed on sol 350 (31 July 2013) [*NASA*, 2013b]. (c) M-100 Mastcam Filter 7 time-lapse views of Phobos passing across the solar disk, observed on sol 369 (20 August 2013) [*NASA*, 2013a].

5.5. Data Products and PDS Archiving

The Mastcam cameras and their DEAs create, store, and transmit to the rover's computer single images or video groups of images that are encoded in five basic formats: raster 8 bits per pixel, raster 16 bits per pixel, losslessly compressed (first difference Huffman encoded) 8 bits, JPEG gray scale, or JPEG color (Table 3). Other kinds of images generated by onboard processing in the Mastcam DEAs (i.e., focus merges, range maps, and 1/64 sampled thumbnails of full-resolution images) are also encoded in these same formats. Additional details on the specific nature of these raw data products, which are eventually downlinked to Earth and archived in the NASA PDS as raw Mastcam EDRs, can be found in the Mastcam Software Interface Specification (SIS) document [*Malin et al*., 2013].

Raw and calibrated Mastcam image products are currently being delivered to the PDS Imaging Node archives in five forms, depending on their level of calibration, if any (Table 13). Unprocessed image data are in CODMAC Level 2 form (Appendix C). These data are delivered to the PDS as files suffixed with XXXX.DAT. Details of the raw data records in a .DAT file are documented in sections 4.4.3 and 4.4.5 of the Mastcam SIS [*Malin et al*., 2013]. Software to extract and decompress all 21 Mastcam EDR data products (Table 3) from raw PDS-archived .DAT files is provided as part of the PDS archive distribution (for example, at http://pds-imaging.jpl.nasa.gov/data/msl/MSLMST_0001/SOFTWARE/). The four derivative forms of processed Mastcam data products (Table 13) are archived in a format recognized by the PDS for representation of image data and suffixed with ".IMG." These Mastcam images have associated detached labels in the PDS3 ODL standard, with the same filenames but suffixed with ".LBL." Details of the file naming scheme used for PDS archived Mastcam images are provided in Appendix D. There are 21 file types (Table 3) and five kinds of delivery products (Table 13) possible in the Mastcam PDS archive. Thus, a total of 105 types of archived files is possible. However, as described in section 4.5 of the Mastcam SIS [*Malin et al*., 2013], only 93 of these 105 possible products are actually being archived in the PDS. Details on the definitions and valid values of the Mastcam PDS label keywords, as well as example Mastcam PDS label file entries, are provided in Appendix A of *Malin et al*. [2013].

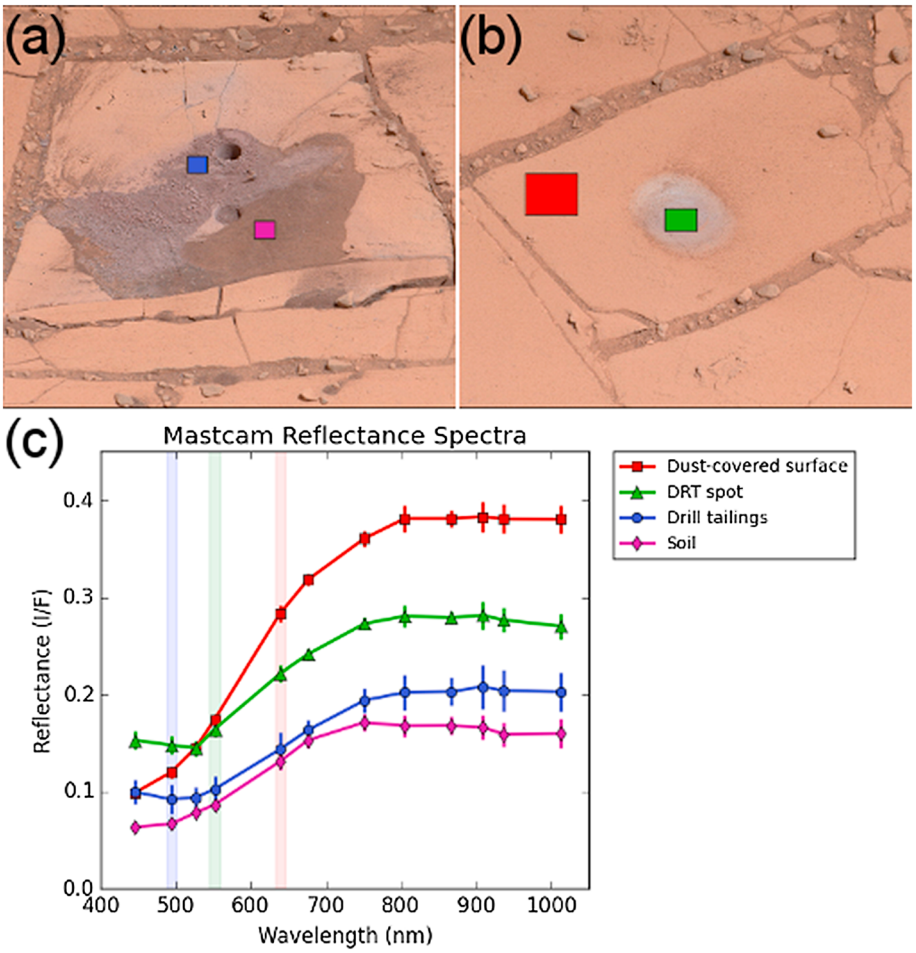

Figure 33. M-100 Bayer broadband RGB color composites of two pointings from a multispectral sequence acquired on sol 762 at the Confidence Hills drill location, targeting (a) the full and mini drill holes and tailings and (b) a nearby surface brushed by the DRT. (c) The colored boxes are the regions from which the spectra shown in matching are derived. The spectra are plotted as the mean *I/F* value as a function of filter band center wavelength, with bars showing the standard deviation of the pixels within each ROI. Spectra from the two cameras have been scaled together and averaged at overlapping filters. Colored stripes show the positions of the broadband Bayer values.

6. Future Work and Lessons Learned

The Mars Science Laboratory Mastcam investigation has been a phenomenal success and has played an important role in the discoveries made by *Curiosity* to date. Still, significant improvements to the calibration of the cameras are in the works, and more are being envisioned, for the future. For example, as described above, opportunities exist to further increase the fidelity of the calibration by completing the development of a rigorous shutter smear subtraction routine, which will be especially important for images taken with very short (several milliseconds) exposure times. Proactively flagging or replacing hot, cold, or gray pixels could be implemented in a future version of the pipeline or at least for images where the best possible analysis of small-scale, high-frequency pixel-to-pixel variations could be deemed important. Finally, additional in-flight monitoring of the performance of the Mastcams during *Curiosity's* continuing mission will result in the collection of additional in-flight assessments of CCD bias and dark current, as well as the flatfield performance of the detectors and optics, over time. Building those occasional monitoring update files into an evolving time-dependent calibration pipeline is an important part of any long-duration flight mission.

Many kinds of scientific analyses of Mastcam data require detailed quantitative image processing and modeling and thus rely on a robust and reliable calibration. Testing and calibration of the MSL Mastcam instruments took advantage of previous preflight and in-flight test and calibration programs

for other Mars imagers [e.g., *Reid et al*., 1999; *Bell et al*., 2003, 2006] as well as the experience gained by MSSS staff and science collaborators in their successful design, fabrication, and testing of more than a dozen other previous flight cameras. In a similar vein, the MSL Mastcam design, fabrication, test, calibration, and operations experience can provide guidance and lessons learned for similar work to be conducted on other future Mars and planetary imaging investigations. For example, much of the test equipment, procedures, and data processing algorithms used for Mastcam calibration will be directly applicable to the testing and calibration of the closely related Mastcam-Z imaging system on the NASA Mars 2020 rover [*Bell et al*., 2014]. Mastcam-Z is a pair of ~3:1 zoom cameras that directly inherit much of the CCD, electronics, and other mechanisms (filter wheel and focus) from MSL Mastcam, and thus, that team will leverage much of the process and experience described here. For example, the Mastcam-Z team is already planning to conduct additional preflight thermal vacuum testing of the geometric properties of the cameras over a wide range of expected Mars operating temperatures, going beyond the testing that was done for the Mastcams based on experience with thermal variations in the geometric properties of the cameras observed on Mars (e.g., section 4.2.5). Other examples of lessons learned include the intention to perform more robust assessments of the in-band and out-of-band performance of the multispectral filters (section 3.2.4) in order to increase the accuracy of the radiometric calibration, redesign of the Mastcam calibration target to better anticipate the effects of air fall (or landing-induced) dust on the target (sections 4.3 and 5.3.6), and closer coordination with other imagers on the Mars 2020 rover in order to try to arrive at a more uniform format for file naming, file labels, and archive formats across the vehicle.

7. Summary

This paper describes the preflight and in-flight calibration of the Mars Science Laboratory *Curiosity* rover's Mastcam fixed focal length, multispectral, stereoscopic imaging system [*Malin et al*., 2017], enabling the conversion of DN values downlinked from Mars into reliable estimates of physical quantities like absolute radiance and radiance factor, or *I/F*. Preflight and in-flight corrections and models have been developed to perform corrections for CCD, optics/filters, and geometric effects and to validate the calibration by comparing our derived estimates of radiance or radiance factor with those from known or assumed standards. The absolute radiometric accuracy and filter-to-filter precision of calibrated images meet or exceed the ±10% and ±7% requirements, respectively, except for the narrowband blue (445 nm) filters, which were not as well characterized prior to launch.

We also describe the tactical and strategically refined Mastcam data reduction and calibration pipelines. The latter include improvements in Bayer filter interpolation, flatfielding, and radiance and relative reflectance (*I/F*) calibrations. As validation of the pipeline processing described here, we also show some examples of science results related to calibration and processing of scenes with wide variations in focus and compression parameters, stereo imaging, sky/astronomical object imaging, and narrowband multispectral data sets. Finally, we describe a variety of details regarding Mastcam file naming, file formats, headers/labels, and PDS archiving that should prove useful to end users of the Mastcam images.

Appendix A: Camera Model Transformations

Rotations are expressed using unit quaternions, which are four tuples in the form

$$(s, v_x, v_y, v_z) = (s, \vec{V}) = \left(\cos\frac{\theta}{2}, \vec{A}_x \sin\frac{\theta}{2}, \vec{A}_y \sin\frac{\theta}{2}, \vec{A}_z \sin\frac{\theta}{2}\right) \tag{A1}$$

where $\vec{A}$ is the unit vector axis of rotation, θ is the angle of rotation, and the magnitude (square root of the sum of each of the four elements squared) is 1. Multiplying two quaternions composites their rotations (rotate by Q_2 first and then Q_1):

$$Q_1 \times Q_2 = \left(s_1 s_2 - \vec{V}_1 \cdot \vec{V}_2,\ s_1 \vec{V}_2 + s_2 \vec{V}_1 + \vec{V}_1 \times \vec{V}_2\right) \tag{A2}$$

which uses vector dot and cross products. Rotating a vector by a quaternion is then

$$Q \times \vec{V} = \text{vector}\left(Q \times (0, \vec{V}) \times Q'\right) \tag{A3}$$

where the central term builds a quaternion with scalar 0 and vector component V, $Q^{'}$ is the quaternion inverse, which negates the vector component, and vector() returns the vector part $\vec{V}$ of the quaternion (simply discarding the scalar).

The calibration camera models shown in Table 8 were taken at a specific pose, defined by **Pcal**, **Qcal** in Table 9 (which were themselves derived using the algorithm below). To construct the new pose P, Q, we extract the ARTICULATION_DEVICE_ANGLE from the RSM_ARTICULATION_STATE_PARMS group in the PDS label. The first two elements are the measured joint angles in the azimuth (az_target) and elevation (el_target) directions. These are *not* azimuth and elevation but are rather angles from the hard stop in the two directions, which are nominally at azimuth 181 and elevation −91°. Thus, az_target = 0 is (approximately) 1° past backward (toward the rear of the rover), and el_target=0 is (approximately) 1° past straight down.

```
az_axis = normalize(az_axis)                    // to unit vector
el_axis = normalize(el_axis)
// Rotate around az_axis to align with az_home, el_home
Quaternion axis_rot = Rotation(az_axis, (az_home - az_elaxis))
Vector el_axis_rot = axis_rot × el_axis         // Rotate a vector
Vector el_point_rot = axis_rot × (el_point - az_point) + az_point
// Transform target angles they're 0 at home
double az_angle = az_target - az_home
double el_angle = el_target - el_home

// Compute reference point position Pr at home configuration
Vector tmp1 = az_axis × el_axis_rot             // Cross product
Vector Pr = el_point_rot
if (magnitude(tmp1) > epsilon) {                // not parallel, safe
    double det = el_axis_rot · az_axis          // dot product
    double alpha = (az_point - el_point_rot) · el_axis_rot +
          ((el_point_rot - az_point) · az_axis) * det
    alpha = alpha / (1.0 - det*det)
    Pr = el_point_rot + el_axis_rot * alpha
}
// Compute final point and rotation as composite of az and el
Quaternion az_rot = Rotation(az_axis, az_angle)
Vector P_rmech = az_rot × (Pr - az_point) + az_point
Vector P = P_rmech + rmech_to_rnav              // Final point, RNAV
Quaternion el_rot = Rotation(el_axis_rot, el_angle)
Quaternion Q = az_rot × el_rot                  // Final quaternion
```

Note that the kinematics algorithm works with the "Rover Mechanical" coordinate frame throughout, converting to "Rover Navigation" frame at the end. The calibration and final camera models are expressed in Rover Navigation frame. A more complete treatment of these and other MSL coordinate frames can be found in *Alexander and Deen* [2015, section 6.3]. Conversion to "Local Level" or "Site" frames (which point north/down instead of relative to the rover) can be accomplished using the parameters in the ROVER_COORDINATE_SYSTEM_PARMS group in the PDS label.

Finally, the camera model extrinsics must be transformed from Pcal, Qcal to P, Q. To do this, we compute a composite transform that "unrotates" by Qcal and then rotates by Q:

$$Q_{\text{rot}} = Q \times Q^{'}_{\text{cal}} \tag{A4}$$

The CAHVOR vectors are then transformed as follows:

$$C' = Q_{rot} \times (C - P_{cal}) + P \tag{A5}$$

$$A' = Q_{rot} \times A \tag{A6}$$

$$H' = Q_{rot} \times H \tag{A7}$$

$$V' = Q_{rot} \times V \tag{A8}$$

$$O' = Q_{rot} \times O \tag{A9}$$

$$R' = R \tag{A10}$$

$$E' = E \tag{A11}$$

The E vector is a part of CAHVORE models [*Gennery*, 2006] that are not used by Mastcam but is included here for completeness.

For full-frame images, this is the final result. However, most Mastcam images are subframed, often to return only the science imaging area, and thumbnail images are additionally downsampled by a factor of 8. These operations affect the camera model as follows.

The starting coordinate of a subframe is found in the IMAGE object in the PDS label. The FIRST_LINE keyword defines the line start (dy), and FIRST_LINE_SAMPLE defines the sample start (dx). Note that the number of lines or samples does not affect the camera model. The model is then transformed as follows:

$$H' = H - (\mathrm{d}x - 1)A \tag{A12}$$

$$V' = V - (\mathrm{d}y - 1)A \tag{A13}$$

The -1 is because the PDS keywords start counting at 1, by definition. If there is no subframe, dx and dy are 1 and there is no modification. The other vectors are unchanged.

Mastcam thumbnail images are downsampled by a factor of 8. This factor is in the PIXEL_AVERAGING_HEIGHT ($h$scale) and PIXEL_AVERAGING_WIDTH ($v$scale) keywords in the IMAGE_PARMS group (which are always eight for Mastcam thumbnails). Downsampling is accomplished via the following:

$$H' = H/h\mathrm{scale} \tag{A14}$$

$$V' = V/v\mathrm{scale} \tag{A15}$$

which should be done after the subframe shift. To be technically correct, one additional step is needed to compensate for integer coordinates being in the center of a pixel: the H and V vectors must be shifted by -0.5 pixel before the scale and +0.5 pixel afterward, making the full equations:

$$H' = (H + A/2)/h\mathrm{scale} - A/2 \tag{A16}$$

$$V' = (V + A/2)/v\mathrm{scale} - A/2 \tag{A17}$$

However, as of this writing, the additional shift is not being performed by the Mastcam software, so the PDS labels reflect the simple scaling of equations (A14) and (A15) instead.

Appendix B: MSL/Mastcam Standard Companding (11 to 8 bit DN) and Decompanding (8 to 11 bit DN) "Look-Up Table 0"

This table provides the conversion to go from 8 bit downlinked Mastcam Data Number (DN) data back to the data's original 11 bit dynamic range.

8 bits	11 bits	8 bits	11 bits	8 bits	11 bits	8 bits	11 bits	8 bits	11 bits	8 bits	11 bits	8 bits	11 bits	8 bits	11 bits
0	0	**32**	46	**64**	150	**96**	316	**128**	542	**160**	829	**192**	1177	**224**	1586
1	2	**33**	48	**65**	154	**97**	322	**129**	550	**161**	839	**193**	1189	**225**	1600
2	3	**34**	50	**66**	159	**98**	328	**130**	558	**162**	849	**194**	1201	**226**	1614
3	3	**35**	53	**67**	163	**99**	334	**131**	566	**163**	859	**195**	1213	**227**	1628
4	4	**36**	55	**68**	168	**100**	341	**132**	575	**164**	869	**196**	1225	**228**	1642
5	5	**37**	58	**69**	172	**101**	347	**133**	583	**165**	880	**197**	1237	**229**	1656
6	5	**38**	61	**70**	177	**102**	354	**134**	591	**166**	890	**198**	1249	**230**	1670
7	6	**39**	63	**71**	181	**103**	360	**135**	600	**167**	900	**199**	1262	**231**	1684
8	7	**40**	66	**72**	186	**104**	367	**136**	608	**168**	911	**200**	1274	**232**	1698
9	8	**41**	69	**73**	191	**105**	373	**137**	617	**169**	921	**201**	1286	**233**	1712
10	9	**42**	72	**74**	196	**106**	380	**138**	626	**170**	932	**202**	1299	**234**	1727
11	10	**43**	75	**75**	201	**107**	387	**139**	634	**171**	942	**203**	1311	**235**	1741
12	11	**44**	78	**76**	206	**108**	394	**140**	643	**172**	953	**204**	1324	**236**	1755
13	12	**45**	81	**77**	211	**109**	401	**141**	652	**173**	964	**205**	1336	**237**	1770
14	14	**46**	84	**78**	216	**110**	408	**142**	661	**174**	974	**206**	1349	**238**	1784
15	15	**47**	87	**79**	221	**111**	415	**143**	670	**175**	985	**207**	1362	**239**	1799
16	16	**48**	90	**80**	226	**112**	422	**144**	679	**176**	996	**208**	1374	**240**	1814
17	18	**49**	94	**81**	231	**113**	429	**145**	688	**177**	1007	**209**	1387	**241**	1828
18	19	**50**	97	**82**	236	**114**	436	**146**	697	**178**	1018	**210**	1400	**242**	1843
19	20	**51**	100	**83**	241	**115**	443	**147**	706	**179**	1029	**211**	1413	**243**	1858
20	22	**52**	104	**84**	247	**116**	450	**148**	715	**180**	1040	**212**	1426	**244**	1873
21	24	**53**	107	**85**	252	**117**	458	**149**	724	**181**	1051	**213**	1439	**245**	1888
22	25	**54**	111	**86**	258	**118**	465	**150**	733	**182**	1062	**214**	1452	**246**	1903
23	27	**55**	115	**87**	263	**119**	472	**151**	743	**183**	1074	**215**	1465	**247**	1918
24	29	**56**	118	**88**	269	**120**	480	**152**	752	**184**	1085	**216**	1479	**248**	1933
25	31	**57**	122	**89**	274	**121**	487	**153**	761	**185**	1096	**217**	1492	**249**	1948
26	33	**58**	126	**90**	280	**122**	495	**154**	771	**186**	1108	**218**	1505	**250**	1963
27	35	**59**	130	**91**	286	**123**	503	**155**	781	**187**	1119	**219**	1519	**251**	1979
28	37	**60**	134	**92**	292	**124**	510	**156**	790	**188**	1131	**220**	1532	**252**	1994
29	39	**61**	138	**93**	298	**125**	518	**157**	800	**189**	1142	**221**	1545	**253**	2009
30	41	**62**	142	**94**	304	**126**	526	**158**	810	**190**	1154	**222**	1559	**254**	2025
31	43	**63**	146	**95**	310	**127**	534	**159**	819	**191**	1166	**223**	1573	**255**	2033

Appendix C: Committee on Data Management and Computation Processing Levels and Their Traditional NASA Equivalents

This table provides definitions of the standard Committee on Data Management and Computation (CODMAC) processing levels used for MSL Mastcam data processing [*CODMAC*, 1982].

CODMAC Level	NASA Level	Type	Description
1	Packet data	Raw data	Telemetry data stream as received at the ground station, with science and engineering data embedded
2	0	Edited data	Instrument science data (e.g., raw voltages and counts) at full resolution, time ordered, with duplicates and transmission errors removed
3	1-A	Calibrated data	Level 0 data that have been located in space and may have been transformed (e.g., calibrated and rearranged) in a reversible manner and packaged with needed ancillary and auxiliary data (e.g., radiances with the calibration equations applied)
4	1-B	Resampled data	Irreversibly transformed (e.g., resampled, remapped, and calibrated) values of the instrument measurements (e.g., radiances and magnetic field strength)
5	2	Derived data	Geophysical parameters, generally derived from Level 1 data, and located in space and time commensurate with instrument location, pointing, and sampling
5	3	Derived data	Geophysical parameters mapped onto uniform space-time grids

Appendix D: Mastcam Archived PDS Data File Names

Mastcam images archived in the NASA Planetary Data System (PDS) follow the file naming convention described here:

SSSSIIFFFFFFLLLXXCCCCCPGV_DXXX.ZZZ

defined as follows:

SSSS: Four-digit sol number after landing day (which was defined as sol 0).
II: Two-digit camera code: "ML" = Mastcam Left (M-34) and "MR" = Mastcam Right (M-100).
FFFFFF: Six-digit sequence number identifier.
LLL: Three-digit command number within the sequence that corresponds to this image.
XX: Two-digit Camera Data Product Identifier (CDPID) counter that records the number of times this CDPID has been used over the lifetime of the mission.
CCCCC: Five-digit CDPID value, uniquely assigned by the camera to an image product.
P: One-letter product type (see Table 3).
G: One-letter Group of Pictures (GOP) hexadecimal counter, for video sequences.
V: One-digit version number.
DXXX: Four-letter data processing code (see Table 13).
ZZZ: Three-letter file extension (typically, "DAT" or "IMG").

The string "FFFFFFLLL" within the filename is also known as the Product Identifier or a numerical identifier assigned to images when they are commanded from the ground. Depending on how the image was commanded, this number contains values related to the sequence used to command the image.

The product identifier is useful to distinguish among groups of images commanded with the same imaging sequence, such as for a panorama, video, or multispectral observation.

Additional details and examples of the Mastcam (as well as MAHLI and MARDI) file naming scheme can be found in section 3.4.1 and Table 3.4-1 of *Malin et al*. [2013].

Acknowledgments
We are indebted to the incredibly talented men and women of NASA, JPL/Caltech, and many other government labs and companies who helped to design, build, and deliver *Curiosity* to Mars and thus to enable the success of the Mastcam investigation. We also acknowledge the broader Mastcam, MAHLI, and MARDI Development and Operations team at Malin Space Science Systems, Inc., for their skill, perseverance, and outstanding attention to quality and detail in the preflight testing and sol-to-sol operations of the cameras on Mars. We thank Peter Smith, Rebecca Greenberger, Ken Edmundson, and Janet Richie for patient and thoughtful reviews and suggestions on an earlier draft of this paper. We acknowledge ASU student helpers Matt Jungers, Julie Mitchell, and Hannah Kerner for assistance with Mastcam calibration pipeline development and data processing. This work was funded by grants and contracts from NASA, JPL/Caltech, and MSSS. The data used here are listed in the references, tables, appendices, and/or the NASA Planetary Data System Imaging Node at, for example, http://pds-imaging.jpl.nasa.gov/data/msl/MSLMST_0001/CALIB/ and http://pds-imaging.jpl.nasa.gov/data/msl/MSLMST_0001/SOFTWARE/.

References

Alexander, D., and R. Deen (2015), Mars Science Laboratory Project Software Interface Specification (SIS): Camera & LIBS Experiment Data Record (EDR) and Reduced Data Record (RDR) Data Products, Version 3.5, December 17. [Available at http://pds-imaging.jpl.nasa.gov/data/msl/MSLNAV_0XXX/DOCUMENT/MSL_CAMERA_SIS.PDF.]

Bell, J. F., III and W. T. Sullivan (2004), The MarsDial: A sundial for the Red Planet, *The Planet*. *Rep*., pp. 6-11, Jan./Feb. 2004.

Bell, J. F., III et al. (2003), The Mars Exploration Rover Athena Panoramic Camera (Pancam) investigation, *J. Geophys. Res.*, *108*(E12), 8063, doi:10.1029/2003JE002070.

Bell, J. F., III, J. Joseph, J. N. Sohl-Dickstein, H. M. Arneson, M. J. Johnson, M. T. Lemmon, and D. Savransky (2006), In-flight calibration and performance of the Mars Exploration Rover Panoramic Camera (Pancam) instruments, *J. Geophys. Res.*, *111*, E02S03, doi:10.1029/2005JE002444.

Bell, J. F., III, M. C. Malin, M. A. Caplinger, M. A. Ravine, A. S. Godber, M. C. Jungers, M. S. Rice, and R. B. Anderson (2012), Mastcam Multispectral Imaging on the Mars Science Laboratory Rover: Wavelength coverage and imaging strategies at the Gale Crater field site, *43rd Lunar & Planetary Science Conf*., Abstract 2541.

Bell, J. F., III et al. (2013), Initial Multispectral Imaging Results from the Mars Science Laboratory Mastcam investigation at the Gale Crater field site, *44th Lunar & Planetary Science Conf*., Abstract 1417.

Bell, J. F., III, J. N. Maki, G. L. Mehall, M. A. Ravine, M. A. Caplinger, and the Mastcam-Z Science Team (2014), Mastcam-Z: A geologic, stereoscopic, and multispectral investigation on the NASA Mars-2020 rover, Abstract 1151, Presented at "International Workshop on Instrumentation for Planetary Missions (IPM-2014)," Greenbelt, Maryland, November 4-7.

Bertelsen, P., et al. (2004), Magnetic properties experiments on the Mars Exploration Rover Spirit at Gusev Crater, *Science*, *305*, 827-829.

Brown, D. C. (1958), A solution to the general problem of multiple station analytical stereotriangulation. *RCA Data Reduction Tech. Rep. No. 43*, Radio Corporation of America.

Burns, P. D. (2015), sfrmat3: SFR evaluation for digital cameras and scanners. [Available at http://losburns.com/imaging/software/SFRedge/sfrmat3_post/index.html.]

Burns, P. D. (2000), Slanted-Edge MTF for digital camera and scanner analysis, Proc. IS&T 2000 PICS Conference, pg. 135-138.

Caplinger, M. A. (2013), MSSS MSL Camera Calibration Summary, 9 May 2013; [Available online from the NASA Planetary Data System's Imaging Node. [Available at http://pds-imaging.jpl.nasa.gov/data/msl/MSLMST_0008/CALIB/MSL_MMM_CAL.TXT.]

CODMAC (1982), *Committee on Data Management and Computation, Data Management and Computation*, Volume 1 *Issues and Recommendations, Space Science Board, Assembly of Mathematical and Physical Processes*, National Research Council; 185 pp., National Academy Press, Washington, D. C.

Colina, L., R. C. Bohlin, and F. Castelli (1996), The 0.12–2.5 micron absolute flux distribution of the Sun for comparison with solar analog stars, *Astron. J.*, *112*, 307-315.

Deen, R., A. Chen, K. Capraro, H. Gengl, S. Algermissen, N. Ruoff, and O. Pariser (2015), Pointing correction for Mars surface mosaics, 2nd Planetary Data Workshop, poster #7055, Flagstaff, AZ, June 8-11.

Di, K., and R. Li (2004), CAHVOR camera model and its photogrammetric conversion for planetary applications, *J. Geophys. Res.*, *109*, E04004, doi:10.1029/2003JE002199.

Eastman Kodak Company (2009), "Kodak Interline Image Smear," Application note 2.0 MTD/PS-0899, March 5, 2009.
Edgett, K. E., et al. (2015), Curiosity's robotic arm-mounted Mars Hand Lens Imager (MAHLI): Characterization and calibration status, MSL MAHLI Technical Report 0001, *ResearchGate Technical Reports*, doi:10.13140/rg.2.1.3798.5447.
Edgett, K. S., et al. (2012), Curiosity's Mars Hand Lens Imager (MAHLI) investigation, *Space Sci. Rev. 170*, 259–317, doi:10.1007/s11214-012-9910-4.
Gennery, D. B. (2001), Least-squares camera calibration including lens distortion and automatic editing of calibration points, in *Calibration and Orientation of Cameras in Computer Vision*, edited by A. Grun and T. Huang, chap. 5, pp. 123–136, Springer, Berlin.
Gennery, D. B. (2006), Generalized camera calibration including fish-eye lenses. *Int. J. Comput. Vis. 68*(3), 239–266, doi:10.1007/s11263-006-5168-1.
Ghaemi, F. T. (2009), Design and fabrication of lenses for the color science cameras aboard the Mars Science Laboratory rover, *Opt. Eng.*, *48*, 103002-1-103002-15.
Grotzinger, J. P., et al. (2012), Mars Science Laboratory mission and science investigation, *Space Science Rev.*, *170*, 5–56.
Grotzinger, J. P., et al. (2014), A habitable fluvio-lacustrine environment at Yellowknife Bay, Gale Crater, Mars, *Science*, *343*, doi:10.1126/science.1242777.
Hapke, B. (1993), *Theory of Reflectance and Emittance Spectroscopy*, 455 pp., Cambridge Univ. Press, New York.
Howell, S. B. (2000), *Handbook of CCD Astronomy*, 184 pp., Cambridge Univ. Press, New York.
Imatest, LLC (2015) SVG Squares and Wedges pattern test chart. [Available at http://www.imatest.com/docs/testcharts_sfr_svg/.]
Janesick, J. R., K. P. Klaasen, and T. Elliott (1987), Charge-coupled-device charge-collection efficiency and the photon-transfer technique, *Opt. Eng.*, *26*, 972–980, doi:10.1117/12.7974183.
Johnson, J. R., et al. (2006), Radiative transfer modeling of dust-coated Pancam calibration target materials: Laboratory visible/near-infrared spectrogoniometry, *J. Geophys. Res.*, *111*, E12S07, doi:10.1029/2005JE002658.
Kinch, K. M., J. Sohl-Dickstein, J. F. I. Bell, J. R. Johnson, W. Goetz, and G. A. Landis (2007), Dust deposition on the Mars Exploration Rover Panoramic Camera (Pancam) calibration targets, *J. Geophys. Res.*, *112*, E06S03, doi:10.1029/2006JE002807.
Kinch, K. M., M. B. Madsen, J. F. Bell III, J. R. Johnson, and W. Goetz (2013), Dust on the Curiosity Mast Camera calibration target, *44th Lunar & Planet. Sci. Conf., Abstr.* 1061, 2013.
Kinch, K. M., J. F. Bell III, W. Goetz, J. R. Johnson, J. Joseph, M. B. Madsen, and J. Sohl-Dickstein (2015), Dust deposition on the decks of the Mars Exploration Rovers: 10 years of dynamics on the panoramic camera calibration targets, *Earth and Space Science*, *2*, 144–172, doi:10.1002/2014EA000073.
Lemmon, M. T., J. F. Bell III, M. C. Malin, K. M. Bean, A. Vasavada, F. J. Martin-Torres, M.-P. Zorzano-Mier, and the MSL Science Team (2013), Astrometric observations of Phobos and Deimos during solar transits imaged by the Curiosity Mastcam, 44th Lunar and Planetary Science Conference, LPI Contribution No. 1719, p. 1787, March 18-22.
Lemmon, M. T. (2014), The Mars Science Laboratory optical depth record, Eighth Int. Conf. on Mars, held July 14-18, 2014 in Pasadena, California, LPI Contribution No. 1791, p.1338.
Lemmon, M. T., J. F. Bell III, M. C. Malin, M. J. Wolff, J. N. Maki, J. Lasue, and S. Le Mouelic (2014), Imaging of comet C/2013 A1 (Siding Spring) from the Martian surface, AGU, Fall Meet., Abstract P42A-08.
Lemmon, M. T. (2015), Martian upper atmospheric aerosol properties from Phobos eclipse observation, American Astronomical Society, DPS meeting #47, id.401.09.
Lemmon, M. T., M. J. Wolff, J. F. Bell III, M. D. Smith, B. A. Cantor, and P. H. Smith (2015), Dust aerosol, clouds, and the atmospheric optical depth record over 5 Mars years of the Mars Exploration Rover mission, *Icarus*, *251*, 96–111, doi:10.1016/j.icarus.2014.03.029.
Lensation GmbH (2013), How to convert pixel size in line pairs per millimeter, *OptoWiki Knowledge Base*. [Available at http://www.optowiki.info/faq/how-to-convert-pixel-size-in-line-pairs-per-millimeter/.]
Madsen, M. B., et al. (2003), The Magnetic Properties Experiments on the Mars Exploration Rover mission, *J. Geophys. Res.*, *108* (E12), 8069, doi:10.1029/2002JE002029.
Maki, J. N., et al. (2003), The Mars Exploration Rover engineering cameras, *J. Geophys. Res.*, *108*(E12), 8071, doi:10.1029/2003JE002077.
Maki, J., D. Thiessen, A. Pourangi, P. Kobzeff, T. Litwin, L. Scherr, S. Elliott, A. Dingizian, and M. Maimone (2012), The Mars Science Laboratory engineering cameras, *Space Sci. Rev.*, *170*, 77–93, doi:10.1007/s11214-012-9882-4.
Malin, M. C., et al. (2010), The Mars Science Laboratory (MSL) Mast-mounted Cameras (Mastcams) flight instruments, Lunar and Planetary Science Conference 41, Lunar and Planetary Institute, Houston, TX, p. 1123, March 1-5.
Malin, M. C., et al. (2017), The Mars Science Laboratory (MSL) Mast cameras and Descent imager: I. Investigation and instrument descriptions, *Earth and Space Science*, *4*, doi:10.1002/2016EA000252.
Malin, M. C., K. Edgett, E. Jensen, and L. Lipkaman (2013), Mars Science Laboratory Project Software Interface Specification (SIS): Mast Camera (Mastcam), Mars Hand Lens Imager (MAHLI), and Mars Descent Imager (MARDI) Experiment Data Record (EDR), Reduced Data Record (RDR), and PDS Data Products," Version 1.2, JPL D-75410, October 29, 2013. [Available at http://pds-imaging.jpl.nasa.gov/data/msl/MSLMST_0001/DOCUMENT/MSL_MMM_EDR_RDR_DPSIS.PDF.]
Malvar, H. S., L. He, and R. Cutler (2004), High-quality linear interpolation for demosaicing of Bayer-patterned color images, *Acoustics, Speech, and Signal Processing*, 2004, *Proceedings*, *IEEE International Conference*, vol. 3, pp. 485–488, IEEE, Piscataway, N. J. [Available at http://www.ipol.im/pub/art/2011/g_mhcd/, 2004.]
Mikhail, E. M., J. S. Bethel, and J. D. McGlone (2001), *Introduction to Modern Photogrammetry*, 496 pp., John Wiley, New York.
Morris, R. V., and T. G. Graff (2002), Athena instrument validation and data library development for the Mars Exploration Rover (MER) mission, *Eos Trans. AGU*, *83(47)*, Fall Meet. Suppl., Abstract P21C-03.
National Aeronautics and Space Administration (NASA) (2013a), Annular eclipse of the Sun by Phobos, as seen by Curiosity, press release, 28 August 2013a. Online at http://mars.nasa.gov/msl/multimedia/images/?ImageID=5536.
National Aeronautics and Space Administration (NASA) (2013b), NASA rover gets movie as a Mars moon passes another, press release, 15 August 2013b. [Available at http://mars.nasa.gov/msl/news/whatsnew/index.cfm?FuseAction=ShowNews&NewsID=1509.]
National Aeronautics and Space Administration (NASA) (2014), Mercury passes in front of the Sun, as seen from Mars, press release, 10 June 2014. [Available at http://mars.nasa.gov/news/whatsnew/index.cfm?FuseAction=ShowNews&NewsID=1647.]
National Aeronautics and Space Administration (NASA) (2015), Sunset in Mars' Gale Crater, press release, 8 May 2015. [Available at http://mars.nasa.gov/msl/multimedia/images/?ImageID=7188.]
ON Semiconductor, Inc. (2014), "Interline CCD Smear," Application Note AND9184/D, online at http://www.onsemi.com/pub_link/Collateral/AND9184-D.PDF.
ON Semiconductor, Inc. (2015), "KAI-2020 1600 (H) x 1200 (V) Interline CCD Image Sensor," Pub. Order No. KAI-2020/D. [Available at http://www.onsemi.com/pub/Collateral/KAI-2020-D.PDF.]

Pennebaker, W. B., and J. L. Mitchell (1992), *JPEG: Still Image Data Compression Standard*, 638 pp., Springer, New York.
Peters, S. (2016), Mars Science Laboratory Pointing, Positioning, Phasing, and Coordinate Systems (PPPCS) document, volume 9: Surface remote sensing and navigation, Version "Public Release", May 13, 2016. [Available at http://pds-imaging.jpl.nasa.gov/data/msl/MSLNAV_0XXX/DOCUMENT/PPPCS_VOL9_MAY_13_2016.PDF.]
Reid, R. J., et al. (1999), Imager for Mars Pathfinder (IMP) image calibration, *J. Geophys. Res.*, *104*, 8907–8926, doi: 10.1029/1998JE900011.
Rice, M. S., J. F. Bell III, E. A. Cloutis, A. Wang, S. Ruff, M. A. Craig, D. T. Bailey, J. R. Johnson, P. A. de Souza Jr., and W. H. Farrand (2010), Silica-rich deposits and hydrated minerals at Gusev Crater, Mars: Vis-NIR spectral characterization and regional mapping, *Icarus*, *205*, 375–395, doi:10.1016/j.icarus.2009.03.035.
Stamnes, K., S. C. Tsay, W. J. Wiscombe, and K. Jayaweera (1988), Numerically stable algorithm for discrete-ordinate-method radiative transfer in multiple scattering and emitting layered media. *Appl. Opt.*, *27*, 2502–2509.
Tomasko, M. G., L. R. Doose, M. T. Lemmon, P. H. Smith, and E. Wegryn (1999), Properties of dust in the martian atmosphere from the Imager for Mars Pathfinder, *J. Geophys. Res.*, *104*, 8987–9007, doi: 10.1029/1998JE900016.
Truesense Imaging, Inc. (2012), "KAI-2020 Imaging Sensor," Device performance specification document, Rev. 1.0, PS-0017, June 22, 2012.
Wellington, D. F., J. F. Bell III, J. R. Johnson, K. M. Kinch, M. S. Rice, A. Godber, B. L. Ehlmann, A. A. Fraeman, C. Hardgrove, S. Le Mouélic, and the MSL Science Team (2017), Visible to near-infrared MSL/Mastcam multispectral imaging: Initial results from select high-interest science targets within Gale Crater, Mars, *Am. Mineral.*, *102*, 1202–1217, doi:10.2138/am-2017-5760CCBY.
Wiens, R. C., et al. (2012), The ChemCam instrument suite on the Mars Science Laboratory (MSL) rover: Body unit and combined system performance, *Space Sci. Rev.*, doi:10.1007/s11214-012-9902-4.
Williams, R. M. E., et al. (2013), Martian fluvial conglomerates at Gale crater, *Science*, *340*, 1068–1072, doi:10.1126/science.1237317.
Wolff, M. J., M. D. Smith, R. T. Clancy, R. Arvidson, M. Kahre, F. Seelos IV, S. Murchie, and H. Savijaärvi (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.
Yakimovsky, Y., and R. Cunningham (1978), A system for extracting three-dimensional measurements from a stereo pair of TV cameras, *Computer Graphics and Image Processing*, *7*, 195–210, doi:10.1016/0146-664X(78)90112-0.
Yotam, E., P. Ephi, and Y. Ami (2007), MTF for Bayer pattern color detector, in *Signal Processing, Sensor Fusion, and Target Recognition XVI*, edited by I. Kadar, Proc. of SPIE, vol. 6567, 65671M, SPIE—The International Society for Optical Engineering, Bellingham, Wash., doi:10.1117/12.723140.
Zdunkowski, W., T. Trautmann, and A. Bott (2007), *Radiation in the Atmosphere*, 1st ed., 496 pp., Cambridge Univ. Press, Cambridge, U. K.
Zhang, D., and Xiaolin Wu (2005), Color demosaicking via directional linear minimum mean square-error estimation, *IEEE Trans. Image Process.*, *14*(12), 2167-2178.