

PhD thesis

Simone F. Juul Jakobsen

Modelling and Correction of an Artifact in the Mars Exploration Rover Panoramic Camera

Advisors: Morten B. Madsen, Kjartan M. Kinch

November 30, 2018

Contents

1 Introduction	15
1.1 Motivation and aim for this thesis work	16
1.2 The structure of this thesis	18
2 Mars: fundamentals and basic geology	19
2.1 Martian surface morphology	20
2.1.1 Craters	21
2.1.2 Martian epochs - the geological timescale	22
2.1.3 Volcanoes	23
2.1.4 Aeolian features	25
2.1.5 The Martian dichotomy	25
2.2 Water on Mars	26
2.2.1 The hydrological history of Mars	27
2.2.2 Past climate conditions	28
3 Mars Exploration Rovers (MER)	31
3.1 MER instrumentation - the Athena payload	32
3.1.1 Pancam	32
3.1.1.1 Filters	33
3.1.1.2 CCD	34
3.1.1.3 Calibration target	37
3.1.2 Pancam in-flight calibration and data products	38
3.2 Overview of mission results	40
4 Pancam R7 effect: Characterisation and simulation	43
4.1 The R7 effect in in-flight geology images	44
4.1.1 The R7 effect indicated in the literature	45
4.2 Characterising the effect	48
4.3 Finding a hypothesis	51
4.4 The model	53
4.4.1 Single wavelength approximation	54
4.4.2 Full band model	55
4.5 Simulation - optimizing the model	57
4.5.1 Results of the simulation	58

5 Pancam R7 effect: Correction	61
5.1 Inverse process - an iterative deconvolution method	61
5.2 Termination criteria and first test	63
5.2.1 The correction algorithm	65
6 Pancam R7 effect: Correction verification	67
6.1 Correction of independent, in-flight, geology images	67
6.2 Correction of in-flight calibration target images	75
6.3 Correction of images used for previously published work	78
7 Discussion	87
7.1 Simulating the effect - the accuracy of the model	87
7.2 Notes on the correction method	88
7.3 Correction Performance on Pre-Flight, Stray-Light Image	89
7.4 Correction Performance on In-Flight Geology Images	90
7.5 Correction performance on cal target offsets	91
7.6 Correction of images used for previous publications	92
8 Conclusion and outlook	93
8.1 Implications	94
8.2 Future work	94

“Look for the valleys, the green places,
and fly through them. There will always be a way through.”
- C. S. Lewis, Narnia

“Yet if there were no hazards there would be no achievement,
no sense of adventure.”
- Arthur C. Clarke, Rendezvous with Rama

Preface

This thesis has been submitted to the PhD School of The Faculty of Science, University of Copenhagen. The work presented here was carried out at the Niels Bohr Institute, University of Copenhagen during the years 2014-2018.

The majority of the content of this thesis is based on a paper that was submitted to the journal *Earth and Space Science* under AGU in September 2018. Two reviewer reports were received in October 2018. Both recommended publication after minor revisions. A revised manuscript was submitted in November 2018 and is now awaiting final approval and publication. Sections [6.3](#) and [7.6](#) describes the preliminary work done towards a second publication, and is currently unsubmitted.

The first part of this thesis (Chapters 1-3) is written as a short popular science introduction to the field of Mars research, introducing some of the basic tools and concepts used throughout the second part of the thesis. The second part (Chapters 4-8) describes the original work done by the author in collaboration with advisor and main advisor, Kjartan M. Kinch and Morten B. Madsen.

Copenhagen, November 2018

Simone F. Juul Jakobsen

Acknowledgements

Many people have supported me through my PhD endeavor.

First and foremost I would like to thank my two advisors Kjartan Kinch and Morten Bo Madsen. Kjartan was quick to catch my interest with an idea he had for a, at the time, small project. What was first intended to constitute a minor part of my PhD ended up being the main focus of my work during the past four years. I have learned so much during my work with you, so thank you for all the engaging science talks and for always being available for questions when I needed it.

I want to thank Morten for his eternal optimism which is quite contagious. At times when I was stuck trying to locate a bug in my IDL code, Morten would look at it without knowing what I was doing and find the bug in a few moments. Thank you for that! You spared me a great many hours of pulling my hair out. Your door has always been open and I've enjoyed all the many stories and anecdotes you know from your many years of participation in missions to Mars.

Thank you both for your support and understanding through these past years.

Thank you to Jim Bell for hosting me during my stay at ASU. It was a great experience and I met a lot of incredibly talented people.

I would also like to thank my husband without whom I wouldn't have come this far. This has been a very tough four years, for more reasons than one, and my husband has been there the whole way through. During the main part of my thesis writing, he has been the main reason that I remembered to eat. You have been my biggest supporter, and I am so happy that we are the kind of couple who enjoy solving integrals while eating dinner, and discuss which coding language is better (you know what I think). You have been the main reason I've kept the first quote on the previous page in mind.

I of course want to thank my parents who have endured my ghostly presence these past couple of months. Thank you for genuinely being interested when I tell you about my project. You have always believed in me. And thank you for letting me occupy your dinner table with my books whenever my husband actually engaged with other people than his stressed out wife.

And finally, an unconventional thanks to my cat who basically wrote my thesis with me while purring my stress away.

Summary

In 2004 two identical rovers touched down on the surface of Mars, on opposing sides of the planet. The two rovers Spirit and Opportunity were both part of the joined mission called Mars Exploration Rovers (MER), and their mission was to characterise their surroundings and identify possible signatures of past liquid water. From the collected data, information about the climate of the planet can be derived and used to shed light on whether life could have existed on the planet at earlier times or whether life could still exist to this day just below the red martian dust. One of the main rover instruments is a panoramic camera called Pancam. It's aim is to obtain high resolution images of Mars that can be used both for navigation purposes as well as scientific analyses. The importance of being able to trust the collected observations cannot be underlined enough and it was therefore important that Pancam could maintain it's calibration throughout the MER mission. In order to calibrate Pancam during the mission a so called calibration target, mounted on the deck of each rover, was utilized. The calibration method consists of having characterised the calibration target to a high degree prior to launch, by measuring the precise reflectance of it's different coloured and grey scale areas, and from that derive a model. Later, during the mission, it would be possible to obtain Pancam images of the cal target and extract spectra from the images of the same coloured and grey scale areas, and scale them to the models. This way the Pancam images, obtained at the same time and place as the cal target images, could be calibrated. Prior to launch, the models were compared to extracted reflectances from Pancam images of the cal target in order to pre-calibrate Pancam, and during these activities, a discrepancy was found at the longest wavelengths let in by the camera, in the filter called the R7 filter. The discrepancy consisted of a decrease in contrast in the R7 filter, that is the brightest area of the cal target was darker in the image than what the model predicted, and the darkest area of the cal target was brighter in the image than what the model predicted. The work presented here focus on describing the effect, both in images of the cal target, but also in images of the martian surface. That the R7 filter artifact is observed in all types of images points towards it being an instrumental effect, and as the work here will show, an effect that most likely arises in the CCD after the light has passed through the camera optics. From this assumption, a hypothesis is formulated saying that the effect arises as a consequence of light being reflected off the backside of the camera CCD, which leads to a signal registered in other pixels than the intended one. Based on this, a mathematical model is developed that describes the effect. The developed model is used to simulate the effect, and it is shown that this can be done to a very high degree of precision. Since the primary aim is to remove the effect from affected R7 images, the model is invoked on the inverse problem through an iterative deconvolution algorithm, where the output is designed to converge towards the corrected version of the image. The work presented here shows that the developed correction removes the effect to a high degree, both from simple, symmetric images, and complex images of the martian surface. Furthermore, it is shown how the effect could have influenced previously published work, where the analysis was based on R7 filter images.

I 2004 landede to identiske rovere på Mars overflade på hver sin side af planeten. De to rovere Spirit og Opportunity var begge en del af en fælles mission kaldet Mars Exploration Rovers (MER) og deres mission var at karakterisere deres omgivelser og identificere mulige tegn på at flydende vand har været til stede på et tidligere tidspunkt i Mars historie. Ud fra de opsamlede data kan information omkring planetens klima gennem historien blive udledt og kaste lys over om der kan have eksisteret liv på planeten, eller om der måske stadig gør det, evt. gemt under planetens røde støv. Et af hovedinstrumenterne ombord på de to rovere, er et panoramisk kamera kaldet Pancam. Dets formål er at tage højopløselige billeder af Mars der kan bruges til både navigeringsformål og til videnskabelig analyse. Vigtigheden af at kunne stole på de indsamlede observationer kan ikke understreges nok og det var derfor vigtigt at Pancam blev kalibreret løbende gennem missionen. For at kunne kalibrere Pancam under missionen bruger man et såkaldt kalibreringstarget monteret på hver rover. Metoden består i at have karakteriseret dette kalibreringstarget så godt som muligt før opsendelsen, dvs. ved at måle den præcise reflektans fra de forskellige farvede og gråskalaområder den består af og ud fra det formuleret en model. Ved senere, under missionen, at tage billeder af kalibreringstarget med Pancam, og derefter udtrække spektre af de samme områder på kalibreringstarget, kan man ved hjælp af skalering til modellerne, kalibrere de Pancambilleder der er taget på samme tid og sted som billedet af kalibreringstarget. Før opsendelsen af de to rovere sammenlignede man netop modellerne med de udtrukke reflektanser fra Pancam billeder af kalibreringstarget for at for-kalibrere Pancam, og der fandt man en afvigelse ved de længste bølgelængder, i det filter der kaldes R7. Her så man at det lyseste område af kalibreringstarget var mørkere i billedet end modellen forudså, og det mørkeste område var lysere i billedet end modellen forudså. Dvs. der var lavere kontrast i R7 filteret i forhold til modellen. Det arbejde der præsenteres her omhandler en beskrivelse af den effekt der ses i R7 filteret, både i billeder af kalibreringstarget og i billeder af Mars overflade. At R7 filtereffekten ses i alle typer af billeder, tyder på at effekten er en instrumentel effekt, og som det vil vise sig i denne afhandling, en effekt der opstår i CCDen efter lyset er gået igennem kameraets optik. Ud fra dette opstilles der en hypotese om at effekten skyldes en refleksion på bagsiden af kameraets CCD, som leder til et ekstra signal registreret i andre pixler end den der burde have registreret lyset. Baseret på dette udvikles en matematisk model til at beskrive effekten på baggrund af hypotesen. Den formulerede model bruges dernæst til at simulere R7 effekten, og det viser sig at dette er muligt til en høj grad af præcision. Da det primære formål er at fjerne effekten fra R7 filter billeder bruges modellen på det inverse problem vha. en iterativ deconvolution algoritme hvor outputtet konvergerer mod det korrigerede billede. Denne afhandling viser hvordan den udviklede korrektion kan fjerne effekten i høj grad, både på simple symmetriske billeder, og på reelle billeder af Mars overflade. Der vises ydermere hvordan effekten kan have haft betydning for tidligere publicerede resultater baseret på R7 billeder.

1

Introduction

As the fourth planet from the Sun, Mars is the outermost terrestrial planet and neighbor to the Earth. With its radius of 3389.5 km it subtends an angle in the sky ranging from 3.5" to 25" depending on its current proximity to Earth. The distance means that the view of Mars from an Earth based telescope does not reveal its topography in detail, only the changing albedo of the surface and the change in atmospheric opacity can be seen at this distance. All topographic and geological data acquired therefore stem from the number of telescopes in orbit around Mars, and from the rovers exploring its surface.

In January 2004, the Spirit rover (MER A) touched down in the Gusev crater on Mars. Spirit was the first half of the NASA lead Mars Exploration Rover mission. The second half, the Opportunity rover (MER B), landed 3 weeks later on the opposing side of the planet on the Meridiani Planum, 2 degrees South of the Martian Equator. Both sites were chosen for the high likelihood of water being present sometime in the past. The two rovers each had a planned mission span of 92 Earth days ahead of them, but both rovers succeeded in surpassing this goal. The Spirit rover continued to be operational for a little over 5 years before getting stuck in soft soil. Still it continued to be functional for approximately one more year before all communication stopped. The Opportunity rover still continued to be operational up until June 3, 2018, when a dust storm cut off the connection to the rover. Efforts are being put into recovering the connection, and will continue into the first month of 2019 [19].

Both rovers are identical in construction and are equipped with the Athena Science payload, with the Panoramic Camera (Pancam) imaging system as one of the main tools available for remote sensing of the nearby terrain. Pancam is a two-camera system mounted on a mast on the deck of each rover, with the two cameras acting as the rover eyes. Mounted in front of each camera is a 8-step filter wheel containing filters L1-L8 for the left eye, and R1-R8 for the right eye. Two of the filters, namely the L8 filter at $\lambda_{eff} = 440$ nm and the R8 filter at $\lambda_{eff} = 880$ nm are used for solar imaging, and the L1 filter at $\lambda_{eff} = 739$ nm is a broadband filter¹. The remaining 13 filters are narrowband "geology" filters, whereas two of the right eye filters are overlapping with two of the left eye filters, providing the capability of stereoscopic imaging. The Pancam system is therefore able to provide 11 point spectra, spanning the full wavelength range of 432 nm-1009 nm.

Besides being a navigation tool used to direct the rover in directions of interest, Pancam is also an important diagnostic system in the search for mineral, and especially hydrological signatures

¹ λ_{eff} here denotes the "effective", or central wavelength, defined as the weighted mean wavelength of the band, where the weighting is by the sensitivity of the whole camera system (filter, optics, detector).

in rocks and soil on the martian surface. By analysing the 11 point spectra, it is possible to detect signatures of hydrated minerals along the traverse, providing insight into present and past history of water on Mars. Therefore, it is of course important that the data collected can be trusted with full confidence across the entire wavelength span.

1.1 Motivation and aim for this thesis work

In order to extract spectra from Pancam images of Mars and being able to convert them into units that makes it possible to do meaningful analysis, a baseline has to be established by photographing an object with a known response. Such an object, called the calibration target (hereafter referred to as the cal target) is mounted on each of the rover decks. It consists of a number of different coloured areas, each giving rise to a very specific, known response signature, characterised and modeled pre launch. The caltarget is by far the most photographed object on Mars by Pancam, since it is photographed several times each day, so that images obtained with Pancam of the martian surface can be calibrated correctly. Further detail on this is provided in Section [3.1.1.3](#) and [3.1.2](#).

During preflight calibration activities of Pancam in 2002, it was however noted that the R7 filter at $\lambda_{eff} = 1009$ nm gave a slightly different output than what was expected based on the model. That is, the spectra obtained from images of the calibration target using Pancam, and the ones obtained using a spectrometer showed a discrepancy in the R7 filter. This fact is depicted in Fig. [1.1](#). The panels (a) and (c) show the expected response based on modeled data for the different coloured areas of the cal target, the panels (b) and (d) show the extraction from Pancam images from the same regions of the cal target. Panel (d) shows clearly how the contrast between the white and black area of the caltarget in the R7 filter at $\lambda_{eff} = 1009$ nm is less than what the model in panel (c) predicts. In more detail, the top curve in panel (d) dips down at 1009 nm, and the bottom curve has an upturn at 1009 nm, which is not seen in the model plot of panel (c). At the time it was believed that the cause for the lower contrast was absorption properties of the cal target material at that specific wavelength (mentioned in [10] in the caption of Figure 20 and on page 24). No further investigations were performed to clear up the question, as the effect was deemed small enough to not affect the image data analysis. However, prior to the onset of this project, an indication of the effect being present in not only images of the cal target, but also in images of the martian terrain, was discovered by Kjartan Kinch. The presence of the effect in in-flight images made it much more likely that the effect is an instrumental inherent effect rather than it being related to properties of the cal target materials. This in turn made it critical to perform further investigations into the nature, and potential cause of the effect, the ultimate goal being to remove the effect from affected images. This idea was what initiated the work presented in this thesis, and it will therefore be discussed in much more detail later on.

Based on the a priori knowledge of the R7 effect, two overall aims for this work were formulated:

1. The primary aim was to establish a successful way of removing the effect from R7 filter images without compromising the scientific merit of the data.
2. The second aim was to get a sense of the cause of the effect.

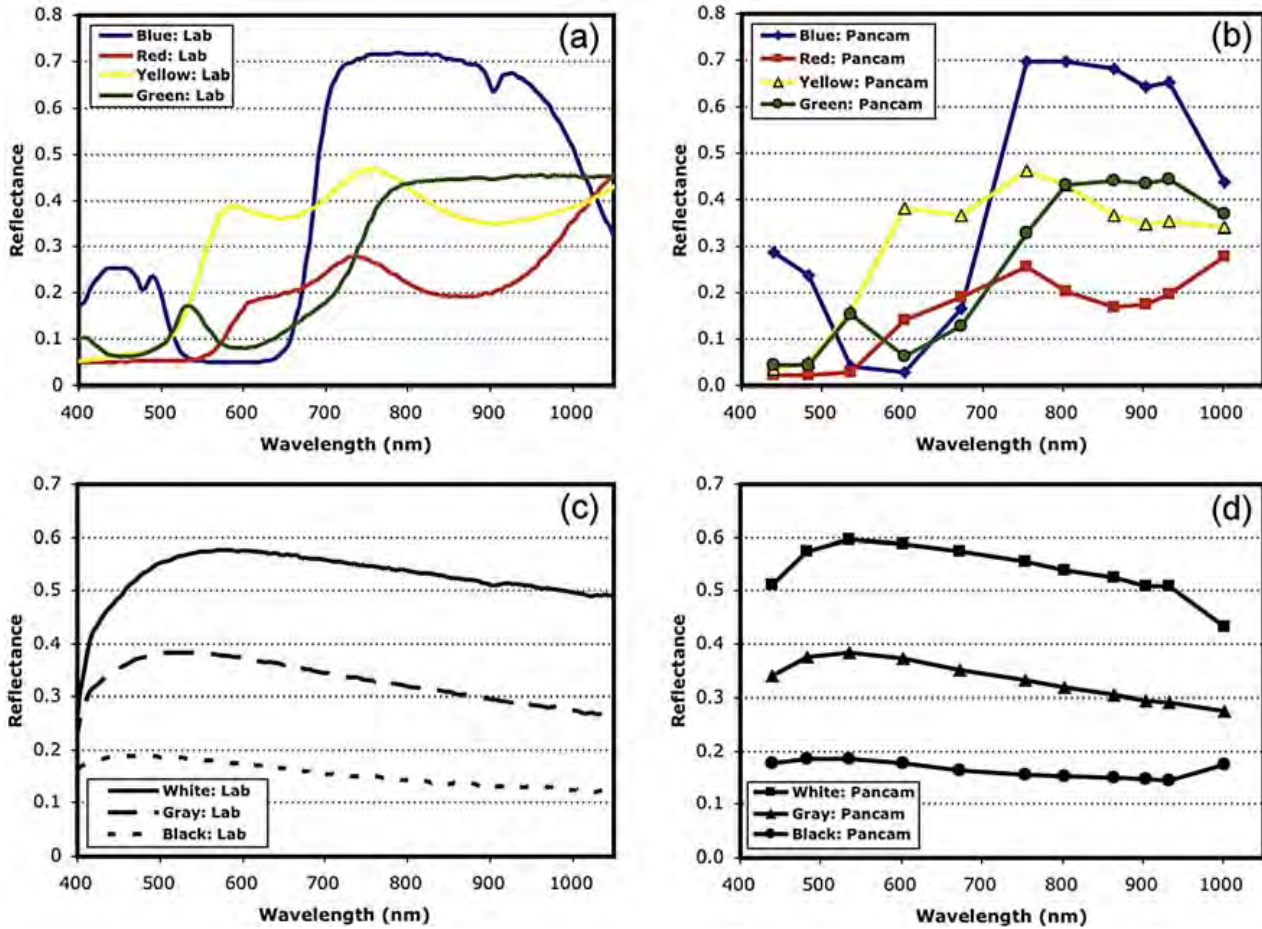


Figure 1.1: A plot of the results from pre-flight calibration of Pancam, where a comparison was made between spectra of the cal target obtained with a spectrometer, and spectra extracted from Pancam images of the cal target. A discrepancy is seen between the modeled spectra and the Pancam spectra for the white and black areas of the cal target, indicating a lower contrast than expected. Figure from [10]

As we will see in Chapter 4, the secondary aim was dependent on a hypothesis being made before any work was done towards fulfilling the first aim, and for that reason the two aims are somewhat intertwined. The key point here is therefore that the work presented in this thesis can be thought of as a success if only the first aim turned out to be fulfilled, with the second aim being merely the "cherry on the top".

The initial part of the work presented here will reveal, that:

- The effect is most likely related to properties of the Pancam instrument itself.
- The effect is more prominent than first thought, and could be significant and potentially affect the outcome of analysis based on R7 images.

The further work presented will show that:

- It was possible to develop a correction with the ability to remove the effect from the images obtained with the R7 filter.

- Not only can the effect be removed to a high degree, the correction algorithm does not introduce any new artifacts, or distort the corrected images in any way. Thereby fulfilling the first aim above.
- It is possible to provide a likely explanation for the cause of the effect using the hypothesis established prior to the actual work done. Thereby fulfilling the second aim above.

1.2 The structure of this thesis

Chapter 2 will provide an overview of Mars as a planet, as well as provide a walk through of why Mars continues to be a corner stone in the search for evidence of life elsewhere in the Universe. The work presented in this thesis is mainly instrumental oriented and will therefore not cover in depth the geology, geography, or mineralogy of Mars. The introductory part of this thesis is meant as a general overview of those subjects, since they are important in regards to the motivation for this work.

Chapter 3 focuses on the instrumental aspects of the MER mission, more specifically, the technical aspects of the Pancam system. This includes a description of the multispectral capabilities of the camera, as well as the detector construction, which both provide the base knowledge needed to establish a hypothesis for how the effect arises, as we will see later. Furthermore, a short overview of the calibration procedure for Pancam images is provided at the end of the chapter, and an overview of some of the key MER results so far.

Chapter 4 presents the R7 effect in more detail than provided above. The methods used to characterise and model the effect are described, and the results are evaluated, since they create the basis for the further work towards fulfilling the first aim written above.

Chapter 5 describes the construction of a correction algorithm capable of removing the R7 effect from affected images. The method used for the correction, an iterative deconvolution algorithm is described with a solid base in the physical understanding of the process.

Chapter 6 demonstrates the performance of the correction algorithm on different types of images.

Chapter 7 evaluates the performance of the correction and provides insight into the interpretation of the corrected images.

Chapter 8 sums up and conclude on the fulfillment of the thesis aims listed above.

2

Mars: fundamentals and basic geology

The planet Mars has been an object of fascination to humans since the earliest of times. Writings mentioning Mars have been found dating as far back as to ancient Egypt. The exact distance between Mars and Earth changes between oppositions, since the Mars orbit is eccentric with an eccentricity of 0.093 (compared to the eccentricity of the terrestrial orbit of only 0.017), but the planet has always been visible to the naked eye, as a glowing, red dot in the night sky. In the 1600s, the first observations of Mars through a telescope revealed structures on the surface that had until then been hidden by the large distance. By tracking one of these structures (which were later identified as Syrtis Major), the rotational speed of the planet was determined. One day on Mars is approximately 40 minutes longer than a day on Earth, at 24 hours and 39.6 minutes, called a Sol. The length of a martian year, that is the time it takes Mars to complete one entire orbit around the Sun, is 687 Earth days (or 669 Sols). The very early observations of Mars revealed that the structures covering the surface were connected with thin lines. At that time it was largely interpreted as canal-like structures build by intelligent beings [20]. Observations of the climate on Mars was done as early as the 1700s, where changes in the polar caps were noticed, and interpreted as seasonal changes on the planet. Today, it is a fact, that since the obliquity is near the value for Earth, Mars experiences similar seasonal changes as we experience on Earth. However, since the orbit eccentricity is larger than for the Earth, the seasonal changes is experienced differently for an observer on the North pole, than for an observer on the South pole. Since the planet is further away from the Sun during winter in the southern hemisphere than it is during winter in the northern hemisphere, the winters are colder and longer in the southern hemisphere. For the same reason, the summers in the southern hemisphere are shorter and warmer. Generally speaking the seasons are more extreme in the southern hemisphere compared to the northern.

More general details of the composition of the planet, determined much later in history, revealed how the harsh conditions on the planet made it unlikely to find any form of life immediately on the surface. Mars is a differentiated planet like the other terrestrial planets. That is, it consists of a core, mantle and crust. The core is expected to be solid, since Mars does not have a magnetic field, but the size is unknown since no seismic data is available, however, the Insight Mars Lander which arrived on Mars on November 26, 2018 brought with it a seismometer, which will finally reveal the insides of the martian planet, as well as determine the amount of seismic activity on Mars today[48]. The mantle of all the terrestrial planets are thought to be very similar in composition, as they are all created from the same cloud of dust. The composition of which is expected to be well represented by chondritic meteorites. Partial melting of chondrites yield basaltic magma, and it is therefore not surprising that the martian crust has a mainly basaltic

composition. From basaltic magma, we get a whole range of different types of rock all under the name basalt. Basalts are fine-grained rock containing equal amounts of plagioclase feldspar¹, pyroxene², and less than 20% of other minerals, mainly olivine³. The Martian meteorites are also almost all basaltic in nature.

The gravity on Mars is lower than on Earth, since the mass of the planet is only 10% of the mass of the Earth, which gives an acceleration of gravity of 3.69m s^{-2} . The atmosphere of Mars is very thin, which means that the atmospheric pressure has an average value of 600 Pa compared to an atmospheric pressure of 101325 Pa on Earth, at sea level. The composition of the martian atmosphere is also different from that of the atmosphere on Earth, since $\sim 96\%$ consists of pure CO_2 , whereas the rest is made up of small amounts of Argon, N_2 , O_2 , and CO[50]. The thin atmosphere makes it impossible for Mars to have a working greenhouse effect, leading to the planet's temperature being completely controlled by the influx from the Sun. This creates a big temperature fluctuation between night and day, since the surface turning away from the Sun cools down fast when it does not receive Sun light. A common air temperature on Mars is -55° , but it can vary greatly. For instance, the air temperature at Gale crater, the landing site of the Curiosity rover, was measured to a record low of -127° and a record high of 20° . Additionally, the thin atmosphere provides almost no protection from Solar radiation, and the absence of a magnetic field makes the planet vulnerable to charged particles.

However, the search for past and present life on Mars is still the main goal set by key missions to the red planet, since much evidence points towards a past climate much more accommodating to the presence of life, as we will see in this chapter. The main ingredient necessary for life, as we know it, to exist, is water. Even extremophiles, which are able to survive in very harsh environments on Earth, need some amount of water to exist. Without water, a key element in the cycle of life would be missing, and for this reason, the search for life is focused on this main criteria. As we will see in this chapter, the existence of water ice just below the surface of Mars allows for a range of possibilities for the past existence of life, that potentially could have survived to this day. This chapter provides first an introduction to the fundamentals of the red planet, and continues by establishing a timeline for the geological history of Mars to be used throughout the rest of the chapter. The chapter furthermore contains a brief walk-through of the most characteristic morphological features of the surface caused by impacts, volcanic activity, and wind erosion, to establish the past and present conditions on the surface. The second to last section of this chapter presents a short introduction to the hydrological history of Mars as established through remote and in situ observations. The last section provides a short overview of the past climate on Mars.

2.1 Martian surface morphology

The change in the surface of a planet and other celestial bodies is dependent on a long list of factors. On Earth, the changes brought on by geological dynamics, such as volcanic activity and plate tectonics, wind- and water erosion, and biology, constantly renews the surface and erase much of the evidence for past climate conditions. However, on Earth it is possible to dig into the ground and search for related signatures. For instance, large ice cores are drilled

¹Plagioclase feldspar - a mineral of aluminosilicates with sodium and/or calcium.

²Pyroxene - a group of silicate minerals with calcium, magnesium, and iron.

³Olivine - a group of silicate minerals containing magnesium, and iron

from the ice sheets from which it is possible to identify and track the climate conditions back in time. This is of course not possible on Mars, since most knowledge is acquired through remote sensing. The knowledge we do acquire from in situ missions, such as the rover missions, are not as extensive as what can be achieved on Earth, for good reason. Despite the wish to do similar investigations on Mars, the constraints set by the mission in terms of weight and volume of the payload, does not yet allow the rovers to investigate more than the upper ten centimeters of the martian crust. Even if it was possible to dig deep into the ground it would be close to, if not impossible, to perform the necessary investigations without having access to Earth based laboratories. Until the plans for a sample return mission becomes a reality, the information we acquire must come from the telescopic data of the surface, and from the in situ observations and measurements of the top layers done by lander missions. However, this collection of data is able to provide us with a vast variety of information which we will only graze the surface of in this chapter.

2.1.1 Craters

One of the most distinguishable features of the martian surface is its craters, and the morphology of a martian crater can tell us a lot about the processes that have been in play since it formed. One distinguishes between two types of craters. The simple craters and the complex craters. Typically the craters become more complex with size, so that those smaller than ~ 3 km will have a simple, deep bowl shape, those with a size of ~ 3 -130 km in diameter will be complex with a central peak and terraced walls⁴, and those larger with a diameter of ~ 130 -1000 km will exhibit an inner ring instead of, or in addition to the central peak. The largest craters have no inner ring.

The craters on Mars are one of a kind in the sense that many of them have very distinct ejecta patterns that generally look fluid. The craters will typically show lobes of ejecta, where each lobe will terminate in a low ridge. The pattern is thought to arise from fluidization of the ejecta because of volatiles in the surface.

Generally one can think of the process of formation of a crater on Mars as the action of a water droplet falling into water. When the impacting projectile hits the surface, two shock waves arise - one traveling downward and into the target, while the other travels upward and through the projectile. Pressure reaches the GPa regime exceeding strength of material resulting in the material acting like a fluid. The material caught and compressed between the two shock waves is spewed out as lateral jets. Much as a droplet of water, a central peak will form containing material from the deepest layers, excavated by the impact. The material will only act as a fluid in a short amount of time, and the peak solidifies and "freezes" into place. A curtain of ejecta is sent outwards after an impact with the coarser and bigger debris traveling slowly and close to the ground, and smaller debris traveling with larger velocity and further away from the crater. This means that the ejecta falling close to the crater will barely disturb the surface, whereas debris falling further away will hit the ground with large velocity and mix with the surface material. These secondary, minor craters formed by falling debris will typically be irregular in shape when falling close to the main rim, whereas ejecta traveling hundreds of kilometers might form circular craters that are difficult to distinguish from primary craters. The importance of being able to identify individual primary craters, and being able to describe their morphology will be evident from the following section.

⁴Terraced crater walls - steep walls of an impact crater with stair steps created by slumping due to gravity and landslides.

2.1.2 Martian epochs - the geological timescale

In general, since only limited surface renewing, geological activity is present on Mars today, the Martian surface holds a vast record of past events in the form of structures and landforms on the surface, the main structure being that of craters as mentioned above. How well the impact crater morphology is preserved, depends on the location of the crater, since other impact events, or surface processes such as wind erosion, water erosion, and lava flows could modify the crater morphology severely. In order to establish a relative timeline for the past martian history, a method was developed where the density of craters on the surface is related to a relative timescale. This is why, as mentioned in the previous section, it is important to be able to distinguish between individual craters, and determine their degree of modification by surface processes.

The surface of all planetary bodies and natural satellites in the Solar System is scarred to some degree. On Earth the scarring has been largely hidden by the renewal of the surface. The surface of the Moon, on the other hand, has a valuable record of impact events. Its lack of an atmosphere and geological activity has preserved all its craters in a pristine condition making it possible to count the individual craters in an area. Furthermore, the sample return from manned missions to the Moon have made it possible to date different parts of the surface and relate that to the number of craters in the same area. Making the reasonable assumption that the rate of impacts and the size distribution of the impacting objects has been largely the same for the Moon and Mars, one can use the relation between the absolute ages and crater density on the Moon, directly on Mars.

One distinguishes between three epochs, or stratigraphic systems for the geological history of Mars. The Noachian age refers to the earliest time in the history of the planet and the absolute age is estimated to stretch from 4.1 Gy to 3.7 Gyr ago, coinciding with the Late Heavy Bombardment, where the terrestrial planets were struck by an abnormally large number of meteorites in a relatively short amount of time. During this epoch, Mars is believed to have been potentially warmer, geologically more active, wetter, with a denser atmosphere, and generally more Earth-like than what it is today [41]. It is worth noting that the Noachian epoch roughly coincides with the period on Earth when the first lifeforms started to appear. Following the Noachian age, the Hesperian age stretches from roughly 3.7 Gyr ago to 2.5 Gyr ago, when the planet became gradually colder, and drier. The Amazonian age is the youngest and started at approximately 2.5 Gyr ago and stretches until present day. This epoch is characterized by the cold, dry climate, and the sparse atmosphere we see today on Mars.

A plot showing the relation between the cumulative number of craters larger than 1 km in diameter and time can be seen in Fig. 2.1. The different stratigraphic epochs are indicated on the plot. It follows that this method of dating relies on the counting strategy being as exact as possible. As mentioned in Section 2.1.1, secondary craters can be difficult to distinguish from primary craters, when the ejecta travels far, and can therefore pose a problem in dating the surface. Furthermore, extensive erosion and weathering of craters can erase their features and make it difficult to determine their size. However, generally speaking, most craters located on the Hesperian and Amazonian plains (the youngest areas) are well preserved, since very low erosion rates are expected for the post-Noachian period. Craters in the Noachian terrain are highly eroded. They have been filled up slowly over time, which has erased the presence of the central peak, and wind erosion has removed the ejecta patterns formerly surrounding the craters.




Figure 2.1: A plot showing the cumulative number of craters larger than 1 km in diameter against age based on the Hartmann and the Neukum models. The labels indicate the epochs. Figure from [41].

2.1.3 Volcanoes

Volcanism on Mars probably started very early in the martian history, but young crater ages, especially in the areas of Cerberus and Amazonis, as well as the recent discovery of a cluster of volcanoes inside the young valley network, Valles Marineris [18], suggest that volcanic activity could have occurred quite recently. The largest volcanoes on Mars lie in the region of Tharsis, but other volcanic regions count Elysium and the area around Hellas crater⁵. The morphology of the majority of the Martian volcanoes resembles terrestrial shield volcanoes, characterized by a gentle slope near the foot of the volcano that continues as a convex shape with a caldera at the top. An excellent example is Olympus Mons, the highest volcano in the Solar System, rising 21 km above the surface of the planet (see Fig. 2.2).

Since there is no immediate evidence of plate tectonics being present on Mars during the past 4 Gyr, the martian volcanoes are most likely of the intraplate type. This type of volcano is not created by plate movement but rather by a mantle plume beneath the surface, which consists of warmer, lighter material that rises towards the surface by convection. Given the highly basaltic composition of the surface of the planet, the volcanoes on Mars are most likely basaltic⁶ in nature, as such volcanoes are on Earth. However, the basaltic volcanism is expected to express itself differently on Mars than on Earth because of the lower gravity, lower atmospheric pressure, and the different structure of the crust. Because of the low gravity, the buoyant forces that help the magma rise through the mantle are smaller and the chambers form deeper in the mantle than on Earth. The size of the magma chambers then needs to be correspondingly bigger for the magma to break through the crust of the planet. This means that the volcanic eruption rate is much smaller on Mars, even though each eruption event is more violent. Another consequence

⁵Hellas crater - a roughly circular crater in the Southern hemisphere of Mars, measuring 2300 m in diameter and reaching a depth of 7152 m.

⁶basalt - the rock type that is left behind after cooling of lava rich in iron and magnesium (mafic) minerals.




Figure 2.2: The Tharsis bulge. The largest volcanic region on Mars. Image from [41].

of the low gravity is that the lava flows a greater distance, up to 6 times longer than lava on Earth. As a result, lava plains cover about 60% of the planet surface[41][50].




Figure 2.3: Lava plains in the Tharsis region: A) A tube-fed flow between two sheet flows nearby Ascreus Mons. B) Relics of lava tubes near Tharsis Tholus. C) Sheet flows from Pavonis Mons. Image from [41].

Plains with many flow patterns are generally located close to the center of the two main

volcanic provinces, Tharsis and Elysium, and two different kinds of lava flows can be distinguished in these areas: sheet flows, and tube-fed flows. Sheet flows are characterized by the lava not being confined into channels, but rather spread out like lobes over a large area (see Figure 2.3). Tube-fed flows refer to the event where the lava flowing closest to the surface solidifies and creates a "roof" over the still moving part of the lava. When the remaining lava leaves the tube, the "roof" collapses to form a channel or line of pits (see Figure 2.3). Such structures are found all across the volcanic regions, and could in some cases be misinterpreted as structures left behind by water erosion, as we will see in Section 2.2.

2.1.4 Aeolian features

Wind is accountable for most of the changes to the surface of Mars today. However, wind erosion is still minor, since the thin atmosphere only allow erosion of the top layer of porous sedimentary material and the colour variations of the surface is therefore mostly caused by relocation of dust. The wind is able to lift dust into the atmosphere and scatter it across great distances during global dust storms, whereas dust devils account for the more local dust transport. On Earth, the dust being transported by the wind will eventually come to rest in the oceans and be incorporated into the rock record. On Mars, some ocean burial might have occurred in the distant past, and it is expected that some dust is being trapped in the northern polar cap each dust storm season. However, much of the dust carried about in the atmosphere could potentially be very old and date back to even the Noachian epoch.

Wind eroded rocks (called ventifacts) are scattered across the surface of Mars, and are recognized by their flat smooth surfaces that meet in sharp edges (see Fig. 2.4). The smooth rock faces are created by long-term sand abrasion. Wind erosion is also responsible for removing the top layer of sedimentary deposits, which in many cases exposes the stratigraphy below. In the bottom image of Fig. 2.4, the stratigraphic layers are just visible in the foreground of the image as thin horizontal lines in the exposed rock. Only minor signs of wind erosion on volcanic plains is seen, where most fine features of volcanic flow are still recognizable. So called mega ripples (with a wavelength of a few centimeters to a few meters) and larger structures that resemble regular sand dunes are visible all over the planet. They are found at the bottom of craters and in other low areas, and an example of such a structure is shown in the bottom image of Fig. 2.4

2.1.5 The Martian dichotomy

Mars has a global dichotomy that can be seen in both elevation, crustal thickness, and crater density (see Fig. 2.5). The southern uplands are on average 5.5 km higher than the northern plains and the elevation difference offsets the planet's center of figure from its center of mass by 2.986 km measured along the polar axis. The crustal thickness also varies from 5.8 km to 102 km, with the thinner crust in the north. The crater density is observed to be higher in the southern regions compared to the generally more smooth plains of the North. This, however might just be a superficial feature, since the northern region is heavily cratered at a shallow depth. Speculations have been made, whether the large depression, that is the northern hemisphere, could have been left behind by a large impact event. The fact that no obvious impact signatures are observed (as compared to what is seen for the Hellas crater), indicates that the impact would have happened very early in the history of Mars. Many features along




Figure 2.4: Top left image: Example of a ventifact on Earth. Top right image: Sand ripples on Mars. Image obtained by the MER Opportunity in the Endurance crater. Ripple area measures approximately 30 m across. Bottom image: Example of wind erosion of top sedimentary layer, exposing the stratigraphy of the bedrock. Image obtained with the THeMal EMission Imaging System on board the Mars Odyssey spacecraft. Image credit: NASA

the dichotomy border have been found that suggest a lacustrine origin, such outflow channels and deltas. The largest canyon system on the planet, called Valles Marineris is found at the border [22]. Based on these findings, it has been proposed that the northern plains is the relic of a former northern ocean [41]. This idea is however the focus of much controversy.

2.2 Water on Mars

Large amounts of water exists on Mars today, but since water cannot exist in liquid form on the surface, it is found in the form of ice[2][41][1]. During summer in the Northern hemisphere, the cap of dry (CO_2) ice evaporates and exposes a cap of water (H_2O) ice. A similar water ice cap is expected in the Southern hemisphere below the CO_2 ice cap, but is never exposed, since the Southern CO_2 ice cap never evaporates entirely. However, recently, a survey using MARSIS (Mars Advanced Radar for Subsurface and Ionosphere Sounding), a low frequency radar on board the Mars Express spacecraft revealed the presence of a stable body of liquid water beneath the South polar cap[6].

Mars has a kilometer-thick cryosphere, which is defined as the part of the crust that continuously has a temperature below the freezing point of water. The thickness of the cryosphere varies with latitude, and reaches its maximum depth (~ 10 km) at the polar regions [1]. It has further been




Figure 2.5: The Martian global dichotomy could have been made by a large impact early in the Martian history. Image from [41].

proposed that beneath this layer of ice, liquid groundwater may be present today. If this was the case, it would be expected that features on the surface of volcanic plains would be left behind from the violent interaction between the ground water and magma traveling through it towards the surface. Such surface features have been proposed several places on Mars, by comparing them to known examples from Earth. However, no definite conclusions have been made, since for instance impacts could cause similar signatures. Efforts towards detecting a global groundwater table using radar has been made[1], but as of yet has not borne fruit. How deep the groundwater table lies is also dependent on the amount of perchlorates present, since these types of salts act as a freezing point depressing agent, and if the depth is large, the detection will be highly dependent on the penetration depth of the radar[1].

2.2.1 The hydrological history of Mars

Both observations of morphological surface features as well as the detection of spectral signatures from what is called hydrated minerals are used in order to identify areas on Mars that have been effected by water in the past. A hydrate is defined as the product left behind when water has been added to the crystal structure of an existing mineral, resulting in a new mineral. For instance, clay minerals are hydrous aluminium phyllosilicates, which means that the aluminium phyllosilicates have been altered by water. The detection of such minerals is therefore a strong argument for water having been present at the point in time, when the clay minerals were formed. The detection of secondary alteration products, such as carbonates and sulfates can also be used as an indicator for water, since they are created as byproducts of weathering or hydrothermal alterations.

Even though liquid water cannot flow on the surface of Mars today, plenty of surface features indicate that liquid water has been present on Mars in the past. South of equator, East of Tharsis Montes, on the dichotomy border, a collection of large canyons called Valles Marineris,

is located. The canyons most likely result from faulting⁷, since they are oriented radially from Tharsis, and therefore is located at a place of high tension caused by the Tharsis bulge. Several of the canyons show signs of water erosion. One of those canyons is called Juventae Chasma, and its morphology points towards water having overflowed the canyon and created a large outflow channel named Maja Vallis (see Figure 2.6). The water could have originated from a ground water eruption in connection with the faulting[41][23]. The elevation of the outlet is ~ 4 km above the canyon floor, which means that a lake would have been left behind. The sediments seen at the bottom of Juventae Chasma are sulfate rich which supports this hypothesis, since sulfates can form by evaporation of saline water. It is however important to note that this is not the only way that sulfates can form. The presence of sulfates is therefore not a direct evidence of water, but when it is held together with other evidence, it can be used to support the argument. The thick sediments seen in many of the canyons have a horizontally layered structure, which is consistent with them being deposited in lakes, and this together with the signs of outflow channels from local depressions, and the detection of sulfates, makes the water hypothesis even more likely. Evidence against the lake hypothesis is that not all sediments are horizontally layered. Some places they bare resemblance to sediments originating from pyroclastic flows, which is when volcanic matter (ash and steam) acts as a fluid.

The biggest outflow channels terminate at Chryse planitia. This region contains several streamlined teardrop-shaped islands, that are suggested to have formed by large amounts of flowing water in violent, large-scale, flood events.

When gullies were found on slopes facing the poles, it was proposed that they could be formed by seeping groundwater, and more interestingly, that they were formed at present day. However, this explanation had some difficulties, since groundwater is unlikely to break through the surface at present conditions on Mars with a mean annual temperature of 215 K (-58°). To accommodate this, another theory was proposed, that the gullies could be formed by melting ice[24]. However, recently, the liquid water interpretation was challenged by [3], who showed that dry transport of sedimentary material is able to create much the same features.

Since life is in need of stable conditions to form, the search for ancient lakes has been of high priority for the planning of rover and lander missions. Gale crater was indeed selected as the landing site of the Mars Space Laboratory mission based on its high likelihood of being the relics of a former lake. Sedimentary structures on the crater floor were consistent with deposition in a standing body of water. Delta formations, which are sedimentary fan deposits originating at the crater rim and terminating in lobe structures on the crater floor are found all along the rim of the crater. The delta formations have been weathered to expose a layered structure of materials which further supports the lacustrine interpretation. The Curiosity rover has later confirmed that the sedimentary deposits on the crater floor contains a wide variety of clays, which, as mentioned above, is the low-temperature hydrothermal alteration products formed by hydration of aluminum phyllosilicates[25][51][53].

2.2.2 Past climate conditions

Most of the valley networks on Mars is dated back to the Noachian, with some continuing to form into the Hesperian and later. If the valley networks and canyons were indeed filled with

⁷A fracture in a surface caused by movement of the rock material




Figure 2.6: An image showing Juventae Chasma. A large outflow channel called Maja Valles can be seen to the North as soft bending elongated structures, indicating that the Canyon was overflowed with water sometime in the past. At the point A, sulfate-rich sediments are observed, supporting the liquid water hypothesis. Image from [41].

water, the question is how they were filled. As mentioned above, eruption or seepage of groundwater is one hypothesis, that is of course dependent on the existence of a ground water table below the cryosphere. Also rainfall and snowfall could explain the origin of the water, however, this explanation requires a vastly different climate than what is observed on Mars today. Water supply by rainfall would indicate that the early epochs on Mars were characterised by terrestrial conditions with a much warmer climate with no cryosphere. As mentioned in the introduction, it is worth noting that the Noachian epoch coincides with the era on Earth where the first life forms started to emerge. Another possibility is that the lakes and other liquid water features formed in a colder climate, as a result of melting of ice on the surface, or melting of ground ice by dikes. The detection of hydrated minerals on the surface, such as clays, does not exclude this possibility, since the clays could have formed below the surface at that time and could have been brought to the surface by later weathering[25]. Whether the Noachian epoch was warm or cold is therefore not completely understood, since it is generally difficult to explain how Mars could

have been warmer. Warming by a $\text{CO}_2\text{-H}_2\text{O}$ greenhouse effect alone is not sufficient, because of the low solar flux in the early age of the planet, but other explanations could be warming by volcanic eruptions, or large impacts[54].

It is difficult to explain how the early Mars could have been warmer, but it is even more difficult to explain how some of the valleys were formed very recently, since the climate became significantly colder after the Noachian. A theory is that the formation could have been facilitated by a temporary global warming caused by a release of CO_2 and possibly SO_2 into the atmosphere[41].

3

Mars Exploration Rovers (MER)

Surface operating missions provide a unique perspective on the hydrological history of Mars compared to what can be achieved by orbiting spacecrafts. Being able to investigate sites of special interest up close means that theories made based on remote sensing by orbiters can be tested in situ. Additionally, the lander missions make it possible to expose layers of soil that otherwise would not have been accessible.

The first lander missions Viking 1 and 2 launched in 1975 and made history as they took the very first images from the martian surface. The landers carried biology experiments with different phases customized to test the martian soil for contents of biology. No organic molecules were detected, however, one of the biology experiments designed to search for signs of metabolism in a sample of soil, showed a positive result. This result was later interpreted as a solely chemical reaction between a strong oxidant in the soil, and the water that was added during the experiment. This hypothesis was supported when the Phoenix lander in 2008 detected the presence of low concentrations of perchlorates, a strong oxidant, in the top layer of soil[5]. Since then a number of other missions have successfully placed landers on Mars. The Sojourner rover became the first rover on Mars when it arrived on board the Pathfinder lander in 1997. The main objective was to determine whether the landing site had been experiencing flood events and with the discovery of a number of large boulders that showed signs of water transport, this idea was supported [5]. A secondary but also important objective was to test new technologies that were later used for the Mars Exploration Rover mission.

In 2004, three weeks apart, the two rovers Spirit (MER A) and Opportunity (MER B) touched down on opposing sides of the red planet. The Spirit landing site at the Gusev crater (location: 14.5684°S 175.472636°E), and the Opportunity landing site at Meridiani Planum (location: 1.9462°S 354.4734°E) were chosen based on both sites showing signs of water alterations. The Gusev crater has a delta-like deposit suggesting an inflow of water from the valley network Ma'adim Vallis located North of the Gusev crater. Furthermore it has several deposits across the crater floor that exhibits a layered structure indicating that it was deposited in standing water. The upper floor plains consists of remains of basaltic lava flows dated back to the early Hesperian [7], which are superimposed on the older Noachian age crater floor.

The Meridiani Planum site does, on the other hand, not exhibit any mentionable morphological features indicative of fluvial activity. The motivation for choosing this site rather stems from the spectroscopic evidence of the mineral grey hematite, which is an iron oxide that is generally formed in aqueous environments on Earth[52].

Other factors for site selection were also taken into account, since it was important to avoid any

largely hazardous terrain in the form of large boulders, high wind activity, and slopes. Furthermore, since both rovers were solar powered, the landing sites were required to be located not far from the Equator. The planned mission span was 90 Sols for each rover, which they both surpassed. After the termination of Spirit's mission in 2010, when the rover got stuck in soft soil, Opportunity continued to collect data up until June 3, 2018, when a dust storm cut off the connection to the rover[19]. Efforts are presently being put into recovering the connection. The Viking 1 lander held the record for the longest working surface mission on Mars, until the record was broken by the Opportunity rover in 2010.

3.1 MER instrumentation - the Athena payload

The two rovers Spirit and Opportunity are identical in construction and design, each carrying the Athena science payload consisting of two main parts, the Panoramic Mast Assembly (PMA) and the Instrument Deployment Device (IDD) (see Fig. 3.1). The PMA makes up the remote sensing part of the payload and count the Panoramic Camera (Pancam) and the Miniature Thermal Emission Spectrometer (Mini-TES). Pancam is capable of obtaining high-resolution, colour, stereo imaging, and Mini-TES collects spectra cubes at mid-infrared wavelengths, and is able to measure the temperature of the different surfaces encountered by the rovers. For calibration of Pancam, a calibration target is mounted on the rover deck. The IDD is a robotic arm holding four additional tools used for in-situ investigations, the Microscopic Imager (MI), and Alpha Particle X-Ray Spectrometer (APXS), a Mössbauer Spectrometer (MB), and an Abrasion Tool (RAT). The MI is used for close up imaging of soil and rocks, the APXS is for elemental analysis, the Mössbauer spectrometer is for mineralogical analysis of iron-bearing minerals, and the RAT is for removing the top layer of soil, in order to expose the hidden unaltered surface for close up imaging with the MI or elemental characterization using the APXS. Additionally, the payload includes a Magnetic Properties Experiment that is designed to collect magnetic particles, which makes it possible to investigate the magnetic properties of the martian dust. Finally, in addition to the scientific instruments, the rover has six engineering cameras: two Navigation Cameras (Navcams), which are two wide-angle monochromatic cameras, and four Hazard Avoidance Cameras (Hazcams), which are fisheye monochromatic, stereo cameras, two pointing forwards and two pointing backwards[12][16][10][8].

In this section we will focus on the Pancam system, and not provide any further detail on any of the other instruments that are part of the Athena payload, since these are not directly relevant for the present work.

3.1.1 Pancam

Pancam is an important tool in the task of identifying and describing the morphology of the landing site, but it is also used for navigation. The Pancam system consists of two cameras separated horizontally by 30 cm, making up the eyes of the rover (see Fig. 3.1). With its high resolution of 0.27 ± 0.02 mrad/pixel, the cameras have a spatial resolution of 2.7 cm/pixel at a range of 100 m, which gives them a three times higher spatial resolution compared to the Pathfinder and Viking Landers, and is equivalent to a 20/20 human vision.

The camera lens is a Cooke triplet that provides an effective focal length of 43 mm and a focal ratio of f/20, giving each camera a Field of View (FOV) of $16^\circ \times 16^\circ$ making it possible for




Figure 3.1: A diagram of the instrumental layout of the two MER rovers Spirit and Opportunity.
Image Credit: NASA/JPL/Cornell

the cameras to maintain focus from infinity to 1.5 meters from the camera. Using its imaging capability of 360° in azimuth and $\pm 90^\circ$ in elevation, it has the ability to acquire a panorama view of its surroundings. From these images it is possible to obtain preliminary information of composition, morphology, and texture of specific targets of soil or rock from a distance before the decision is made to direct the rover to the location. On location, other instruments can be used for in-situ observations and investigations. Also, using the high spatial resolution, and multispectral capability of Pancam, it is possible to acquire information about the mineralogical composition of the surrounding terrain from a distance [12].

3.1.1.1 Filters

Each camera has a small 8-position filter wheel attached in front of the optics. The 8 filters per camera cover the wavelength range 432 nm-1009 nm with the filters denoted R1-R8 for the right eye and L1-L8 for the left eye (see Table 3.1 for an overview). Two of the filters (L8 at $\lambda_{eff} = 440$ nm and R8 at $\lambda_{eff} = 880$ nm) are used for solar imaging, and one is a broadband filter (L1 at $\lambda_{eff} = 739$ nm). λ_{eff} here denotes the "effective" or central wavelength, defined as the weighted mean wavelength of the band, where the weighting is by the sensitivity of the whole camera system (filter, optics, detector). The remaining 13 filters are narrowband "geology" filters. Two of these are overlapping bands for the right and left eye, making it possible to obtain stereoscopic images at these specific wavelengths. This leaves the possibility of obtaining 11 point spectra spanning the full wavelength range [10]. Generally, the filters are arranged so that the shorter wavelength filters (< 750 nm) are positioned in front of the left eye camera, and the longer wavelength filters (> 750 nm) are positioned in front of the right eye camera. An image of the normalised response in all filters is shown in Fig. 3.2. The Pancam filters

have been chosen specifically to detect absorption from hydrated minerals, as well as other key minerals, such as iron-bearing, primary basaltic minerals, such as pyroxene and olivine.

Table 3.1:
Pancam Filters

Filter name ¹	λ_{eff} [nm]	Band Pass [nm]	Type
L1	739	338	no filter
L2	753	20	red stereo
L3	673	16	geology
L4	601	17	geology
L5	535	20	geology
L6	482	30	geology
L7	432	32	blue stereo
L8	440	20	solar
R1	436	37	blue stereo
R2	754	20	red stereo
R3	803	20	geology
R4	864	17	geology
R5	904	26	geology
R6	934	25	geology
R7	1009	38	geology
R8	880	20	solar

¹The filter name is the camera eye combined with the filter position number. Table values from [10]

During pre-flight calibration, testing showed that the effective wavelength of the Pancam filters would increase a small amount of 1.5 ± 0.5 nm with increasing temperature from -25° to 50° . The filters showed no change in bandwidth. When increasing the angle of incidence from 0° to 9° , the filters showed a decrease in effective wavelength of 1.5 ± 1 nm. Both shifts were deemed so small that they should not affect the scientific analyses [10].

The blocking performance of the geology filters, meaning their ability to block out radiation outside the passband was determined acceptable and within or close to the rejection-band requirement of 1%. The average rejection band level was calculated as the total measured transmissivity outside the $\pm 1\%$ pass-band limit of the filter, divided by the total transmissivity of the filter between 400 and 1100 nm. The average rejection band level for the narrow-band geology filter was found to be $1.1 \pm 0.7\%$, that is within the requirement. The short-pass (L7) and long-pass (R7) geology filters exhibited a slightly larger rejection band level of 1.5 – 5.0%, but were still judged as acceptable, since the out of band transmission was found to be smooth and monotonic and no leakage was deemed meaningful in any way.

3.1.1.2 CCD

The Pancam detectors are standard front illuminated, buried-channel, frame transfer Charge Coupled Device (CCD) systems. One system connected to each of the Pancam eyes.

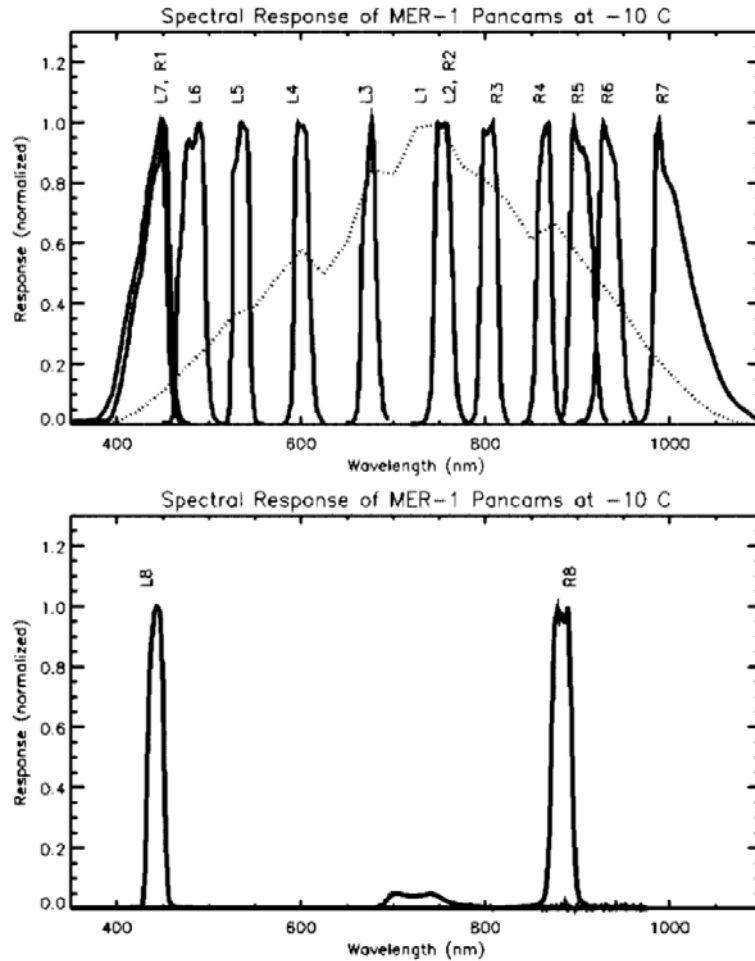


Figure 3.2: The top most figure shows the normalized spectral response for all Pancam geology filters. The Quantum Efficiency (QE) is indicated in the background as a dotted line (the QE will be discussed more in Section 3.1.1.2). The bottom figure shows the spectral response of the two solar filters. Plot from [10]

When a photon is incident on a CCD, the photon will interact with the bulk silicon of the CCD and create one or more charge carriers (ie. electrons) depending on the energy of the photon via the photoelectric effect. The charge generation efficiency, that is the fraction of photons that produces a charge, is described by the Quantum Efficiency (QE) of the detector. The QE is highly dependent on wavelength, and falls off steeply at the Near InfraRed (NIR) and Ultra Violet (UV) end of the spectrum. An example of the QE of one of the Pancam detectors, measured prior to installing is shown in Figure 3.3. The plot shows the QE across the wavelength range of the Pancam at temperatures $+5^\circ$ and -55° with a QE difference of $\pm 3\text{-}5\%$ between the two. Since the UV response of a CCD is affected by reflection losses at around 400 nm, and by the short penetration depth at UV wavelengths and beyond (see Figure 4.7 in Section 4.3), the QE is relatively low at the high energy end of the wavelength range, as can be seen in Figure 3.3. Furthermore, a photon will only generate a charge if the energy is high enough to do so, which is generally above the silicon band gap at ~ 1.14 eV (~ 1090 nm). Most photons with less energy will pass through the detector without interacting, whereas some will interact anyway but will not create charge carriers (called free-carrier absorption), which again means that the QE will be relatively low at the silicon band gap and beyond. Some of these issues can be counteracted by adding coatings of various materials to the CCD surface, however

no UV-enhancing or antireflection coatings have been added to the Pancam CCD[57]. Instead the relatively low QE at either end of the spectrum is counteracted partly by giving the shortest wavelength filter and longest wavelength filter a broader band pass, which was already discussed in Section [3.1.1.1](#)




Figure 3.3: The plot shows the QE of the detector for the two temperatures tested. Plot from [10]

Since each detector system is of the frame-transfer type, it consists of two 1024x1024 pixel CCDs, where one frame is used as the active, collecting area and the other frame is used as the read out area. Each pixel is of the size 12 by 12 μm , and each pixel has a full-well capacity of $170,000 \pm 20,000$ electrons (~ 4095 Data Numbers (DN)), and the exposure time can be set between 0 msec to 335s in steps of 5.12 msec. When an exposure is finished, the collected charge is shifted rapidly (5.2 ms) from the active frame to the readout frame, and the active frame can start a new exposure immediately, avoiding a delay between consecutive exposures. The charge transfer efficiency (meaning the time it takes for the collected charges to be read out of the read out frame) is significantly better when using a buried channel CCD, rather than a surface channel CCD, since the charge is transported in a potential well beneath the surface, partly avoiding that some of the charge is trapped in the Si-SiO₂ layer intersection. For this reason, CCDs used for scientific instruments are generally of the buried-channel type, as is the Pancam detector. Furthermore, using a frame-transfer CCD means that vertical smear is largely avoided since the readout frame has an opaque cover that prevents more than 99% of light at the full wavelength band to be registered. When the charge has been shifted to the read-out frame, the charges are "clocked" out row by row with a read out rate of 200,000 pixels/s, which adds up to 5 s per full frame image. Both the read-out frame and the active frame of the CCD are subject to thermally induced dark currents, the shielded storage area 4-5 times less than the active area, though. The dark current was modeled prior to launch using an exponential relationship between the rate of the accumulation of dark current in one pixel with temperature, and the resulting model, one for the shielded area and one for the active charge collecting area were stored in each camera, so that a dark current image can be generated for any temperature.

Blooming is a general issue in CCD detectors, and describes the charge overflow that happens when a pixel reaches saturation, and the extra charge flows into adjacent pixels. The behaviour

causes unwanted streaks in the resulting image and different approaches can be made to avoid this effect. For the Pancam detector, the method used is called clocked anti-blooming. Here the charge is moved back and forth between two phases of the detector in the same pixel just before blooming occurs in the pixel by inverting the gates of one phase, one at a time. The extra charge (electrons) that would cause blooming will then be eliminated by recombining with the holes trapped in the Si-SiO₂ interface. Just before blooming starts in the phase where the charge was moved, the charge is lead back into the first phase, and again the extra charge is eliminated by recombination at the Si-SiO₂ interface. The charge is lead back and forth at a rate shorter than the blooming rate.

The responsivity of the CCD was well described during pre-flight calibration procedures of Pancam. The responsivity in each geology filter was determined by obtaining images of an integrating sphere, creating a uniform diffuse light source at different temperatures. Based on these images, a responsivity coefficient was determined for each geology filter to be used for calibration of Pancam images during the mission. Flatfields, that is images of the pixel-to-pixel variations in responsivity across the detector, were acquired by pointing the camera towards a field known to be spatially flat within 1-2%, meaning that the response in each pixel of the detector is expected to be the same. Spatial variations could be caused by intrinsic variations in the sensitivity of the individual pixels, it could be caused by the \cos^4 falloff with distance from the optical axis, contamination or faults in fabrication of the filters or optics, or it could be related to undesired stray-light or scattering in the camera system or test equipment. In the case of Pancam, the spatial variation across the field in all filters seemed to be dominated by the \cos^4 falloff component, which was found to be responsible for a 5-10% decrease in responsivity along the edges of the images, since all other possibilities were ruled out by extensive inspection of the CCD, filters, and optics. All were found to live up to the required standards [10]. In order to make sure that the Pancam spatial response is stable, images of the Martian sky are used as a good approximation of a flatfield image during the mission.

3.1.1.3 Calibration target

Letting Pancam routinely take images of an object with a well characterised, known response, makes it possible to maintain the calibration of Pancam throughout the mission. Such an object, called the calibration target (hereafter referred to as the cal target) is mounted on the deck of each rover. The target measures 8 x 8 cm and is shaped as a sundial, with a 6 cm, vertical pole in the middle meant to cast a shadow across the target. The cal target consists of three grey scale annuli shaped regions, centered around the central pole. A "black" region with 20% reflectivity, a "grey" region with 40% reflectivity, and a "white" region with 60% reflectivity, with the white being the outermost annuli. In the corners of the square plate in which the target materials are mounted, 4 coloured regions are placed, one in each corner. Two plates of polished aluminum is placed on the target as well, which are responsible for reflecting the colour of the martian atmosphere. The material of the cal target coloured regions are silicone pigmented with either titanium, or carbon black to generate the gray scale, and submicron powder of either hematite, goethite, chromium oxide, or cobalt aluminate to create the red, yellow, green, and blue coloured regions.

During pre-flight calibration of Pancam, a complete characterisation of the cal target was conducted. The response of the individual calibration target areas was carefully measured using a spectrometer resulting in a good determination of a Bidirectional Reflectance Distribution Function (BRDF) model for all target areas, across the Pancam wavelength range and for relevant

temperatures. The BRDF is the radiance of the surface (measured by the spectrometer) divided by the incident irradiance (or power per surface area) from the irradiating source (in this case the laboratory light source), at all possible viewing geometries. The BRDF model of the cal target returns a reflectance coefficient $R_{pre-flight}(\lambda, az, el, az', el')$ for a particular wavelength and viewing geometry, where az, el and az', el' denotes the azimuth angle and elevation of the incident and emitted light, respectively.

3.1.2 Pancam in-flight calibration and data products

Here is provided a short overview of the Pancam data products available through the Planetary Data System (PDS) [49], which were used throughout this work. Hereunder, some of the key points related to the pre-flight and in-flight calibration procedures resulting in these products. Further details on the calibration of Pancam is given in [10] and [8].

When data from Pancam is downlinked, the raw data is converted into Experimental Data Records (EDR) which contains the raw Data Numbers (DN) for all pixels. These numbers are directly proportional to the number of photons incident on the pixel of the CCD.

In order to restore the true radiance values of the images, the images are processed through the MER calibration pipeline, where they are calibrated using data achieved during pre-flight and in-flight calibration tests. The bias noise and the dark currents are removed, the images are corrected for bad pixels and shutter smear, and they are flat-field corrected. From this step in the calibration pipeline, several data products are created, all going under the umbrella term Science Reduced Data Records (Science RDRs).

- Using the radiometric conversion coefficients determined in pre-flight and in-flight calibration of Pancam, the images are first converted into units of spectral radiance ($\text{W}/\text{m}^2/\text{nm}/\text{sr}$). This data product is referred to as RAD RDRs and includes the acronym RAD (for Radiance) in the image filename.
- The images are alternatively converted into a relative reflectance product, which makes it easier to compare to laboratory measurements, because it places the observed data in the same geometry as the laboratory data. These data products are denoted IOF (for I Over F), since the relative reflectance (or radiance factor as it is sometimes called) is defined as:

$$IOF = \frac{I}{F}, \quad (3.1)$$

where I is the measured radiance of the surface and πF is the collimated incident irradiance (power per area perpendicular to the direction of travel)[58]. The data product is referred to as IOF RDRs and is recognised by the acronym IOF in the image filename.

- An alternative and related relative reflectance product is called R^* (or reflectance factor) which is the bidirectional reflectance of the surface divided by the reflectance of Phyllosilicates on Mars and implications for early Martian climate a perfectly diffuse (Lambertian)

reflector under the same viewing and illumination geometry as the observation. R^* is therefore related to IOF in the following way:

$$R^* = \frac{IOF}{\cos(i)}, \quad (3.2)$$

where i is the solar irradiance incidence angle. Assuming unchanged illumination geometry, the reflectance factor and the radiance factor therefore only differs by a constant, the shape of the spectra are the same.

The IOF data product is produced using Pancam images of the cal target obtained during the mission and close in time to the actual scene observation to be calibrated. Several images of the calibration target was taken each day of the mission, making the caltarget the most photographed object on Mars. The Pancam image of the cal target in the relevant filter for the observation is run through the calibration pipeline to produce a RAD file, and an average radiance value for each cal target area is extracted by marking the Regions Of Interest (ROIs) in the image by hand. Making use of the pre-flight determined BRDF model for the cal target, one can calculate predicted reflectance values for the different cal target areas and compare them to the extracted values from the Pancam images. As mentioned in Section 3.1.1.3, the BRDF model of the cal target returns a reflectance coefficient $R_{pre-flight}(\lambda, az, el, az', el')$. The observer geometry (Pancam) is predetermined by the rover design, and the incident illumination is calculated from the known position of the sun at the time of the observation. Using the output reflectance coefficient in combination with the responsivity of the system $R_k(\lambda)$, determined pre-flight, a predicted naïve reflectance value can be calculated for each area of the caltarget at the specific Pancam bandpass of interest:

$$R_{predicted} = \frac{\int R_k(\lambda) \cdot R_{pre-flight}(\lambda, az, el, az', el') \cdot d\lambda}{\int R_k(\lambda) \cdot d\lambda} \quad (3.3)$$

The extracted RAD values from the Pancam image of the cal target are then plotted against the predicted reflectance values $R_{predicted}$ from the reflectance model, and a straight line is fitted to the data. Since ideally a zero input should yield a zero output, when no other effects are involved, the fit is constrained to go through (0,0). The slope m of this best fit line is then used in the conversion of Pancam observations of the martian surface (obtained close in time to the cal target observation) from units of radiance (RAD products) into relative reflectance products (IOF), since the fitted line can be described by the equation:

$$RAD = R_{predicted} \cdot m, \quad (3.4)$$

since $R_{predicted}$ is in units of R^* , and $R^* = \frac{IOF}{\cos(i)}$ as mentioned above, this gives:

$$RAD = \frac{IOF}{\cos(i_{target})} \cdot m = \frac{RAD}{m} \cdot \cos(i_{target}), \quad (3.5)$$

with $i_{caltarget}$ as the solar irradiance incidence angle. This way, Eq. 3.5 provides a way to reflectance calibrate images of the martian surface during the mission by exploiting the already well described reflectance properties of the cal target.

The above is a simplified version of the reflectance calibration, since Eq. 3.5 is only valid as long as no other effects are in play. However, a large factor needs to be accounted for, which is the accumulation of dust on the cal target during the mission. A dust model is included in

the calculation of $R_{predicted}$, which is described thoroughly elsewhere [9]. The above exercise is performed for every cal target image, and a calibration coefficient is calculated and added to the header of the relevant images, to use for later conversion into IOF.

3.2 Overview of mission results

The goal and aim of the Mars Exploration Rover mission was to search for, and characterize a vast number of rocks and soil in order to uncover information about the past and present climate conditions, and water activity on the planet. Especially water activity local to the landing sites, which included the search for sedimentary deposits (such as clays) and evaporites, which are minerals left behind by water evaporation processes, or by hydrothermal processes. Furthermore, the rovers were to provide a good description and overlook of the immediate surroundings of their respective landing sites, including a description of the geologic processes that shaped the terrain, such as wind erosion, volcanism, cratering, and so on. Many of these goals would also serve to verify information already acquired by orbiting instruments. Here we highlight a few of the interesting discoveries made by the two rovers during the mission.

Opportunity landed at the Meridiani Planum, on January 25th, 2004. The rover touched down in a small crater named Eagle crater. Immediately after landing a layered outcrop was seen in the crater wall (later dubbed "El Capitan") (see Fig. 3.4). Characteristic features called "vugs" were seen in the surface, which generally occur when crystals form inside a rock and are later eroded away by for instance water. The rock also featured fine parallel layers that were consistent with sediment deposition. Such finely layered rocks were encountered numerous places throughout the mission. Investigation of the rock layers encountered at the bottom of the Endurance crater revealed the presence of the iron mineral jarosite, which is only formed in the presence of acidic water, and the mineral hematite in the form of spherules, which are now known as "blueberries". On Earth, hematite is generally formed as water evaporates, though not exclusively. So called crossbeds were encountered, which are rock layers that tilt at an angle compared to the horizontal layers. The exposed bedrock had a rippled surface structure, which is a characteristic sign, known from Earth, of water flowing back and forth (see Fig. 3.4). All of these findings were clear indication of a watery past on Mars, but not only that, the discovery of the many layered rocks, ripples and evaporites (hematite and jarosite) indicated that the water would have been acidic, would have flown across the surface, and would have been present for a long time[37][38][56].

Spirit touched down at the Gusev crater landing site, named Columbia Memorial Station, on January 4th 2004. The rover landed in dusty soil, covered with rocks. The rocks were quickly identified as basaltic, containing olivine, pyroxene and plagioclase. Soon after landing, the Spirit rover traversed to the Columbia hills, a hill-like formation protruding from the crater floor, rising to an elevation of ~ 130 m above the plain. From orbital observations, the formation had a visibly layered structure characteristic of deposition in a standing body of water. During the traverse across the plain towards Columbia Hills, the investigated soil and rock targets, were all basaltic in nature with very little or no sign of water alterations. When the rover reached Columbia Hills, however, the first rock target named "Pot of Gold" showed signs of the mineral hematite, which can be formed in connection with water. Soon hereafter, several rocks in the area were found to contain minerals indicative of an aqueous environment. All the rocks of




Figure 3.4: The left image shows a close up of the outcrop rock dubbed "El Capitan" exposed in the Eagle crater wall. The martian hematite spherules (or "blueberries") can be seen in the top right corner of the image. The right image shows the best obtained image of the ripples (called "festoons") indicative of flowing water. Image Credit: NASA/JPL/Cornell

Columbia Hills contained varying amounts of phosphor, sulfur, chlorine, and bromine, which are all capable of being transported in water solutions. The rocks furthermore contained inversely proportional amounts of olivine and sulfates, which is expected in an aqueous environment, since water alter the olivine, but creates sulfates. The most exciting find, however, was of the iron bearing mineral goethite, which is only formed in aqueous environments, and was therefore a direct evidence of water[33][35][34][5]. For the 2020 rover mission, Gusev crater was among the final four sites considered, but was ultimately not chosen.

4

Pancam R7 effect: Characterisation and simulation

As mentioned in Section 1.1, the preliminary calibration of Pancam was done pre-flight by comparing the measured reflectance of the cal target (the BRDF) to the extracted relative reflectance from Pancam images of the cal target. Images of the cal target in all geology filters, using all four flight Pancams, were obtained. The Pancam images were calibrated for instrumental effects and converted into relative reflectance, before the 11-point reflectance spectra were extracted from each cal target area. The extracted data was scaled to match the laboratory measurements of the grey ring of the cal target, making it possible to directly compare the shape of the response profiles obtained from the Pancam data to those of the laboratory response profiles. Fig. 1.1 already presented in Section 1.1 shows the comparison between the two data sets. Panel (a) and (b) show the reflectance of the 4 coloured regions (blue, red, yellow, green) on the cal target from the laboratory data, and the Pancam spectral extraction, respectively. Panel (c) and (d) show the reflectance from the grey scale (white, grey, black) regions from laboratory, and Pancam data, respectively. The Pancam data and the laboratory data show an overall consistency across the wavelength range with similar reflectance levels and spectral shape, however in the R7 filter at $\lambda_{eff} = 1009$ nm, a discrepancy between the laboratory data and extracted Pancam spectra is observed. In panel (d) when comparing to panel (c), it is evident that the contrast between the white and black region reflectances becomes less for the Pancam data, compared to the laboratory data. We further observe that all coloured curves of panel (a) and (b) in Fig. 1.1 lie above the 0.3 reflectance value threshold in the R7 filter, which makes them "brighter" than the grey middle (grey region) curve of panel (c) and (d), and comparable in "brightness" to the top (white region) curve. When examining the trends of the coloured curves of panel (b), a similar tendency to the downturn in R7 of the white region curve can be seen. All the Pancam extracted R7 reflectance values for all the coloured curves are slightly, and in some cases (for the yellow and green curves), dramatically less than the models in panel (a) predicts.

The explanation for the discrepancy was reported in [10] as being related to a higher transparency of the caltarget materials at the longest wavelengths: *"For example, the relatively large discrepancy between the Pancam and laboratory measurements for 1000 nm band (filter R7) are known to be due to the higher transparency of the silicone rubber calibration target materials at these long wavelengths, allowing the reflectance "signature" of the underlying aluminum substrate to exert an influence on the detected signal. These and other viewing geometry effects are the subject of continuing analysis"* (citation from [10] on page 24 line 20). It was mentioned again in the caption of Fig. 20 in [10]: *"Small remaining differences between the laboratory data and Pancam measurements are due to residual uncorrected differences in illumination and viewing geometries between the measurements and, for the 1000 nm band (filter R7), to uncorrected*

effects related to the higher transparency of the RTV silicone calibration target materials at these long wavelengths. All of these effects will be incorporated in the calibration target model to be used during Mars surface operations". At the time, this conclusion was based on the fact that the anomaly was observed in images of the cal target, and since the effect was deemed small enough to not affect the calibration procedure, it was not investigated further. However, as we will see in the following section, the effect is not just seen in images of the cal target, but is observed to occur even in the images of the Martian surface. In some extreme cases where the R7 filter data is used as a key part of an analysis, and spectra are extracted from regions of generally high contrast, the effect can potentially affect the analysis and in some cases make it difficult to draw conclusions.

In this Chapter an effort will be made to describe and characterise the effect to as high a degree as possible, a hypothesis will be proposed for how the effect arises, and that hypothesis will be used to develop a model with the ability to simulate the effect.

4.1 The R7 effect in in-flight geology images

The left panel of Fig. 4.1 shows a radiance calibrated image of an outcrop on Mars, taken by the Opportunity rover on Sol 954 (image ID: 1P212872229RAD76EVP2586R7C1) during its traverse along the rim of the Victoria crater. The cracks and crevices in the rock creates a large number of shadowed areas, where many of them are positioned next to a bright rock face, creating sharp transitions from very dark to very bright. This image was chosen based on that morphology, since this would be the ideal place to investigate the R7 filter effect, as the effect was observed to create an artificial decrease in contrast in the longest wavelength band. A calibrated radiance image in each filter was downloaded from the PDS archive [49] and a number of ROIs were selected by hand (indicated by the coloured shapes in the images of Fig. 4.1) using merspect available through the mertools package. Care was taken in choosing ROIs that represented a broad selection of the different, interesting parts of the image, before a spectrum of each region in each filter was extracted. The radiance spectrum was converted into I/F using the calibration coefficient of the observation given in the header of the corresponding I/F image (the procedure is described further in Section 3.1.2). The resulting I/F spectra are shown in the right panel of Fig. 4.1. It is evident from the plot that the spectrum of the darkest shadows in the image (red curve and green curve) show a substantial increase in the R7 filter compared to the neighboring filter R6 (at $\lambda_{eff} = 934$ nm). Before conversion into I/F, the red curve showed a R7/R6 radiance ratio of 1.39, and the green curve showed a R7/R6 radiance ratio of 1.33. After conversion into I/F, the ratio for the red curve is at 1.31 and the green curve is at 1.26. Still significant increases, and even visible by eye in the plot. That these ratios should in fact reflect real wavelength-dependent differences between the two filters is not likely, since the only significance of the chosen extraction regions are that they are shadows and are therefore not expected to have a different mineralogical composition from the other materials in the scene that show different behaviors in the R7 filter. If the upturn in the spectra was caused by a mineralogical signature, the likelihood of other areas showing the same tendency would be high, and this is not what is observed. On the contrary, when looking at the spectra of the brightest areas of the image (the light blue curve), a decrease in radiance is observed for the R7 filter. Again these brightest areas are not expected to be significant in any other way than them being bright. Areas that are neither among the brightest or the darkest (the dark blue curve) does not show any noticeable upturns or downturns.

From the spectra it is evident that an effect of some kind is causing a decrease in contrast of the image in the R7 filter just as what was observed in images of the cal target during pre-flight calibration. That the effect is observed in images of Mars as well, indicates that the effect is not related to properties of the cal target, but is rather an effect related to the instrument itself.

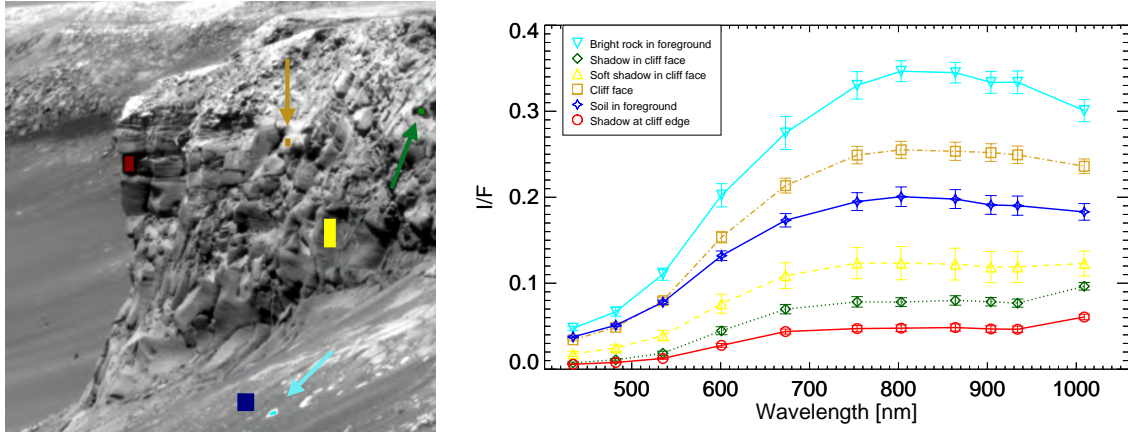


Figure 4.1: An example of the effect seen in images obtained with the R7 filter. The R7 image to the left was obtained on Sol 954 by the Opportunity rover. The extracted spectra, after conversion into I/F values can be seen in the plot to the right. Spectra from very dark areas exhibit a dramatic upturn at R7 ($\lambda_{eff} = 1009$ nm), the downturn at R7 for the brightest areas is more subtle but yet noticeable. Image retrieved from the PDS archive for MER data [49]. Image ID: 1P212872229RAD76EVP2586R7C1.

4.1.1 The R7 effect indicated in the literature

Acknowledging that the R7 effect is observed in in-flight geology images makes it critical to investigate how this might have affected previous analysis based on R7 filter data. It was found that several previous works point out anomalies in R7 images relative to other filters of the Pancam system.

Several papers report the use of R7 data for analysis[26][27][28][29][30]. The focus in this thesis will be on studies done in [9] and [32], since these two studies link to the work presented here.

In a study of dust deposition on the Pancam radiometric calibration target [9], it was found that derived dust albedo values showed a dramatic upturn at the longest wavelength (Fig. 8 and Section 4.1.2 of [9]). The authors also found that the ratio between the direct and diffuse illumination (sunlight exposed and shadowed) showed a dramatic downturn at the longest wavelength (Fig. 10 and Section 4.2.2 of [9]). The authors hypothesize that the observations are caused at least partly by an instrumental artifact, and mention under "future work" that this issue is in need of further characterization (Section 5.4 of [9]).

This discovery was what instigated the investigation into the effect and the work presented here.

In [31], the authors reported the existence of a characteristic spectral downturn from filter R6 ($\lambda_{eff} = 934$ nm) to R7 ($\lambda_{eff} = 1009$ nm) for regions on the Martian surface known to be silica rich. The authors of [32] use this idea to look for a correlation between the magnitude of

the spectral slope and the amount of silica in the soil, in order to look for hydrated minerals. They use spectral features at ~ 1009 nm in combination with the definition of other spectral parameters as a hydration signature and produce maps of hydrated materials along the Spirit traverse across the Columbia Hills. Since hydrated silica forms as a hydrothermal alteration product, the mapping is important in the understanding of how the silica formed, that is, whether it might have been formed by a fumarolic vent (if the silica is found to be clustered in local areas), has been transported by hydrothermal fluids (if the silica rich soil is more wide spread), or perhaps formed in a standing body of water (if found mainly in topographic depressions).

To calibrate the study and define the spectral parameters characterising the silica-rich soil, the authors extract spectra from Pancam images of already confirmed silica-rich materials, based on APXS and/or Mini-Tes analyses. The resulting spectra of the silica-rich soil is shown in Fig. 4.2 together with their corresponding Regions Of Interest (ROIs) marked out in the top image of the figure.

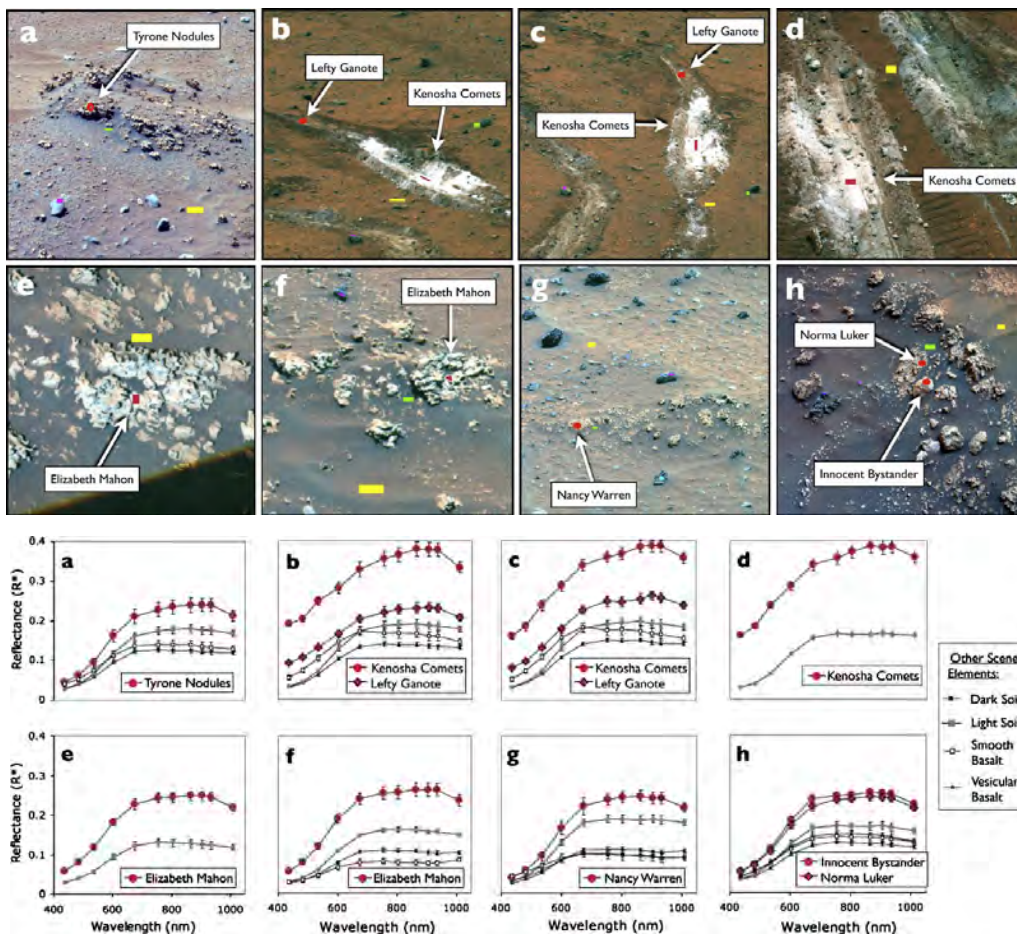


Figure 4.2: Top image: Pancam images of the silica-rich materials, as confirmed by APXS and/or Mini-Tes. The images are false colour with blue = 432 nm, green = 545 nm, red = 753 nm. Image a) Sol: 1101, Sequence: P2552, b) Sol: 1158, Sequence: P2581, c) Sol: 1187, Sequence: P2533, d) Sol: 1198, Sequence: P2539, e) Sol: 1160, Sequence: P2582, f) Sol: 1174, Sequence: P2588, g) Sol: 1190, Sequence: P2534, h) Sol: 1294, Sequence: P2581. Coloured regions are ROIs for spectral extraction: red = silica-rich targets, yellow = light soil, green = dark soil, purple and blue = two types of basalt. Bottom plots: Corresponding extracted spectra for the ROIs marked in the top images. Figures from [32]

In order to characterise the spectral shape and quantify the downturn from R6 to R7, they

measure the spectral slope between two Pancam filters i and j in the following way:

$$S_{ij} = \frac{R_j^* - R_i^*}{\lambda_j - \lambda_i}, \quad (4.1)$$

where R^* is the relative reflectance in each filter and λ is the effective band of the filter. Based on these spectra, the authors find a correlation between the magnitude of the R6-R7 slopes and the abundance of silica measured by the rover instruments used for in-situ analyses. They find other spectral characteristics and use those characteristics to produce parameter-space plots of the extracted relative reflectance values (see Fig. 4.3).

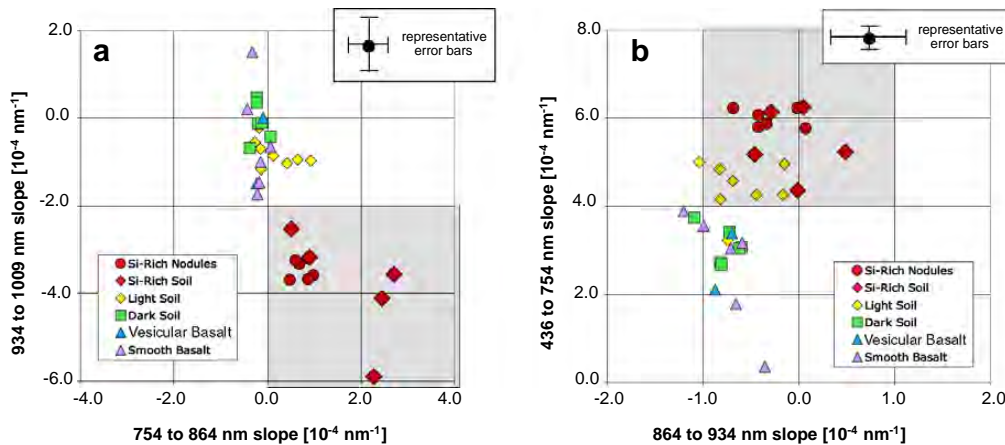


Figure 4.3: The reflectance values extracted from the Pancam images shown in Fig. 4.2 are here plotted in two parameter space plots. The silica-rich soil reflectance values (the red data points) seem to cluster together separate from the rest of the scene material. The limits set by these plots are used in the calculation of hydration maps. Figures from [32]

Based on those plots, they conclude that the silica-rich materials are spectrally distinct from the other elements in the scene. The spectral signature parameters were defined as follows: (1) 934-1009 nm (R6-R7) slopes less than $-2.0 \cdot 10^{-4} \text{ nm}^{-1}$, (2) 754-864 nm (R2-R4) slopes greater than 0.0 nm^{-1} , (3) 436-754 nm (R1-R2) slopes greater than $4.0 \cdot 10^{-4} \text{ nm}^{-1}$ and (4) 864-934 nm (R4-R6) slopes between -1.0 and $1.0 \cdot 10^{-4} \text{ nm}^{-1}$.

The R6-R7 spectral slope is reported as the strongest feature of the silica-rich soil in [32], but they note that the spectral feature cannot be a direct evidence of silica-rich material, since Si-bearing phases are not very spectrally active, hence it must be a signature perhaps related to the formation of the silica-rich material. The details of the experiments designed to determine the origin of the spectral feature is not described further here, but the subject is covered in [32]. They conclude that the spectral feature observed as a negative R6-R7 slope is not necessarily a direct representation of water-bearing silica, but could be a signature of other mineral phases which form together with the silica material. Based on the requirements set by the defined spectral parameters mentioned above, they isolate a number of 5 mineral candidates, all exhibiting a strong negative spectral slope across the Pancam wavelength band 934 nm (R6) -1009 nm (R7). Four out of the five candidate minerals exhibit structurally bonded H_2O or OH , the only candidate that does not is excluded based on its very unlikely presence in martian soil. They conclude that the negative R6-R7 slope observed in Pancam images must be due to some kind of hydrated mineral.

Based on the above conclusions the authors use the defined spectral parameters to map the

potential distribution of hydrated soil in the Columbia Hills area. They use the spectral parameters defined based on the spectra of the confirmed silica-rich materials mentioned above to remotely identify similar areas. Maps are generated for the 8 images that was used to calibrate and define the spectral parameters, and they are presented here in Fig. 4.4.




Figure 4.4: Hydration maps for the 8 images of Figure 4.2 are shown above. The coloured areas are pixels that fulfill the criteria set by the four spectral parameters given in the text. The colours scale with the magnitude of the negative R6-R7 slope in the spectra. The maps are overlain on 754 nm (R2) images. Figures from [32]

They argue that there is good correlation between the mapping method and the location of silica-rich deposits, and they therefore apply the method to a number of images along the Spirit traverse in order to map out the location of hydration signatures. During this analysis, they report the detection of false hydration signatures at very low sun angles, and attribute it to inaccuracies in the cal target calibration model (p. 384, second column [32]), and they report a detection of hydrated minerals on surfaces at high local emission angles, and argue that this signal could be caused by the Martian dust itself (p. 393, first column in [32]). Both of these anomalies, could potentially be attributed to the R7 filter effect.

4.2 Characterising the effect

Characterising the effect involves being able to identify a collection of simple images that maximizes the manifestation of the anomaly. Using a simple image eliminates the interference from other effects that could occur in a complex geology image, and makes it possible to isolate the R7 effect.

A systematic search was done of the large set of images acquired during the various calibration tests of the Pancam system. In this search we came across an image sequence acquired in the beginning of a stray-light test of the Spirit Pancam system. A number of images were acquired of a bright (effective signal level $> 2 \cdot 10^6$ DN¹) fiber optics light source on-axis and off-axis

¹DN stands for Data Numbers, which is

at different intensities. The stray-light test was done to determine how internal and external scattering of light affects the performance of the Pancam system. Most of the images were overexposed since this was part of the test setup, however some sequences were not overexposed and are of a single symmetric, on-axis light source on a very dark background. As described in Section [1](#), [4](#) and [4.1](#), the effect is observed to decrease the contrast between the brightest and darkest regions of an image, therefore it was expected that the effect should be maximized in such images where the transition between light and dark is abrupt and the contrast is high. Since the effect was expected to be a large scale effect that has no or very little influence on pixel to pixel variations, the size of the source made these images a better choice than for instance the PSF images obtained with Pancam. We looked for similar images for the Opportunity Pancam system, but was unsuccessful in finding a sequence that was not over-exposed. However, since the right eye cameras on the two rovers are designed to be identical, it is argued that any results acquired from analysis done using the Spirit images can be applied to the Opportunity images directly, as long as all images have been calibrated according to the official procedures stated in [\[10\]](#) and [\[8\]](#) and described in Section [3.1.2](#) of this thesis.

The image sequence selected provided us with 13 radiance calibrated images of the same on-axis source, one image for each spectral filter R1-R7 and L2-L7. This made it possible to compare the images directly, and that way identify how the R7 filter image deviates from the other filter images (R2-R6). Prior to any analysis, the source in each filter image was fitted with a 2 dimensional Gaussian function in order to determine the center of each source. Each image was then background subtracted using the average count per pixel of an annulus centered on the source with an outer ring radius of 260 pixels and a thickness of 2 pixels, well outside the source area. Each filter image was then subsequently cut out to a size of 481x481 pixels centered on the source and normalised to an integrated value of 1 by dividing all pixel values with the sum of all pixels in the cropped image. This made sure that the light source spectrum would not influence the comparison between filters. The assumption for comparison is that in this simple image of a symmetric, on-axis, brought-spectrum, bright source on a dark background, there is no realistic physical effect of the source radiation that can cause the light distribution to change drastically from filter to filter, hence any change in distribution must therefore be instrumental. This is discussed further in Section [4.3](#). The most interesting part of the images is the transition zone from the bright source to the dark background, since it is expected that the anomaly will cause the transition from bright to dark to be less sharp in the R7 filter, compared to that in the other filters. To investigate the light distribution in the different filters a number of radial profiles centered on the source were extracted, one for each filter image. The profiles were generated by binning the pixel values in 60 concentric annuli centered on the source, each with a width of 2 pixels and the largest with a maximum radius of 120 pixels. The resulting profiles are plotted together and shown in Fig. [4.5](#). Each data point is the average pixel value in a bin.

Panel a) shows the profiles on a linear scale where it is already noticeable that the light distribution profile extracted from the R7 filter image (black curve) deviates from the shape of the R2-R6 profiles. Panel b) shows a zoom-in of the tail part of the profiles where it is even more obvious that the transition from light to dark is more gradual for the R7 profile, since the darkest area of the image is artificially brightened (indicated by the excess signal in the tail part of the R7 profile) and the brightest area is artificially darkened (indicated by the lower signal at the peak of the R7 profiles), compared to the R2-R6 images. Panel c) shows a logarithmic plot of all profiles including error bars for the R7 filter profile and for the neighboring R6 filter profile. The error bars were generated as the standard deviation in each bin using the standard

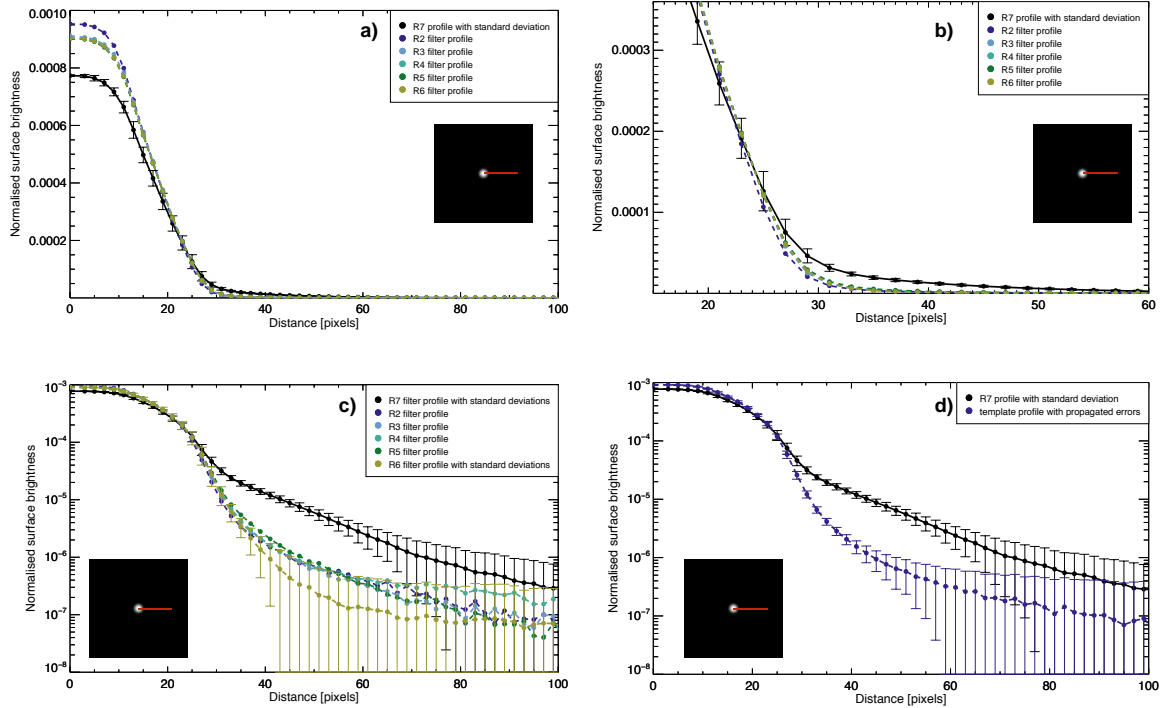


Figure 4.5: The plot shows 6 different radial profiles, each representing the light distribution of the source in a filter mounted in front of the right eye of the Pancam on board the Spirit rover. Black curve: R7, coloured curves: R2-R6. Note how the R2-R6 profiles are generally very similar and have a distinctively different shape than the R7 filter profile. a) linear representation of all filter profiles, b) linear zoom-in of the region where the difference between the R2-R6 profiles and the R7 profile is largest, c) logarithmic representation of all filter profiles, d) the template profile taken as the average of the R2-R6 profiles, and the R7 profile. Image ID: 020715222717_0001989_R7_103_rad.

formula $\sigma_x = \sqrt{\frac{\sum_{i=1}^N (x_i - \bar{x})^2}{N-1}}$. The plot shows very clearly that the R7 profile is well outside the individual errors for each filter profile at the transition region. What is also evident is that the filter profiles R2-R5 have very similar distributions. The R6 profile also shows a very similar distribution to the R2-R5 profiles except for a small deviation beyond a distance of 40 pixels visible on a logarithmic scale. However, at this distance the signal in the filters R2-R6 is generally very low and it is believed that the difference in the R6 distribution is related to background subtraction. Additionally, it is stressed that the difference between the R6 distribution and the R2-R5 distributions is much smaller than that between the R7 distribution and the rest of the filters. The R6 profile only deviates from the R2-R5 profiles by a value comparable to the standard deviation in each bin of the R6 filter image, and is only visible because of the logarithmic scale of the plot (see Fig. 4.5 panel c) for comparison).

In order to be able to compare the R7 light distribution to that of the other filters, the normalized, background subtracted R2-R6 images were averaged and a new profile was generated using the same binning procedure as for the individual filter images. The resulting image is hereafter referred to as the template image, and the extracted profile as the template profile. This image and profile is used throughout this work as a representation of how the light is expected to be

distributed when entering the camera system. Panel d) of Fig. 4.5 show a comparison between the R7 profile and the template profile. The errorbars of the template profile are calculated as the Root Mean Square (RMS) of the individual errors in each extraction bin for each filter. From this plot it is obvious that the most dramatic difference between the two light distributions is found between a distance of ~ 30 and ~ 60 pixels from the source center, which is just on the dark side of the transition zone between the bright source and the dark background. The maximum difference is at a distance of 49 pixels from the source center, where the R7 profile is a factor of ~ 10 larger than the template value. In other words, ~ 25 pixels away from the edge of a bright light source, the R7 image shows a signal about 10 times larger than the signal recorded in all the other filters imaging the same scene.

In general, assuming that the discrepancy between the R7 profile and the template profile is due to the R7 filter anomaly, the effect can be described as a blurring effect that causes photons in the brightest region to bleed into the darker region.

4.3 Finding a hypothesis

As described in Section 1, we wish to not only characterise the effect but also be able to get a sense of why and how it arises. What we are after in general is a model capable of taking the true light distribution (the template profile) and transform it into the light distribution seen in the R7 image (the R7 profile). Formulating a hypothesis about the physical transformation process, and basing the development of a model on that, prior to testing which functions would be the best to use, gives us the possibility of discussing the validity of the hypothesis based on whether or not the model turns out to be a good description of the manifestation of the effect. In order to establish such a hypothesis, we need to evaluate the different explanations for the effect. As mentioned in Section 4.2, no physical effects should be able to cause such a dramatic change in the light distribution from filter to filter, and the effect is therefore expected to be instrumental. Possibilities are therefore, that the effect must be inherent to either the R7 filter itself, the optics, or the detector of the Pancam system.

Characteristics of the effect:

- Symmetric - Using images of a symmetric source for characterization of the effect has the advantage that we can get a sense of the spatial manifestation of the effect. Since the source does not lose its symmetry in the R7 filter image, as established by the ability to easily fit a 2 dimensional symmetric Gaussian function to the source in each filter, we can deduce that the effect must as well be symmetric in nature.
- Wavelength dependent - Since the effect only occurs in the R7 filter, it must be wavelength (or at least filter) dependent.

Possible causes:

- Filter errors - Manufacturing errors, such as kinks and holes in the filter itself would most likely manifest itself asymmetrically and not cause the symmetric broadening of the light distribution we observe in Fig. 4.5.
- Inaccuracies in the optics lens - Inaccuracies in the lens of the camera optics would most likely also cause asymmetrical effects, since most aberration effects are not symmetrical

even for an on-axis source. Spherical aberration is symmetrical for an on-axis source, but would not be wavelength dependent, and should therefore manifest itself in all filters, and not only the R7 filter. Chromatic aberration is as the name suggests, wavelength dependent and caused by dispersion. The result could be a symmetric blur for an on-axis source as the one seen in the R7 filter stray-light image, however, the Pancam lens is a triplet Cooke lens which is specifically designed to rule out such an effect. Furthermore, chromatic aberration would cause off-axis sources to be asymmetrically blurred, which is not what is observed in Pancam R7 filter images (see Fig. 4.1). Also, you would expect the wavelength dependency to be more gradual in this case rather than a sudden onset in the longest-wavelength filter.

- Other manufacturing errors - The most straight forward argument for the effect not being due to defects during manufacturing is however that the effect is present in the R7 filter of both the Spirit and the Opportunity rover. That the same defects should be present in both systems would be very unlikely.

Ruling out effects that have to do with the construction of the Pancam R7 filter and the optics, leaves one possibility behind, the CCD. We therefore assume that the effect occurs after the light is incident on the CCD. Furthermore, the R7 filter is, unlike the filters R2-R6, a longpass filter, meaning that the filter relies on the rapidly declining sensitivity of the detector to provide a long-wavelength cut-off. For comparison, on the similar Mastcam system on board the Mars Laboratory (MSL), all filters are of the bandpass type [14], and since the effect is not seen for MSL, it is believed that the effect arises somewhere close to the silicon band-gap at 1100 nm, above the MSL Mastcam long-wavelength cut-off boundary at $\lambda_{eff} = 1012 \pm 21$ nm and $\lambda_{eff} = 1013 \pm 21$ for the Left and Right camera, respectively [14]. However, since the CCDs are different devices, it is difficult to make a direct comparison between the two.

One idea for the cause of the excess signal in the CCD was previously tested by [43]. The idea was that light let in by the R7 long-pass filter above 1200 nm would lead to the excess signal in the R7 filter band. The author conducted an investigation of the CCD response above 1200 nm, and found that the response was proportional to the temperature of the sensor, and the integrated response >1200 nm was .8% of the response within the R7 band ($\lambda_{eff} = 1009 \pm 38$ nm). This contribution is far too small to explain the effect, but could be a small contributing factor.

In the present work, it is proposed that the vast majority of the effect is rather related to the long penetration depth in silicon at wavelengths close to 1100 nm where the response in the R7 filter is still significant (see Fig. 3.2). It is proposed that light at the Near InfraRed (NIR), let in by the R7 filter will be able to penetrate the thick, bulk layer of silicon of the front-illuminated CCD ($> 400 \mu\text{m}$ [15] and see also [40] for standard silicon thickness for a front-illuminated CCD) and reflect off the backside of the bulk material. As the light is reflected, it spreads out diffusely and travel back through the bulk silicon layer, eventually reaching the active layer of the CCD for the second time, where part of it is registered in pixels nearby to the one receiving the incident radiation, giving rise to an excess signal. The amount of light registered at these nearby pixels is determined by the distance x from the pixel the light was initially incident on. This results in a symmetric broadening of the signal for this one pixel, and assuming that the effect will be the same for all pixels of the active layer, the large scale effect for a symmetric on-axis source is exactly what we see manifested in the stray-light R7 test image; a broadening

of the light distribution (see Fig. 4.5). The process is shown in the schematics of Fig. 4.6 and illustrates the path we expect some of the light to follow when incident on one single pixel of the CCD. This process explains why a broadening of the signal is only seen at the longest wavelengths let in by the R7 filter, since only light closest to the silicon band gap will be able to penetrate through the rather thick layer of silicon. The penetration depth in silicon increases quickly at increasing wavelength in the NIR part of the spectrum (see Fig. 4.7). Very close to the band-gap at $\lambda = 1100 \text{ nm}$, the penetration depth approaches $1000 \mu\text{m}$, and since the thickness of the silicon substrate of the CCD is $\sim 400 \mu\text{m}$, the light will be able to easily make one round trip through the bulk silicon.




Figure 4.6: The figure shows a schematic representation of our hypothesis for the nature of the observed effect and how it affects the registration of radiation in one pixel. Part of the light incident on the CCD penetrates the pixel layer without giving rise to a signal, and instead reflects off the backside of the CCD eventually resulting in an additional signal being registered at a distance x from the pixel receiving the incident light. C is the bulk thickness of the CCD.

It is important to note that the hypothesis proposed above is similar but not equal to etaloning seen in back-side illuminated CCDs. In back-side illuminated CCDs, the silicon layer has been edged away and left very thin, and light at the NIR end of the spectrum is known to be able to reflect back and forth giving rise to constructive interference which leads to a sometimes large excess signal. The silicon layer of the front-illuminated CCD is too thick for several reflections back and forth to occur. What is proposed here is that a single reflection is enough to produce the small effect seen in the R7 images.

4.4 The model

As mentioned in Chapter 1, our main goal for the present work is to eventually be able to remove the effect from the affected images. Since the reality of this task is that of an ill posed problem where we only have knowledge of the outcome of the effect in the form of the R7 filter images, we need a method capable of restoring the images under these conditions. Such a method is the iterative deconvolution technique, as we will see in Chapter 5. However, the problem will be much more approachable if we have a good model capable of describing the effect. This is the purpose of this Chapter, to establish a good, simple model capable of simulating the effect using a few well defined parameters, with well defined physical meaning rooted in our hypothesis.

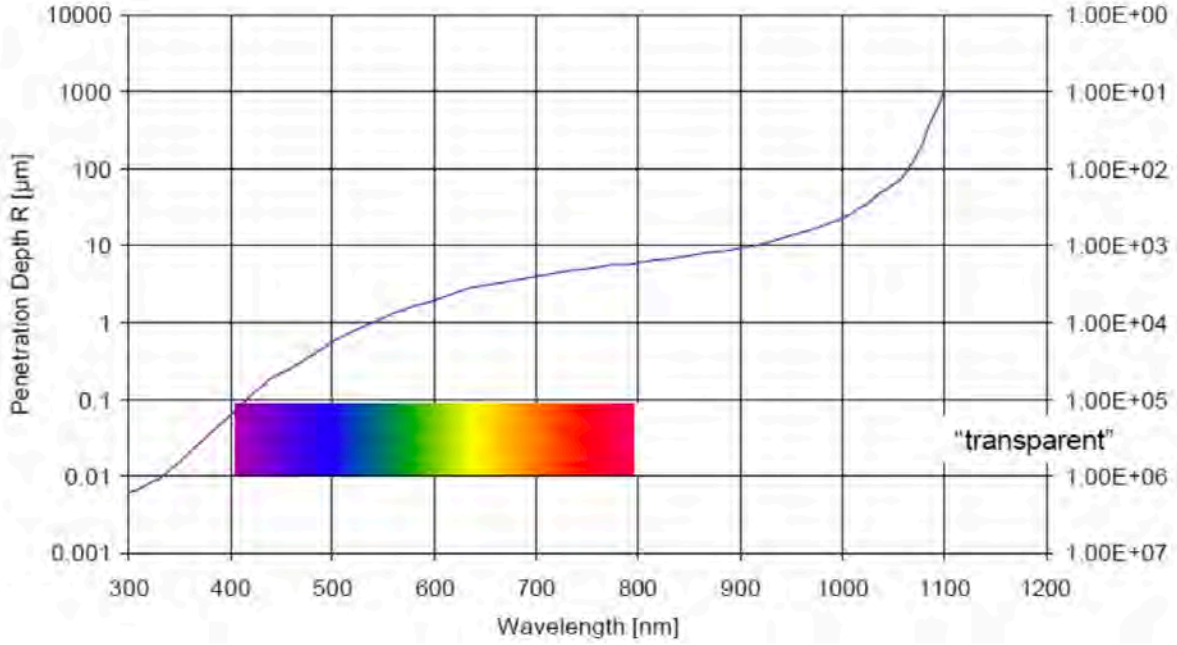


Figure 4.7: Penetration depth in Si.

4.4.1 Single wavelength approximation

The total distance travelled by light hitting normal to the surface of the CCD, is given by: $L = C + \sqrt{C^2 + x^2}$, where C is the thickness of the CCD (as indicated in Fig. 4.6) and x is the distance from the pixel receiving the incident light and the pixel registering the scattered light. The second term of the equation is the length travelled back through the bulk silicon given by the pythagoras theorem. All values are measured in units of pixels for convenience (1 pixel = 12 μm [10]). The light will be affected by attenuation on its way through the silicon. The attenuation is, according to the Beer-Lambert law, given by: $I(L) = I_{in} \cdot e^{-B(\lambda) \cdot L}$, where $I(L)$ is the intensity of the light after traveling a distance L through the medium, I_{in} is the intensity of the light incident on the CCD, and $B(\lambda)$ is the absorption coefficient of the bulk silicon, which is wavelength dependent and equal to the inverse of the penetration depth (see Fig. 4.7 for the wavelength dependency of the penetration depth in silicon).

To simplify the task, we assume that the scattering of the light off the backside of the CCD is perfectly diffuse, and that the surface is therefore perfectly Lambertian. This means that we include a cosine factor into the equation leading to the relation: $I(x) \propto I_{in} \cdot e^{-B(\lambda) \cdot L} \cdot \cos(\theta)$, where $I(x)$ is now the intensity of the backscattered light reaching a pixel at a distance of x from the pixel receiving the incident light. θ is the angle between the incoming light normal to the plane of the backside of the CCD, and the path the light travels after being backscattered.

Writing the expression in terms of the parameters specific to the hypothesis, we get: $\cos(\theta) = \frac{C}{\sqrt{C^2 + x^2}}$, and in turn: $I(x) \propto I_{in} \cdot e^{-B(\lambda) \cdot (C + \sqrt{C^2 + x^2})} \cdot \frac{C}{\sqrt{C^2 + x^2}}$, where we have substituted for L . After scattering, the light will no longer be collimated, and the intensity reaching a pixel at distance x (where x is measured as defined in the schematics of Fig. 4.6) can be approximated using the inverse square law. This gives us the final form of a function describing the backscattering of light at one specific wavelength:

$$I(x) \propto I_{in} \cdot e^{-B(\lambda) \cdot (C + \sqrt{C^2 + x^2})} \cdot \frac{C}{\sqrt{C^2 + x^2}} \cdot \frac{1}{C^2 + x^2} = I_{in} \cdot e^{-B(\lambda) \cdot (C + \sqrt{C^2 + x^2})} \cdot \frac{C}{(C^2 + x^2)^{3/2}} \cdot (4.2)$$

It is assumed that the attenuation is dominated by the absorption caused by electrons being lifted into the conduction band. By ignoring other causes of absorption, we can make an approximation stating that the chance of a signal being generated at a given wavelength (within the R7 wavelength band), is inversely proportional to the penetration depth. The assumption is likely to be good over most of the R7 wavelengths, but is expected to break down close to the silicon band gap at 1100 nm, where the probability of excitation becomes very small compared to other causes of absorption.

Additionally, by introducing an overall scaling factor A_{scale} which contains the intensity of light incident on the CCD, the efficiency of the backscattering, as well as the registration probability of the backscattered light in a pixel, we get the final expression describing the signal generated at a specific distance x from the pixel receiving the initial light:

$$f(x)_\lambda = A_{scale} \cdot B(\lambda) \cdot e^{-B(\lambda) \cdot L} \cdot \frac{C}{(C^2 + x^2)^{3/2}}, \quad (4.3)$$

with $B(\lambda) = \frac{1}{\delta(\lambda)}$ where $\delta(\lambda)$ is the penetration depth. As usual we have $L = C + \sqrt{C^2 + x^2}$.

Eq. 4.3 expresses how the artificial signal caused by the R7 filter effect varies as a function of the distance x from the pixel where the light was initially incident on the CCD. It describes how part of the incident light (the amount included in the scaling parameter A_{scale}) penetrates the CCD bulk silicon, gets backscattered diffusely off the backside of the CCD (described by the last term of the equation), travels back through the silicon, again affected by attenuation according to the Beer-Lambert law, and eventually being registered in pixels across an area of the CCD centered on the pixel receiving the initial light. The absorption coefficient falls off exponentially with wavelength, as indicated by inverting the graph in Fig. 4.7. At the longest wavelengths (where $B(\lambda)$ is low), the signal will get suppressed by the low sensitivity of the CCD, and the signal will die out. At the shortest wavelengths (where the absorption coefficient $B(\lambda)$ is high), the signal will disappear because too much of the light is absorbed on its way through the CCD.

This model is a single wavelength approximation, meaning that it is assumed that there is one dominating wavelength responsible for the entire extent of the R7 effect. This might be a good approximation for a narrowband geology filter to some extent, however since the R7 filter is a longpass filter relying solely on the decreasing sensitivity of the CCD for the long-wavelength cut-off, and since we expect the effect to arise close to the silicon band-gap at 1100 nm, where the absorption coefficient $B(\lambda)$ falls off quickly with wavelength, Eq. 4.3 is not sufficient in describing the observed effect. Therefore, we have extended the model to include the entire wavelength band of the R7 filter. The expansion is covered in the following section.

4.4.2 Full band model

In order to formulate a model capable of describing the full band effect, we need to describe how the penetration depth in silicon changes with wavelength, so that we can integrate over the full wavelength band. In Fig. 4.8 we show a visual representation of the variation of the penetration depth with wavelength in the long wavelength region (~ 900 nm - 1100 nm) up to the silicon band-gap. From the plot we see that the penetration depth is increasing at a close to exponential rate at these wavelengths, which is further supported by an exponential fit to the data, represented by the solid black line. The plot was generated using table values listed in [40] (page 172 Table 3.1). The exponential fit is used as an approximation for the trend of

the penetration depth in silicon at wavelengths covered by the R7 filter band. Since our model is written in terms of the absorption coefficient, which is the inverse of the penetration depth, we use $B(\lambda) = B_0 \cdot e^{-\alpha \cdot \lambda}$ as an approximation.




Figure 4.8: The plot shows the penetration depth in silicon as a function of wavelength in the R7 filter wavelength range. The black line is an exponential fit to the data. The data was adopted from Table 3.1 on page 172 in [40].

Eq. 4.3 is then integrated with the lower limit of the integration λ_{min} assumed to be at a wavelength where the effect disappears because the light gets fully absorbed before penetrating the entire CCD, that is where $e^{-B(\lambda_{min}) \cdot L} = 0$. The lower limit of the integration can be thought of as the lower wavelength cut-off of the R7 filter (~ 960 nm), below which we expect the R7 filter to block all light. Additionally, since the CCD is at least $400 \mu\text{m}$ thick [15] and the penetration depth is $\sim 100 \mu\text{m}$ at this lower boundary, we do not expect wavelengths below this limit to contribute to the effect anyway. The upper wavelength limit of the integration λ_{max} can be thought of as the point where the silicon becomes almost transparent to incoming light, and where absorption is dominated by other processes that do not excite electrons into the conduction band, and do therefore not contribute to a signal in the CCD. This limit is expected to lie close to the band-gap at 1100 nm and by keeping this upper limit a free parameter in our model, we have the means of comparing the model parameter of the optimized model, to the expected value of 1100 nm. This will be discussed further in Section 7.1.

Using all the assumptions mentioned above, the signal generated in one pixel can then be expressed as:

$$f(x) = \int_{\lambda_{min}}^{\lambda_{max}} f(x, \lambda) d\lambda = A_{scale} \cdot \frac{C}{(C^2 + x^2)^{3/2}} \int_{\lambda_{min}}^{\lambda_{max}} B(\lambda) \cdot e^{-B(\lambda) \cdot L} d\lambda \quad (4.4)$$

The integration is done by substitution in which we assume that A_{scale} can be treated as

wavelength-independent. The exponential function describing the behaviour of the absorption coefficient across the full R7 band $B(\lambda) = B_0 \cdot e^{-\alpha \cdot \lambda}$ is differentiated with respect to λ , which gives us $\frac{dB(\lambda)}{d\lambda} = -\alpha \cdot B(\lambda)$. Additionally, changing the limits of the integration gives us:

$$f(x) = A_{scale} \cdot \frac{C}{(C^2 + x^2)^{3/2}} \cdot \frac{1}{\alpha} \int_{B(\lambda_{max})}^{B(\lambda_{min})} e^{-B(\lambda) \cdot L} dB . \quad (4.5)$$

Carrying out the integration and assuming $e^{-B(\lambda_{min}) \cdot L} = 0$ as stated above, we get:

$$f(x) = A_{scale} \cdot \frac{C}{(C^2 + x^2)^{3/2}} \cdot \frac{1}{\alpha} \cdot \frac{1}{L} \cdot e^{-B(\lambda_{max}) \cdot L} . \quad (4.6)$$

Simplifying the expression by introducing $A = A_{scale}/\alpha$, writing B as shorthand for $B(\lambda_{max})$, and inserting $L = C + \sqrt{C^2 + x^2}$, we finally get:

$$f(x) = A \cdot \frac{1}{C + \sqrt{C^2 + x^2}} \cdot e^{-B \cdot (C + \sqrt{C^2 + x^2})} \cdot \frac{C}{(C^2 + x^2)^{3/2}} . \quad (4.7)$$

Eq. [4.7](#) is the final expression for the total contribution to the R7 signal in one given pixel at a distance x from the pixel receiving the incident light. The expression contains three free parameters: A is the overall scaling factor governing the strength of the effect, B is the absorption coefficient at the upper wavelength limit where the silicon becomes close to transparent to the radiation, and where the absorption becomes dominated by other processes that do not excite electrons into the conduction band. C is the physical thickness of the silicon wafer of the CCD. Eq. [4.7](#), when applied to each pixel in an image of the true light distribution, should be able to recreate (or simulate) the R7 effect, resulting in an image matching the light distribution we see in the R7 filter images. Given of course that our hypothesis and model are correct.

4.5 Simulation - optimizing the model

As mentioned in the previous section, Eq. [4.7](#) should, when applied to each pixel in an image of the true light distribution, result in an image of the light distribution we see in the R7 filter. Since we were able to identify a set of simple images of a standard light source on a very dark background as described in more detail in Section [4.2](#), we have the means of testing and optimizing our model. Using the template image as an approximation of the true light distribution where the effect is not present, we can apply the model (Eq. [4.7](#)) to each pixel in the image, and in that way simulate the R7 effect. This relies on the assumption that the large scale effect observed in the stray-light image is a representation of the pixel scale effect described by the hypothesis (Fig. [4.6](#)). Additionally, since we have an R7 filter image of the same source, we can optimize the model against that image and achieve estimated values for the model parameters.

If we use the representation X_{R7} for the R7 image, and X_T for the template image, we can write up the transformation process as follows:

$$X_{R7}(i, j) = X_{Template}(i, j) + X_{Template}(i, j) \cdot D + \sum_{k \neq i}^K \sum_{l \neq j}^L X_{Template}(k, l) \cdot f(x) , \quad (4.8)$$

Basically, the above equation should map the template image into the R7 filter image by simulating the effect pixel by pixel. In more detail, Eq. 4.8 defines how a specific pixel value in the R7 image ($X_{R7}(i, j)$) is simulated by adding an incremental value determined by our model $f(x)$ from Eq. 4.7 to the pixel at the same position in the template image $X_{Template}(i, j)$. The first term in the transformation is therefore the template image pixel value $X_{Template}(i, j)$ itself, the second term introduces a negative parameter D , which is responsible for the preservation of the normalisation by multiplying it with the pixel value $X_{Template}(i, j)$ and adding that (negative) product to the initial pixel value $X_{Template}(i, j)$. The third term is a double sum running over all pixels (k, l) of the template image with the size (K, L) , and is responsible for adding up the total contribution from these pixels. The size of the individual pixel contribution is determined by the product of the pixel value $X_{Template}(k, l)$ and our model $f(x)$.

In order to be able to perform the above transformation of the entire template image, the model (Eq. 4.7) was written up as a function using the Interactive Data Language (IDL). Additionally, an IDL procedure was constructed, capable of applying the model to each pixel of the template image through a nested loop running across the entire image. After one full transformation, a radial profile was extracted from the resulting image and compared directly to the radial profile extracted from the R7 filter stray-light image. This direct comparison was possible solely because of the symmetric and simple nature of the stray-light images, as mentioned earlier. The comparison of the two radial profiles provided us with the possibility of fitting the model parameters through a 1 dimensional Levenberg-Marquardt weighted least-squares fitting routine (MPFIT by Craig Markwardt [42]). The fit included distances of $x < 120$ pixels, meaning that the double sum of Eq. 4.8 did not run over all pixels of the template image. This truncation at $x=120$ pixels made the fit much faster, and since the function $f(x)$ is generally very small at $x > 120$ pixels, it did not influence the result of the fit.

At the initiation of the fit, a set of start values for the parameters A , B , C and D were chosen. The first guess for the parameter values was based on their expected physical interpretation according to the hypothesis that the model was constructed from (as defined in Fig. 4.6). This was done to avoid the fit being stuck in a local minimum when running the fit with most of the parameters free. A more fine tuned guess for the parameter values were determined by running the IDL procedure with slightly different values for the parameters without invoking the fit, and evaluate the fit by eye. When the fit was eventually turned on with the fine tuned parameter guess, it proved rather difficult to achieve a good fit when all 4 parameters (A , B , C , D) were kept free. However, when fixing the parameter C at the value 33, we were able to achieve an excellent fit. The physical interpretation of the parameter C in our model is the thickness of the CCD, and since this value was the best known quantity out of the 4 parameters contained in our model (the minimum thickness of the CCD reported by the manufacturer is $400 \mu\text{m}$, [15]), this was the parameter we chose to fix at 33, corresponding to 33 pixels = $396 \mu\text{m}$. The fit was run until a best fit result for the parameters were achieved.

4.5.1 Results of the simulation

As described in Section 4.5, a one dimensional fitting routine was used to achieve best fit values for the 3 model parameters that were kept free. The fit resulted in parameter values: $A = 96.2$, $B = 0.0388$, and $D = -0.211$. The visual result of the fit can be seen in Fig. 4.9. The template profile is shown in blue, the R7 profile is shown in black, and the fitted (simulated) profile is shown in dark red. From the figure it is evident that a very good fit was achieved, as

the simulated profile fits the R7 profile almost perfectly. The goodness of fit was reported by MPFIT as $\chi^2 = 0.83$, and with degree of freedom, d.o.f. = 56 (number of bins - free parameters of model), but should not be used as a direct indication of the quality of the fit (a short discussion of this is provided in Section [7.1](#)).




Figure 4.9: This figure shows a logarithmic plot of the fit used to determine the parameter values for the optimised model. The dark blue curve is the template profile, the black curve is the R7 profile to be simulated, the dark red curve is the best fit, simulated profile. The uncertainties for the R7 profile is the standard deviation in each bin and the errorbars of the template profile show the RMS of the error on the individual profiles R2-R6 in each bin.

The fit in Fig. [4.9](#) shows that it is possible to reproduce the R7 filter effect to a very good accuracy using our developed model. The excellent fit achieved indicates that our mathematical description is good enough to be used for the development of a correction capable of removing the effect from the affected images, which is the ultimate goal for this work.

5

Pancam R7 effect: Correction

In Chapter 4 a hypothesis was formulated for the physical origin of the effect seen in the Pancam R7 filter images, and the development of a model capable of simulating the R7 filter effect, was described in detail. In the following chapter the optimised model is invoked on the inverse problem, where we attempt to remove the R7 filter effect. A correction is formulated and tested on the same stray-light images used in the previous chapter, as this gives us a way to quality check the outcome of the correction. Since we are using the template image as a representation of the "true" light distribution, we can use the correction to remove the effect from the R7 stray-light image, and directly compare the result to the template image, and in that way determine whether the correction was successful or not.

The task of removing the R7 filter effect from the images can evidently not be approached the same way as the simulation of the effect, since for all other independent, more complex landscape R7 filter images of Mars, we do not have the luxury of a simple template image representing the "true" light distribution. To overcome this difficulty, an iterative deconvolution algorithm is constructed, capable of retrieving an approximation of the original "true" light distribution of the image being corrected, by letting the output converge towards the true light distribution. In order to use this approach, a good description of the R7 image degradation is needed. The identification of a good model describing the degradation of the R7 filter images, was exactly the point of Chapter 4, and the constructed model (Eq. 4.7) is used as an estimate of the blurring kernel of the convolution.

The stray-light images from the development of the model are used as a first test of the correction, since this makes it possible to evaluate the performance of the algorithm, and what will be evident below is that it makes it possible to determine a termination criteria for the iterative process to use in the correction of other independent R7 images.

5.1 Inverse process - an iterative deconvolution method

In this section we describe the general method used to write up a correction capable of removing the R7 filter effect from R7 filter images. That the task of correcting for the effect is not as straight forward as simulating it, but has the form of an ill-posed problem, was mentioned above. It can also be shown mathematically by rearranging Eq. 4.8:

$$X_{True}(i, j) = X_{R7}(i, j) - X_{True}(i, j) \cdot D - \sum_{k \neq i}^K \sum_{l \neq j}^L X_{True}(k, l) \cdot f(x) . \quad (5.1)$$

Eq. 5.1 shows how to derive the best estimate for the "true" light distribution from a R7 image, i.e. how to correct the image for the R7 filter effect. As is evident from Eq. 5.1, the term X_{True} appears on both sides of the equation, and more importantly, it appears inside the sum that is responsible for implementing our model. This means that in order to estimate the size of the correction to each pixel, we need to know the "true" light distribution, which we evidently do not since this is exactly what we are after. In order to overcome this obstacle, the correction is formulated as an iterative process, so that the output of the correction slowly converges towards the true light distribution. The task is not treated as a *blind* deconvolution process, since we use the model we have developed as a good estimate for the blurring function[63][64][46].

The process of correcting one R7 filter image is illustrated by the schematics of Fig. 5.1. The entire correction process consists of a number of iterations N , each using the best fit parameter values determined in the fitting routine described in Section 4.5.1 (i.e. $f(x)$ is unchanged).

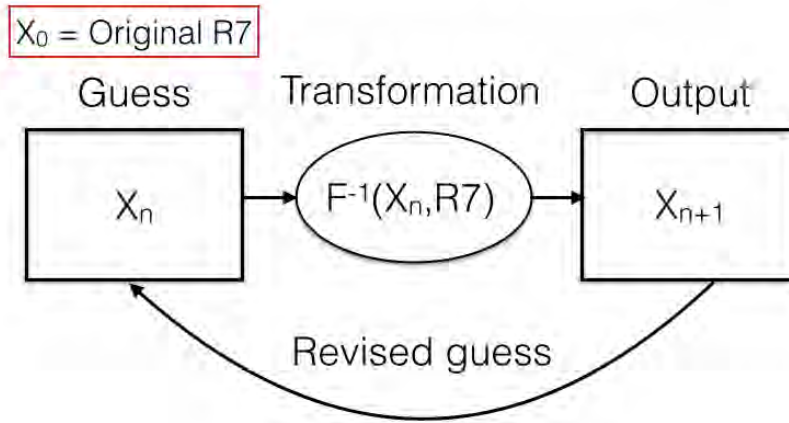


Figure 5.1: This schematic shows how the R7 images are corrected using an iterative process. Each iteration consists of a full circle in the schematic and results in an output which is a fully transformed image one pixel at a time. This output is then used as the new guess in the following iteration. The output will converge, and when a certain pre-defined cut-off value is reached, the iterative process will stop.

Each iteration starts out with a guess X_n for the corrected version of the image X_{True} . The very first iteration $n = 0$ uses the original R7 image itself as the guess for the "true" light distribution. Since the R7 effect is generally very small, this is an appropriate first guess. The output image from the first iteration $X_{(0+1)}$ is used as the guess for the second iteration $n=1$, and so on. In each iteration, the image being modified (corrected) is the original R7 image.

Writing up the process as illustrated in Fig. 5.1 using the terminology of the schematic, we get the correction for each pixel in the R7 image:

$$X_{n+1}(i, j) = X_{R7}(i, j) - X_n(i, j) \cdot D - \sum_{k \neq i}^K \sum_{l \neq j}^L X_n(k, l) \cdot f(x) . \quad (5.2)$$

Eq. 5.2 describes how the correction to the (i, j) pixel of the original R7 filter image is the summed correction caused by the blurring of light in all neighboring pixels (k, l) within the distance of 120 pixels, remembering that the model is truncated at the distance $x = 120$

pixels. Meaning that if all pixels in the region have been affected by the R7 filter effect (which we assume is the case), then they will all have contributed a small amount of light to the (i, j) pixel which is not supposed to be there. The amount of which is determined by the distance x between the (i, j) pixel and the (k, l) pixel, through the model in Eq. 4.7. All those corrections are collectively summed up and subtracted from the (i, j) pixel of the original R7 image, that is the amount of light added to the (i, j) pixel caused by the effect in the neighboring pixels, is removed by subtraction. This action is repeated for each pixel (i, j) of the original R7 image, resulting in an output array X_{n+1} , which is used as the new guess for the "true" light distribution in the next iteration.

To avoid artificial dimming or brightening of the edge pixels (that is, pixels within a distance of 120 pixels from the edge), a weighting function was applied to Eq. 5.2, which takes into account the fact that the pixels closest to the edge of the image will be affected only by the pixels inside the image. The weighting is done by dividing the summed correction by the fraction of the pixels inside the image.

5.2 Termination criteria and first test

In general it is not possible to evaluate the accuracy of the iterative deconvolution approach by examining the output itself, since we do not know how it is supposed to look after correction. The above described process could potentially continue forever, since the process will continue to move incremental amounts of light around the image. If the image is very noisy, the noise will eventually make the output diverge from the corrected version of the image. For this reason it is necessary to have a pre-determined termination criteria that decides whether a convergence of the output has been reached, and whether the iterative process should be stopped. This is a tricky task, since as mentioned, we do not have a way to determine whether or not we have reached the "true" light distribution when working with complex, landscape images of the Martian surface and no universal termination criteria exists which works for all types of inverse problems. Given this special case of an effect that is isolated to the R7 filter, we have the luxury of knowing the "true" light distribution of the R7 stray-light image, namely the template image, which means we have the opportunity to use the stray-light images to determine a general termination criteria to use on other R7 images, where the "true" light distribution is not known a priori.

In order to examine the convergence of the output of the correction of the R7 stray-light image, the image was cropped the same way as for the simulation, meaning a 481 by 481 pixel image with the source in the center, and the correction was setup to run over 20 iterations. The output of each iteration was saved to a separate file and a test-value was calculated, capable of relating the output of one iteration to that of the previous iteration. The test-value was defined as follows:

$$test - value \equiv \sum (X_N - X_{N-1})^2 / (K \cdot L) , \quad (5.3)$$

where X_N is the output image from the current iteration, X_{N-1} is the output image from the previous iteration, and $K \cdot L$ is the size of the image in units of pixels. Following each individual iteration, the output is evaluated by comparing it to the output from the previous iteration. This way, by determination of a cut-off value, the process can be terminated at convergence. The value for the cut-off was selected based on the 20 iterations of the R7 stray-light image.

A radial profile was extracted from each of the saved outputs and compared to the template profile, and the result was examined by eye. A selection of the iteration outputs is shown in Fig. 5.2.




Figure 5.2: Convergence of the output towards a good correction over 10 iterations. An acceptable convergence of the output is reached already after 3-4 iterations.

The radial profiles were generated the same way as the described procedure in Section 4.5. The black curve depicts the light-distribution in the uncorrected R7 image, the dark blue curve is the template profile, which in this case indicates the "true" light distribution of the image. The remainder of the coloured curves represent the radial profile extracted from the output image from the first five, and the tenth iteration. It is evident from the plot that the R7 profile converges towards the "true" light distribution after only a few iterations. It is important to note that the plot in Fig. 5.2 is logarithmic, meaning that the difference between the R7 profile and the template profile is much larger at distances of ~ 40 pixels, compared to distances > 60 pixels. At distances > 60 pixels, the background noise will dominate over the signal in the template image (remember that the template profile is generated as the average over profiles R2-R6, of which the signal at > 60 pixels is essentially zero), the R7 profile on the other hand has significant signal at > 60 pixels which leads to bigger variation in each bin, and a greater potential for overcorrection. Furthermore, realising that the dip of the iteration profiles at ~ 70 pixels below the template profile only looks drastic because of the logarithmic scale, if the size of the dip (that is the difference between the iteration profiles and the template profile at ~ 70 pixels) was inverted to be positive it would be comparable in size to the discrepancy seen at ~ 90 pixels.

No noticeably different results were achieved from iteration 5 to iteration 10. The correction seems to converge already after 4 iterations, and since the aim was to construct an algorithm that is efficient and fast without any large compromises with the quality of the correction, the test-value from the 4th iteration as our cut-off value in our termination criteria. In Fig. 5.3 the convergence of the output of the correction is represented by the test-value for each iteration.

The horizontal line marks the chosen cut-off value at 10^{-14} , which is the rounded off value from the 4th iteration. The termination criteria now takes the form:

$$\sum (X_N - X_{N-1})^2 / (K \cdot L) \leq 10^{-14} , \quad (5.4)$$

In praxis, this means that when the criteria of [5.4](#) is fulfilled, the iterations stop and the output of the current iteration is used as the final corrected image. We make use of such a cut-off value, since the number of iterations necessary to get a completely converged output could vary from image to image, therefore it is not useful to dictate the number of iterations needed on beforehand, but rather a value describing how much the output changes from one iteration to the next.




Figure 5.3: The convergence of the output represented by the test-value for the 10 first iterations. The horizontal line indicates the chosen cut-off value of 10^{-14} .

5.2.1 The correction algorithm

The correction algorithm was implemented using Interactive Data Language, and the final version of the correction algorithm works over three stages: remapping of the correction array, bin by bin correction of each array element, evaluating and resetting. One run-through of the three steps amounts to one iteration.

Prior to the onset of the algorithm, an array holding the correction values based on the model was generated. The correction array is an array with nonzero elements in a circular area forming the center of the array, surrounded by null elements.

Stage 1: Remapping of the correction array

In this first stage of the algorithm, the correction array is translated across the R7 image, remapped to center around the pixel (i, j) being corrected. The remapped correction array is multiplied with the guess image X_n to determine the incremental correction for each (k, l) pixel

to the (i, j) pixel in the original R7 image.

Stage 2: Bin by bin correction of each array element

The correction array is binned in consecutive annuli corresponding to the annuli used to extract the radial profiles from the stray-light images. The summed, and weighted correction in each bin is subtracted from the (i, j) pixel of the original R7 image, one bin at a time (starting with the annulus with the largest radius). Stage 1 and 2 are repeated for each pixel in the original R7 image.

Stage 3: Evaluating and resetting

When all pixels of the original R7 image has been corrected by repeating Stage 1 and 2, the resulting array is evaluated against the previous output with the calculation of the test-value of Eq. 5.3. If the termination criteria is not fulfilled, the correction array for the present iteration is reset, and the output array is used as a new guess for the "true" light distribution when the next iteration starts (stage 1 and 2 are repeated with the new guess). If the termination criteria is fulfilled, the output of the iteration is saved as the final correction of the image.

6

Pancam R7 effect: Correction verification

In this chapter, we demonstrate the performance of the correction on images obtained during the mission of both the Spirit rover and the Opportunity rover, meaning images other than the stray-light images from the Spirit rover. It is important to test the correction on independent images, that is images which were not used in the development of the model, since this will make sure that the correction does not distort the images in any way, and that it does not introduce any unwanted artifacts. The efficiency of the correction is also important to test, since a more complex image could require more extensive computer time, which we seek to minimize as much as possible.

The images presented in this chapter are images of the martian surface from both the Spirit rover, as well as the Opportunity rover. This makes it possible to test the assumption that since the Pancam systems on the two rovers are designed to be identical, and since the effect is present in both Spirit and Opportunity images, a correction developed based on images from the one, should be general enough to be applicable to the other. Furthermore, the correction of a set of calibration target images obtained during the mission by both rovers, is presented in Section [6.2](#).

All images in this Chapter was downloaded from the PDS archive and radiance calibrated according to the established procedures described in [8] and section [3.1.2](#) of this thesis. All images were kept at their original size, and un-normalised.

6.1 Correction of independent, in-flight, geology images

The two geology images used in this section were selected based on their complexity, and on the fact that they exhibit many high contrast areas, meaning very dark shadows next to very bright surfaces. The Spirit image selected depicts a collection of boulders and regolith that were encountered by the Spirit rover on its way from the Bonneville crater to the Missoula crater on Sol 77 (Image ID: 2P133203792RAD2224P2572R7C5). The many rocks covering the surface create many shadowed areas with sharp transition from dark to bright. The Spirit image is shown in Fig. [6.1](#) as the top left image. The Opportunity image is of a cliff-like structure exhibiting many rock protrusions, creating shadowed regions underneath as well as in crevices in the rock faces. The image was obtained on Sol 954 on the Opportunity rover's traverse along the rim of the Victoria crater (Image ID: 1P212872229RAD76EVP2586R7C1). The Opportunity image is presented in Fig. [6.1](#) as the top right image.




Figure 6.1: Improvement in contrast after correction of two radiance calibrated R7 images. The left column shows the correction of an image obtained by the Spirit rover on Sol 77. Image ID: 2P133203792RAD2224P2572R7C5. From top, down: Uncorrected, corrected, and the corrected image divided by the uncorrected. The uncorrected and corrected R7 images have been stretched to match the overall minimum and overall maximum of the two images. Min: $8.80 \text{ mW}/(\text{m}^2 \cdot \text{nm} \cdot \text{sr})$, Max: $34.2 \text{ mW}/(\text{m}^2 \cdot \text{nm} \cdot \text{sr})$. The ratio image was stretched according to its own minimum and maximum value: Min: 0.771, Max: 1.06. Right column shows the correction of an image obtained by the Opportunity rover on Sol 954. From top, down: Uncorrected, corrected, and the corrected image divided by the uncorrected one. Again, the uncorrected and corrected R7 images have been stretched between the overall minimum and maximum values of the images. Min: $0.0 \text{ mW}/(\text{m}^2 \cdot \text{nm} \cdot \text{sr})$, Max: $26.7 \text{ mW}/(\text{m}^2 \cdot \text{nm} \cdot \text{sr})$. The ratio image was again stretched to its own minimum and maximum values. Min: 0.736, Max: 1.09. Image ID: 1P212872229RAD76EVP2586R7C1.

The Figure shows how the correction improves the contrast in the two images. The images at the top of each column are the uncorrected R7 images, and the middle row contains the corrected images. The correction of both images were obtained after 7 iterations, at which point the termination criteria was fulfilled and the iterative process terminated. When comparing row one and two, it is marginally noticeable by eye, that the darkest shadows become more distinct in the corrected images in row two, and the transition zone between dark and bright regions, becomes sharper. When looking at the brightest areas of the images, a similar improvement is seen, where the bright areas become brighter and the transitions becomes sharper. Row three of the figure shows the corrected image divided by the uncorrected. The ratio images show that no additional information is added to the image during correction, no artificial brightening or dimming of the edges is seen, and no artifacts or distortions are added. In general, the correction provides a sharpening of the image, and for this reason, the ratio images of the (sharper) corrected image to the (less sharp) uncorrected image shows a reasonable representation of the terrain. This is an important quality check of the correction, since no visible extreme changes from the uncorrected to corrected images were expected, since the effect is generally very small as mentioned previously. To give a sense of the size of the correction, the brightest pixel, and a random pixel were selected, prior to correction, and the value of the two were logged for each iteration. The fluctuations of those pixel values are presented in Table [6.1](#).

Table 6.1:
Pixel fluctuations

Iteration #	Random pixel [$W/(m^2 \cdot nm \cdot sr)$]	Brightest pixel [$W/(m^2 \cdot nm \cdot sr)$]
Spirit image		
1	0.033749275	0.036148906
2	0.034182043	0.037252270
3	0.034198940	0.037412721
4	0.034202003	0.037446240
5	0.034202666	0.037453373
6	0.034202806	0.037454885
7	0.034202835	0.037455204
Opportunity image		
1	0.018142220	0.029299956
2	0.018260203	0.032095784
3	0.018248261	0.032649725
4	0.018246649	0.032766022
5	0.018246534	0.032790449
6	0.018246545	0.032795573
7	0.018246553	0.032796647

Note that the brightest pixel gets significantly brighter from first to the 7th iteration, as expected, but generally the correction is incremental.

In Section [4.1](#) where we introduced and characterised the effect, we saw how the spectra of a geology image, namely the selected Opportunity image was affected by the R7 filter effect. The decrease in contrast caused by the effect was made obvious by the extraction of spectra

from a selection of areas across the image (see the plot of Fig. 4.1). The areas were marked out by hand using the merspect tool from the mertools package, and were selected so that they represented the complexity of the image the best. This means that besides selecting the darkest shadows and the brightest surfaces, the selection also included regions representing smooth areas where no sudden transitions from bright to dark occur. Here is presented the same spectra, but including the spectral extractions from the corrected version of the R7 image. The spectra was extracted from the radiance calibrated images and converted to I/F Using the calibration coefficient derived during the standard IOF-calibration procedure from the calibration target image acquired close in time. Same procedure was performed on the selected Spirit image, and in the Figures 6.2 and 6.3, the resulting spectra are shown. The error bars represent the pixel to pixel variation in each extraction bin. The regions used for spectral extraction are shown in the top image of each figure, and the extracted spectra are shown in the bottom plot. The I/F values for the uncorrected images are represented by unfilled symbols, the I/F values extracted from the corrected images are represented by filled symbols. In the plot for each rover we observe a change in contrast as expected. The darkest shadows in each image becomes darker, and the brightest areas become brighter, and just as importantly, the spectra from the smooth areas that do not exhibit any sharp transitions, do not change noticeably, which can be seen as a quality check of the correction. The correction removes the dramatic upturns and downturns at the R7 filter wavelength, resulting in much flatter, smoother and less surprising spectra without adding any artificial signals.

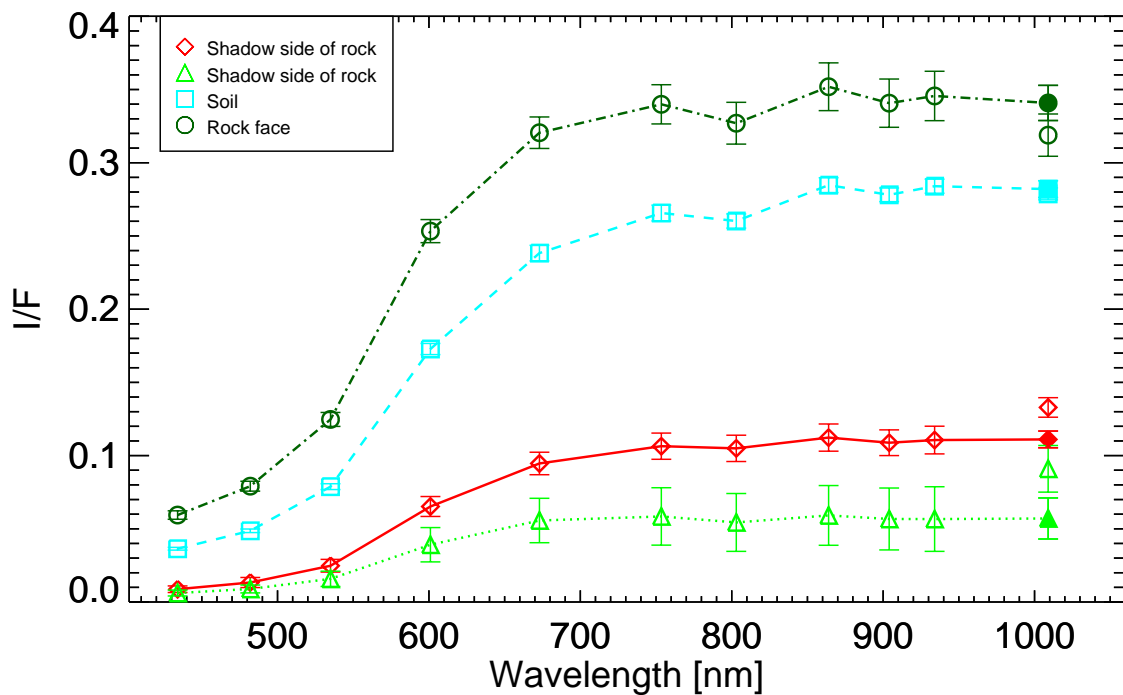



Figure 6.2: I/F spectra extracted from regions of interest marked on the image on top. For R7 both uncorrected (open symbols) and corrected (filled symbols) values are shown. The error bars represents the standard deviation in each ROI.

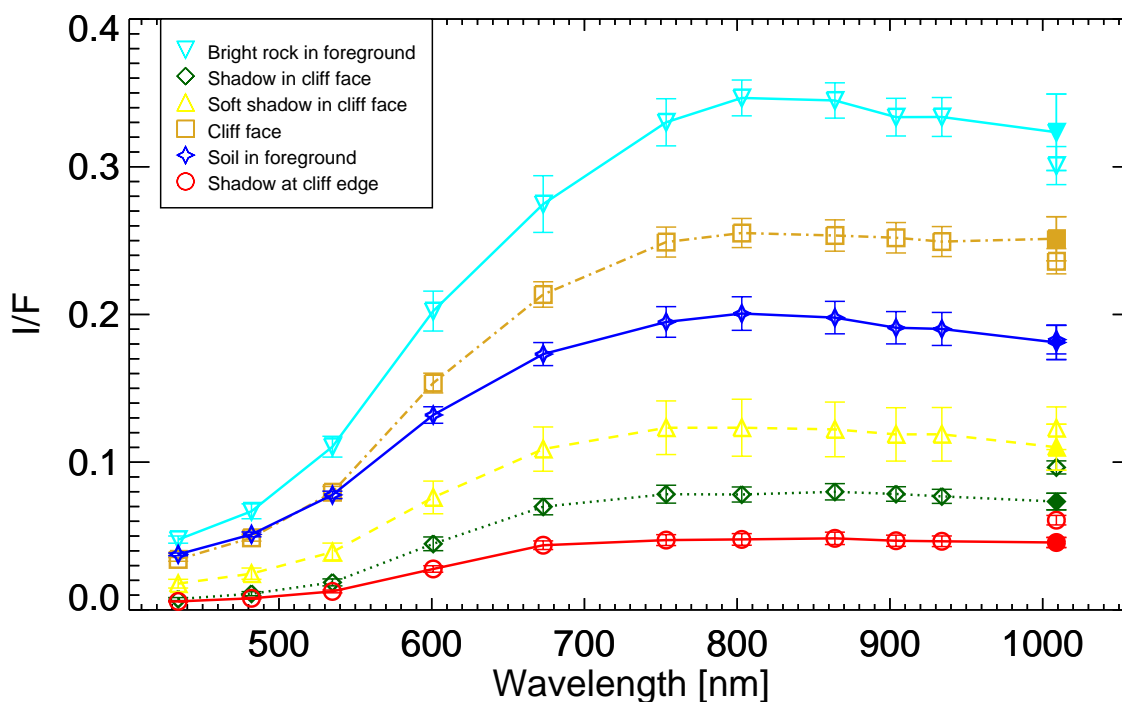



Figure 6.3: I/F spectra extracted from regions of interest marked on the image. For R7 both uncorrected (open symbols) and corrected (filled symbols) values are shown. The error bars represents the standard deviation in each ROI.

In Fig. [6.1](#) it is just possible to distinguish the increase in contrast for the corrected R7 image when investigating the darkest and brightest areas of the image, and in Figures [6.2](#) and

6.3 the increase in contrast is shown through extraction of spectra. However, by examining the ratio between the shortest wavelength image (R2 image at $\lambda_{eff} = 754$) and the R6 filter image at $\lambda_{eff} = 934$ nm, and comparing that to the ratio between the same R2 image and the uncorrected R7 image, and then again comparing that to the ratio between the same R2 image and the corrected R7 image, it makes for a good visualization of the effect and the capabilities of the correction. In Fig. 6.4, the ratio images for the Spirit rover are shown in the left column, and the same ratio images for the Opportunity rover are shown in the right column. The first row shows the R2 image divided by the R6 image, the second row shows the R2 image divided by the uncorrected R7 image, and the third row shows the R2 image divided by the corrected R7 image. The overall maximum and minimum pixel value of the three images in a column determines the grey scale stretch of each image in that same column (see caption of Fig. 6.4 for numbers).

In Section 4.1, we argued that no physical effect can explain the downturns and upturns in the spectra in the longest wavelength filter R7, and that the shape of the spectra extracted based on uncorrected R7 images are not consistent with what would be expected for spectral responses of geological materials. The same argument applies in this investigation of the residual radiation between the shortest wavelength filter and the two neighboring filters R6 and R7. There is no immediate reason why the residual images R2/R6 and R2/R7 should be dramatically different. However, what we observe is indeed a dramatic difference between the images in the first row of Fig. 6.4 and the images in the second row of the figure, where the bright and dark regions are almost inverted. The R2/R7 image shows a fair representation of the terrain for the same reason that the ratio image of the corrected R7 divided by the uncorrected R7 did (Fig. 6.1, bottom row), namely that the R7 image is significantly less sharp and has lower contrast than the other images. No mineralogical presence in the scenery could be countable for the systematic lowering of contrast across the entire image.

The residual images in the third row of the figure, which are the R2 image divided by the corrected version of the R7 images, immediately show how well the correction removes the effect. The R2/R7 corrected images are almost identical to the R2/R6 images, with only subtle wavelength-dependent differences.




Figure 6.4: Left column: Spirit image. Right column: Opportunity image. From top to bottom: R2 divided by R6, R2 divided by the uncorrected R7, R2 divided by the corrected R7. Both images are radiance calibrated, and all images have been scaled to match the overall minimum and overall maximum of the images for each rover. Spirit colour stretch: Min: 1.06 , max: 1.70. Spirit sequence: P2572 of Spirit Sol 77. Opportunity colour stretch: Min: 1.04 , max: 1.80. Opportunity sequence: P2586 of Opportunity Sol 954.

6.2 Correction of in-flight calibration target images

As described in detail in Section 3.1.2, Pancam relies on the calibration target for converting collected images from radiance to radiance factor I/F [10][8]. Images of the cal target are routinely collected by Pancam, and radiance values from the 7 regions of the cal target are extracted and plotted against the reflection values determined pre-flight in order to obtain calibration coefficients to use for conversion from RAD to I/F through Eq. 3.5. Ideally, the observed radiance should be proportional to the reflectance of the surface, so when plotting observed radiances against known reflectances the data should fall on a straight line through the origin. However, effects like dust accumulation on top of the target will cause an offset from (0,0). To correct for this, a dust model is included in the calibration procedure, which corrects for the change in measured radiance caused by the varying layer of dust on top of the target.

Early in the MER mission, the cal target was largely dust free, and data would therefore be expected to fall on a straight line as described above. This was generally the case for all filters - except the R7 filter. In order to see if the correction is able to better this strange behaviour, we examined the cal target images acquired by both rovers during the first 5 Sols of the mission. We chose Sol 5 as the cut-off to include as much data as possible without including images that were largely affected by dust accumulated on the cal target. This left us with a reasonably sized data set of 92 cal target observations from Spirit and 172 observations from Opportunity, when counting all images obtained in any of the 13 "geology" filters of the Pancam system. For each image, a radiance value from each of the a) sunlit surfaces on the cal target, and b) shadowed surfaces were extracted and plotted against the model reflectance factor values. As mentioned in Section 3.1.2, the reflectance factor R^* is defined as: $R^* = \frac{I/F}{\cos(i)}$, where i is the angle of incidence.

The two sets of data a) and b) were simultaneously fitted with a simple linear model, two straight lines constrained to intercept the y (radiance) axis at the same point, but with the offset from (0,0) kept as a free parameter of the fit. The offset value was then back-converted from radiance into data numbers (DNs) by multiplying the offset value with the exposure time and dividing the result with the standard pre-flight determined conversion factors mentioned in Section 3.1.2 and described further in [10]. The conversion was done in order to be able to make a direct comparison of the offset value between filters, since the exposure time could be different for each image. Furthermore, the offset value in DN can be directly compared to the maximum value of the 12-bit digitization (0-4095 DN empty to full well), which could give an idea of the size of the offset relative to the average signal value in the image. The average signal value of the image will generally take a significant fraction of the maximum value, since the exposure is controlled by an auto-exposure algorithm.

An example of a plot like the one described above is shown in Fig. 6.5. The plot was generated based on data extracted from a R7 filter image acquired by the Spirit rover on Sol 3. The 10 regions of interest (ROIs) in the image were marked by hand using the merspect tool from the mertools package. The regions count the 4 corner colour chips, as well as the 3 grayscale rings plus shadowed regions of the grayscale rings. The mean radiance value for the different ROI were extracted and plotted against the already known reflectance factor values. The reflectance factor values were collected from the header of each individual image. This entire procedure was repeated using the corrected version of the same R7 filter image, and the exact same regions of interest. A linear fit was performed in the way described above for the uncorrected R7 image data and for the corrected R7 image data, and the resulting plot is presented in Fig. 6.5. The

solid lines in the plot represent the best fit to the data extracted from the uncorrected R7 image (the black circular symbols), and the dashed lines represent the best fit to the data extracted from the corrected version of the R7 image (the red, square symbols).

The plot reveals an offset of 225 DN before correction, and an offset of 51 DN after correction. This definitely shows in a quantifiable way that the correction reduces the anomalous effect in the R7 filter to a very high degree. This plot is just an example of the general trend, and the image was selected randomly from a set of images where all regions of the cal target were visible. Producing a similar plot to the one in Fig. 6.5 for each of the 92 cal target images from the Spirit rover, and the 172 cal target images from the Opportunity rover, provided an offset value for each individual cal target image obtained during the first 5 Sols of the mission. Additionally the offset value for each of the corrected R7 filter cal target images, was calculated. The plot in Fig. 6.6 shows all the calculated offset values for each filter specific wavelength, for each rover. The black coloured data represents the offset values for all images in all filters, including the uncorrected R7 filter images. The red data represents the offset values extracted from the corrected R7 filter images. The small dots of both colours represent the individual offset values, the filled square is the mean offset value, and the error bars are $\pm 1\sigma$.

Looking at the black data only, it is evident from the plot that the individual offset values for each filter other than the R7 filter, scatter close to zero, with a mean offset within 60 DN of zero. The uncorrected R7 offsets are noticeably different from the rest, and hold mean values of 107 DN for Opportunity and 220 DN for Spirit. In particular, the R7 offsets are obviously different from the offsets of the other right-eye filters at the longest wavelengths ($\lambda_{eff} > 700$ nm) that show mean values within 30 DN of zero.

Looking at the red data showing the offset values extracted from the corrected R7 images, the mean offsets are brought much closer to zero for both rovers. On Opportunity, the correction brings the mean offset from 107 DN to -29 DN, in line with the observations in the other longer wavelength filters. On Spirit, the correction brings the mean offset from 220 DN and down to 54 DN, a value that is still higher than the rest of the filters, but the difference is definitely less dramatic than before.

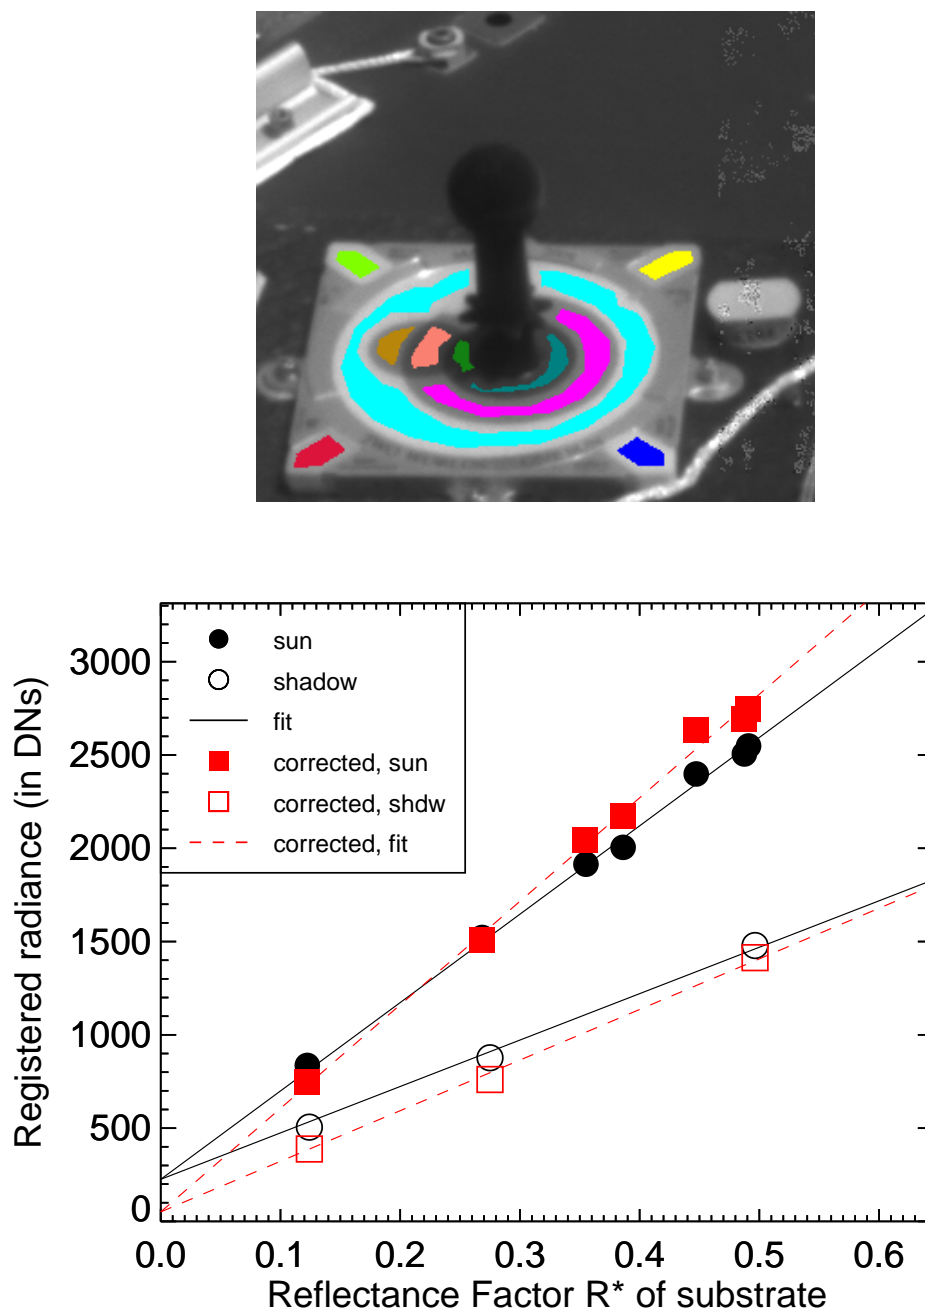


Figure 6.5: Example of a calibration target image and radiance-reflectance plot with demonstration of determination of the radiance offset. Image ID: 2P126802681ESF0200P2110R7M1.IMG was acquired on Spirits sol 5 through the R7 (1009 nm) filter. The black circles (filled and open) show the mean radiance values extracted from 10 regions on the uncorrected calibration target image. The filled circles correspond to sunlit regions and the open circles correspond to shadowed regions. The solid black lines show the fit with the two different slopes representing the different solar irradiances, one for the sunlit- and one for the shadowed regions. The red squares (filled and open) and red dotted lines represent the data and fit for the corrected image. The radiance values are back-converted from radiance to data numbers (DN's) i.e., they are not obtained from raw image data. The images was dark current subtracted, flatfield corrected etc. prior to analysis.

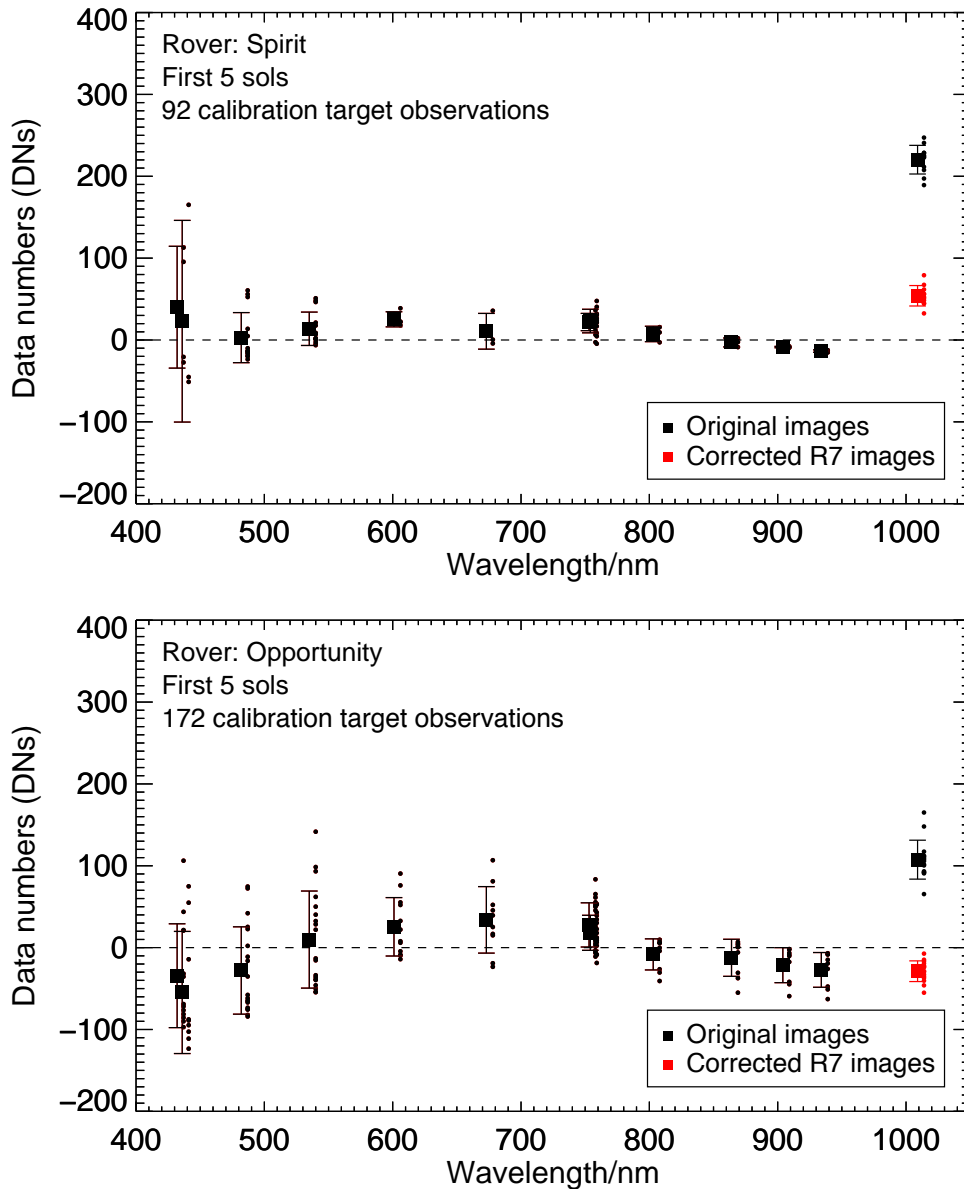


Figure 6.6: Offsets in radiance-reflectance plots for all calibration target images obtained during the first 5 Sols on Spirit (top) and Opportunity (bottom). As mentioned in the text, the offsets are found by fitting two straight lines with a common intercept to the radiance-reflectance plot and recording how far from zero they intercept the y-axis. The offsets were converted back to DN before plotting. Each small filled circle is an offset found from a single cal target image. For each filter, the mean offset is shown as a large filled square with error bars showing $\pm 1\sigma$. The black data points are obtained from original images. Red data points are R7 images corrected using the developed correction algorithm.

6.3 Correction of images used for previously published work

As described in Section [4.1.1](#) the existence of the R7 filter effect was hinted at in previously published papers, and R7 filter images have been widely used in previous publications [26][27][28][29][30][9][32]. Especially in [32], described in more detail in Section [4.1.1](#), the analysis was highly dependent on the value of the spectral slope between the R6 and the R7 filter, since the focus on the paper was to show that a correlation exists between the magnitude of

the negative spectral slope from 934 nm to 1009 nm (R6 to R7) and the amount of silica in a silica-rich target. Based on the idea that spectral characteristics of the known silica-rich materials could be used to calculate hydration maps along the Spirit traverse, a set of such parameter definitions were defined by the authors of [32]. The four chosen parameters are listed in Table 6.2.

Table 6.2:
Spectral parameters characterising silica-rich soil (from [32])

Parameter #	Slope [$R\#$ to $R\#/\text{nm}$]	Range [10^{-4}nm^{-1}]
(1)	R6 to R7/934-1009	< -2
(2)	R2 to R4/754-864	> 0.0
(3)	R1 to R2/436-754	> 4.0
(4)	R4 to R6/864-934	> -1.0 < 1.0

In order to examine how the correction for the R7 filter effect could potentially affect the results given in [32], some of the main plots and figures of the publication (the plots and figures presented in Section 4.1.1) were reproduced using the corrected version of the R7 images. As a preliminary investigation, four sets of the 8 sets of images presented in [32] were used in the analysis presented here (corresponding to the images a), b), e), and f) of Fig. 4.2), and the R7 images were corrected using the correction algorithm. All four images reached convergence (the termination criteria was fulfilled) after 8 iterations. Each set of images were read into merspect and regions as similar as possible to the ones marked out in the top image of Fig. 4.2, were selected in all images, including the corrected version of the R7 image. Fig. 6.7 shows the resulting spectra. The open symbols represent the extracted spectra based on all filter images, including the uncorrected R7 images, the filled out symbols represent the extracted values based on the corrected R7 filter images. At overlapping/stereo wavelengths, the I/F value from the right eye was chosen, since this was the method used in [32]. The y-axis is I/F, whereas the unit used in [32] was R^* , however, they are close to directly comparable, since the shape would be exactly the same, given that the only difference between radiance factor, I/F and reflectance factor R^* , is the constant $\cos(i)$, where i is the solar incidence angle on the surface at the time of the observation. Furthermore, the slopes are calculated based on relative values between the individual data points of the spectra, which amounts to the same result for both cases. The spectra presented here do not distinguish between the different kinds of silica-rich targets (red symbols and curves), or between the different curves representing other scene material (black curves and symbols). This is for convenience, since the focus is on the removal of the R7 filter effect, which only affects the slope between the R6 and R7 filter, and the point of these plots is to examine whether or not the spectra from the known silica-rich material can still be distinguished from the other scene material as a whole, after correction. Notice how the correction removes a big part of the slope from the R6 filter to the R7 filter in all spectra of silica-rich material.

The authors of [32] proceeded to plot the extracted data in spectral-parameter space plots based on spectral characteristics of the silica-rich materials examined. The plots showed that data from the silica-rich material plotted in a different area than the rest of the scene material examined. In order to determine whether the general trend seen in those plots could be reproduced using the four out of eight image sequences selected for the preliminary investigation, similar plots were generated using the uncorrected R7 filter images. The result can be seen in

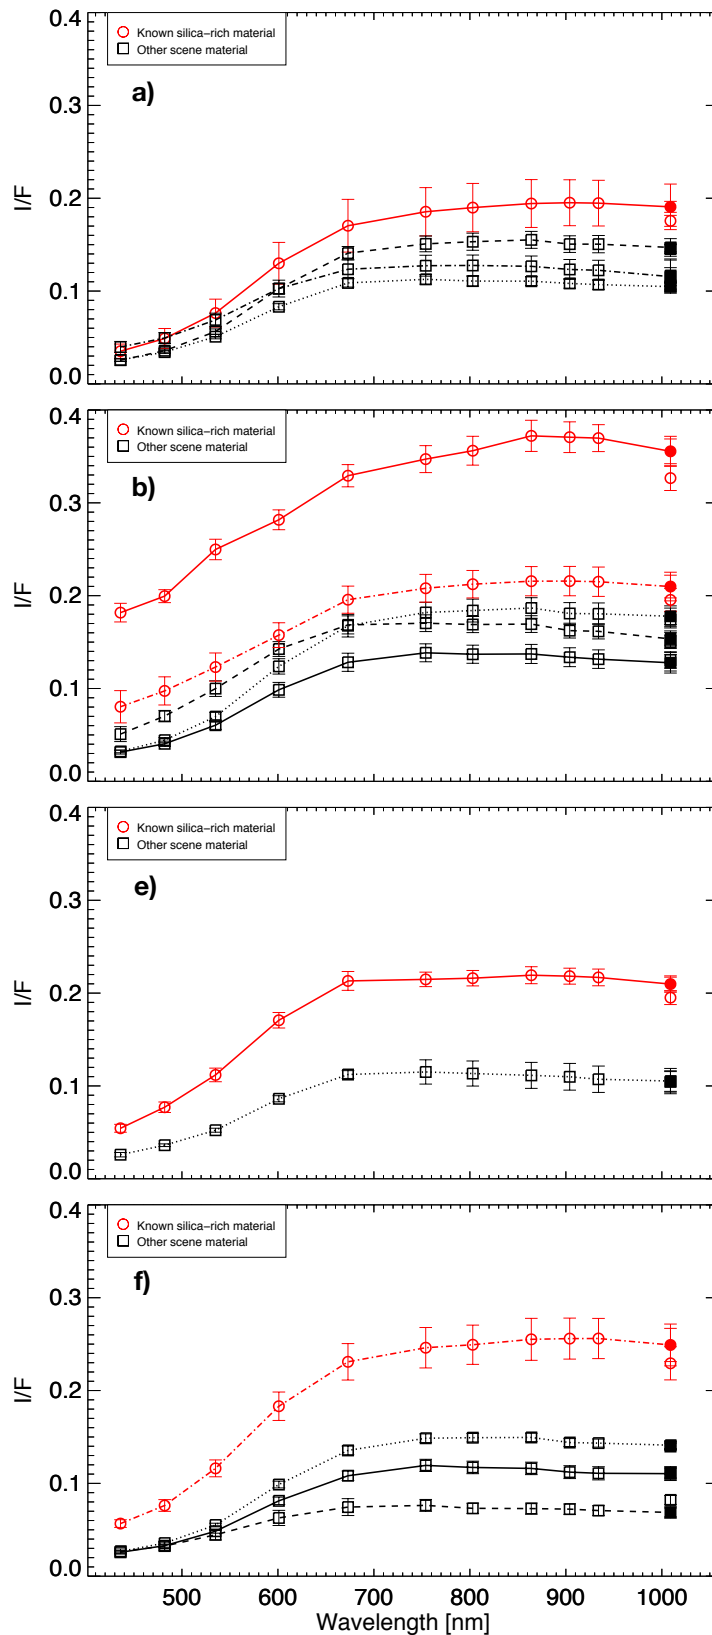


Figure 6.7: The four panels above are the reproduced spectra extracted from four of the eight images used in [32]. The letter of each panel corresponds to the lettering of Fig. 4.2. The extraction regions were chosen to be as similar to the ones in [32] as possible, but keep in mind that they are not identical. The open symbols are the original spectra acquired from all filters including the uncorrected R7 filter. The filled out symbols are extracted from the corrected R7 images. The errorbars corresponds to 1σ .




Figure 6.8: The two plots are the reproduced plots from Section 4.1.1 using four of the eight images: a, b, e, and f and using the uncorrected R7 images. These plots are provided to show that the four images are a good representation of the complete set of eight, since the general trends seen in the plots given in Section 4.1.1 are well represented by the four data sets shown here.

Fig. 4.3. When comparing these plots to the ones from the publication presented in Section 4.1.1, it is evident that the four sequences selected are a good representation of the published figures. It is still obvious that the silica-rich material (red symbols) group together separate from the other scene materials. Remember that the ROIs selected for extraction were chosen to be as similar to the ones used to produce the plots of Fig. 4.3, but are not identical, which means that some differences when comparing the plots presented here and the plots presented in 4.1.1 should be expected. The plot of Fig. 6.9 shows the same parameter-space plot as the bottom plot of Fig. 6.8, but was generated based on the corrected R7 images. It becomes evident when comparing the two plots, that the spectral characteristics of the silica-rich material become harder to distinguish from the rest of the material (represented by the green, yellow and purple symbols). The previous spectral parameter definitions given in Table 6.2 says that the silica-rich material will plot below $-2 \cdot 10^{-4} \text{nm}^{-1}$ on the y-axes, however when substituting the uncorrected R7 images with the corrected ones, this parameter definition actually excludes all the red data points in the plot showing the R7 corrected data.

The spectra from the silica-rich material (red symbols in plots) are still somewhat distinguishable from the rest of the data points in respects to the R2 to R4 slope, but in order to include potential silica-rich material when calculating the hydration maps, the spectral parameter definition for the R6 to R7 slope must be redefined. An obvious choice would be to change parameter definition (1) to include all R6 to R7 slopes below 0.0nm^{-1} . However, this would potentially include some of the bright soil of the scene (yellow symbols in plot), which are not

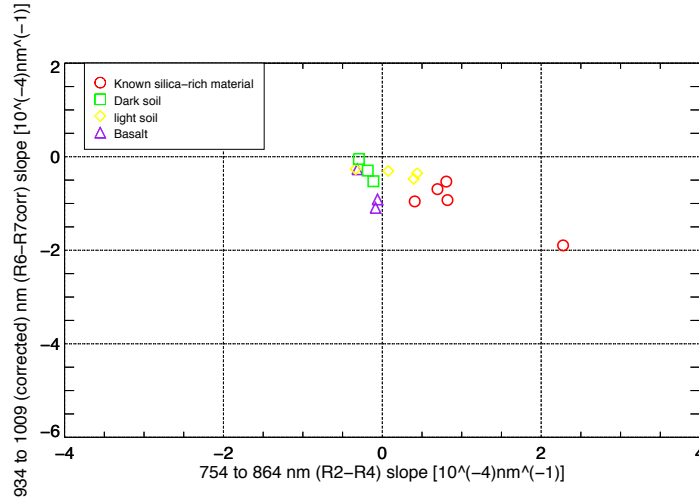


Figure 6.9: The plot above is the same as the bottom plot of Fig. 6.8 using the corrected R7 images instead of the uncorrected original R7 images. Note that the data still clumps together according to the nature of the soil, but occupy a smaller subset of the parameter-space.

expected to be silica-rich. A compromise could be to redefine the spectral parameter definition (1) to include all R6 to R7 spectral slopes less than $0.5 \cdot 10^{-4} \text{ nm}^{-1}$ or maybe less than $-1.5 \cdot 10^{-4} \text{ nm}^{-1}$ to be sure to exclude as much of the bright soil as possible. A possibility could also be to change the spectral-parameter definition (2), so that the new definition includes all R2-R4 slopes larger than $0.5 \cdot 10^{-4} \text{ nm}^{-1}$ in order to again further exclude as much of the bright soil as possible. The parameter definitions can be tweaked in many ways and several hydration maps were therefore calculated for each of the possibilities above for one of the image sets (Sol 1158, P). This made it possible to evaluate the different versions of the parameter-definitions before calculating the rest of the hydration maps. The calculated hydration maps for the image are shown in Fig. 6.10 and the changes to the two spectral parameters (1) and (2), the slopes between the R6 and R7 filter and the slopes between the R2 and R4 filters, is listed in Table 6.3. Image A) of Fig. 6.10 is a reproduced version of image b) in Fig. 4.4. It is calculated based on the uncorrected R7 image and the original spectral parameters as they are given in Table 6.2, and is provided here for comparison purposes. Image B) is calculated based on the corrected R7 image using the original spectral parameters and the same colour scale and range. It is clear that the map still picks out the main feature of the image, the Kenosha Comets, which is known to be silica-rich. However, the slope magnitudes are much less than they are in image A). Images D, G and H all have a lot of noise from other materials of the scene. Maps C, E and F are generally the best at both picking up on the main feature known to be silica-rich, and at the same time limiting the noise from other scene material. The F) map has the tightest constrain on the selection parameters of the three. The upper bounds on the colour scales (black/blue) are determined by the maximum pixel value in the map, which also corresponds to the upper bound on the spectral parameter (1), as given in Table 6.3. The lower bound on the colour scale was determined by the maximum slope value plotted for a silica-rich material (red symbols) in Fig. 6.9.

For the calculation of hydration maps for the three additional image sets, the parameter constraints from C), E) and F) were used, as they were deemed the best in terms of balancing detection significance and noise level from other scene materials that are not expected to be silica-rich. The resulting hydration maps are depicted in Fig. 6.11.




Figure 6.10: This figure shows the hydration maps calculated for different spectral parameter definitions. The lettering corresponds to the lettering in Table 6.3, where an overview of the different parameters are given. Only parameters (1) and (2) of Table 6.2 has been changed.

Table 6.3:
Spectral parameters for the hydration maps in Fig. 6.10)

Parameter #	Slope [$R\#$ to $R\#/\text{nm}$]	Range [10^{-4}nm^{-1}]
Image A		
Uncorrected R7 image and unchanged spectral parameters		
Image B		
Corrected R7 image and unchanged spectral parameters		
Image C		
(1)	R6 to R7/934-1009	< -1.0
(2)	R2 to R4/754-864	> 0.5
Image D		
(1)	R6 to R7/934-1009	< -1.0
(2)	R2 to R4/754-864	> 0.0
Image E		
(1)	R6 to R7/934-1009	< -1.5
(2)	R2 to R4/754-864	> 0.0
Image F		
(1)	R6 to R7/934-1009	< -1.5
(2)	R2 to R4/754-864	> 0.5
Image G		
(1)	R6 to R7/934-1009	< 0.0
(2)	R2 to R4/754-864	> 0.5
Image H		
(1)	R6 to R7/934-1009	< -0.5
(2)	R2 to R4/754-864	> 0.0

The table only lists the changed spectral parameters (1) and (2).
Parameters (3) and (4) are kept as listed in Table 6.2

It is clear that the calculated hydration maps do not have the strong signature of the maps presented in [32], however, the general features in the images are still being picked up by the map, even though the R6 to R7 slope magnitudes are less significant. The tightest constrained parameters used on the image in the F column seem to constrain the selection criteria a bit too much, whereas the middle column E picks out the interesting features in the image pretty well with next to no signal from the surrounding soil. Even a subtle clustering is seen at the positions indicated by the red arrows.




Figure 6.11: The above images show the resulting hydration map calculation based on the three parameter constraints C, E and F. The letters a), e) and f) to the left refer back to the lettering used in [32] and Section [4.1.1](#) for comparison.

The previous three chapters provided a detailed account of how the R7 filter artifact was characterised through the development of a model capable of simulating the effect in an image of the "true" light distribution (Chapter 4), how the developed model was implemented in a correction capable of removing the unwanted effect from affected images (Chapter 5), and how the correction performed on images dependent, as well as images independent of the development of the model (Chapter 6). It was shown how the correction performs on a simple image, such as the Spirit stray-light R7 image, and on complex images, such as the geology images from the Spirit rover as well as from the Opportunity rover. It was further shown that the correction output converges quickly, and that no artifacts were introduced during the process. In this chapter, a discussion of the different steps of the development of the correction is presented, as well as of the validity and implications of the correction.

7.1 Simulating the effect - the accuracy of the model

The most dramatic difference between the R7 profile and the template profile is found between a distance of ~ 30 pixels and ~ 60 pixels from the source center (see Fig. 4.9 for reference), which is just on the dark side of the transition region between the bright source to the dark background of the stray-light image. The maximum difference is found at a distance of 49 pixels, where the R7 value is a factor of ~ 10 larger than the template value. Extraction of spectra from specific high contrast regions can therefore be affected by areas well outside the region and cause erroneous results.

In Chapter 4 several possibilities for the nature and origin of the effect were discussed. It was ruled out that the culprit could be the cal target materials themselves, based on the realisation that the R7 filter effect is present even in in-flight geology images obtained by both rovers. The mere fact that the effect is present on both rovers makes it highly unlikely that the effect would be caused by manufacturing errors of the filter, and/or optics. Furthermore, no manufacturing errors seem to be able to create such a wavelength-dependent, spatially symmetric effect as the one seen in the R7 filter images. The size, symmetry, and wavelength dependency of the effect made us conclude that the effect is an inherent, instrumental effect of the CCD. Based on this assumption 4 parameter, simple model capable of simulating the effect to a high degree, was developed. The model was optimised through a one dimensional fitting procedure of the radial profile of a simple bright source on a dark background, providing us with approximated values for the 4 model parameters. The fit statistics were returned by the MPFITFUN routine as

$\chi^2 = 0.83$ and a *d.o.f.* = 56, resulting in a reduced χ^2 value of $\bar{\chi}^2 = 0.015$, which indicates severe overfitting. The error bars in each bin were generated as one standard deviation, which makes sure that potential unknown uncertainties are covered by the rather large errorbars. This however, builds on the assumption that the data in each bin fits a Gaussian distribution, which it most likely does not. The goodness of fit is determined based on the fact that when plotted, the model is judged by eye to fit the data very well (refer back to [4.9](#)).

The physical interpretation of the model parameter values is rooted in the hypothesis on which the model was developed. Of the 4 parameter values of the model, two of them, A and D, are scaling parameters, which makes them difficult to interpret per definition. The two others, B and C, however, have both a very clear physical interpretation. C denotes the thickness of the bulk silicon layer of the CCD, and was held fixed at 33 pixels, equivalent to 396 microns during the fitting procedure. The selected value was based on the 400 micron minimum thickness reported by the manufacturer, and was rounded to the nearest whole number of pixels for ease. The mere fact that the model fits the data so well using this physically reasonable value for the parameter C, may be taken as evidence in favor of the hypothesis. The physical interpretation of the parameter B of the model is the minimum absorption coefficient in silicon in the wavelength range covered by the model, hence we would expect a value of the parameter B close to the absorption coefficient at the silicon band-gap at 1100 nm, at which wavelength the model breaks down (a more detailed description of the wavelength dependency of our model is given in Section [4.4](#)).

Converting the parameter B to a penetration depth, for easier interpretation, gives us a value of $1/0.0388 \text{ pixels}^{-1} * 12 \text{ microns pixel}^{-1} = 309 \text{ microns}$, where 12 microns is the physical size of a pixel [10]. This value can be compared to the penetration depth values for pure silicon listed in Table 3.1 on page 172 of [40]. The table reports penetration depth values of 250 microns at 1060 nm, 370 microns at 1080 nm and 582 microns at 1100 nm. Comparing the value for B to this list it is found that while the value is in the right ball park, the value would be expected to be higher and closer to the 582 microns listed for the 1100 nm wavelength, since this is the point where the effect is expected to break down and the signal to be dominated by other effects. However, the penetration depths given in the table are for pure silicon, and the silicon of the CCD is most likely not pure, but probably doped with some other materials. Unfortunately, we don't have access to values for the penetration depth in the doped silicon material of the CCD. Also, the model neglects loss processes due to impurities, and any thermal losses etc. that do not add to the signal. These processes may however dominate at 1100 nm and the best fit value for B probably reflects the highest penetration depth that contributes significantly to the signal, which is most likely not all the way at 1100 nm. Taking these uncertainties into consideration, the value of 309 microns must be considered fully consistent with the model.

7.2 Notes on the correction method

Deconvolution is a very common inverse problem in imaging, and a lot of already developed algorithms exist for this type of task[44][47][45][46][60][62]. It is quite possible that a more efficient correction could be achieved using one of the existing deconvolution methods. However, in the effort of avoiding a *blind* deconvolution process[59], an estimate of the "kernel" function (not to be confused with the PSF, since that would imply that the effect was a result of the light-path through the optics, which is argued here that it is not) was found from the development of a mathematical model capable of simulating the effect in a set of images depicting a very

simple, symmetric source. The simulation was performed on the two dimensional image with subsequent fitting of a one dimensional representation of the effect (the radial profile of the source in different filters). To complete this task, an algorithm was written for the purpose, and it therefore made sense to use the already developed algorithm on the inverse problem as well. The available sequence of images of a simple, symmetric source also made it possible to determine a regularization parameter value (a cut-off) determining when the iterative process should be stopped. No singular method described in literature exists that dictates how best to generate new estimates of the solution to use as the guess for the next iteration. The method used here (shown in the schematics of Fig. 5.1) where the output of the correction X_{n+1} is evaluated against the output from the previous iteration X_n , has the advantage that it gives a sense of the stability of the approximate solution, but the disadvantage that it loses the possibility of determining how close to the solution it is, meaning how well it mimics the original R7 image when the kernel (our model of Eq. 4.7) is applied. For an ideal noiseless case it would be expected that after infinite iterations, the difference between consecutive outputs will tend to zero. In the case of a noisy image, the output could potentially approach the correct solution, and then go away when the noise begins to dominate, since the noise will be amplified and cause the difference between consecutive iterations to be very small, even though the image has lost its structure. For this reason a so called regularization parameter is used, which determines when the iterations should stop, preferably before information about the object is lost in noise. The method used here was based on the availability of the image sequence of a simple, symmetric source, where the termination criteria was determined based on the convergence of the one dimensional representation of the light distribution (radial profile). An alternative method of producing new iterates is to use the "kernel" (the model) to blur or convolve the approximate solution, resulting in the return of a blurred image directly comparable to the original blurred image. This way the discrepancy between the blurred image output, and the original blurred image can be minimized until a best approximate solution has been found.

These two methods of regularization arguably leads to the same result for an ideal situation of a noiseless image, since in the case of the method used throughout this work, the "perfect" solution will be perfectly stable, meaning that the difference between the last iteration and the one just before would essentially be zero, and if the process is truly converging towards the corrected image, simulating the effect using the "perfectly" corrected image would produce the original blurred R7 image. In the case of the second method, the "perfect" output would be identical to that original blurred image as well. This is of course a hypothetical case since the reality is that the noise level will play a large role and make it impossible to ever reach a "perfect" solution.

7.3 Correction Performance on Pre-Flight, Stray-Light Image

In order to test the performance of the correction before introducing it on actual scientific geology images, the correction was implemented on an image where the effect was ensured to be maximized. The stray-light R7 filter image from Spirit was a natural choice for the first test because of its simplicity, symmetry, and high contrast. Additionally, given the fact that this made it possible to directly compare the output of the correction to the template image, a termination value for the iteration of other more complex images could be defined. It was expected that the correction of this simple image would be faster and reach convergence quicker

than it would for a more complex geology image. Accordingly, the output was observed to converge after only 4 iterations, making the rounded up test-value (10^{-14}) for this iteration the selection for the termination value. Given the expectancy that it would be much easier to correct a simple image than a more complex one, the selected termination value provided a strict limit on an acceptable correction of the more complex geology images. It was important to select a good termination value, since the iterations of one image could per definition continue indefinitely while the output converges towards the "true" light distribution. If the corrected image is very noisy, the output could potentially approach the "true" light distribution, and then go away again. To avoid this situation, a conservative termination value is favorable. This must be weighted against the time expense of correcting a large image, since the processing time will be proportional to the number of iterations.

7.4 Correction Performance on In-Flight Geology Images

The selected images used for a secondary test and validation of the correction were, as mentioned, selected based on their complexity and high occurrence of high contrast areas, meaning very dark shadows, next to very bright regions. These selection criteria for the first complex geology image corrections serve the purpose of making sure that the correction is first implemented on images where the effect is expected to be maximized, since this makes sure that no artifacts are added to the images during correction and that it is safe to draw conclusions from analysis performed on corrected images. Generally, the effect corresponds to a smooth filter with same strength across the entire image, as the results from our characterisation of the effect (described in detail in section 4.2) also indicated. This additionally rules out the possibility that the R7 filter effect should be caused by chromatic aberration effects, since that would manifest itself as a blurring effect that would worsen the more off-axis the source is located. This is not what is seen in the second row of Fig. 6.4 or third row of Fig. 6.1.

Spectra extracted from the corrected R7 filter images do not show the dramatic upturn seen in the uncorrected images and the long wavelength end of the spectra are generally more in line with what would be expected from spectra of geological scenes. The contrast between the darkest and brightest areas on the R7 filter image increases, as expected. Examining the ratio between the corrected R7 image and the uncorrected R7 image, the pixel values span the range 0.771 - 1.06 for the Spirit example, and 0.736 - 1.09 for Opportunity. In other words, for these specific examples, some dark pixels get $\sim 25\%$ darker after correction, and some bright pixels get 5-10 % brighter after correction. Inverting these percentages indicates that the R7 effect inflates the observed radiances from some dark shadows up to $\sim 33\%$.

Examining the images directly revealed that when taking the ratio between the shortest wavelength filter R2 and the R6 filter and comparing that with the ratio between the R2 filter and the uncorrected R7 filter, there is a noticeable difference between the two (see Fig. 6.4). This difference is attributed to the R7 filter effect since, again, the light distribution is not expected to change drastically from the R6 filter image to the R7 filter image. Especially not when the light distribution in filters R2-R6 are so similar as indicated by Fig. 4.5. When removing the R7 filter effect, i.e. invoking the correction on the image, the resulting ratio image is much more in line with the expected. Based on this direct observation, and based on the results from the spectral extraction, it is safe to claim that the correction is efficient and that

the model on which it is based is sound.

7.5 Correction performance on cal target offsets

The Pancam calibration target consists of a number of surfaces that span relatively few pixels in Pancam images but have dramatically different reflectance properties in close proximity to each other. As such it would be expected that images of these targets are exactly the type of images affected by the R7 effect, since the R7 effect is a blurring effect causing a broadening of the light distribution, and thereby a decrease in the contrast of the image. It was shown in Section 6.2 in Fig. 6.5, that when plotting observed radiances against known reflectances of the R7 filter images, the low-reflectance data points will be offset upwards in radiance, while the high-reflectance data points will be offset downward, ultimately causing the fitted straight line to intercept the y-axis above zero. Performing this analysis on all cal target observations for the first 5 Sols on each rover revealed that this is a general trend. By invoking the correction on this collection of cal target images, it is shown in Fig. 6.6 that the correction satisfyingly removes a great part of this trend on both rovers. However, as is evident from the figure, the correction works best for the Opportunity data, which is curious since the correction was developed based on Spirit images. A couple of approaches were made to clear up why the correction works better on data from Opportunity than data from Spirit. The first idea was to test whether the region selection files used for the extraction of radiance data were incomplete. The selection files were all downloaded from the same database as the images themselves. As the whole process of extraction of data was done using a script, only random sampling was done to check that the selection files were complete. To make sure that incomplete selection files were not the reason for the discrepancy, all selection files for both Spirit and Opportunity were checked one by one, and missing regions were drawn in and added by hand. This exercise proved to be unfruitful in the respects of explaining why the correction worked best for the Spirit data. All that was gained was better statistics since more region extractions were possible.

The second idea was to examine whether the difference was due to saturated pixels in the Spirit image. Overexposure of the CCD could cause certain pixels to be saturated, such that the pixels next to the overexposed pixels would be exposed to the full R7 effect from the total intensity, while the value of the overexposed pixel in the image would not reflect the value of the total intensity received, resulting in an under-correction of that part of the image. This idea was tested by manually adding to the pixel value of the 5% brightest pixels of the image. By multiplying the brightest pixels with a factor of 2 prior to correction, the correction was forced into weighting these pixels higher. However, this only improved the Spirit mean offset by 2.5%, and could therefore not explain the discrepancy between the results for the two rovers.

Based on the above investigations, it was concluded that the discrepancy between the corrected offset values for Spirit and Opportunity was not due to saturated pixels, nor missing data, and remains unexplained. However, the 54 DN offset for Spirit after correction is less than 1.5% of the full well level (4095 DN), and the discrepancy is therefore minor and could be explained by systematic uncertainties. Some systematic uncertainties will be related to the calibration target reflectance model, especially in the longest and shortest wavelength filters that both have wider bands than the rest of the filters, as well as lower Signal to Noise (S/N). This is in particular reflected in the Spirit plot of Fig. 6.6, by the substantially larger errorbars for the shortest wavelength filters and the longest wavelength filters, compared to the rest. It is also reflected by the filter-to-filter variation in mean offsets for Spirit (Fig. 6.6). Other uncertainties, such as

small differences in the filter material, optics and CCD chip could lead to a difference in how the R7 artifact manifests itself on Spirit and Opportunity. One would expect, however, that since the model was developed based on Spirit data, the correction should also work best on Spirit data.

7.6 Correction of images used for previous publications

A preliminary investigation was presented in Section 6.3 of how the R7 filter artifact could potentially influence the results of previously published studies using the uncorrected R7 filter images, such as the results given in [32]. The investigation showed that the R7 filter effect removes a large part of the R6 to R7 filter slope in the spectra of known silica-rich materials (based on APXS/Mini-Tes measurements), thereby compromising the spectral characteristics of those spectra when comparing them to the spectra of other scene materials. It was shown that when plotting the data from the corrected set of images in a spectral parameter-space plot, the previously well-defined clustering of the data points representing the silica-rich material was hard to distinguish from the data points representing the rest of the scene elements. It was attempted to change two of the spectral parameters, in order to exclude the interference from bright soil, that tends to have steeper negative R6 to R7 slopes as well, especially if they are near a shadow, since this is where the R7 filter artifact is maximized. Several attempts were made to make the silica-rich material more distinguishable, which resulted in the calculation of 6 different hydration maps for an image containing a known silica-rich feature called Kenosha Comets, which is a bright patch of dirt exposed by one of the wheels on the Spirit rover as it was dragged across the ground. All the calculated maps picked up on the signature from the bright patch, and especially 3 of the maps (C, E and F) were able to pick up on the feature as well as minimize the signal from other scene elements. Using the spectral parameters defined for the calculation of the C, E and F maps, the hydration maps for the three remaining images of the set showed that it is still possible to obtain a map capable of selecting the interesting features in the frame, even though the maps do not show the same clear clustering around the silica-rich targets as the maps presented in [32]. It is in fact quite possible that the redefined parameter constraints are valid and better at picking out the hydration signatures with the number of false positives reduced. Some subtle clustering is even noticeable in a few frames at the areas indicated by the red arrows in Fig. 6.11. That the hydration maps do pick up on the Kenosha Comets feature in all frames of Fig. 6.10 is probably a consequence of it being the strongest silica-rich target measured by Spirit with a concentration of 90.1 ± 0.83 wt.% SiO_2 [36]. The magnitude of the negative spectral slope between R6 and R7 was by far the largest[32], which is also evident from the spectra in Fig. 6.7.

8

Conclusion and outlook

The previous chapters looked at the characterisation of an instrumental artifact observed in Pancam images obtained using the longest wavelength filter R7 (1009 nm). It was shown that the artifact is present in both pre-flight images obtained during calibration procedures, and in-flight terrain images from Mars, which ruled out the initial argument in [10] that the effect would be caused by transparency of the cal target materials. It was shown in Chapter 4 that the effect manifests itself as a long, faint tail in the image of a symmetric, simple source on a dark background. By assuming that this manifestation can effectively be seen as the large scale effect of an underlying pixel scale effect, a hypothesis was proposed that provides a physical explanation for why and how the effect arises. Namely that photons close to the silicon band gap at 1100 nm are let in by the longpass R7 filter, penetrate the CCD, scatter from the backside, goes back through the CCD, and are registered in a different pixel than the one they entered through initially. A model was developed based on this hypothesis, which was subsequently implemented into a correction, which largely removes the effect from the affected images. The preliminary verification of the correction on actual terrain images from Mars, shows that the correction removes the effect to a good degree, and does not introduce any new artifacts into the images. The correction was furthermore used to correct a number of cal target images, that showed an improvement of the mean offset values calculated based on the corrected R7 cal target images (see Fig. 6.6). The correction was furthermore implemented on R7 images used for previously published analyses. The previously published results consisted of hydration maps calculated along the Spirit traverse, based on an assumed correlation between the magnitude of the spectral slope between the R6 and R7 filter and the silica-richness of a target. The correction was found to remove most of the observed negative slope, thereby making it difficult to isolate the silica-rich targets from the rest of the scene elements.

The work presented throughout this thesis was focused on the fulfillment of the two aims listed in the Introduction under Section 1.1.

1. The primary aim was to establish a successful way of removing the effect from R7 filter images without compromising the scientific merit of the data.

This aim was fulfilled to a high degree, since it was shown that the developed correction performs well on all images presented in Chapter 6.

2. The second aim was to get a sense of the cause of the effect.

There is a good indication that this aim has been fulfilled. The mere fact that the correction, which is build on the base of the hypothesis about the origin of the effect, performs so well on all tested data, supports the conclusion that the R7 filter effect indeed is a manifestation of light at the longest wavelengths scattering off the backside of the CCD leading to an excess signal.

8.1 Implications

The artifact was found to be fairly weak in most cases, but significant in certain observations containing sharp edges with high contrast. It was found that the effect could cause the signal in some dark shadows in terrain images to be inflated by up to $\sim 33\%$. Analyses of early mission cal target images further showed that the reduced contrast was > 100 DN (full well = 4095 DN for comparison) for a hypothetical, perfectly dark pixel. For this reason it is possible that scientific conclusions based on un-corrected R7 images could be affected. This possibility was explored through the correction of R7 images used for previously published work regarding the calculation of hydration maps along the Spirit traverse. As mentioned it was found that the correction removes a large part of the negative slope between the R6 and R7 filter in all spectra tested, but a change in the parameter definitions used for the selection criteria in calculation of the hydration maps lead to a revised version of the hydration maps for a subset of selected images, showing that the method could potentially still be good for strong hydration targets. Potentially other previously published results could be affected as well.

8.2 Future work

All images presented in this thesis were radiance calibrated (RAD) images, since this made it possible to exclude any unwanted interference of the cal target calibration model, motivated by initial argument in [10], that the effect was due to transparency of the cal target materials. When extracting spectra from the images, the spectra were subsequently converted into I/F, since the physical interpretation of the radiance spectra would be difficult. It is worth noting that the R7 cal target images used in the calculation of the calibration coefficients were not corrected for the R7 effect on beforehand. However, it is expected that part of the R7 effect is unintentionally removed through the application of a dust model which purpose is to account for the accumulating layer of dust on top of the cal target. The accumulating dust layer will cause an offset in the fit of the modeled reflectances and the extracted radiances very similar to the example plot in Fig. 6.5. The dust model corrects for this offset, that way unintentionally correcting for a part of the R7 effect at the same time. it would still be very interesting to correct all cal target images, and see if this will affect the extracted spectra.

For further verification that the correction algorithm is sound, the plan is to implement the correction on a larger dataset, particularly on a number of other specific R7 observations that have previously been used for scientific analysis in published work. This work would make a natural follow-up publication to the present work. Provided that the correction algorithm is fully verified, the goal is to implement it on the full dataset of acquired R7 images eventually providing a recalibrated database of corrected R7 images.

Bibliography

- [1] Clifford, Stephen M. (2010). Depth of the Martian cryosphere: Revised estimates and implications for the existence and detection of subpermafrost groundwater. *JGR* 115, E7.
- [2] Dundas, Colin M., et al. (2018). Exposed subsurface ice sheets in the Martian mid-latitudes. *Science* 359 (6372), pp. 199-201.
- [3] Raack, Jan (2017). Water induced sediment levitation enhances downslope transport on Mars. *Nature Communications* 8, 1151.
- [4] Rice, J.W. (2016). Morphologic and Topographic Analysis of Geologic Features in the Columbia Hills, Gusev Crater, Mars. *47th Lunar and Planetary Science Conference*
- [5] Arvidson, Raymond E. (2016). Aqueous history of Mars as inferred from landed mission measurements of rocks, soils, and water ice. *J. Geophys. Res. Planets*, 121, 1602–1626, doi:10.1002/2016JE005079.
- [6] Orosei, R. (2018). Radar evidence of subglacial liquid water on Mars. *Science*,
- [7] Rice, J. W. (2016). Morphologic and Topographic Analysis of Geologic Features in the Columbia Hills, Gusev Crater, Mars. *47th Lunar and Planetary Science Conference*
- [8] Bell III, J. F., et al. (2006). In-flight calibration and performance of the Mars Exploration Rover Panoramic Camera (Pancam) instruments. *JGR*, 111, E2.
- [9] Kinch, M. K., et al. (2015). Dust deposition on the decks of the Mars Exploration Rovers: 10years of dust dynamics on the Panoramic Camera calibration targets. *ESS*, 2, 5.
- [10] Bell III, J. F., et al. (2003). Mars Exploration Rover Athena Panoramic Camera (Pancam) investigation. *JGR*, 108, E12.
- [11] Squyres, S. W., et al. (2004). The Opportunity Rover’s Athena Science Investigation at Meridiani Planum, Mars. *Science*, 306, 5702, p1698-1703.
- [12] Squyres, Steven W., et al. (2003), Athena Mars rover science investigation. *JGR*, 108, E12.
- [13] Squyres, S. W., et al. (2004). The Spirit Rover’s Athena Science Investigation at Gusev Crater, Mars. *Science*, 305, 5685, pp. 794-799.
- [14] Bell III, J. F., et al. (2017). The Mars Science Laboratory Curiosity rover Mastcam instruments: Preflight and in-flight calibration, validation, and data archiving. *ESS*, 4, 7.
- [15] J. Maki, pers. comm.

- [16] Maki, J. N., et al. (2003). Mars Exploration Rover Engineering Cameras. *JGR*, vol. 108, E12
- [17] Kahn, M. K., et al. (2013). Iterative methods of Richardson-Lucy-type for image deblurring. *Numer. Math. Theor. Meth. Appl.*, 6(1).
- [18] Broz P., Hauber E., Wray J.J., Michael G. (2017). *Earth and Planetary Science Letters*, 473 , pp. 122-130.
- [19] NASA Press Release October 29, 2018, URL: <https://www.nasa.gov/feature/update-on-opportunity-rover-recovery-efforts>
- [20] NASA Mars Exploration, URL: <https://mars.nasa.gov/allaboutmars/mystique/history/>
- [21] Parker, Timothy J., et al. (2010). The northern plains: A Martian oceanic basin?. Lakes on Mars, Edited by N.A. Cabrol and E.A. Grin. Elsevier B.V. ISBN 978-0444528544, 2010, p. 249-274.
- [22] Parker, Timothy J., et al. (1993). Coastal geomorphology of the Martian northern plains. *Journal of Geophysical Research*. 98, E6, pp. 11061-1107.
- [23] Bertaux, J.-L. , et al. (2007). Conversations on the Habitability of Worlds: The Importance of Volatiles. *Space Science Reviews*. 129. 123-165.
- [24] R Christensen, Philip. (2003). Formation of recent Martian gullies through melting of extensive water-rich snow deposits. *Nature*. 422. 45-8.
- [25] Bristow, Thomas F. et al. (2018). Clay mineral diversity and abundance in sedimentary rocks of Gale crater, Mars. *Science Advances*, vol. 4, 6.
- [26] M. T. Lemmon et al. (2004). Atmospheric Imaging Results from the Mars Exploration Rovers: Spirit and Opportunity. *Science*, vol. 306, Issue 5702.
- [27] J. F. Bell III et al. (2004). Pancam Multispectral Imaging Results from the Opportunity Rover at Meridiani Planum. *Science*, vol. 306, Issue 5702.
- [28] J. F. Bell III et al. (2004). Pancam Multispectral Imaging Results from the Spirit Rover at Gusev Crater. *Science*, vol. 305, Issue 5685.
- [29] Jeffrey R. Johnson (2015). Spectrophotometric properties of materials observed by Pancam on the Mars Exploration Rovers: 3. Sols 500–1525. *Icarus*, vol. 248, pp. 25-71.
- [30] William H. Farrand (2016). VNIR multispectral observations of aqueous alteration materials by the Pancams on the Spirit and Opportunity Mars Exploration Rovers. *American Mineralogist*, vol. 101, number 9.
- [31] Wang, A., et al. (2008). Light-toned salty soils and coexisting Si-rich species discovered by the Mars Exploration Rover Spirit in Columbia Hills. *JGR*, vol. 113, E12S40.
- [32] Rice, M. S., et al. (2010). Silica-rich deposits and hydrated minerals at Gusev Crater, Mars: Vis-NIR spectral characterization and regional mapping. *Icarus*, 205, 2, pp. 375-395.

- [33] Arvidson, R. E., et al. (2006). Overview of the Spirit Mars Exploration Rover Mission to Gusev Crater: Landing site to Backstay Rock in the Columbia Hills. *JGR*, 111, E2.
- [34] Arvidson, R. E., et al. (2010), Spirit Mars Rover Mission: Overview and selected results from the northern Home Plate Winter Haven to the side of Scamander crater. *JGR*, vol. 115, E7
- [35] Arvidson, R. E., et al. (2008). Spirit Mars Rover Mission to the Columbia Hills, Gusev Crater: Mission overview and selected results from the Cumberland Ridge to Home Plate. *JGR*, vol. 113, E12.
- [36] Ming, D. W., et al. (2008). Geochemical properties of rocks and soils in Gusev Crater, Mars: Results of the Alpha Particle X-Ray Spectrometer from Cumberland Ridge to Home Plate. *JGR Planets*, vol. 113, E12.
- [37] Squyres, S. W., et al., Overview of the Opportunity Mars Exploration Rover Mission to Meridiani Planum: Ragle crater to Purgatory Ripple. *JGR*, 111, E12.
- [38] Arvidson, R. E., et al. (2011). Opportunity Mars Rover mission: Overview and selected results from Purgatory ripple to traverses to Endeavour crater. *JGR*, vol. 116, E7
- [39] Squyres, S. W., et al. (2006). Rocks of the Columbia Hills. *JGR*, 111, E2.
- [40] Janesick, J. R. (2001). Scientific Charge-coupled Devices. *SPIE - The International Society for Optical Engineering*.
- [41] Carr, M. (2007). The Surface of Mars (Cambridge Planetary Science). Cambridge: Cambridge University Press. doi:10.1017/CBO9780511536007
- [42] Markwardt, C. B., Non-linear least squares fitting in IDL with MPFIT (2008), *Astronomical Data Analysis Software and Systems XVIII*, vol. 411, edited by D. Bohlender, P. Dowler, and D. Durand, pp. 251-254, Astronomical Society of the Pacific, Quebec, Canada.
- [43] Alizai, Khaled (2015). Master thesis, University of Copenhagen, Niels Bohr Institute, Mars Group.
- [44] Vuik, C. (2017). Iterative Solution Methods. *Research School for Fluid Mechanics, Delft Institute of Applied Mathematics, Delft, The Netherlands*.
- [45] Hansen, P. C. (2006). Deblurring Images, Matrices, Spectra, and Filtering. *Society for Industrial and Applied Mathematics*.
- [46] Hansen, P. C. (2010). Discrete Inverse Problems, Insight and Algorithms. *Society for Industrial and Applied Mathematics*.
- [47] Bertero, M. & Boccacci, P. (1998). Introduction to Inverse Problems in Imaging. *IOP Publishing*.
- [48] Folkner, W. M. et al. (2018). The Rotation and Interior Structure Experiment on the InSight Mission to Mars. *Space Science Reviews*. 214(5).
- [49] PDS archive, URL: <http://pds-geosciences.wustl.edu/missions/mer/index.htm>

- [50] Carr, M. H. & Head, J. W. (2010). Geologic history of Mars. *Earth and Planetary Science Letters*, vol. 294, issue 3-4, pp. 185-203.
- [51] Ehlman, B. L. & Buz, Jennifer (2014), Mineralogy and Fluvial History of the Watersheds of Gale, Knobel, and Sharp craters: A Regional Context for MSL Curiosity's Exploration. *Geophysical Research Letters*, vol. 42, issue 2.
- [52] Farrand, W. H., et al. (2014). Observations of rock spectral classes by the Opportunity rover's Pancam on northern Cape York and on Matijevic Hill, Endeavour Crater, Mars. *JGR*, 119, pp. 2349-2369.
- [53] Miliken, R. E., et al. (2009). Clay and Sulfate-bearing Rocks in a Stratigraphic Sequence in Gale Crater. *40th LPSC*
- [54] Miliken, R. E., et al. (2010). Paleoclimate of Mars as captured by the stratigraphic record in Gale Crater. *JGR*, vol 37, L04201.
- [55] Recktenwald, Gerald (2015). Stopping Criteria for Iterative Solution Methods. *Lecture Notes for ME 448/548*
- [56] Cabrol, Nathalie A. (2014). Sands at Gusev Crater, Mars. *JGR Planets*, 119, pp. 941-967.
- [57] Jerram, Paul A., et al. (2010). Back-thinned CMOS sensor optimization. *Proceedings of SPIE - The International Society for Optical Engineering*.
- [58] Shepard, M. K. (2017). Introduction to Planetary Photometry. *Cambridge University Press*.
- [59] Patrizio Campisi & Karen Egiazarian (2017). Blind image deconvolution: theory and applications. *CRC Press*.
- [60] A. Anusha & Dr. P. Govardhan (2017). Improvement of Image Deblurring Through Different Methods. *IOSR-JCE*, vol. 12, issue 6, pp. 79-82.
- [61] Board of Studies (NSW 1995). Methods of Algorithm Description, *Board of Studies NSW, Australia*
- [62] Armand Wirgin (2004). The inverse crime. *Math Phys*
- [63] Samuli Siltanen (2018). MAT-52506 Inverse Problems. *Lecture notes, University of Helsinki*
- [64] Prato, M, et al. (2012). Efficient deconvolution methods for astronomical imaging: Algorithms and IDL-GPU codes. *A&A*, vol. 539, A133.

1 **An Instrument Anomaly in the Mars Exploration**
2 **Rover Panoramic Camera (Pancam) 1009 nm Filter**
3 **(R7): Characterisation, Simulation, Correction and**
4 **Preliminary Verification**

5 **Simone J. Jakobsen¹, Kjartan M. Kinch¹, Morten Bo Madsen¹, James F. Bell**
6 **III², Danika Wellington², Lorinda Dajose³, Khaled Alizai⁴**

7 ¹Astrophysics and Planetary Science, Niels Bohr Institute, University of Copenhagen, Copenhagen,
8 Denmark

9 ²School of Earth and Space Exploration, Arizona State University, Phoenix, AZ, USA

10 ³California Institute of Technology, Pasadena, CA, USA

11 ⁴Astrophysics and Atmospheric Physics, Danish Technical University Space, Kgs. Lyngby, Denmark

12 **Key Points:**

- 13 • CCD
14 • Multispectral
15 • Calibration

Corresponding author: Simone Jakobsen, simonejj@nbi.ku.dk, simonejuuljakobsen@gmail.com

Abstract

During pre-flight calibration of the Panoramic Camera (Pancam) instrument on board the Mars Exploration Rovers MER A (Spirit) and MER B (Opportunity), a discrepancy was noted between 11-band spectra extracted from Pancam images of the camera's radiometric calibration target and reflectance spectra obtained with a spectrometer. This discrepancy was observed in the longest-wavelength filter of the camera (the longpass R7 filter with system $\lambda_{eff} = 1009$ nm), and consisted of a reduction in contrast between bright and dark regions. Here, we describe and characterise this effect. We propose that the effect arises because long-wavelength photons close to the silicon band-gap at 1100 nm, are allowed through the R7 filter, pass through the bulk CCD, scatter from the back-side, pass through the CCD again, and are registered in a pixel other than the pixel through which they originally entered. Based on this hypothesis we develop a model capable of accurately simulating the effect, and correct for it. We present preliminary results from testing this correction on pre-flight, as well as in-flight, images. The effect is small, but in some specific cases in small regions of high contrast, the effect is significant. In in-flight images of martian terrain we observed the signal in dark shadows to be artificially inflated by up to $\sim 33\%$ and analysis of early-mission cal target images indicated that the reduced contrast due to the artifact is equivalent to >100 DN (full well = 4095 DN) for a hypothetical perfectly dark pixel.

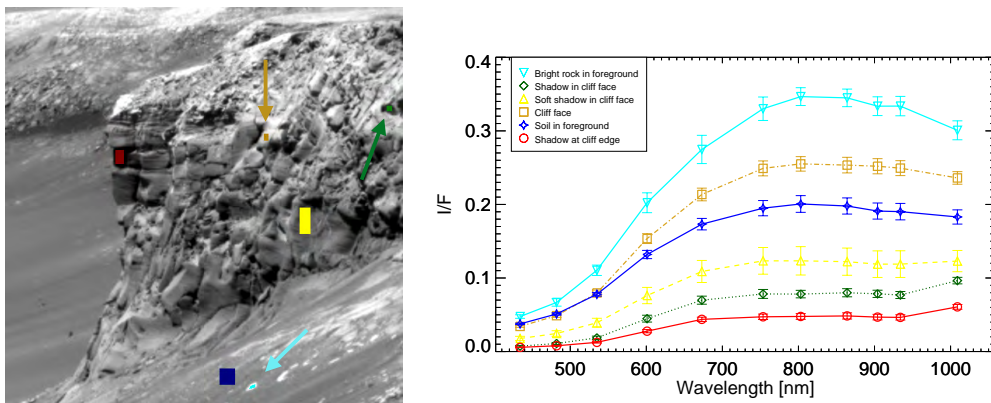
1 Introduction

The two rovers Spirit (MER A) and Opportunity (MER B) of the Mars Exploration Rover mission, landed on Mars in 2004 where they began studying the surface using the remote sensing and in-situ instruments of the Athena Science Payload [Squyres et al., 2004][Squyres et al., 2004][Arvidson et al., 2006][Squyres et al., 2006][Arvidson et al., 2008]. From their respective landing sites on opposing sides of the planet, they have contributed to the increasing knowledge of the history of water on the Martian surface. After the termination of Spirit's mission in 2010, Opportunity continued to collect data up until June 3, 2018, when a dust storm cut off the connection to the rover. Efforts are presently being put into recovering the connection.

One of the main tools available on both rovers is the panoramic camera imaging system (Pancam) [Bell et al., 2003], [Bell et al., 2004a], [Bell et al., 2004b], [Bell et al., 2006]. The system consists of two cameras on each rover, each equipped with an 8-position filter wheel. The 8 filters per camera cover the wavelength range (432 nm-1009 nm) with the filters denoted R1-R8 for the right eye and L1-L8 for the left eye. Two of the filters (L8 at $\lambda_{eff} = 440$ nm and R8 at $\lambda_{eff} = 880$ nm) are used for solar imaging, and one is a broadband filter (L1 at $\lambda_{eff} = 739$ nm). λ_{eff} here denotes the "effective" or central wavelength, defined as the weighted mean wavelength of the band, where the weighting is by the sensitivity of the whole camera system (filter, optics, detector). The remaining 13 filters are narrowband "geology" filters with band widths of 17-38 nm. Two of these are overlapping bands for the right and left eye, making it possible to obtain stereoscopic images at these specific wavelengths. This leaves the possibility of obtaining 11 point spectra spanning the full wavelength range [Bell et al., 2003].

In addition to being a navigation and assessment tool used as a way to identify interesting sites encountered along the traverse, the Pancam system provides image data used for a variety of scientific data analyses. In order to most accurately convert acquired images from units of radiance to units of radiance factor I/F (where I is the measured scene radiance at the sensor and πF is the incident solar spectral irradiance at the top of the martian atmosphere at the time of the observation), it is necessary to take several images of the onboard radiometric calibration target (hereafter referred to as the cal target) each Sol (a Sol is a day on Mars, which is approximately 24 hours and 37 minutes long). The calibration process and the conversion from radiance to I/F is described more carefully in [Kinch et al., 2015], and [Bell et al., 2006].

68 During preflight calibration activities in 2002, a discrepancy was noted between I/F
 69 spectra extracted from Pancam images of the cal target, and equivalent spectra obtained
 70 with a spectrometer [Bell *et al.*, 2006]. The discrepancy was observed as a small, yet notice-
 71 able, decrease in contrast between the brightest areas and the darkest areas of the
 72 cal target at the R7 filter wavelength ($\lambda_{eff} = 1009$ nm). At the time, this observation
 73 was hypothesized to be related to absorption properties of the cal target material (men-
 74 tioned in [Bell *et al.*, 2003] on page 24 line 20, and again in the caption of Figure 20),
 75 but a full explanation was not found, and the effect was deemed small enough to not af-
 76 fect the image data analysis on the tactical timescales needed for planned rover oper-
 77 ations. We present here evidence of the same effect seen in in-flight geology images ob-
 78 tained in the R7 filter, and argue that the effect can be significant in images covering
 79 high contrast regions, and that the effect may influence the spectra extracted from such
 80 regions. The image shown in the left panel of Fig. 1 was obtained by the Opportunity
 81 rover on Sol 954, and the I/F spectra presented in the right panel show an example of
 82 the R7 filter effect. The I/F values from the darkest shadow in the R7 image show sub-
 83 stantial increases relative to the neighboring filter (R6 at $\lambda_{eff} = 934$ nm). One shadow
 84 (shadow at cliff edge - dark red) shows an R7/R6 radiance ratio of 1.39 and another (shadow
 85 in cliff face - dark green) shows an R7/R6 radiance ratio of 1.33. After conversion to I/F
 86 the ratios are 1.31 and 1.26, respectively. These ratios are very large and unlikely to re-
 87 flect real wavelength-dependent differences in the radiance from these regions that are
 88 just shadows but not otherwise expected to be significantly different in composition or
 89 mineralogy from other materials in the scene. That this dramatic effect is observed in
 90 images of Mars indicates that it is in fact not related to properties of the cal target, but
 91 is rather a general effect related to the instrument itself.



92 **Figure 1.** An example of the effect seen in images obtained with the R7 filter. The R7
 93 image to the left was obtained on Sol 954 by the Opportunity rover. The extracted spectra,
 94 after conversion into I/F values are in the plot to the right. Spectra from very dark areas ex-
 95 hibit a dramatic upturn at R7 ($\lambda_{eff} = 1009$ nm). Image retrieved from the PDS archive for
 96 MER data. No scaling was performed between the right and left eye values. Values at overlap-
 97 ping wavelengths (L7R1 and L2R2) are presented as an average of the two eyes. The errorbars
 98 show the standard deviation between individual pixels within each region of interest. Image ID:
 99 1P212872229RAD76EVP2586R7C1.

100 Several previous works point out anomalies in R7 images relative to other filters
 101 of the Pancam system. In a study of dust deposition on the Pancam radiometric cali-
 102 bration target [Kinch *et al.*, 2015], it was found that derived dust albedo values showed
 103 a dramatic upturn at the longest wavelength (Fig. 8 and section 4.1.2 of [Kinch *et al.*,
 104 2015]). The authors also find that the ratio between direct and diffuse illumination showed

105 a dramatic downturn at the longest wavelength (Fig. 10 and section 4.2.2 of [Kinch *et al.*,
106 2015]). The authors hypothesize that these observations are caused at least partly by
107 an instrumental artifact, and mention under "future work" that this issue is in need of
108 further characterization (section 5.4 of [Kinch *et al.*, 2015]).

109 In [Wang *et al.*, 2008], the authors reported the existence of a characteristic spec-
110 tral downturn from filter R6 to R7 for regions on the Martian surface known to be sil-
111 ica rich. The authors of [Rice *et al.*, 2010] use this idea to look for a correlation between
112 the magnitude of the spectral slope and the amount of silica in the soil, in order to look
113 for hydrated minerals. They report a detection of false hydration signatures at very low
114 sun angles, and attribute it to inaccuracies in the cal target calibration model (p. 384,
115 second column [Rice *et al.*, 2010]). They further report a detection of hydrated miner-
116 als on surfaces at high local emission angles, and argue that this signal could be caused
117 by the Martian dust itself (p. 393, first column in [Rice *et al.*, 2010]).

118 In this paper, we describe and characterise the R7 anomaly, discuss a possible ex-
119 planation for it, as well as develop a model to use in an automated algorithm for the pur-
120 pose of correction. The focus here is on *showing* that an instrumental artifact exists, ex-
121 plaining the *cause* of this artifact, developing a *correction* for this artifact and *demon-*
122 *strating* the performance of this correction on a small dataset, specifically chosen to show
123 most clearly the effect of this artifact. It is an obvious implication of the present work
124 that a number of previous studies, including those mentioned above, are in need of re-
125 evaluation in light of the results presented here. We are committed to fully reprocess the
126 relevant data and re-evaluate the conclusions of these previous studies in light of the present
127 work. This work is ongoing, and will be presented in a follow-up paper, but is not cov-
128 ered here.

129 In Section 2 we describe the method of developing a model and algorithm that sat-
130 isfactorily simulates the R7 anomaly, in Section 3 we present the results for the model
131 parameters, and their interpretation, as well as examples of in-flight images before and
132 after correction. In Section 4 we discuss the model as a physical description of the R7
133 effect as well as the performance of the algorithm in the process of correcting data.

134 2 Methods

135 In order to develop a correction for the R7 images, it was necessary to characterise
136 the nature of the effect, and propose a hypothesis for how it arises. When a model had
137 been constructed that successfully simulated the effect, the model was applied to the in-
138 verse problem in order to correct for the effect. The hypothesised model was tested un-
139 der known and well defined conditions before it was introduced on independent image
140 data from the two rovers Spirit and Opportunity. In this section we provide an overview
141 of the entire process of developing a correction for the R7 filter images.

142 This section is divided into five separate parts. Section 2.1 describes and charac-
143 terises the observed effect, Section 2.2 treats the development of a model that describes
144 the effect in a satisfactory way, Section 2.3 describes how the developed model was im-
145 plemented into an algorithm that can be applied to image data, Section 2.4 looks at the
146 inverse problem of actually correcting for the effect, and finally Section 2.5 introduces
147 the correction on in-flight image data, giving us the opportunity to test its performance.

148 2.1 Characterisation of the Anomaly

149 We performed a systematic search of the large data set of images acquired during
150 preflight calibration of the Pancam system in order to identify image sequences that most
151 clearly display this subtle artifact in the R7 filter. We identified a sequence of images
152 acquired using the Spirit Pancam at the beginning of a stray light test as the ideal data
153 set for visualizing and characterizing the effect.

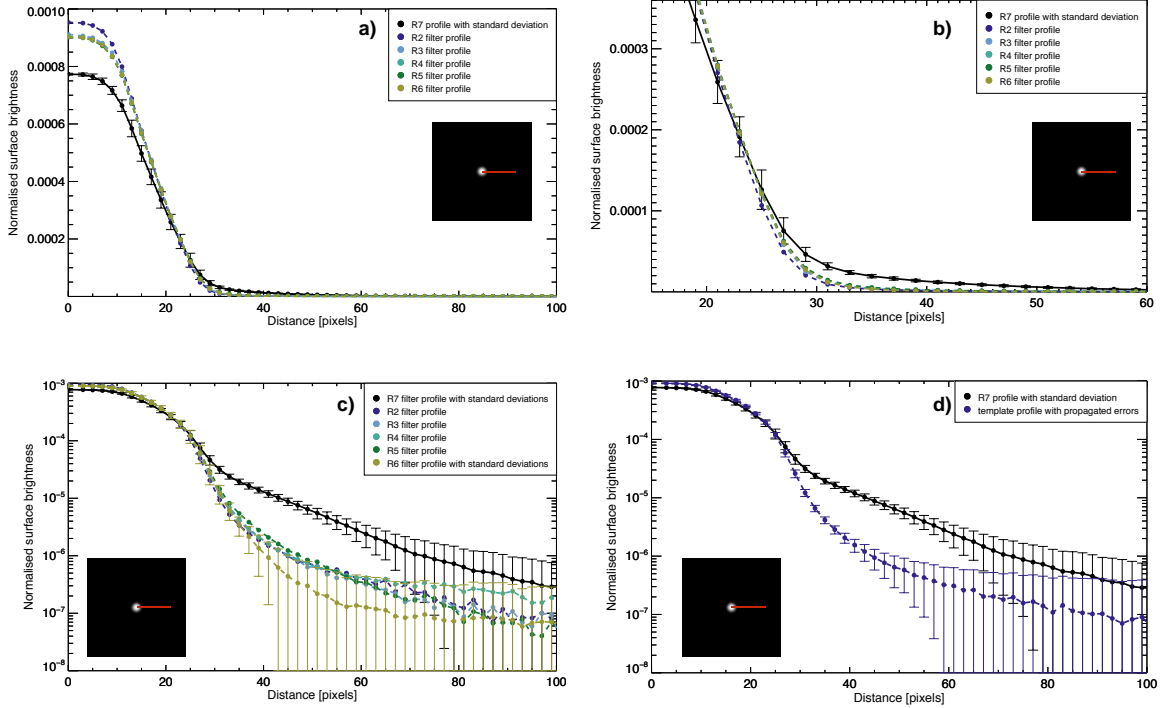
154 During the test, a series of images of a broad spectrum light source on a dark back-
 155 ground were collected using all 13 spectral filters R1-R7 and L2-L7. These test images
 156 were obtained in order to look for ghost images and stray-light effects arising as a con-
 157 sequence of CCD overexposure. However, a few images acquired at the beginning of the
 158 test sequence only showed the bright, circular light source on a dark background, with
 159 no overexposed pixels. The high symmetry, sharp edge and high contrast between the
 160 bright source and the extremely dark background in these images provided the ideal base
 161 for characterisation of the anomalous effect. These images gave us the means of com-
 162 paring images between filters, by relying on the sensible assumption that when normal-
 163 ising the images by dividing each image by the sum of all pixels, the light distribution
 164 in the image across all filters should be the same. The assumption is that there is no phys-
 165 ical effect that would make the light distribution substantially different from filter to fil-
 166 ter in this simple image of a bright, broad-spectrum source on a dark background. Any
 167 difference in light distribution between filters would be due to properties of the instru-
 168 ment. Thus, after normalisation, the image should be very similar in all filters. This is
 169 indeed the case for the images acquired in the filters R2-R6, whereas R7 images look very
 170 different (see Fig. 2 panels a), b), and c)).

171 We made use of these non-overexposed, radiance calibrated pre-flight, stray-light
 172 images from the Spirit rover for characterisation of the effect, and subsequent modeling
 173 and simulation. We looked for similar images from the Opportunity rover, but were un-
 174 fortunately not able to locate any non-overexposed images that were suitable for use (the
 175 point of the stray-light test, for which these images were acquired, is overexposure). How-
 176 ever, since the right eye cameras on the two rovers are designed to be identical, and in-
 177 deed are effectively identical in all respects as borne out by published pre-flight calibra-
 178 tion parameters [Bell *et al.*, 2003], we argue that a correction developed based on im-
 179 ages from the Spirit camera can be applied to Opportunity images, as long as the im-
 180 ages have been calibrated according to established procedures.

181 The source center in each filter image was identified by fitting the source with a
 182 2 dimensional Gaussian. An estimate of the background in each image was calculated
 183 as the average radiance value in an annulus with an inner ring radius of 258 and a width
 184 of 2 pixels, well outside the source area. Each background-subtracted image was then
 185 cropped to a size of 481 pixels by 481 pixels with the source centered, and normalised
 186 by dividing each pixel value with the pixel sum of the entire (cropped) image. Lastly,
 187 each image was binned by applying 60 concentric annuli centered on the source, each with
 188 a width of 2 pixels, and the largest annulus stretching a distance of 120 pixels from the
 189 source center. By determining the average radiance in each bin, it was possible to ex-
 190 tract a radial profile representing the light distribution in each filter. A plot showing the
 191 resulting profiles is presented in Figure 2.

200 Comparing the profiles for the right eye of Pancam (filters R2 at 754 nm, R3 at
 201 803 nm, R4 at 864 nm, R5 at 904 nm, R6 at 934 nm, and R7 at 1009 nm) reveals that
 202 the filters R2-R5 show a very similar distribution. The R6 profile also show a very sim-
 203 ilar distribution except for some deviation visible on a logarithmic scale at distances be-
 204 yond 40 pixels from the center, where all these filters show very low signal. We believe
 205 that this small difference between the R6-profile and R2-R5 is related to the background
 206 subtraction, and we stress that although distinctly visible on a logarithmic plot the dif-
 207 ference in absolute values is much smaller than the difference between the R7-profile and
 208 the rest of the filters (see Figure 2 panel b). The R6-profile is lower than R2-R5 only by
 209 a value comparable to the standard deviation of pixels in each annulus of the R6 image
 210 (see Figure 2 panel c)).

211 In order to compare the distribution of the R7 profile to that of the other filters,
 212 we averaged the normalised, stray-light images from the filters R2-R6, and extracted a
 213 new profile using the same binning procedure. This averaged image is hereafter referred
 214 to as the template image, and the profile extracted from the image is hereafter referred
 215 to as the template profile (see Figure 2 panel d)). This template profile is used through-
 216 out this work as a representation of how the light is expected to be distributed when in-



192 **Figure 2.** The plot shows 6 different radial profiles, each representing the light distribution of
 193 the source in a filter mounted in front of the right eye of the Pancam on board the Spirit rover.
 194 Black curve: R7, coloured curves: R2-R6. Note how the R2-R6 profiles are generally very similar
 195 and have a distinctively different shape than the R7 filter profile. a) linear representation of all
 196 filter profiles, b) linear zoom-in of the region where the difference between the R2-R6 profiles and
 197 the R7 profile is largest, c) logarithmic representation of all filter profiles, d) the template profile
 198 taken as the average of the R2-R6 profiles, and the R7 profile. Image ID: 020715222717.0001989_
 199 R7_103_rad.

217 cident on the CCD. The error bars of the template profile are the Root Mean Square (RMS)
 218 of the individual errors for the profiles R2-R6 in each extraction bin. By treating the tem-
 219 plate profile as "truth", we can think of the deviation of the R7 profile as solely a conse-
 220 quence of the instrumental effect, making it possible to model it. This is only possi-
 221 ble because of the simplicity, and symmetry of the stray-light images.

222 The most dramatic difference between the R7 profile and the template profile is
 223 found between pixel ~ 30 and pixel ~ 60 from the source center, just on the dark side
 224 of the transition between the bright and the dark area of the image. In relative terms,
 225 the difference is at a maximum at 49 pixels from the source center (see Figure 2 panel
 226 d)), where the R7 value is a factor of ~ 10 larger than the template value.

227 2.2 The Model

228 In this section we propose a detailed hypothesis capable of describing the instru-
 229 mental artifact by analysing the difference between the R7 filter image and the rest of
 230 the filters. Unlike the R2-R6 filters, the Pancam R7 filter is a longpass filter, relying on
 231 the rapidly declining sensitivity of the CCD itself to provide the long-wavelength cut-

off. In contrast, the conceptually similar Mastcam system on board the Mars Science Laboratory (MSL), all filters are bandpass filters [Bell *et al.*, 2017]. Since the described effect is only observed on MER, not on MSL, and only in the R7 filter, this leads us to believe that the effect arises from light at wavelengths very close to the edge of the silicon band gap at 1100 nm. Perhaps the fact that this effect is not observed on MSL Mastcam is related to the long-wavelength cutoff of the longest-wavelength filter on the MSL Mastcam ($\lambda_{eff} = 1035 \pm 50 \text{ nm}$). However, the CCDs are different devices, and other differences in the construction and material of the two systems may also be important, and therefore it may be difficult to directly compare the two.

In any case, assuming that the effect is an intrinsic occurrence of the CCD, we propose that the effect is related to the very long penetration depth in silicon attained by light with wavelengths close to 1100 nm. We propose that the artifact arises from light penetrating the CCD and reflecting off the backside of the CCD bulk material at the longest wavelengths let in by the R7 filter. Figure 3 shows how instead of being completely absorbed, some of the light hitting a pixel of the CCD will be transmitted through the bulk silicon layer, scatter from the back edge, travel back through the CCD and reach the active layer at a different, nearby pixel location. Some of the radiation will be registered in this nearby pixel, with the amount determined by the distance x from the pixel receiving the incident light. This ultimately results in a broadening of the signal that was intended to be registered in one single pixel. In effect, this small artifact will manifest itself as an additional very long, very faint, tail in the point spread function for the camera when using the R7 filter. The result in the case of a symmetric source, such as the stray-light image source, will be a broadening of the light distribution across the image as seen in Fig. 2.




Figure 3. The figure shows a schematic representation of our hypothesis for the nature of the observed effect and how it affects the registration of radiation in one pixel. Part of the light incident on the CCD penetrates the pixel layer without giving rise to a signal, and instead reflects off the backside of the CCD eventually resulting in an additional signal being registered at a distance x from the pixel receiving the incident light. C is the bulk thickness of the CCD.

In order to correct for the effect we first develop a model capable of simulating the transformation of the template profile into the R7 profile, using our hypothesis as a good description of the behavior of the light let in by the R7 filter ($\lambda_{eff} = 1009 \text{ nm}$).

The total distance traveled by light hitting normal to the surface of the CCD, penetrating the bulk silicon layer, getting backscattered and eventually reaching the active layer is given by: $L = C + \sqrt{C^2 + x^2}$, where C is the thickness of the CCD and x is the distance between the pixel receiving the initial radiation and the pixel registering

the backscattered radiation. For convenience we measure distances in units of pixels (1 pixel = 12 microns).

The light will be attenuated on its path through the CCD according to the Beer-Lambert law: $I(r) = I_{in} * e^{-B(\lambda) * r}$, where $I(r)$ is the intensity of the light after traveling a distance r through the CCD chip, I_{in} is the intensity of the light incident on the CCD, and $B(\lambda)$ is the absorption coefficient of the CCD silicon, which is wavelength dependent.

For simplicity we assume the scattering off the backside of the CCD to be Lambertian and we therefore introduce a cosine factor into the model: $I(x) \propto I_{in} * e^{-B(\lambda) * L} * \cos(\theta)$, where θ is the angle between the incoming light normal to the plane of the backside of the CCD, and the scattered light, which gives us: $\cos(\theta) = \frac{C}{\sqrt{C^2 + x^2}}$, and in turn: $I(x) \propto I_{in} * e^{-B(\lambda) * (C + \sqrt{C^2 + x^2})} * \frac{C}{\sqrt{C^2 + x^2}}$.

After scattering, the light is no longer collimated and the amount of light reaching a pixel at a particular distance x is approximated by the inverse-square law, giving us the final form of the function describing the scattering of light from the backside of the CCD for one specific wavelength λ :

$$I(x) \propto I_{in} * e^{-B(\lambda) * (C + \sqrt{C^2 + x^2})} * \frac{C}{\sqrt{C^2 + x^2}} * \frac{1}{C^2 + x^2} = I_{in} * e^{-B(\lambda) * L} * \frac{C}{(C^2 + x^2)^{3/2}} \quad (1)$$

We now assume that the chance of generating a signal in the CCD at a given wavelength within the R7 band can be treated as inversely proportional to the penetration depth - in effect this means that we assume that attenuation of the light as it penetrates the bulk layer of the CCD is dominated by electrons lifted into the conduction band and ignore other causes of absorption. This assumption is likely to be good over much of the R7 band but will break down very close to the band gap at 1100 nm, where the probability of excitation becomes very small and absorption begins to be dominated by other processes.

Making this assumption and introducing an overall scaling factor A_{scale} that contains the intensity of light incident on the CCD, the efficiency of the reflection off the CCD back surface as well as some factor for the conversion efficiency in a pixel, we can write the signal generated in a given pixel by this process due to light that originally reached the CCD at a pixel a distance of x away:

$$f(x)_\lambda = A_{scale} * B(\lambda) * e^{-B(\lambda) * L} * \frac{C}{(C^2 + x^2)^{3/2}} \quad (2)$$


with $L = C + \sqrt{C^2 + x^2}$.

Eq. 2 expresses how the artificial signal varies as a function of distance x from the pixel where the light originally reached the CCD. It also describes the dependence on wavelength. At shorter wavelengths (high $B(\lambda)$) the signal will disappear because of the exponential term as the absorption is too high for the light to penetrate the CCD. At very long wavelengths the signal is depressed by the low sensitivity of the CCD (low $B(\lambda)$).

Since $B(\lambda)$ varies dramatically over the R7 filter band, we do not expect a single wavelength approximation to be sufficient to describe the observed effect. We therefore expand the model to describe the total effect across the whole wavelength band of the R7 filter. We argue that we can use an exponential profile as an approximation of how the silicon absorption coefficient decreases with increasing wavelength at the wavelengths of interest (~ 900 nm - 1100 nm). By integrating Eq. 2, we find a model describing the integrated effect across the R7 filter band.

The validity of such an approximation is supported by the plot seen in Fig. 4. The plot has been generated based on a table listing the penetration depths in silicon across a broad wavelength range (page 172 Table 3.1 in [Jamesick et al., 2001]).

We use the approximation $B(\lambda) = B_0 * e^{-\alpha * \lambda}$, and integrate Eq. 2. The lower limit of the integration, λ_{min} is assumed to be at a wavelength where the effect disappears because the light does not penetrate far enough (i.e., $e^{-B(\lambda_{min}) * L} = 0$). This can



313 **Figure 4.** The plot shows the penetration depth in silicon as a function of wavelength in
 314 the R7 filter wavelength range. The black line is an exponential fit to the data. The data was
 315 adopted from Table 3.1 on page 172 in [Janesick *et al.*, 2001].

319 be thought of as 960 nm at the approximate lower cut-off of the R7 filter, where the pen-
 320 etration depth is approximately ($\sim 100\mu\text{m}$) whereas the CCD thickness is at least 400
 321 microns [Makin pers. comm.] (see [Janesick *et al.*, 2001] for standard silicon thickness
 322 for a front illuminated CCD). The upper wavelength limit, λ_{max} , will be used as a free
 323 parameter and can be thought of as the point close to 1100 nm where silicon becomes
 324 almost transparent to the radiation and absorption is dominated by other processes that
 325 do not generate electrons in the conduction band and thus do not generate signal in the
 326 CCD. Thus, the signal generated in a given pixel can then be expressed as:

$$f(x) = \int_{\lambda_{min}}^{\lambda_{max}} f(x, \lambda) d\lambda = A_{scale} * \frac{C}{(C^2 + x^2)^{3/2}} \int_{\lambda_{min}}^{\lambda_{max}} B(\lambda) * e^{-B(\lambda)*L} d\lambda \quad (3)$$

327 in which we assume that A_{scale} can be treated as wavelength-independent. We integrate
 328 by substitution. Differentiating $B(\lambda) = B_0 * e^{-\alpha*\lambda}$ with respect to λ gives us $\frac{dB(\lambda)}{d\lambda} =$
 329 $-\alpha * B(\lambda)$ and by changing limits in our expression we get:

$$f(x) = A_{scale} * \frac{C}{(C^2 + x^2)^{3/2}} * \frac{1}{\alpha} \int_{B(\lambda_{max})}^{B(\lambda_{min})} e^{-B(\lambda)*L} dB \quad (4)$$

330 Carrying out the integration and assuming $e^{-B(\lambda_{min})*L} = 0$, we get:

$$f(x) = A_{scale} * \frac{C}{(C^2 + x^2)^{3/2}} * \frac{1}{\alpha} * \frac{1}{L} * e^{-B(\lambda_{max})*L} \quad (5)$$

331 Introducing $A = A_{scale}/\alpha$, writing B as shorthand for $B(\lambda_{max})$, and inserting $L =$
 332 $C + \sqrt{C^2 + x^2}$, we finally get:

$$f(x) = A * \frac{1}{C + \sqrt{C^2 + x^2}} * e^{-B*(C + \sqrt{C^2 + x^2})} * \frac{C}{(C^2 + x^2)^{3/2}} \quad (6)$$

333 Eq. 6 expresses the total contribution to the R7 signal in a given pixel from light that
 334 entered the CCD at another pixel at a distance x. The equation has three free param-
 335 eters: A is an overall scaling factor governing the strength of the effect, B is the absorp-
 336 tion coefficient at the point close to 1100 nm where silicon becomes almost transparent

337 to the radiation and absorption is dominated by other processes that do not generate
 338 electrons in the conduction band and thus do not generate signal in the CCD and C is
 339 the physical thickness of the CCD wafer.

340 The above exercise provides us with an expression that should be able to describe
 341 the transformation of the template profile into the R7 profile, as long as our hypothe-
 342 sis is a good approximation for the nature of the effect. Applying the above model pixel
 343 by pixel to the template image and extracting a profile from the resulting image should
 344 result in a profile that is comparable to the R7 profile. By running this simulation as a
 345 fitting routine we are able to optimize the model, and attain approximated values for
 346 the model parameters, that can then later be used in the correction.

347 2.3 Simulating the Effect

348 Assuming our model is a good approximation of the effect, and arguing that the
 349 template image should be a good representation of how the light would be distributed
 350 in the R7 image if the effect was not present, we can recreate the R7 image from the tem-
 351 plate image. As mentioned in Section 2.2 the large scale effect seen in the stray-light im-
 352 ages (shown in Fig. 2) is assumed to be a representation of the pixel scale effect described
 353 by the hypothesis (Fig. 3). Hence, by applying equation 6 to each pixel of the template
 354 image, it should be possible to simulate the behavior of the radiation incident on the CCD,
 355 and recreate the effect.

356 If we define the R7 image as X_{R7} and the template image as X_T , we can write up
 357 the process as follows:

$$X_{R7}(i, j) = X_{Template}(i, j) + X_{Template}(i, j) * D + \sum_{k \neq i}^K \sum_{l \neq j}^L X_{Template}(k, l) * f(x) \quad (7)$$

358 where the parameter D is negative and introduced to keep the normalisation of the im-
 359 age, and where (K,L) is the size of the image. This assures that the transformation does
 360 not change the total signal in the image (the sum of all pixels will be the same). The
 361 variable x still defines the distance from the pixel (i, j) to the pixel (k, l) and $f(x)$ is as
 362 given in Eq. 6. The double sum runs over all pixels within 120 pixels of pixel (i, j) . The
 363 function $f(x)$ is very small at $x > 120$ (see Eq. 6) so this truncation at $x = 120$ does not in-
 364 fluence the final result significantly.

365 The result of this simulation will be an image that is directly comparable to the
 366 R7 image through its radial profile. The opportunity of comparing the simulated image
 367 directly to the R7 image through their radial profiles makes it possible to fit the model
 368 parameters. This in turn provides us with what is needed to develop a correction with
 369 the purpose of removing the effect from R7 filter images.

370 In order to be able to apply the model to each pixel of the template image through
 371 equation 7, an algorithm containing the model was formulated using the Interactive Data
 372 Language (IDL), and a Levenberg-Marquardt weighted least-squares fitting routine (MP-
 373 FIT by Craig Markwardt [*Markwardt et al.*, 2008]) was used in order to determine the
 374 best fit parameter values for the model. The script reads in the template profile, and the
 375 function containing the model is called by determining a set of start values for the model
 376 parameters. A radial profile is extracted from the output and evaluated against the R7
 377 filter profile, until an acceptable fit has been achieved.

378 The first guess for the parameter values was based on their expected physical in-
 379 terpretation according to our model and hypothesis (see Section 2.2). This was done to
 380 avoid being stuck in a local minimum when running the fit with most of the parameters
 381 free. An acceptable first guess for the parameters was determined by running the sim-
 382 ulation with different values of the parameters without invoking the fit, and evaluating
 383 the result by eye. It proved difficult to get convergence to a stable result when all 4 pa-
 384 rameters (A, B, C, D) were kept free. For this reason the parameter C was fixed at 33,
 385 corresponding to 33 pixels = 396 microns (the minimum thickness of the CCD reported

386 by the manufacturer is 400 microns, [Maki pers. comm.] (see [Janesick *et al.*, 2001] for
 387 standard silicon thickness for a front illuminated CCD). When this was done the result
 388 was a very satisfactory fit as can be seen in Section 3.1.

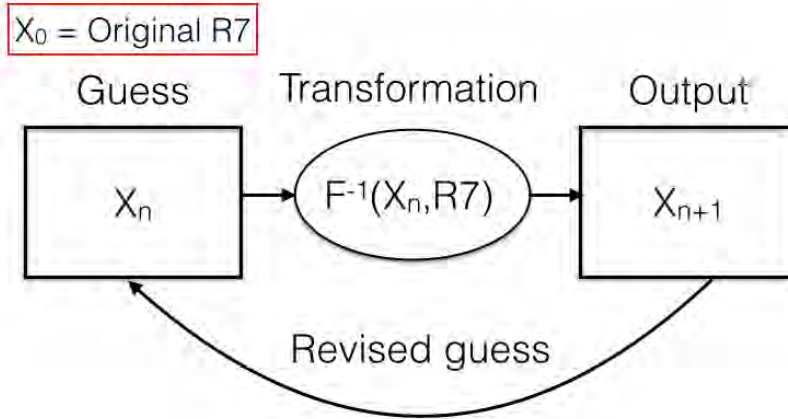
389 2.4 The Inverse Process - The Correction

390 Equation 6 is only strictly valid for pixels that are more than 120 pixels from the
 391 edge of the image. For pixels closer to the edge, a weighting function is applied to Eq.
 392 7, which takes into account the fact that the pixels closest to the edge of the image will
 393 be affected only by the pixels inside the image and therefore will be artificially dimmed
 394 if no weighting is included.

395 By simulating the effect as described in Section 2.3, we found an estimation of the
 396 parameter values in our model. This gave us the means of reversing the transformation,
 397 so that we can actually correct for the effect we have so far been simulating. The inverse
 398 problem is not as simple as the direct simulation, as can be seen from rearranging Eq.
 399 7:

$$X_{True}(i, j) = X_{R7}(i, j) - X_{True}(i, j) * D - \sum_{k \neq i}^K \sum_{l \neq j}^L X_{True}(k, l) * f(x) . \quad (8)$$

400 Eq. 8 describes how to derive the best estimate of the "true" light distribution from a
 401 recorded R7 image, i.e. how to correct the image. Since X_{True} appears on both sides of
 402 the equation and particularly also inside the sum we have to perform the correction through
 403 an iterative process. We can achieve a corrected version of the R7 image by letting the
 404 corrected output converge towards the true light distribution. Each iteration towards
 405 convergence uses the best fit parameter values determined in the simulation of the ef-
 406 fect (i.e. $f(x)$ is unchanged). The model is applied to the original R7 image by making
 407 a guess for the corrected version of the image. The first guess is the original R7 image
 408 itself, which is an appropriate initial guess, since the effect is generally very small. Fol-
 409 lowing each iteration, the output image is used as a new guess until the output converges
 410 towards a corrected result. A schematic showing the procedure is provided in figure 5.



411 **Figure 5.** This schematic shows how the R7 images are corrected using an iterative process.
 412 Each iteration consists of a full circle in the schematic and results in an output which is a fully
 413 transformed image one pixel at a time. This output is then used as the new guess in the follow-
 414 ing iteration. The output will converge, and when a certain pre-defined cut-off value is reached,
 415 the iterative process will stop.

416 Writing it up in a similar way to the one used in Eq. 8, and using the terminol-
 417 ogy from the schematic in Fig. 5, we get a correction for each pixel:

$$X_{n+1}(i, j) = X_{R7}(i, j) - X_n(i, j) * D - \sum_{k \neq i}^K \sum_{l \neq j}^L X_n(k, l) * f(x) . \quad (9)$$

418 The correction is performed pixel by pixel, as the simulation was done pixel by pixel.
 419 The correction to the (i, j) pixel of the R7 filter image is the summed correction from
 420 all neighboring pixels within a distance of 120 pixels from the (i, j) pixel, which is the
 421 maximum distance we include in the correction.

422 The correction from each bin is weighted before it is added to the (i, j) pixel to avoid
 423 an artificial brightening of the edge pixels in the corrected image, since these pixels have
 424 fewer neighboring pixels. The corrected output from the iteration is used as a new guess
 425 for the corrected image. This process is repeated until a convergence is reached. The num-
 426 ber of iterations needed to obtain sufficient convergence is determined by the change in
 427 the mean of the pixel values in the image:

$$test - value = \sum (X_N - X_{N-1})^2 / (K * L) , \quad (10)$$

428 where X_N is the output image from the current iteration, X_{N-1} is the output image from
 429 the previous iteration, and $K \times L$ is the size of the image in units of pixels. The correc-
 430 tion was invoked 20 times on the stray-light R7 image, and the test-value for each iter-
 431 ation was logged and evaluated.

432 2.5 Correction of in-flight Images

433 As a first check of the correction method, the pre-flight, stray-light R7 image was
 434 corrected and the result inspected. This way we were able to choose a cut-off value to
 435 use for consecutive corrections as described in Section 2.4. However, since the model and
 436 therefore the correction was developed using these same pre-flight images, it is perhaps
 437 not surprising that the correction is successful in this case.

438 2.5.1 Verification Using in-flight Geology Images

439 As a first independent verification of the correction we picked two examples of in-
 440 flight R7 images to correct and evaluate. A complex geology image from the Spirit rover
 441 was chosen based on it exhibiting high contrast regions. Additionally a similarly com-
 442 plex and high-contrast image from the Opportunity rover was chosen to further ensure
 443 that the algorithm is general enough to correct images from both rovers. Spectra from
 444 both images before correction and after correction were extracted and results are pre-
 445 sented in Section 3. The regions chosen for spectral extraction were regions exhibiting
 446 particular high contrast to their surroundings, that is regions that are very bright and
 447 surrounded by a very dark area as well as very dark regions surrounded by a very bright
 448 area. The effect of the correction was inspected by producing images for each rover show-
 449 ing the ratio between the uncorrected and corrected R7 image, the ratio between the R2
 450 filter image and all adjacent filters R3-R7, as well as the corrected version of the R7 fil-
 451 ter image. All in-flight images from both rovers used for this test were I/F calibrated
 452 using the standard pipeline for calibration of MER images [Bell *et al.*, 2006] and the re-
 453 sulting images are presented in Section 3.

455 2.5.2 Verification Using in-flight Calibration Target Images

456 The MER panoramic cameras rely on an external calibration target (cal target)
 457 for converting in-flight images of the martian terrain from radiance to radiance factor

I/F [Bell *et al.*, 2003]. Radiance factor is π times the bidirectional reflectance of the surface and can be easily compared to laboratory reflectance data. The calibration target contains 7 different surfaces with well-characterized reflectance properties. During the reflectance calibration process [Bell *et al.*, 2006] radiance values are routinely extracted from these 7 regions of a calibration target image. The observed radiances are then plotted against reflectances that are known from pre-flight characterization. Ideally, the observed radiance should be proportional to the reflectance of the surface, so when plotting observed radiances against known reflectances the data should fall on a straight line through the origin. The slope of this line is a measure of the solar irradiance and this quantity is used to convert images of the martian terrain obtained at the same time from radiance to I/F.

Later in the mission dust deposition on the cal target must be taken into account and corrected for [Kimch *et al.*, 2015], but early in the mission the cal targets were largely dust-free and the observed radiances are generally observed to fall close to straight lines through the origin for all filters - except R7. We obtained data from cal target observations acquired through the first 5 sols of the mission. Sol 5 was chosen as a cutoff from a balance between wanting only very early images to minimize the dust influence but also desiring an adequate number of calibration target observations.

There are a total of 92 cal target observations from Spirit and 172 observations from Opportunity on the first 5 sols when counting all full-resolution images using any one of the 13 narrow-band "geology" filters of the Pancam system. For each cal target observation we performed a fit of observed radiances to model reflectances using a simple linear model with two straight lines representing a) regions of the cal target in direct sunlight and b) regions in the shadow of the shadow post. The two lines were constrained to intercept the y (radiance) axis at the same point, but the offset from zero was a free parameter of the fit. Ideally this offset should be zero as the radiance emitted by the calibration target is proportional to the reflectance of the surface. The recorded offset value was back-converted from radiance to data numbers (DN's) by multiplying with the exposure time and dividing by standard pre-flight-determined conversion factors [Bell *et al.*, 2003]. This was done in order to make a better direct comparison of the offset value between filters. The offset in DN can be directly compared to the maximum value of the 12-bit digitization (0-4095 DN) to get an idea of the size of the offset relative to the average signal level in the image, which will generally be set by the auto-exposure algorithm to be a significant fraction of the maximum value. The results of this analysis are shown in Section 3.2.3.


3 Results

This section is divided into two parts, Section 3.1 presents and discusses the results achieved during the development of the algorithm, that is the simulation of the R7 anomaly based on the stray-light images from preflight calibration of the Pancam for the Spirit rover. Section 3.2 treats the performance of the algorithm when applied to the inverse problem, that is the actual correction of the R7 filter images. This include both the pre-flight, stray-light images as well as the in-flight geology and cal target images.

3.1 Simulation

Here we present the results for the development of the algorithm. As described in Section 2.3, we used a fitting routine to determine the best fit values for the model parameters. The result of the fit can be seen in Fig. 6. The template profile is shown in blue, the R7 profile is shown in black, and the fitted, simulated model is shown as the dark red curve. The fit follows the measured R7 profile almost perfectly, indicating a very good fit. The fit was run with the parameter C held at $C = 33.0$, equivalent to a bulk thickness of the CCD of ~ 396 microns. The fit resulted in parameter values of: $A =$

508 96.2, $B = 0.0388$, and $D = -0.211$. A is a scaling factor determining the overall strength
 509 of the effect, B is the absorption coefficient of the radiation at the longest wavelengths,
 510 equivalent to a penetration depth of ~ 309 microns. D is an overall scaling factor ensur-
 511 ing the normalisation, i.e. that the sum of all pixels in the image is unchanged by the
 512 transformation.



513 **Figure 6.** This figure shows a logarithmic plot of the fit used to determine the parameter val-
 514 ues for the optimised model. The dark blue curve is the template profile, the black curve is the
 515 R7 profile to be simulated, the dark red curve is the best fit, simulated profile. The uncertainties
 516 for the R7 profile is the standard deviation in each bin.

517 The main goal of this investigation was to reproduce the observed R7 filter effect
 518 in order to be able to remove it from the affected images. With the fit in Fig. 6 we have
 519 shown that the effect can be simulated to a good accuracy, which gives us confidence that
 520 we have a mathematical description of the effect that is accurate enough to be used for
 521 development of a useful correction. Secondly, since our mathematical description is based
 522 directly on a specific physical hypothesis for the origin of the effect, the good fit makes
 523 us believe that our physical hypothesis is, in fact, correct, and that the instrumental arti-
 524 fact arises from photons close to 1100 nm that penetrate the CCD, scatter from the
 525 backside, pass through the CCD chip again and are registered on the way out in a dif-
 526 ferent pixel from the one through which they originally entered the CCD. Nonetheless
 527 we point out that the main requirement for the correction is that our model describes
 528 the effect well mathematically as demonstrated by Fig. 6. The correction may be per-
 529 fectly useful even if our physical hypothesis is not entirely correct.


530 3.2 Correction

531 In the following we look at the correction of the pre-flight images and the in-flight
 532 images in separate sections. The pre-flight, stray-light images from Spirit were used to
 533 develop the algorithm, and the results achieved from the correction of these images is
 534 presented in Section 3.2.1. Section 3.2.2 presents the results achieved based on correc-
 535 tion of images independent of the development of the algorithm, that is, in-flight, geol-

536 ogy images from both Spirit and Opportunity, as well as cal target images from both rovers
 537 (3.2.3).

538 *3.2.1 Performance of the Correction on Stray Light Image Data from* 539 *the Spirit Rover (pre-flight)*

540 The results of the correction of the Spirit stray-light images is presented in Fig. 7.
 541 The plot shows the convergence of the output after each iteration. The black curve is
 542 the light distribution in the uncorrected R7 image, the dark blue curve indicates how the
 543 light should in fact be distributed as expected from the template profile. The rest of the
 544 coloured curves represent the first five, and the tenth iterations, showing that the R7 pro-
 545 file converges towards the correct distribution. It was observed by eye that the profile
 546 converged after approximately four iterations. The test value of the fourth iteration (Eq.
 547 10) were of the magnitude 10^{-14} , making this our choice for the cut-off value to be used
 548 for all other corrections, determining when the iterative process should be stopped. The
 549 number of iterations needed for correcting an image could vary, therefore the cut-off value
 550 described in Section 2.4 is introduced to ensure that a sufficient convergence is reached.



551 **Figure 7.** Convergence of the output towards a good correction over 10 iterations. A conver-
 552 gence of the output is reached already after 3-4 iterations. After 4 iterations the corrected profile
 553 does not change much with more iterations, and it matches the template profile within the errors.
 554

555 *3.2.2 Performance of the Correction on some Examples of Image Data* 556 *from Spirit and Opportunity (in-flight)*

557 In this section we present the results of the correction of two geology images, one
 558 from the Spirit rover, and one from the Opportunity rover. Both images are independ-
 559 ent of the development of the correction.

560 In Fig. 8 we show the uncorrected and corrected images side by side, comparing
 561 them for each rover. The figure also shows the ratio between the uncorrected and cor-
 562 rected image for each rover.

578 Both images were selected because they both contain a large number of high con-
 579 trast areas, created by clusters of small and large rocks on the Martian surface for the
 580 Spirit image, and by protrusions in the cliff walls in the case of the Opportunity image.
 581 Looking at some of the darkest shadows it is marginally noticeable by eye that they be-
 582 come more distinct in the corrected version of the images, and the edge between the dark-
 583 est pixels and the brighter pixels outside the shadowed area becomes sharper. The same
 584 tendency is observable for the brightest areas. The ratio between the corrected and un-
 585 corrected R7 image shows that no artifacts develop during the correction, and no sys-
 586 tematic effects are observed to distort the images. The primary effect of the correction
 587 is a sharpening of the corrected image relative to the uncorrected image. For this rea-
 588 son the ratio of the (sharper) corrected image to the (less sharp) uncorrected image shows
 589 a reasonable representation of the terrain.

590 In Fig. 9 and Fig. 10 we show examples of spectra extracted from the Spirit and
 591 Opportunity images. The I/F values in the uncorrected, and corrected R7 filter are rep-
 592 resented by unfilled and filled symbols, respectively. We observe that the correction makes
 593 the darkest shadows darker, and the brightest regions brighter, as expected, and creates
 594 much flatter and unsurprising spectra, removing dramatic upturns and downturns in R7.
 595 As expected, the areas that do not exhibit high contrast do not change much.

606 In Fig. 11 we show ratio images for the Spirit rover (left column) and the Oppor-
 607 tunity rover (right column). Row 1) shows the ratio between the shortest wavelength fil-
 608 ter R2 image at $\lambda_{eff} = 754$ nm and the R6 filter image at $\lambda_{eff} = 934$ nm, Row 2) shows
 609 the ratio between the R2 filter image and the uncorrected R7 filter image, Row 3) shows
 610 the ratio between the R2 filter image and the corrected R7 filter image. The R2/R7 im-
 611 age is dramatically different from the R2/R6 image, almost with bright and dark regions
 612 inverted. The R2/R7 image shows a fair representation of the terrain for the same rea-
 613 son that the ratio image of the corrected R7 divided by uncorrected R7 did (Fig. 8, bot-
 614 tom row), namely that the R7 image is significantly less sharp and has lower contrast
 615 than the other images. The R2/R7 corrected is almost identical to the R2/R6 image and
 616 dominated by more subtle wavelength-dependent differences.

623 **3.2.3 Performance of the Algorithm on in-flight Calibration Target Im-** 624 **ages (Spirit and Opportunity)**

625 As described in Section 2.5.1 above, the panoramic camera cal target is routinely
 626 imaged in all filters of the camera system and these images are used in reflectance-calibration
 627 of images of the martian terrain. Here we show how the radiance offset in linear fits on
 628 the calibration target radiance-reflectance plots demonstrate the artifact in the R7 fil-
 629 ter on the close to dust-free calibration targets of the early mission. We also show how
 630 the correction improves the situation significantly.

631 Figure 12 demonstrates the approach using a typical calibration target image in
 632 the R7 filter acquired on sol 3 of Spirit's mission. 10 regions of interest (ROI's) were marked
 633 by hand on the calibration target image including the 4 corner color chips and 3 grayscale
 634 rings plus shadowed regions of the three grayscale rings. Mean radiance values were ex-
 635 tracted from each ROI and plotted against reflectance factor values known from pre-flight
 636 optical characterization of the cal target. The reflectance factor, or R^* , is defined as: $R^* = (I/F) / \cos(i)$, where i is the angle of incidence. For this image (ID: 2P126802681ESF0200P2110R7M1.IMG)
 637 the recorded offset is 225 DN before correction, 51 DN after correction. This shows that
 638 the correction, which has been implemented, is to a very high degree able to reduce the
 639 anomalous effect.
 640

652 Fig. 13 shows all the observed offsets in all filters for all calibration target obser-
 653 vations for the first 5 sols of the mission on both rovers. The offsets are back-converted

654 from radiance to data numbers (DN's) as described in Section 2.5.1 for better compar-
655 ison between filters.

656 The figure shows very clearly on both rovers that offsets in almost all other filters
657 scatter close to zero (mean offset within 60 DN of zero in all cases) but the offsets in the
658 R7 filter are very clearly above zero with mean values of 107 DN on Opportunity and
659 220 DN on Spirit. In particular the R7 data are obviously different from the other right-
660 eye longwave filters (within 30 DN of zero for all filters > 700 nm).

661 In red are shown the same analysis performed on corrected R7-images. On Oppor-
662 tunity the correction brings the mean offset from 107 DN to -29 DN, bringing the R7-
663 offsets in line with the observations in the other longwave filters. On Spirit the correc-
664 tion brings the mean offset down from 220 DN to 54 DN. This value is still higher than
665 all other filters, but much less dramatically than before. In Section 4.2.3 we further dis-
666 cuss the observed difference between the two rovers.

677 4 Discussion

678 In the following section we discuss the results from development of the correction,
679 as well as the results from the implementation of the correction on in-flight images from
680 the Spirit and Opportunity rovers.

681 4.1 Simulating the Effect - the Accuracy of the Model

682 Of the 4 parameter values used in the model, two of them, A and D, are scaling
683 parameters that are hard to directly compare to independent information. The two oth-
684 ers, however, B and C, have very clear physical interpretations. C, as already described,
685 is the bulk thickness of the CCD, and was fixed at 33, equivalent to 396 microns. This
686 value was based on the 400 micron minimum thickness reported by the manufacturer (rounded
687 to the nearest whole number of pixels). Thus, the fact that the model fits the data us-
688 ing this specific, physically reasonable value for C, may be taken as further evidence in
689 favor of our hypothesis. The physical interpretation of the parameter B in the model is
690 the minimum absorption coefficient in silicon in the wavelength range covered by the model.
691 Converting the parameter B to a penetration depth, for easier interpretation, gives us
692 a value of $1/0.0388 \text{ pixels}^{-1} * 12 \text{ microns pixel}^{-1} = 309 \text{ microns}$, where 12 microns is
693 the physical size of a pixel [*Bell et al.*, 2003]. Consulting table 3.1 on page 172 of ([*Janesick*
694 *et al.*, 2001]) (see fig 4) we find listed values for the penetration depth in silicon of 250
695 microns at 1060 nm, 370 microns at 1080 nm and 582 microns at 1100 nm. Thus, we would
696 perhaps have expected a somewhat higher value than we get. However, 1) the tabulated
697 values are for pure silicon, whereas we don't have access to measurements of the exact
698 penetration depths in the bulk material of this specific (doped-silicon) CCD chip, and
699 2) our model neglects loss processes (impurities, thermal loss, etc) that do not lead to
700 the generation of signal in the CCD and these processes may well dominate close to 1100
701 nm. Thus the fit value we get probably reflects the highest penetration depth that con-
702 tributes significantly to the signal, and this may not be all the way at 1100 nm. Given
703 these uncertainties, we consider the fit value of 309 microns to be fully consistent with
704 our hypothesis.

705 4.2 Corrected Images

706 4.2.1 Pre-flight Images

707 The correction of the stray-light image from Spirit was done as a first check of the
708 iterative method. The effect is expected to be most pronounced in areas of high contrast,
709 and since the stray-light images are indeed high in contrast, correcting these images was
710 a good first test. Furthermore, in correcting these images we were able to directly com-
711 pare the output to a "true" value, since we were able to infer how the light should be

712 distributed before it was affected by the R7 filter effect from a template profile calcu-
 713 lated using the R2-R6 images. By comparing directly to this template, we could ensure
 714 that no inherent issues were present that would introduce other artifacts in the image.
 715 As discussed in Section 3.1 we made sure that no artifacts were created due to the cor-
 716 rection by first correcting a simple stray-light image of which we had a good represen-
 717 tation of the "true" light distribution in the form of the template profile. Using this, we
 718 found the correction to be efficient as the result converged after only four iterations. A
 719 test-value cut-off of 10^{-14} was chosen based on the initial correction, and ensures that
 720 this very simplistic image, which we would expect to be most simple to correct, sets a
 721 strict limit for an acceptable correction of the more complex geology images. It was im-
 722 portant to choose a cut-off value, since iterating a large number of times will eventually
 723 result in a too noisy image.

724 **4.2.2 In-flight Images**

725 From the correction of in-flight images, we saw an improvement in contrast in spec-
 726 tra extracted from both Spirit and Opportunity collected images. Generally the corrected
 727 spectral values do not exhibit the dramatic changes from the R6 filter at $\lambda_{eff} = 934$
 728 nm to the R7 filter at $\lambda_{eff} = 1009$ nm that were observed for the uncorrected images.
 729 We observe an increase in contrast, as expected. When taking the ratio of the corrected
 730 R7 image divided by the uncorrected R7 image, the ratio ranges from 0.771 - 1.06 for
 731 the Spirit example. For Opportunity the ratio ranges from 0.736-1.09. In other words,
 732 in these specific examples, some (dark) pixels get $\sim 25\%$ darker after correction and some
 733 (bright) pixels get 5-10 % brighter after correction. Inverting these numbers, this indi-
 734 cates that the R7 artifact inflates the recorded radiances from some dark shadows by up
 735 to $\sim 33\%$.

736 When examining the images directly, we observed that taking the ratio between
 737 different infrared filters, there is a noticeable difference when including the uncorrected
 738 R7 filter image. We interpret this difference as due to the instrumental artifact in that
 739 the distribution of light is not expected to change drastically from the R6 filter to the
 740 R7 filter, especially not when the images obtained with the filters R2-R6 are so similar.
 741 The observed difference, that is the instrumental artifact, was efficiently removed after
 742 correction, resulting in an R7 image much more similar to the R2-R6 images.

743 **4.2.3 Cal Target Radiance-Reflectance Offsets**

744 The Pancam external calibration targets contain a number of surfaces that span
 745 relatively few pixels in Pancam images but have dramatically different reflectance prop-
 746 erties in close proximity to each other. As such we would expect images of these targets
 747 to be exactly the type of images affected by the R7-artifact as photons from high-reflectance
 748 cal target areas pass through the CCD, scatter from the backside and are redistributed
 749 to nearby dark regions. The net effect would be to reduce observed radiances from bright
 750 regions and increase observed radiances from dark regions. This is what was observed
 751 in pre-flight cal target images (Fig. 20 in [Bell et al., 2003]) and at the time hypothe-
 752 sized to be due to partial transparency of cal target materials at 1009 nm. However, as
 753 we have shown, the R7 effect is present also in images of the Martian surface, proving
 754 that the effect is an instrumental effect unrelated to the cal target. When plotting ob-
 755 served radiances against known reflectances, the low-reflectance points will be offset up-
 756 wards in radiance while the high-reflectance points will be offset downwards (See Fig.
 757 12). Thus, the straight-line radiance-reflectance fit will intercept the y-axis above zero,
 758 as observed in Figures 12 and 13.

759 Fig. 13 demonstrates that the correction performs very satisfyingly on the Oppor-
 760 tunity early-mission calibration target observations, so that after correction no signif-
 761 icant difference in offset is observed between R7 and the other infrared filters. On Spirit,
 762 before correction, the observed offsets are more than double what they are on Oppor-

763 tunity (220 vs 107) and the correction appears not fully successful in removing the off-
764 set - although the situation is much improved after correction. The observed 54 DN off-
765 set on Spirit represents less than 1.5% of the full well level (4095) and we point out that
766 there is some non-vanishing systematic uncertainty inherent in the calibration target re-
767 flectance model - particularly in the longest and shortest wavelength filters that have wider
768 bands and lower S/N. This can also be seen from the filter-to-filter variation in mean off-
769 sets in Fig. 13. It is of course also quite possible that small differences in the filter, op-
770 tics and particularly the CCD chip between Spirit and Opportunity make the R7-effect
771 have different strength between the two rovers. Given that the correction algorithm was
772 developed based on Spirit pre-flight images one would a priori expect it to perform bet-
773 ter on Spirit data, though.

774 5 Conclusions

775 In this section we will summarize the results presented in the preceding sections
776 and discuss implications for researchers working with archived Pancam data and our plans
777 for future work.

778 5.1 Summary

779 We have demonstrated the existence of an instrumental artifact in pre-flight and
780 in-flight Pancam images taken through the R7 (1009 nm) narrow-band filter and shown
781 that the effect is an inherent instrumental effect, rather than an effect related to prop-
782 erties of cal target materials. We have described the effect, which manifests itself as ef-
783 fectively a very long, very faint tail in the point-spread function in images acquired through
784 this filter. We have proposed a physical explanation, namely that the effect arises from
785 photons close to the silicon band gap that penetrate the CCD chip, scatter from the back-
786 side, penetrate the CCD again, and are registered on the way back in a different pixel
787 from the one through which they originally entered the CCD. We have presented an al-
788 gorithm for the correction of such images in order to (largely) remove the artifact and
789 have presented preliminary verification that this correction algorithm 1) largely does re-
790 move the artifact and 2) does not introduce other artifacts.

791 5.2 Implications

792 The artifact in R7 images is in most cases fairly weak, but can be significant in cer-
793 tain observations, particularly at sharp edges with high contrast. In some dark shadows
794 we observed the signal to be artificially inflated by up to $\sim 33\%$ and analysis of early-
795 mission cal target images indicated that the reduced contrast due to the artifact is equiv-
796 alent to >100 DN (full well = 4095 DN) for a hypothetical perfectly dark pixel. Accord-
797 ingly, it is possible that scientific conclusions relying on such high-contrast R7 images
798 may be affected. For scientists currently working on archived or newly-acquired Pancam
799 R7 data we recommend caution in drawing any firm conclusions. Additionally, previously
800 published analyses based on R7 filter images are in need of re-evaluation, in order to make
801 sure that the conclusions drawn are left unchanged by the correction.

802 5.3 Future work

803 For further verification that our correction algorithm is sound, we are working to
804 implement the correction on a larger dataset, particularly on a number of specific R7 ob-
805 servations that have previously been used as the basis of analysis in published work. In
806 section 1 we briefly mentioned some examples of such work. Correction of the data used
807 for those and other publications will give us a better understanding of the implication
808 of this instrumental artifact on higher-order conclusions derived from Pancam observa-
809 tions. The results from this investigation will make a natural follow-up publication to

810 the present work. It is our hope and belief that this work will demonstrate beyond doubt
 811 that the correction algorithm is stable, useful and can be sensibly implemented on the
 812 full dataset of acquired R7 images. Provided that the correction algorithm can be thus
 813 fully verified we will implement it in a recalibration of the full R7 dataset for archiving
 814 of corrected R7 images.

815 **Acknowledgments**


816 This work was funded by the Danish Council for Independent Research in the Natural
 817 Sciences, grant 4002-00292 and by the Carlsberg Foundation, grant CF16-0981 and by
 818 the NASA Mars Exploration Program and Mars Data Analysis Program.

819 The data forming the foundation of this work is the radiance- and reflectance-calibrated
 820 images from the Pancams on the two Mars Exploration Rovers. These images are avail-
 821 able through NASA's Planetary Data System at the geosciences node hosted by Wash-
 822 ington University in St. Louis (<http://pds-geosciences.wustl.edu/>). Preflight calibration
 823 images and derived data products, such as corrected images as well as code used in the
 824 analysis are available by direct communication with the first author (simonejj@nbi.ku.dk).
 825 We would be happy to work with any interested parties to provide corrected R7 images
 826 for analysis. When the algorithm is fully verified, work will be put into archiving cor-
 827 rected R7 images.



828 **References**

- 829 Squyres, S. W. (2004). The Spirit Rover's Athena Science Investigation at Gusev
 830 Crater, Mars. *Science* 305, 1025-1035.
- 831 Squyres, S. W., et al. (2004). The Opportunity Rover's Athena Science Investigation
 832 at Meridiani Planum, Mars. *Science*, 306, 1698-1703.
- 833 Arvidson, R. E., et al. (2006). Overview of the Spirit Mars Exploration Rover Mis-
 834 sion to Gusev Crater: Landing site to Backstay Rock in the Columbia Hills. *JGR*
 835 111, E2.
- 836 Squyres, S. W., et al. (2006). Rocks of the Columbia Hills. *JGR* 111, E2.
- 837 Arvidson, R. E., et al. (2008). Spirit Mars Rover Mission to the Columbia Hills, Gu-
 838 sev Crater: Mission overview and selected results from the Cumberland Ridge to
 839 Home Plate. *JGR* 113, E12.
- 840 Bell III, J. F., et al. (2003). Mars Exploration Rover Athena Panoramic Camera
 841 (Pancam) investigation. *JGR* 108, E12.
- 842 Bell III, J. F., et al. (2004). Pancam Multispectral Imaging Results from the Spirit
 843 Rover at Gusev Crater. *Science* 305, 800-806.
- 844 Bell III, J. F., et al. (2004). Pancam Multispectral Imaging Results from the Oppor-
 845 tunity Rover at Meridiani Planum. *Science* 306, 1703-1709.
- 846 Bell III, J. F., et al. (2006). In-flight calibration and performance of the Mars Explo-
 847 ration Rover Panoramic Camera (Pancam) instruments. *JGR* 111, E2.
- 848 Bell III, J. F., et al. (2017). The Mars Science Laboratory Curiosity rover Mastcam
 849 instruments: Preflight and in-flight calibration, validation, and data archiving.
 850 *ESS* 4, 7.
- 851 Kinch, M. K., et al. (2015). Dust deposition on the decks of the Mars Exploration
 852 Rovers: 10years of dust dynamics on the Panoramic Camera calibration targets.
 853 *ESS* 2, 5.
- 854 Rice, M. S., et al. (2010). Silica-rich deposits and hydrated minerals at Gusev
 855 Crater, Mars: Vis-NIR spectral characterization and regional mapping. *Icarus*
 856 205, 2, 375-395.
- 857 Wang, A., et al. (2008). Light-toned salty soils and coexisting Si-rich species discov-
 858 ered by the Mars Exploration Rover Spirit in Columbia Hills. *JGR* 113, E12S40



- 859 Markwardt, C. B. (2008). Non-linear least squares fitting in IDL with MPFIT, *As-*
860 *tronomical Data Analysis Software and Systems XVIII*, 411. edited by D. Bohlen-
861 der, P. Dowler, and D. Durand, pp. 251-254, Astronomical Society of the Pacific,
862 Quebec, Canada.
- 863 Janesick, J. R. (2001). Scientific Charge-coupled Devices. *SPIE - The International*
864 *Society for Optical Engineering*.




563 **Figure 8.** An example of how the correction can be used to improve the contrast in
 564 an R7 image. The left column shows the improvement to a selected Spirit R7 image. The
 565 image is a radiance calibrated image obtained by the Spirit rover on Sol 77. Image ID:
 566 2P133203792RAD2224P2572R7C5. From top, down: Uncorrected R7 image, corrected R7 im-
 567 age, the corrected image divided by the uncorrected one. The uncorrected and corrected images
 568 have been stretched to the same scale, which is the overall minimum and overall maximum of
 569 the two images. Min: $8.80 \text{ mW}/(\text{m}^2 * \text{nm} * \text{sr})$, Max: $34.2 \text{ mW}/(\text{m}^2 * \text{nm} * \text{sr})$. The ratio im-
 570 age was stretched according to its minimum and maximum value: Min: 0.771, Max: 1.06. Right
 571 column shows the improvement to a selected Opportunity R7 image. The image is a radiance cal-
 572 ibrated R7 image obtained by the Opportunity rover on Sol 954. From top, down: Uncorrected
 573 R7 image, corrected R7 image, the corrected image divided by the uncorrected one. Again, the
 574 uncorrected and corrected images have been stretched between the overall minimum and maxi-
 575 mum values of the images. Min: $0.0 \text{ mW}/(\text{m}^2 * \text{nm} * \text{sr})$, Max: $26.7 \text{ mW}/(\text{m}^2 * \text{nm} * \text{sr})$. The ratio
 576 image was again stretched to its minimum and maximum values. Min: 0.736, Max: 1.09. Image
 577 ID: 1P212872229RAD76EVP2586R7C1.



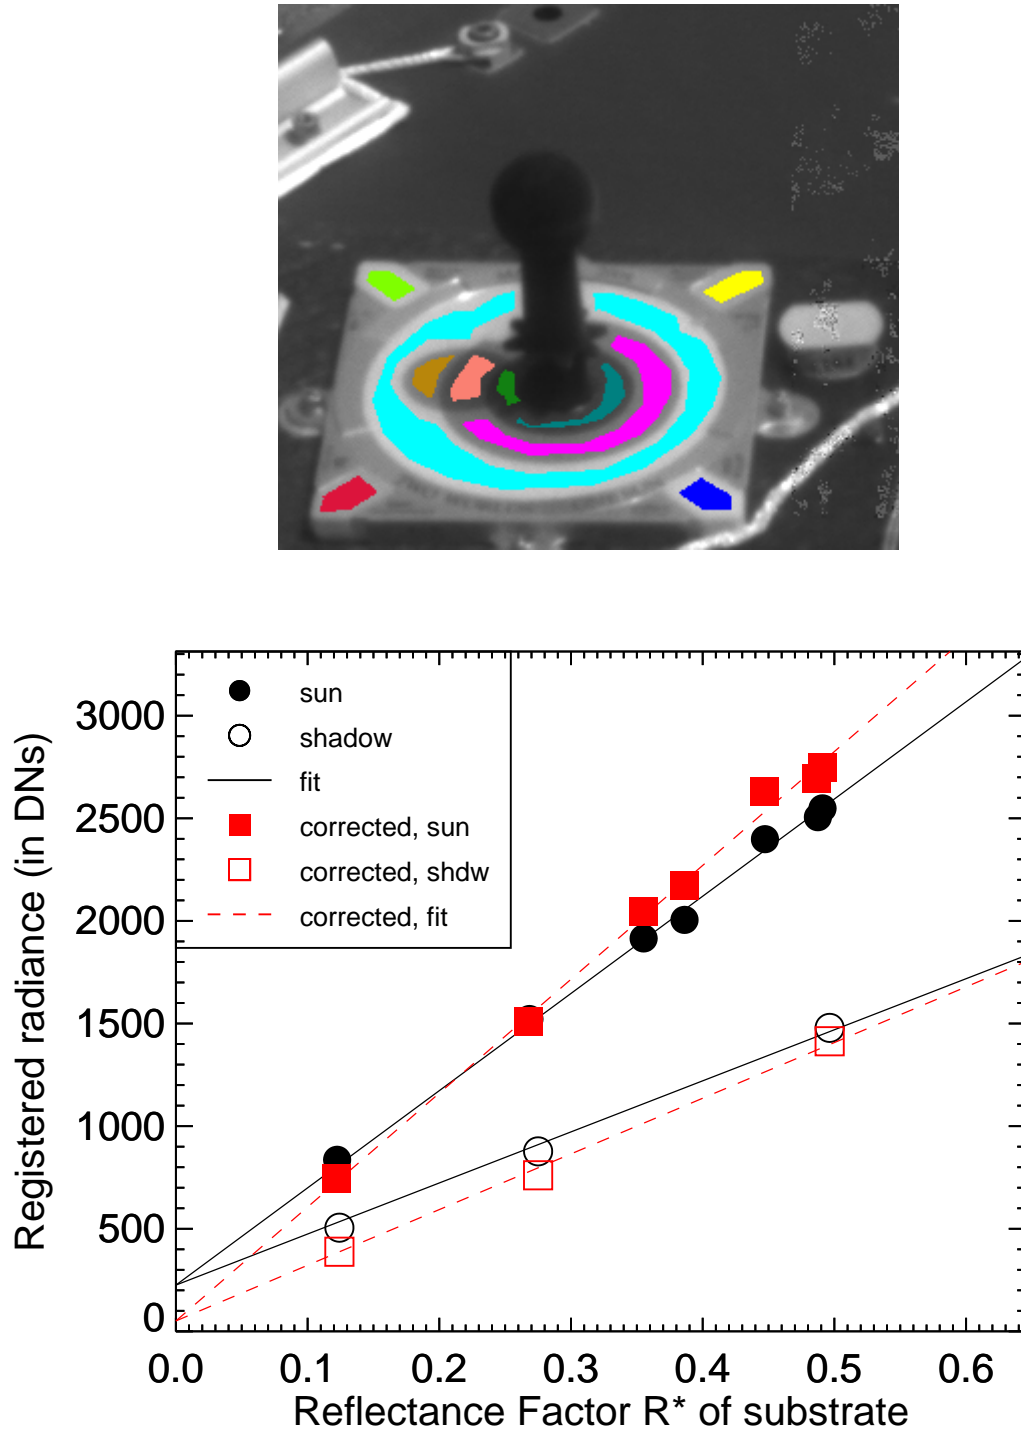
596 **Figure 9.** I/F spectra extracted from regions of interest marked on the image on top. For R7
 597 both uncorrected (open symbols) and corrected (filled symbols) values are shown. No scaling was
 598 performed between the right and left eye values. Values at overlapping wavelengths (L7R1 and
 599 L2R2) are presented as an average of the two eyes. The errorbars show the standard deviation
 600 between individual pixels within each region of interest.



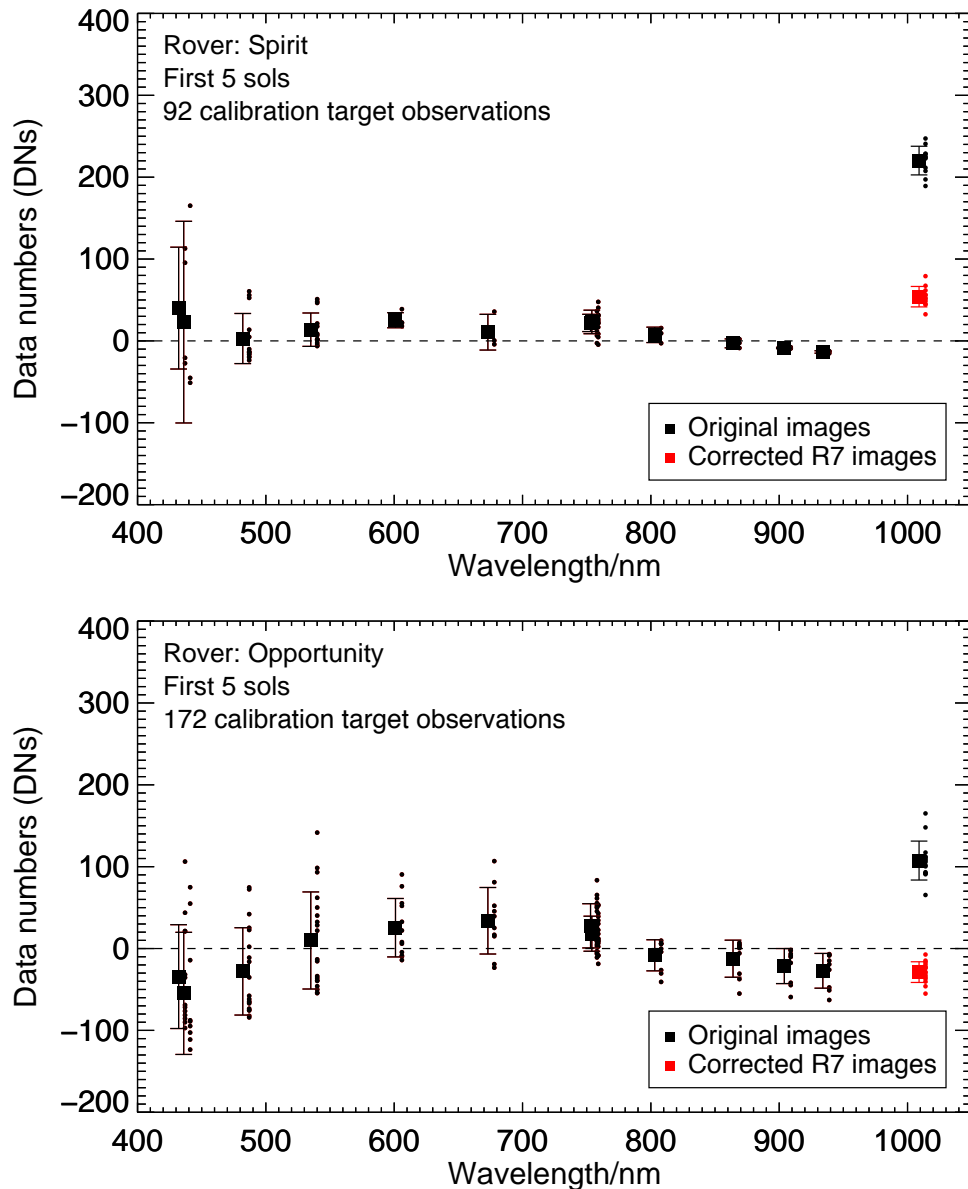
601 **Figure 10.** I/F spectra extracted from regions of interest marked on the image. For R7 both
 602 uncorrected (open symbols) and corrected (filled symbols) values are shown. No scaling was
 603 performed between the right and left eye values. Values at overlapping wavelengths (L7R1 and
 604 L2R2) are presented as an average of the two eyes. The errorbars show the standard deviation
 605 between individual pixels within each region of interest.



617 **Figure 11.** Left column: Spirit image. Right column: Opportunity image. From top to bot-
618 tom: R2 (753 nm) divided by R6 (934 nm), R2 divided by the uncorrected R7 (1009 nm), R2
619 divided by the corrected R7. All images are radiance calibrated. Spirit colour stretch (all three
620 images): Min: 1.06 , max: 1.70. Spirit sequence: P2572 of Spirit Sol 77. Opportunity colour
621 stretch (all three images): Min: 1.04 , max: 1.80. Opportunity sequence: P2586 of Opportunity
622 Sol 954.



641 **Figure 12.** Example calibration target image and radiance-reflectance plot with
 642 demonstration of determination of the radiance offset. This image has image ID:
 643 2P126802681ESF0200P2110R7M1.IMG acquired on Spirits sol 5 through the R7 (1009 nm)
 644 filter. In black circles (filled and open) are shown mean radiance values extracted from 10 regions
 645 on the uncorrected calibration target image. Filled circles are sunlit regions and open circles are
 646 shadowed regions. Solid black lines show the fit allowing two different slopes representing dif-
 647 ferent solar irradiances in sunlit and shadowed regions. In red squares (filled and open) and red
 648 dotted lines are shown data and fit for the corrected image. The radiances are back-converted
 649 from radiance to data numbers (DN's) by multiplying with the exposure time and dividing by
 650 standard pre-flight-determined conversion factor, i.e., even though given in DN, these are not
 651 raw image data, the images have been dark-subtracted, flatfielded etc.



667 **Figure 13.** Offsets in radiance-reflectance plots for all calibration targets during the first
 668 sols on Spirit (top) and Opportunity (bottom). As described in 2.5.1 the offsets are found by
 669 fitting two straight lines with a common intercept to the radiance-reflectance plot and recording
 670 this y-axis intercept. The offsets are back-converted from radiance to data numbers (DN's) by
 671 multiplying with the exposure time and dividing by standard pre-flight-determined conversion
 672 factors. This was done in order to make a better direct comparison of the offset value between fil-
 673 ters. Each small filled circle is an offset found from a single cal target image. For each filter, the
 674 mean offset is shown as a large filled square with error bars showing +/- one standard deviation.
 675 The black data points are original images. Red data points are R7 images corrected using our
 676 correction algorithm as described in section 2.4.