
**Grant agreement no. 607379**

SPA.2013.2.1-01 - Analysis of Mars Multi-Resolution Images using Auto-Coregistration, Data Mining and Crowd Source Techniques

- Collaborative project -

# D4.2
# Completion of CTX DTMs and ORIs

WP 4 – Global DTM/ORI production & validation

Due date of deliverable: month 36 – December 2016

Actual submission date: 06 / 03/ 2017* *(*) EC approval pending*

Start date of project: January 1st 2014 Duration: 39 months

Lead beneficiary for this deliverable: UCL

Last editor: Yu Tao, Jan-Peter Muller (UCL)

Contributors: UCL, UoS

| Project co-funded by the European Commission within the Seventh Framework Programme (2007-2013) | | |
|---|---|---|
| **Dissemination Level** | | |
| **PU** | Public | x |
| **PP** | Restricted to other programme participants (including the Commission Services) | |
| **RE** | Restricted to a group specified by the consortium (including the Commission Services) | |
| **CO** | Confidential, only for members of the consortium (including the Commission Services) | |

# History table

| Version | Date | Released by | Comments |
|---|---|---|---|
| 0.1 | 07.02.17 | UCL | Draft |
| 0.2 | 02.03.17 | UCL | Edited version with updated figures |

# Executive Summary

This report summarises the fully automated multi-resolution DTM processing chain, called CASP-GO, based on the open source NASA ASP software, tie-point based multi-resolution image co-registration, Gotcha sub-pixel refinement, and co-kriging method. The implemented system guarantees global geo-referencing compliance with respect to HRSC and MOLA, provides refined stereo matching completeness and accuracy from the ASP normalised cross-correlation. The development history and processing set-up of CASP-GO are also introduced.

Processing results are shown of the g;lobal CTX DTMs & ORIs in the last section along with an initial visual assessment of their quality.

# Table of contents

# 1. Introduction

## 1.1. Background and Context

Understanding the role of different planetary surface formation processes within our Solar System is one of the fundamental goals of planetary science research. There has been a revolution in planetary surface observations over the last 15 years, especially in 3D imaging of surface shape. This has led to the ability to be able to overlay different time epochs back to the mid 1970's, to examine time-varying changes, such as the recent discovery on Mars of mass (e.g. boulder) movement, tracking inter-year seasonal changes and looking for fresh craters.

To track these changes whether manually or automatically using data mining on the planet Mars, it is important to be able to be able to process data from different sensors and therefore address issues of processing large datasets with different image resolution, lighting conditions, coverage and locational accuracy. The goal of this work is to be able to maximize the exploitation of the available planetary datasets to enable improved understanding of the geology and geomorphology of the Martian surface through the generation of high quality Digital Terrain Models (DTM) and corresponding terrain-corrected images, OrthoRectified Images (ORI), using data from different NASA instruments and co-register these to a global set of accurately areo-referenced HRSC DTMs and ORIs to enable change detection.

## 1.2. Overview of Datasets

The global reference for our multi-resolution DTM production is the MOLA laser altimeter-derived data, which is considered to be the best global Mars 3D reference model to date. Individual MOLA tracks have been interpolated and extrapolated to yield a global MOLA DTM with a spatial resolution of up to 128 grid-points per degree (≈463m/pixel) with a vertical precision of 2-13m except for the polar region (down to 75º of latitude where the many repeat tracks allow DTMs of 256 and 512 grid-points/degree (≈112m/pixel) to be generated.

The HRSC on Mars Express comprises nine channels/looks that together in a single pass collect multi-angular and multi-colour images of the Martian surface, allowing stereo colour images to be produced from single orbit observations. DLR have generated along-track orbital strip DTMs (with a grid-spacing from 50-150m) and ORIs (up to 12.5m/pixel) on sinusoidal map projection system by processing the raw HRSC data using radiometric de-calibration, noise removal, image matching, geo-referencing and photogrammetric processing, and where these have been employed along-track Bundle Adjustment (BA), they are then labelled as "Level-4 Version 50+" when the products reach a satisfactory level of quality. The v50+ HRSC DTMs use the MOLA reference sphere with a radius of 3396.0 km. On iMars, the v50 HRSC ORI/DTMs have been employed for most of the sites as a reference base map for subsequent cascaded CTX/HiRISE DTM production. With increasingly dense HRSC coverage, DLR has

generated a large area mosaic over the US Geological Survey's map quadrant MC11 area (~100 stereo pairs), which covers two of the four proposed ExoMars 2020 landing site candidates. The MC11 HRSC DTM mosaics uses Equidistant Cylindical projection with a grid spacing of 50m for the DTM and colour mosaics, and 12.5m for the panchromatic image mosaic. DLR's MC11-E DTM and ORI are based on the same set of procedures for image filtering and rectification, least squares matching, strip BA, and calculation of 3D points applied for single strip data products, but additionally with bundle block adjustment, joint interpolation of multi-scale 3D point data sets, and photometric correction and image normalisation to an external brightness standard (Thermal Emission Spectrometer albedo). See D3.2 HRSC Preliminary Multi-orbit DTM.

The MRO CTX currently captures single panchromatic grey-scale images at ~6m/pixel over a swath-width of 30km. The CTX images are usually acquired at the same time as HiRISE so the stereo coverage is very limited also, albeit over a wider swath than HiRISE. UCL have processed CTX stereo pairs for MER-A, MER-B, MSL, Viking-1, Viking-2, MPF, and Phoenix to derive ORI (6m) and DTMs of 18m/grid point resolution using the ASP software in early 2013 within the EU FP-7 Planetary Robotics Vision Data Exploitation (PRoViDE) project. CTX ORI and DTM are essential for accurate HiRISE to HRSC co-registration. In late 2015, within the iMars project, CTX stereo pairs for MER-A, MER-B, MSL have been reprocessed at UCL using the optimised CASP-GO processing chain, which brings CTX ORI and DTM to a higher level of quality. In the experimentation stage, a total of 69 CTX stereo pairs over the MC11-E area were originally processed at UCL using the CASP-GO system.

On the other hand, the NASA MRO HiRISE camera is designed to acquire very detailed orbital images of Mars. HiRISE uses 14 CCDs including 10 red channels, 2 blue-green channels, and 2 NIR channels. The nominal maximum size of the red images is about 20,000 x 126,000 pixels and 4,000 x 126,000 pixels for the narrower B-G and NIR bands. To facilitate the mapping of landing sites, HiRISE produces stereo pairs of images from which the topography can be measured to an accuracy of 0.25 metres. By the time of writing up this report, UCL has started processing HiRISE DTMs for the test landing sites and MC11-E area and some areas of special scientific interest.

## 1.3. Pipeline Overview

In iMars, a fully automated multi-resolution DTM processing chain has been developed, called the Co-registration ASP-Gotcha Optimised (CASP-GO), based on the open source NASA Ames Stereo Pipeline (ASP) (Moratto et al., 2010) and (Broxton et al., 2008), tie-point based multi-resolution image co-registration (Sidiropoulos & Muller, 2015), and Gotcha (Shin & Muller, 2012) sub-pixel refinement method. The implemented system guarantees global geo-referencing compliance with respect to High Resolution Stereo Colour imaging (HRSC), and hence to the Mars Orbiter Laser Altimeter (MOLA), providing refined stereo matching completeness and accuracy from the ASP normalised cross-correlation.

### 1.4. Processing Overview

The CASP-GO processing chain has been tested/applied to stereo Mars Reconnaissance Orbiter (MRO) Context Camera (CTX) imagery (6m) over the Mars Exploration Rover (MER-A, B), Mars Science Laboratory (MSL) and a large area mosaic over the US Geological Survey's MC11-E/W area (~100 stereo pairs) and then applied to the production of planet-wide DTMs and ORIs from CTX (~1700+3800 stereo pairs) and MRO High Resolution Image Science Experiment (HiRISE) 25cm NASA images (~441 stereo pairs).

# 2. Method

### 2.1. USGS-ISIS and NASA Ames Stereo Pipeline workflow

The NASA Ames Stereo Pipeline (ASP) is a suite of automated geodesy and stereo-photogrammetry tools designed for processing planetary imagery captured from orbiting and landed robotic explorers on other planets or Earth. It was designed to process stereo imagery captured by NASA and commercial spacecraft and produce cartographic products including digital elevation models (DEMs), ortho-rectified imagery, and 3D models. The original Ames Stereo Pipeline (ASP) was developed by the Intelligent Robotics Group (IRG) of the Intelligent Systems Division at the NASA Ames Research Center. It builds on over ten years of IRG experience developing surface reconstruction tools for terrestrial robotic field tests and planetary exploration (*Moratto et al., 2010*). Figure 1 shows a schematic flowchart of the key processes involved.

ASP takes ISIS formatted "left" and "right" images as input and starts stereo processing by (a) pre-processing, including least squares bundle adjustment (BA), left-right image alignment, map projection that eliminates some of the perspective differences leaving only small perspective differences in the images, image normalisation to bring the two images into the same dynamic range, and image filtering to reduce noise and extract edges; (b) disparity map initialisation that uses a pyramid tiled, integer based correlation approach to find correspondences between pixels in the left image and pixels in the right image; (c) sub-pixel refinement to obtain sub-pixel correlation from their integer estimates; (d) triangulation that uses the geometric camera models stored in ISIS cub files to find the closest point of "intersection" of the two camera rays from disparity map and finally (e) DTM and ORI generation.

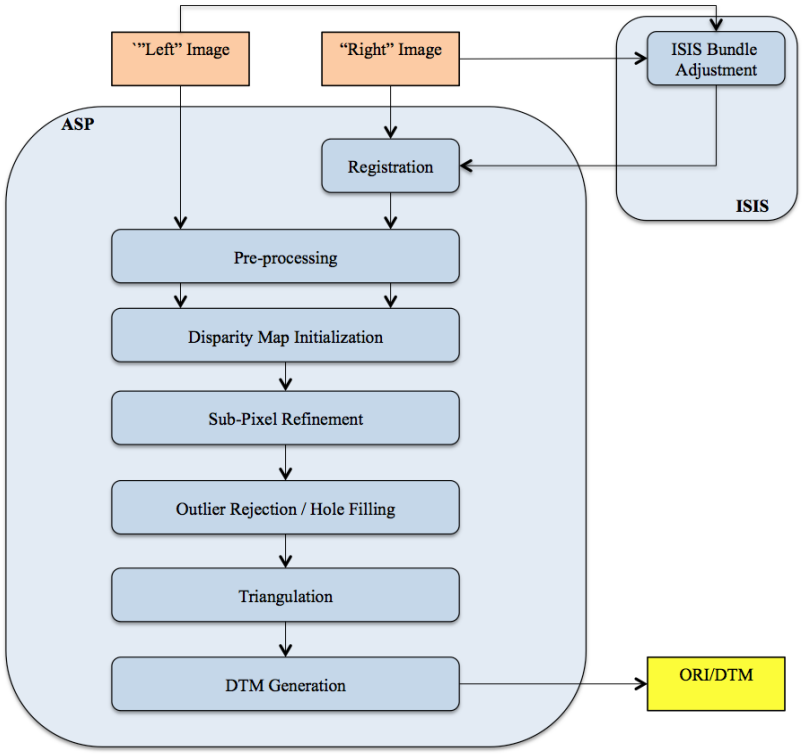


Figure 1. Flow diagram of the original ASP processing chain.

The original ASP pipeline was investigated in the early stages of the iMars project. A quality assessment of the processed results were made by comparing the output DTMs with those produced from the BAE Socet System (SS) and the University of Seoul (UoS) pipeline applied to the same regions on Mars. For this purpose input CTX and HiRISE stereo images of the three most observed sites on MER and MSL were chosen. The supporting (reference) DTMs and images were taken from the HRSC products overlapping with the CTX images. See iMars D1.2 Software for auto-high-res DTMs and its validation.

The averaged and standard deviation of the differences between the HRSC, ASP, SS, and UoS DTMs for CTX was +1.4 ± 84.2m, −2.1 ± 84.4m and −2.7 ± 84.9m, respectively. The same differences for the HiRISE instrument were −13.3 ± 19.7m, +4.2 ± 19.7m and +2.3 ± 37.2m. The large dispersion of the differences is due to a larger number of surface features over a larger area for CTX and a smaller number of features for a smaller area covered by the HiRISE instrument. The results were considered to require improvement in global consistency, completeness and robustness.

Generally speaking different DTMs from different pipelines have different kinds of artefact. For example, the OSU HiRISE DTM employed within the PRoViDE project has mis-matched outliers in some crater regions, the ASP processing chain has a quilting artefact caused by the initial integer base cross-correlation, failed matching areas for texture-less places, and there are known artefacts in the UoA/USGS products, such as "Boxes", CCD seams, faceted areas and manually interpolated areas. These artefacts are generally minor for GIS and visualisation after setting the DTM spacing (down-sampling) to a ratio of 3:1 from full pixel level resolution. However, for a detailed geological study of selected sites, a DTM at a higher level of quality and lower

level of artefacts is always more desirable. Therefore, we have taken advantages of the open source ASP pipeline and further extended/modified several key components to specifically address issues found from the experimental products, in order to develop a more optimal processing chain, called CASP-GO, to provide co-registered geo-spatial coordinates w.r.t HRSC (and MOLA) data, improved DTM completeness, reduced DTM artefacts, and improved DTM accuracy.

## 2.2. CASP-GO Workflow

The CASP-GO pipeline is shown in Figure 2. Apart from the ASP pre-processing, cross-correlation matching, triangulation, and DTM/ORI generation, five additional workflows are introduced to further improve the ASP results. These included (a) a fast Maximum likelihood sub-pixel refinement method to build a floating-point initial disparity map; (b) an outlier rejection and erosion scheme to define and eliminate mis-matches; (c) an ALSC and region growing (Gotcha) based refinement and densification method to refine the disparity value and match un-matched and/or mis-matched area; (d) co-kriging grid-point interpolation to generate the final DTM as well as height uncertainties for each DTM point; (e) ORI co-registration w.r.t. HRSC. Each of these function extensions and modifications (labelled below as the UCL pipeline) will be introduced in detail in this section.

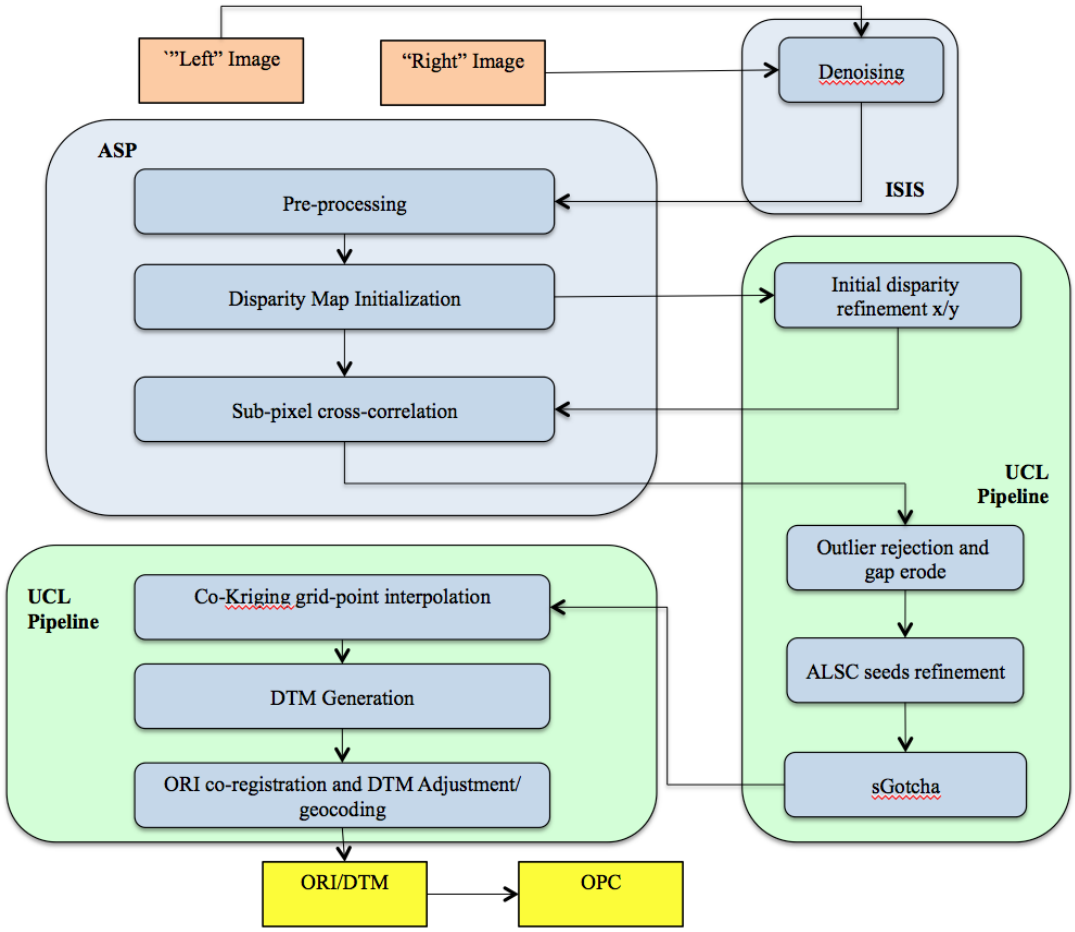


Figure 2. Flow diagram of UCL-Ames CASP-GO processing chain.

### 2.2.1. Initial Disparity Refinement

During our experiments on the CTX DTMs for MER-A, MER-B, and MSL, one of the artefacts we found is a "staircase" effect in the final DTM. This artefact can be observed clearly after a hill-shading process with an elevation of 30º and azimuth of 330º for the light source. In Figure 3, the quilting pattern repeats at the same resolution of the first level of pyramid, which equals to the resolution of the initial disparity map. We found that even though ASP Bayes EM weighted affine adapted correlation exhibits a high degree of immunity to image noise, refining the lower resolution integer disparity map generates severe artefacts resulting in this quilting pattern. Pixel locking effects appeared on both the faster ASP sub-pixel correlation solution, i.e. parabola fitting, and the slower sub-pixel solution, i.e. Bayes EM weighted affine correlation. This is because the sub-disparities tend toward their integer estimates and when using a lower resolution integer estimates, the algorithm doesn't guarantee continuity between sub-pixel disparities from different adjacent integer estimates. The "staircase" becomes most obvious in feature rich areas, e.g. crater edge, because the difference between two adjacent integer estimates is higher in these areas.

In order to reduce the "staircase" artefact, we can either generate integer disparity estimates at the same resolution of the final sub-pixel disparities, or generate float disparity estimates at lower resolution. We applied a fast Maximum Likelihood image matcher (*Olson, 2002)* to generate float disparity estimates on the same resolution as the ASP integer based correlation. The probabilistic formulation for lower resolution matching uses an arbitrary likelihood function for the matching error between edge or image features that eliminate the sharp distinction between matched and unmatched templates. In this approach, we search for a maximum likelihood estimation of template positions, i.e. the joint probability density function (PDF) for the distances. The joint density is modelled as the sum of the error density when an edge pixel is an inlier and a probability density of the distances when the edge is an outlier. (*Olson, 2002*) described a multi-resolution search strategy that examines a hierarchical cell decomposition of the space of possible template positions, which divides the space of template positions into rectilinear cells and determines which cells could contain a position satisfying the acceptance criterion recursively. In our case, given that we have already obtained integer distance estimation and to compute the probability density function, only the magnitude of distances is required, the searching strategy can be simplified to finding template positions in 8 directions with a threshold equal to half of the difference of integer disparity values in each directions.

The maximum-likelihood measure gains robustness by explicitly modeling the possibility of outliers and allowing matches against pixels that do not precisely overlap the template pixel from cross-correlation. It is a fast and robust approach to turn ASP integer disparity to float initial disparity maps at a lower resolution, which can be used to seed disparity in ASP sub-pixel cross-correlation, i.e. the 4$^{th}$ step in the CASP-GO system. The results are shown in the right hand image of Figure 3.

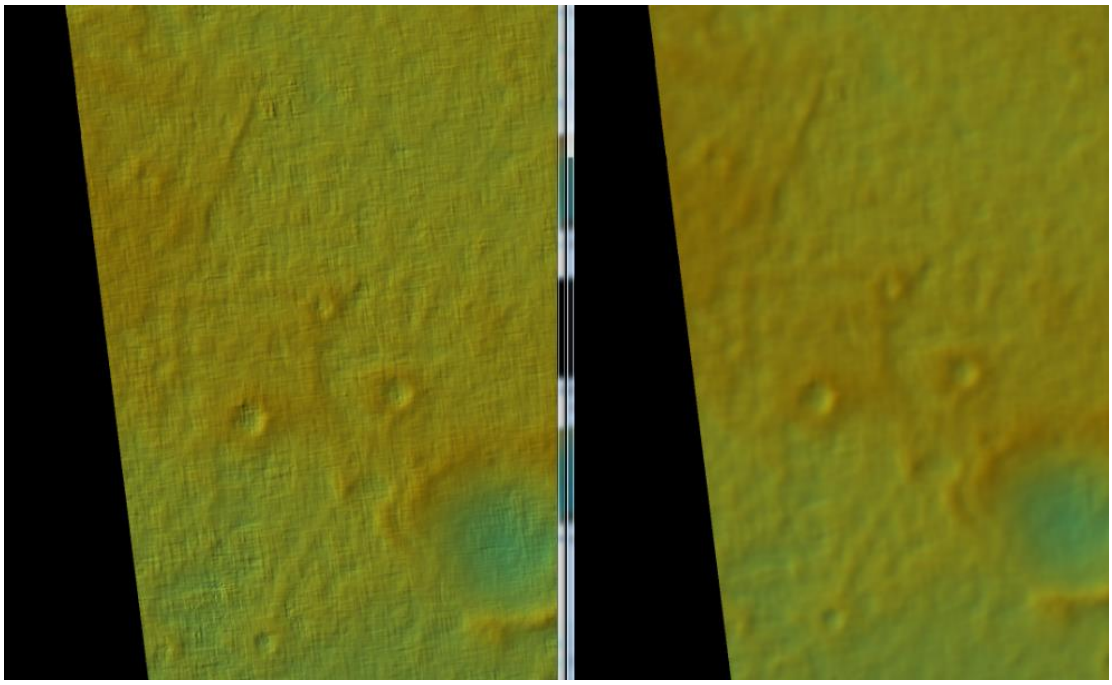

Figure 3. An example of the MER-A DTM, showing quilting artefact from ASP processing (left) and DTM after initial disparity refinement in CASP-GO (right).

### 2.2.2. Outlier Rejection and Erosion Scheme

The final sub-pixel disparity maps obtained from the refined initial disparity map still have several different problems for different data. For example, from our experiments on the CTX DTMs for MER-B and MSL, two obvious problems are unmatched areas (no disparity available, i.e. errors of omission) and mismatched areas (wrong disparity, i.e. errors of commission). The first issue can be triggered in various situations: (a) the regions of matching may have different lighting conditions/contrast/specular properties of the surface; (b) the regions of matching have very little texture or extremely low contrast such that there is insufficient signal to noise ratio and thence may be rejected by the correlation; (c) the regions of matching are highly distorted due to different image perspectives and the steep slopes of the surface, such as crater and canyon walls.

The ASP correlation process attempts to find a match for every pixel using a window that evaluates the lowest cost compared to all the other search locations. Generally, if we use a smaller search kernel and a smaller search range, there will be more unmatched areas, given more chances that the template in the left image cannot find any matched template in the limited search area in the right image. However, if we set the kernel larger, then more areas can be matched but at the expense of losing more sharp features, because the disparity tends to be too smooth. If we set the search range larger, chances are that the template in the left image can find a "matched" template far beyond the true position with a lower correlation cost because the true position meets one of the situations listed above. The first issue hence becomes the second issue, i.e. mis-matches. It is hard to eliminate all mis-matches even though we use a more accurate pre-allocated search range. However, this is not always possible, due to the native weakness of cross-correlation. Also we do not want too many unmatched area, in which case the follow-on densification and co-kriging steps will take too many computing hours.

Therefore, in CASP-GO we normally set a slightly larger correlation kernel, which will minimise un-matched area, but with a smaller search range, which will minimise mis-matches. Then we reject mis-matched areas, and erode on the border of mis-matched pixels and un-matched pixels, since generally the matching quality is much lower in the neighbourhood of "bad" disparities. This will leave "only" un-matched areas, which can be addressed from follow-on Gotcha densification and co-kriging interpolation, on ASP sub-pixel disparity output.

The outlier rejection schemes are: (a) a disparity value differs from a threshold by a percentage of pixels in a kernel; (b) a kernel with standard deviation higher than a threshold; (c) the difference of the mean value of a kernel and neighbouring kernel is higher than a threshold; (d) a kernel with a neighbouring kernel being rejected by a threshold percentage; (e) adjacent disparity values from (a), (b) and (c). Note that the outlier rejection schemes may also remove some disparities that are actually correct. However this can be easily and more precisely re-matched in the next step (Gotcha densification).

### 2.2.3. ALSC Refinement and Gotcha Densification

Based upon the "cleaned" sub-pixel disparity map, an Adaptive Least Squares Correlation (ALSC) refinement is performed on all the remaining disparity values iteratively. These refined disparity values are used as seed points for Gotcha (Grün-Otto-Chau) densification (*Shin & Muller, 2012*). The Gotcha matcher is based on ALSC and region growing. It is very robust and accurate, but very slow for large-scale image matching since it tries to match every point iteratively and re-sort all seed points according to a "best first" strategy when a new point is matched. However, given sufficient number of sub-pixel disparities that pass the outlier rejection schemes, small difficult regions can be matched with Gotcha accurately. For example, the geometrical distortion generated by different viewing angles can be addressed with Gotcha by modifying the shape of the ALSC window iteratively, albeit only a parallelepiped using an affine transformation is currently employed.

The Gotcha algorithm applied in this work can be summarised as (a) with given sub-pixel disparity values, retrieve seed tie-points (point correspondences) on the border (within 5-11 pixel width) according to the x and y translation (disparity); (b) run ALSC on seed tie-points and store similarity value; (c) sort seed tie-points by similarity value; (d) a new matching is derived from the adjacent neighbours of the initial tie-point with the highest similarity value; (e) if a new match is verified by ALSC then it is considered as seed tie-points for the next iteration of region-growing; (f) this region growing process repeats from (c) to (e) until there are no more acceptable matches; (g) retrieve final disparity map after densification.

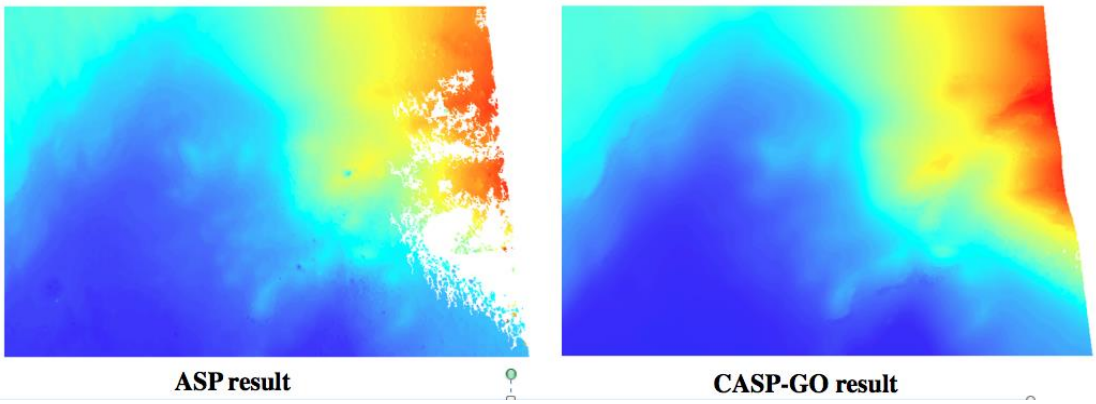


Figure 4. An example of MSL DTMs derived from CTX stereo, showing un-matched areas from ASP processing (left) and improved completeness using Gotcha densification in CASP-GO (right).

With Gotcha densification, we can achieve improved completeness for the final DTM without significantly smoothing out sharp features from a large matching kernel. Generally a larger ALSC window and higher maximum eigenvalue yields better completeness in disparity and hence DTM. The similarity values from the Gotcha matcher are then used together with co-kriging parameters to produce DTM uncertainty values. See Figure 4 for an example of the impact of CASP-GO on completeness.

### 2.2.4. Co-kriging Interpolation

After Gotcha densification, ideally we should get disparities for most of the pixels. However, for some textureless or extremely low contrast areas, where the matching is rejected by both cross-correlation and Gotcha, the final DTM may still have "holes" remaining. In this work, we use the co-kriging method to interpolate onto a gridded DTM. Kriging is a robust technique that uses a spatial model to bias the interpolation process. The basic idea of Kriging in DTM interpolation applications is to use a weighted average, which depends on both the distance of point pairs and spatial variation, of neighbouring known elevation values to predict a missing elevation value. A detailed description of the Kriging method can be found in (*Stein, 1999).*

In CASP-GO, a co-kriging method has been integrated using source code from the Geostatistics Template Library (GsTL)'s co-kriging implementation. Unfortunately, Co-kriging is a computationally expensive algorithm. It is impractical to repeatedly solve the Kriging equations using all observations available. Therefore, a fixed search radius was used for determining neighbouring points in CTX DTM interpolation. Also, the total size of the interpolation needs to be minimised from stereo matching and Gotcha densification steps. Nevertheless, co-kriging greatly improves the accuracy over linear facet interpolation of a Delauney triangulation of the data and provides an accuracy estimate for each point interpolated based on the estimated quality of the data, which can be indicated from the ALSC similarity, and the spatial variation of the terrain.

Each elevation output from stereo matching has associated with it various parameters, which may correlate with the elevation accuracy. These include the maximum eigenvalue of the variance-covariance matrix that reflects uncertainty in the positioning

of the matching window, a measure of local consistency of line disparities, and optionally parameters of radiometric gain and shift, as well as the so-called skewness from the camera model.
For this skewness, the two rays from the camera projection centre to a pixel location never intersect perfectly at a 3D point in practice. This is because any slight error in camera position or orientation information will affect the camera rays' positions. When taking the closest point of intersection of the two rays as the location of 3D from ASP disparity to point clouds triangulation, the actual distance between the camera rays at the point becomes an important elevation accuracy parameter. The perpendicular distance between the two rays at their closet intersection is also a weighted term of the final uncertainty value.

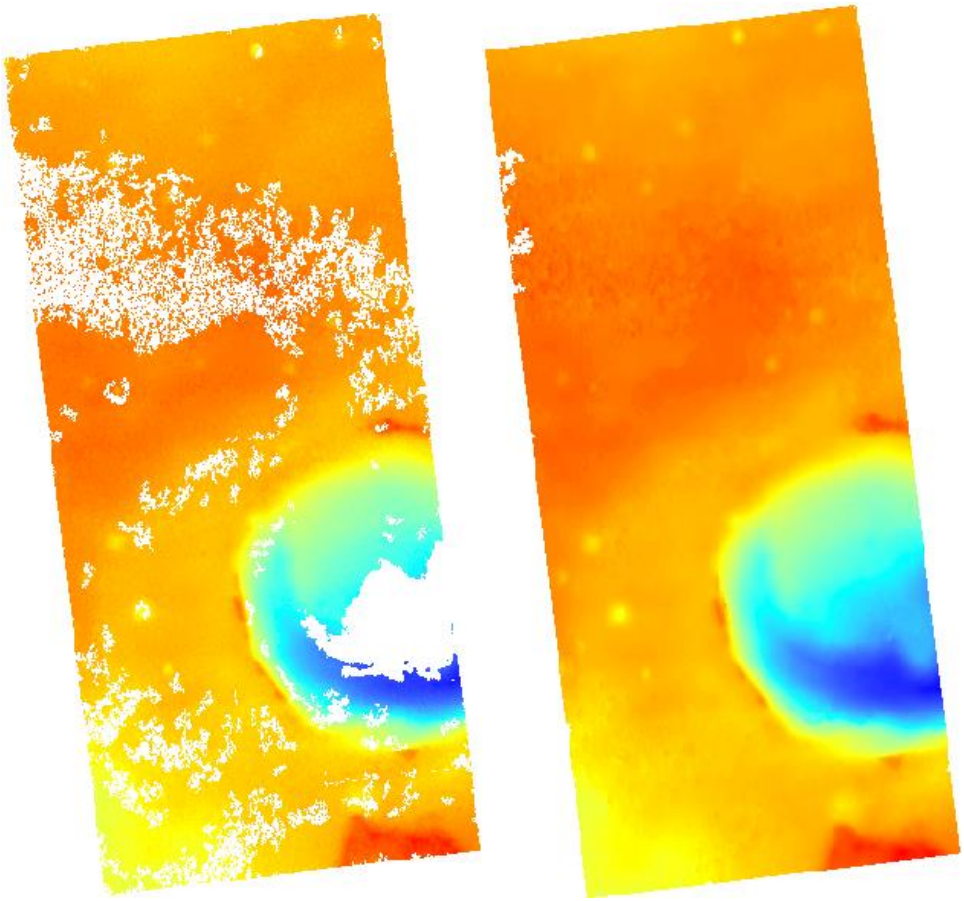

Figure 5. An example of MER-B DTM after bad matching rejection in CASP-GO (left) and final DTM using Gotcha densification and co-kriging.

### 2.2.5. ORI Co-registration and DTM Adjustment

The refined stereo matching workflow has brought the CTX (and coincidentally the HiRISE) DTM production to a new level of automation and accuracy. However, the HiRISE and CTX datasets are generally not co-registered with the HRSC ORI/DTM (DLR processed v50 products) and MOLA dataset. This was reported in a previous EU project – PRoVisG and follow-on HiRISE-CTX-HRSC co-registration work in (Tao, 2015). These mis-registrations are about 100-200m between HiRISE/CTX and HRSC

for the MER and MSL areas, according to manually selected control points on obvious landmark features, such as crater edges.

In this work, we have added a Mutual Shape Adapted Scale Invariant Feature Transform (MSA-SIFT) based co-registration workflow to the final ORI and DTM product. Technical details of MSA-SIFT and the evaluation on ORI/DTM co-registration results for MER and MSL are described in (Tao & Muller, 2015). In this ORI co-registration and DTM adjustment work, we take HRSC ORI as the reference image for CTX ORI co-registration and subsequent shift of the corresponding CTX DTM according to the CTX ORI to HRSC ORI transformation. For the HiRISE co-registration, we take co-registered CTX ORI as the reference image for HiRISE ORI co-registration and subsequent transformation of HiRISE DTM.

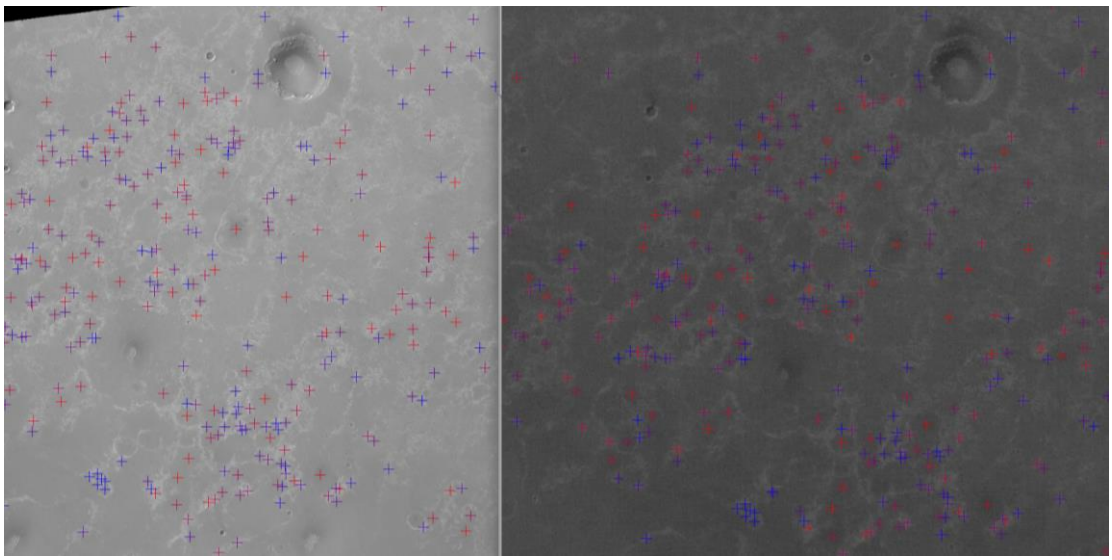

Figure 6. An example of MSA-SIFT tie-points from MER-B CTX ORI (left) and HRSC ORI (right), showing increasing uncertainty value from blue to red.

Bringing all CASP-GO products into an unique geo-spatial context with respect to HRSC and MOLA data is important for data exploitation such as visualisation, GIS and change detection, which are the main goals in iMars.

# 3. Development and Set-up

## 3.1. CASP-GO Development History

The CASP-GO system has been developed simultaneously with the test data validation performed by DLR (see D4.5 Validation) and web-GIS interface tests performed by FUB. Several key changes in the pipeline were executed according to quality assessment, processing time estimation, and web-GIS interfacing requirements made by different partners in iMars. A summary of the development history can be found in (Table 1). The CASP-GO outputs include a DTM, ORI, Gotcha “mask”, co-kriging “mask”, uncertainty map and hill-shaded colourised browser products, which can be directly ingested into the iMars web-GIS server.

*Table 1 Software development timetable.*

| Time | Software version | Development history |
|---|---|---|
| Early 2015 | v0.0 | USGS-ISIS+ASP based processing chain |
| Aug. 2015 | v0.5 | Integrated ACRO and Gotcha densification |
| Sep. 2015 | v1.0 | Added initial disparity refinement and outlier rejection scheme |
| Oct. 2015 | v1.2 | Integrated GsTL co-kriging interpolation |
| Nov. 2015 | v2.0 | Replaced Gotcha with sGotcha, moved ACRO to the last step, added USGS-ISIS denoising |
| Early 2016 | v2.2 | Replaced ACRO with CSA-SIFT, added batch processing wrapper, agreed with JR for OPC FTP processing |
| Jun. 2016 | v2.5 | New metadata designed to adapt to the ISIS PVL format |

## 3.2. Software Set-up and Processing Resources

At UCL-MSSL, we have mirrored the HRSC, CTX, HiRISE PDS data volumes from JPL in a local shared storage system in order to speed up the production process with an option such that if data is unreachable it can be read from the original source again. At UCL-MSSL, the developed software is installed in a shared directory, which is accessible from 14 Linux processing blades (10 with 16 cores and 48GB RAM; 4 with 24 cores and 96GB RAM). Jobs are controlled via a local desktop machine and distributed to the 14 processing blades with multiple sessions of multi-threaded processing. Processed results are stored in several 1TB RAID storage disk partitions and logged back to the local controlling desktop. Failed jobs can be examined through detailed log files and in the future will be reprocessed automatically with different processing parameters.

In the meantime, UCL has been successfully rewarded with free access to $20,000 computing resources from Microsoft® Azure Cloud for Research. UCL worked on the virtual machine set-up, software integration, and test processing at Microsoft® Azure cloud computing since early 06/2016 and started batch processing of 1540 CTX stereo pairs (the original published list of CTX stereo-pairs from the NASA Ames group published online in 2012) in 07/2016 which completed in mid 12/2016. UCL was recently awarded an additional $20,000 computing resources from Microsoft® Azure Cloud for Research in early 01/2017 for finishing up the rest of the ≈4,000 CTX stereo pairs, some of which overlap with the previous set. This new set was derived using a novel algorithm based on the one described by Sidiropoulos & Muller (2015)

# 4. Processing Results

## 4.1. Summary of 3 Experimental Rover Sites

The MER-A, MER-B, and MSL have been used as experimental sites during the CASP-GO development period. An example of improvement has been introduced in the section 2 Method. The analysis of these datasets is described in D4.5.

## 4.2. Processing Result Over MC11E and MC11W

Following on from the original experiments over the MER-A,B and MSL areas, CASP-GO was applied to all the CTX images which lay within a large mosaic over the U.S. Geological Survey's MC11-E/-W [Figure 7-9]. For MC11-E, 78 identified stereo pairs were processed. 22 (out of 78) stereo pairs were re-processed with different processing parameters to produce DTMs with better quality. 6 (out of 78, ≈8%) stereo pairs failed and were removed from the processing list because of bad quality input images. 42 (out of 78) non-repeat stereo pairs have been integrated to the iMars web-GIS website developed by FUB. Processing of CTX ORI and DTMs over the MC11-E area took ~1.5 weeks time on 10 Linux processing blades at UCL-MSSL excluding follow-on reprocessing and metadata format changing required by FUB for web-GIS integration. In late 2016, 39 CTX DTM/ORIs over MC11-W were processed using latest version of the CASP-GO pipeline. These are shown together with the MC11E+W DTM mosaics in Figure 7.

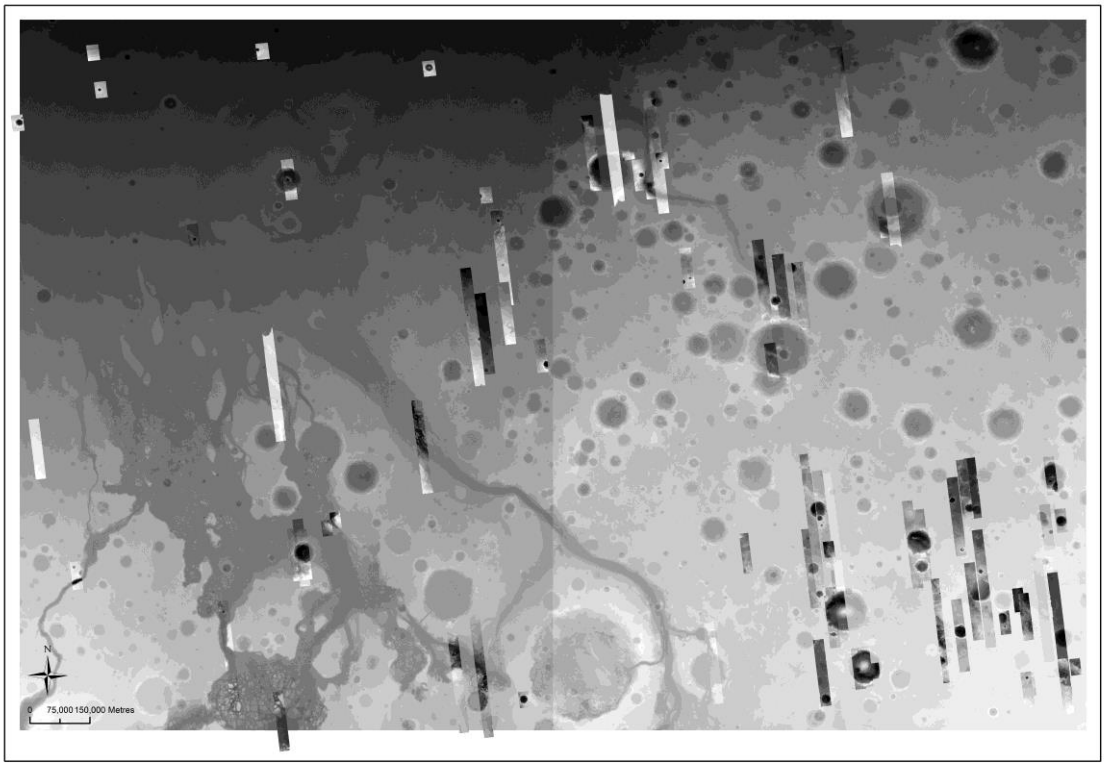


Figure 7. An example of processed CTX DTMs over MC11-E/-W area superimposed on the HRSC DTM mosaic basemap using QGIS.

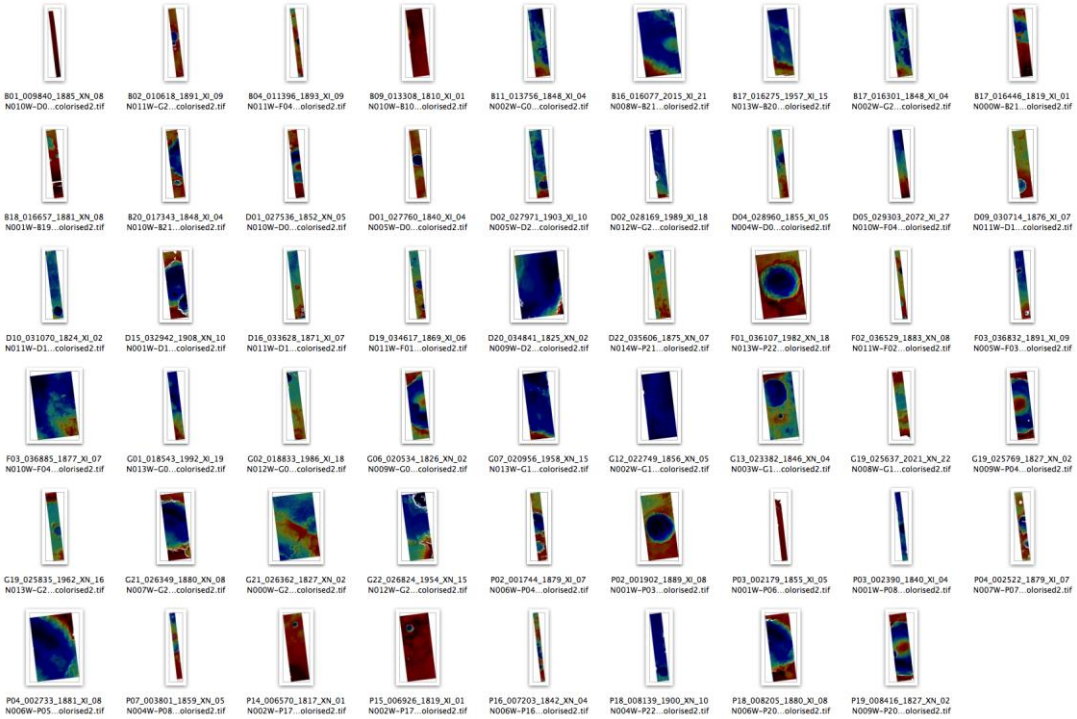

Figure 8. Examples of MC11-E CTX DTM browser products.

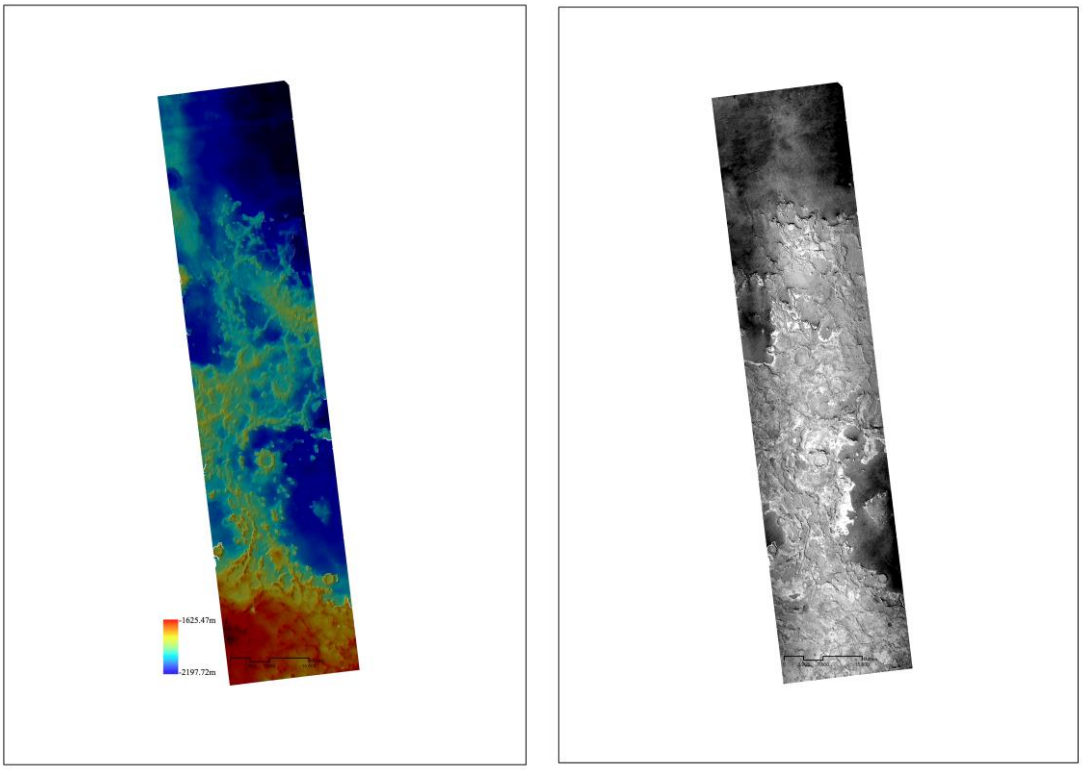


Figure 9. An example of one of the MC11-E CTX DTM and ORI products from CASP-GO pipeline (5th from the left on the top row of Figure 8).

## 4.3. Summary of Global CTX DTM Production

In mid 2016, a total number of 1540 CTX stereo pairs with a global coverage of Martian surface were started to be processed on the Microsoft® Azure® cloud computing resource for massive scale batch processing. In late February 2017, 3820 more up to date stereo pairs were defined (Sidiropoulos & Muller, 2015b) and set-up for batch processing on a very large number of VMs (150) on the Microsoft® Azure® Cloud computing. These DTMs will be ported into the iMars webGIS system as well as PDS4 archive later on.

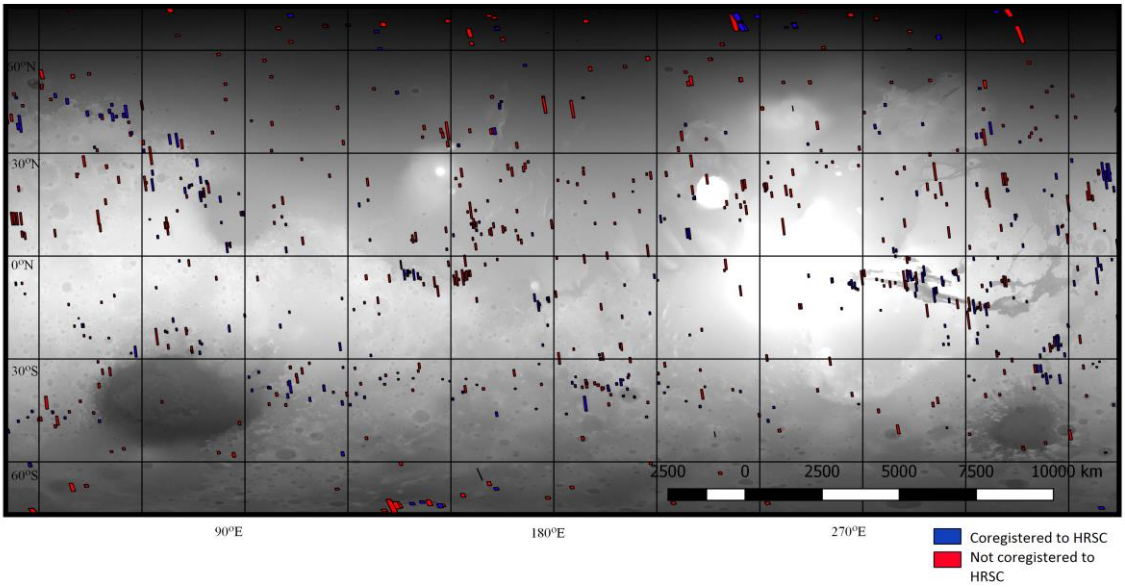


Figure 10. Location of all 1540 18m CTX DTMs processed by the end of 2016. Where no level-4 HRSC DTM exists, the CTX products are not co-registered to a global co-ordinate system.

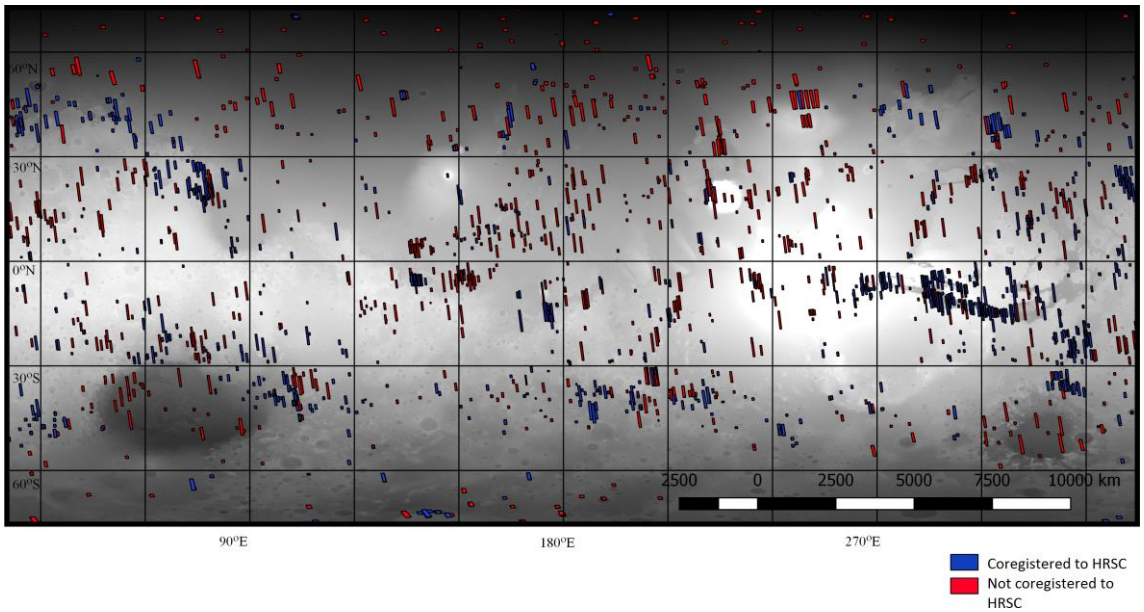


Figure 11. Locations of all the ≈4000 CTX stereo-pairs being currently processed on Microsoft Azure® including overlaps. It is expected that this larger dataset will be completed for the 1st stage processing by mid to end March 2017.

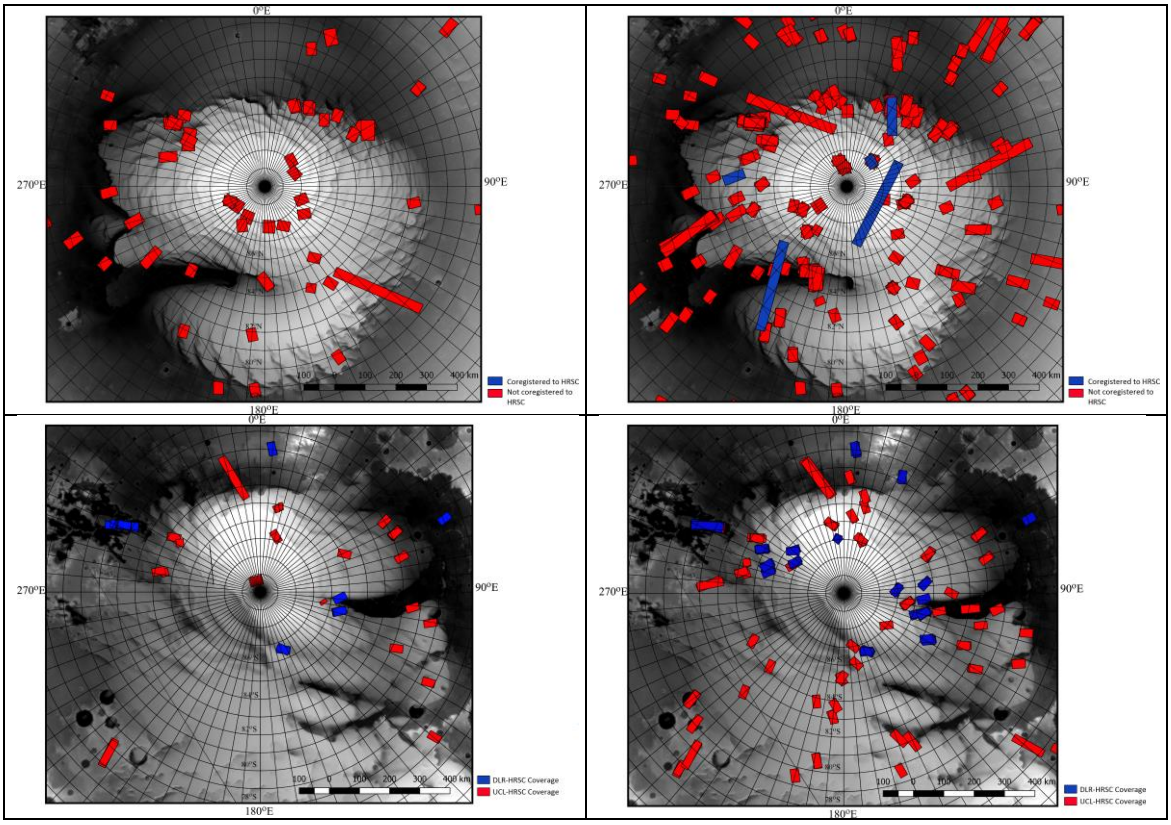


Figure 12. North Pole (upper row) and Sout Pole (lower row) overage of CTX in the original 1540 set (left column) and 4000 set (right column)

The overall statistics for CTX are 1915 pairs have HRSC out of all. Out of 1540 882 pairs don't have HRSC. Out of 3963 2491 pairs don't have HRSC. However, for the polar regions, there is now a complete set of level-4 data so for Figure 12 (SP) these figures are slightly misleading.

A large number of processed DTMs have been assessed using a 5 star rating scheme, i.e. 1 – failed, 2 – major problem, 3 – minor problem, 4 – good quality, 5 – very good quality. 503 out of 620 DTMs have been rated as 3+, i.e. 316 good and very good [Figure 13b-e], 187 has minor problem [Figure 13a]. 31 out of 620 failed all due to bad input images (5%). [Figure 13] shows examples of these failed input images. They are impossible for DTM processing.

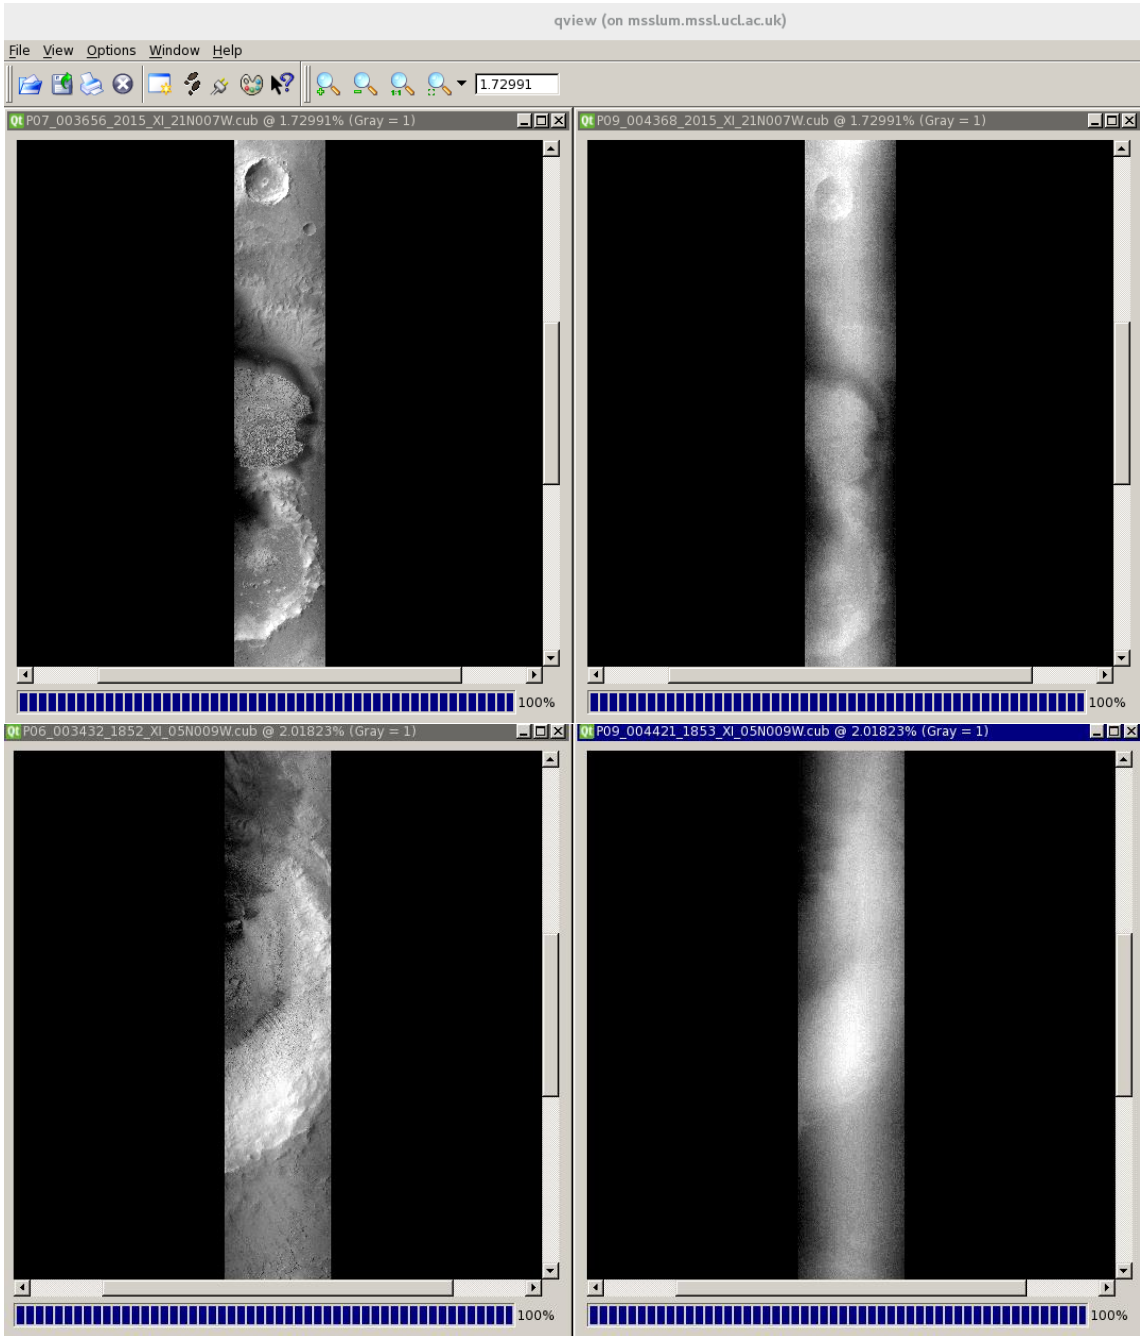
qview (on msslum.mssl.ucl.ac.uk)
File View Options Window Help
1.72991
P07_003656_2015_XI_21N007W.cub @ 1.72991% (Gray = 1)
P09_004368_2015_XI_21N007W.cub @ 1.72991% (Gray = 1)
P06_003432_1852_XI_05N009W.cub @ 2.01823% (Gray = 1)
P09_004421_1853_XI_05N009W.cub @ 2.01823% (Gray = 1)
100%
100%
100%
100%

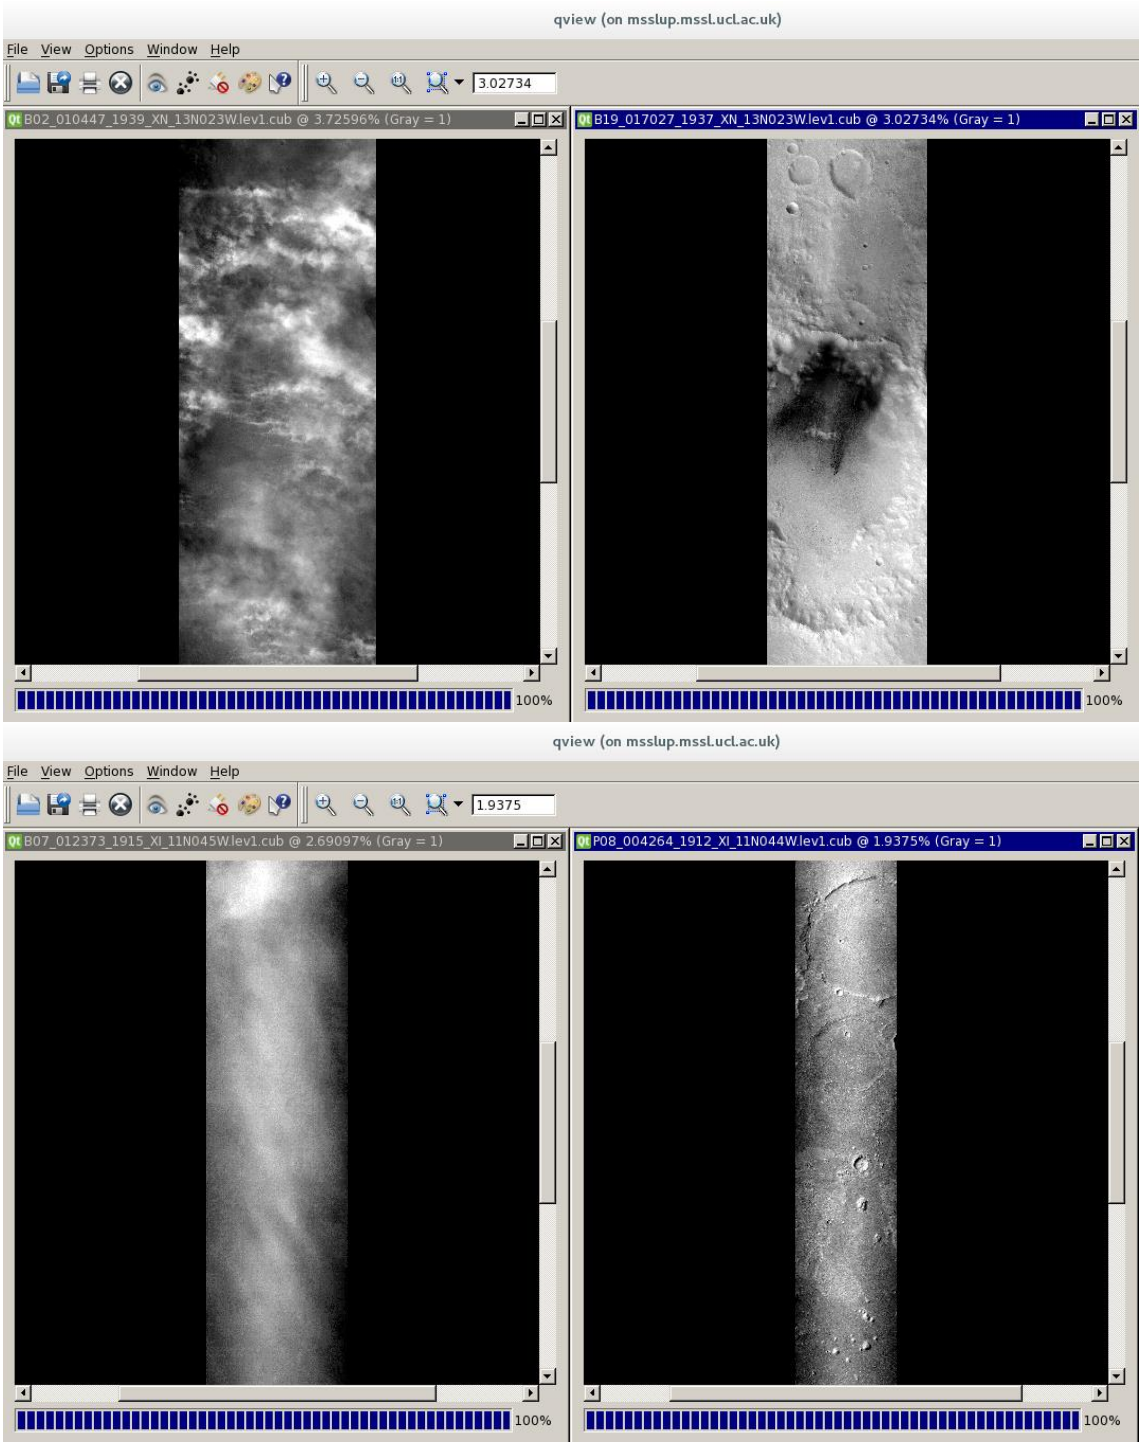


Figure 13. Examples of failed stereo pairs (DTM rating 1).

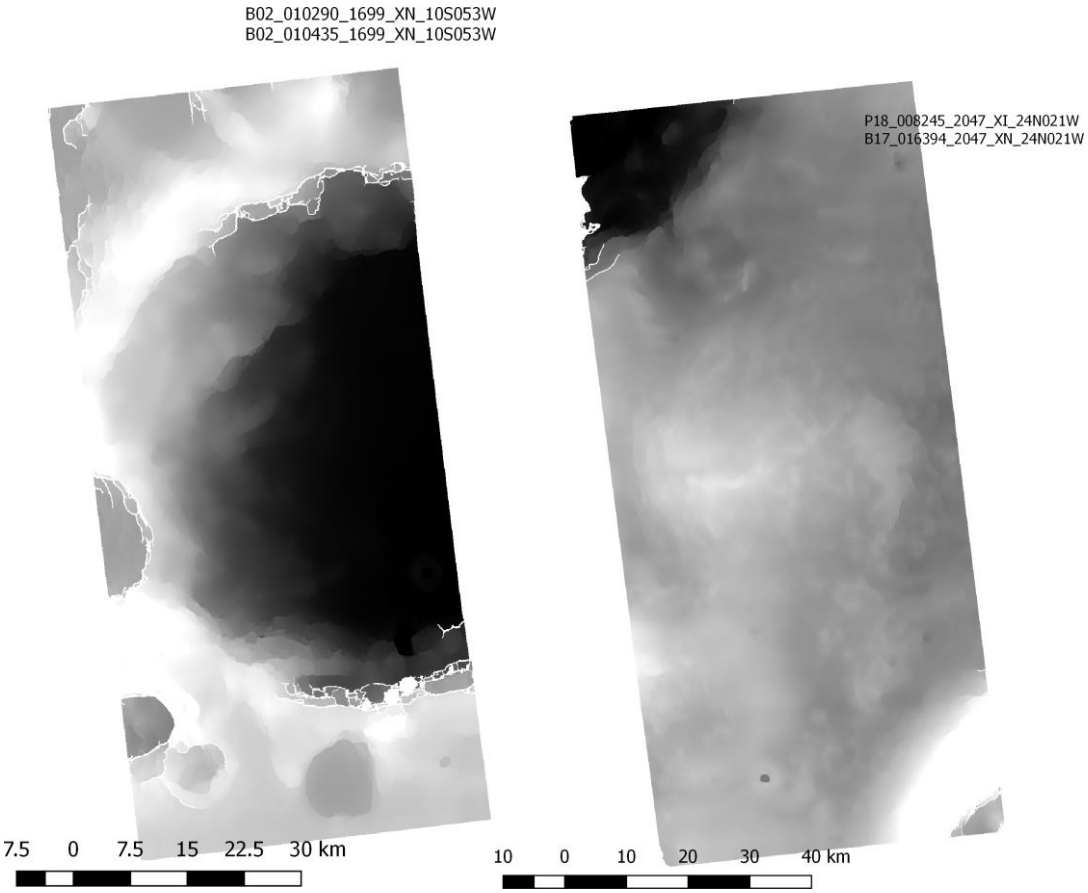


Figure 13a. Example of DTMs rated as 3 with minor gaps in shadow/crater area.

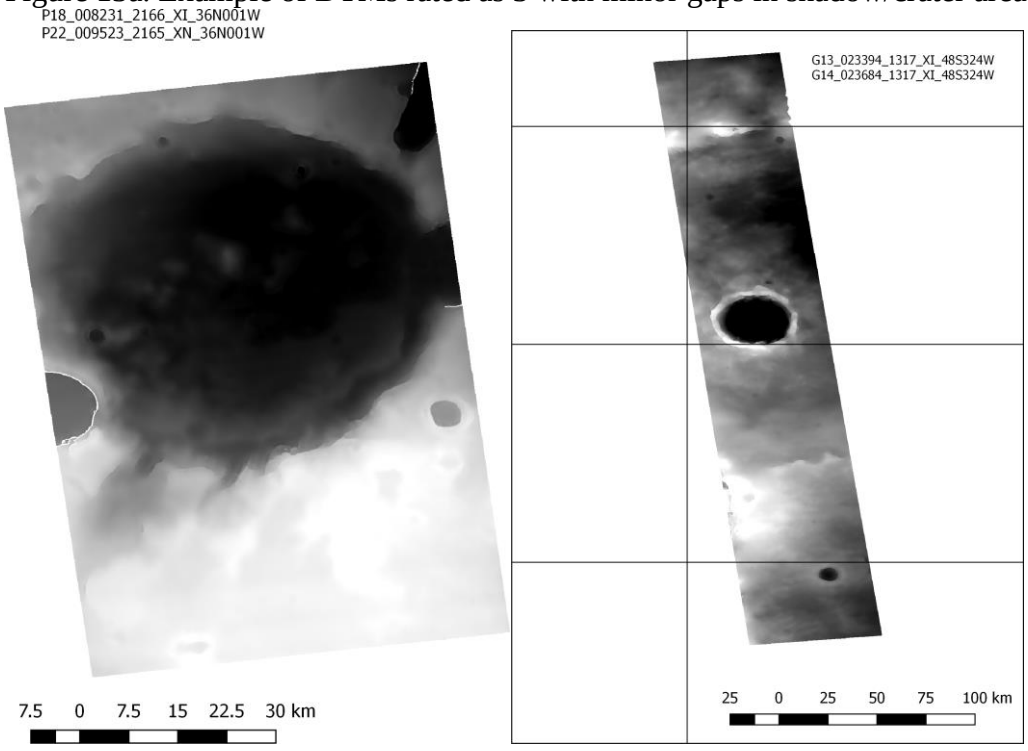


Figure 13b. Example of DTMs rated as 4 and 3 good and minor gaps.

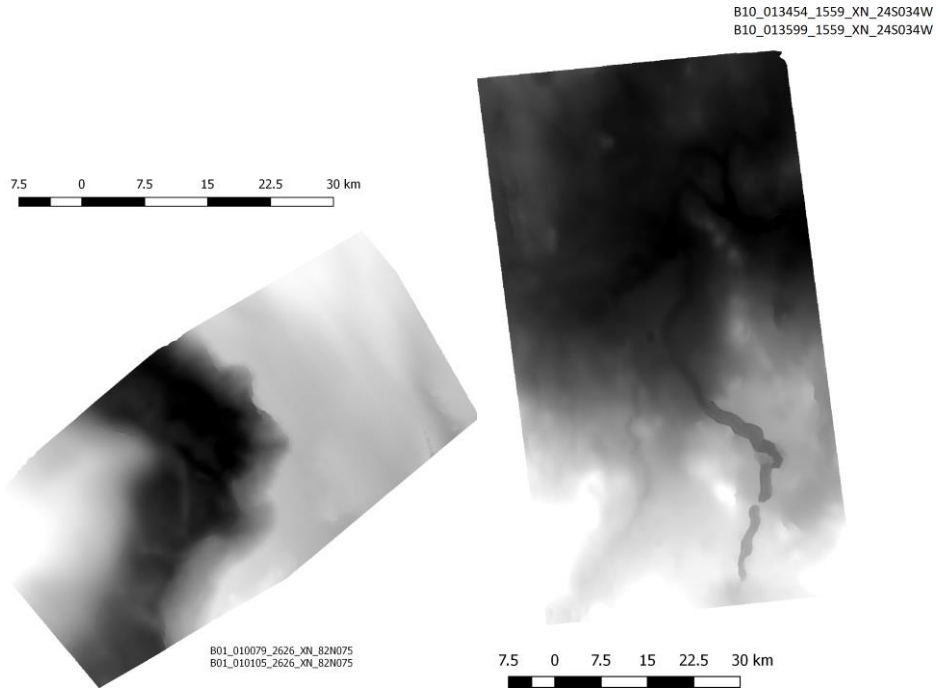


Figure 13c. Example of DTMs rated as 4 and 5 with good qualities.

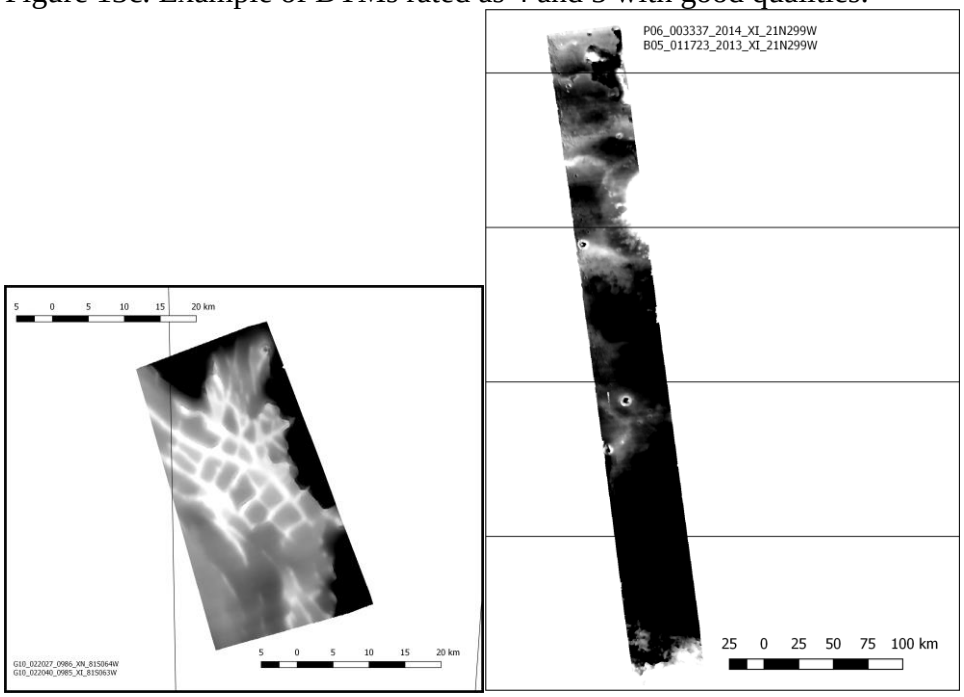

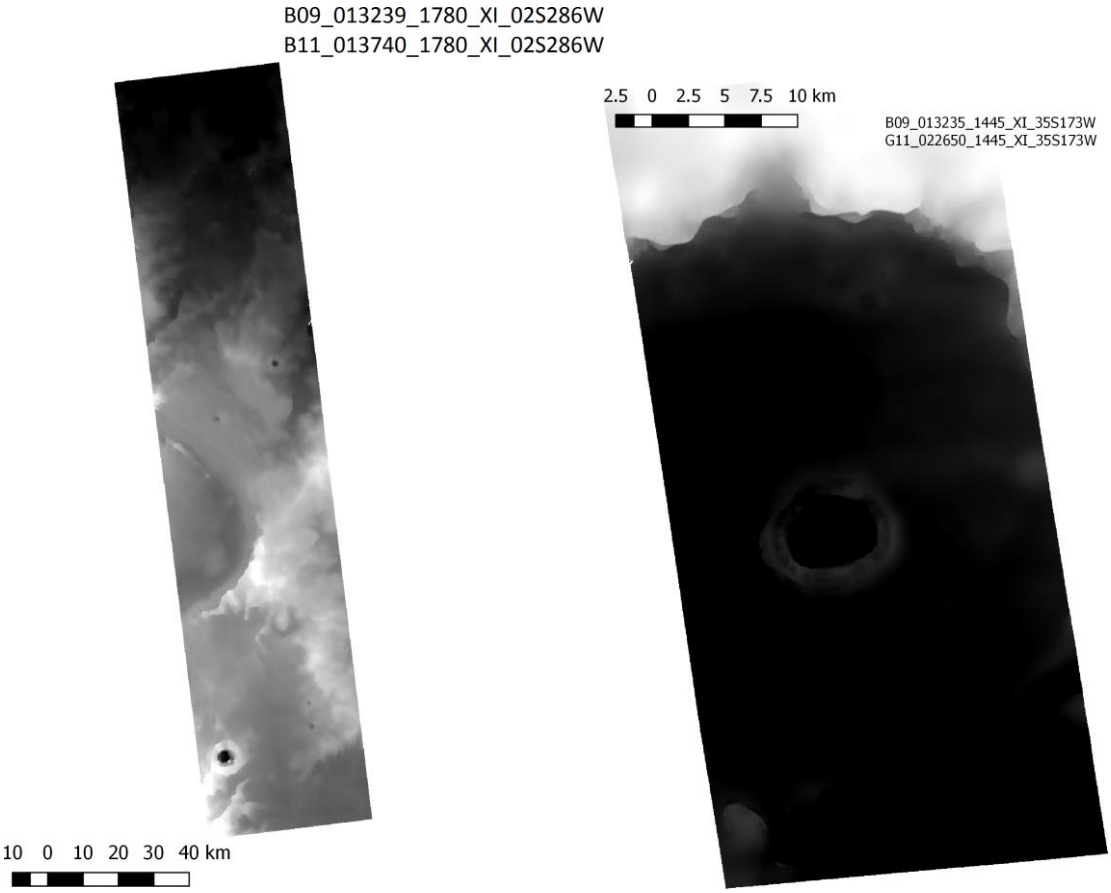


Figure 13d. Example of DTMs rated as 4 or 5 with very good qualities.

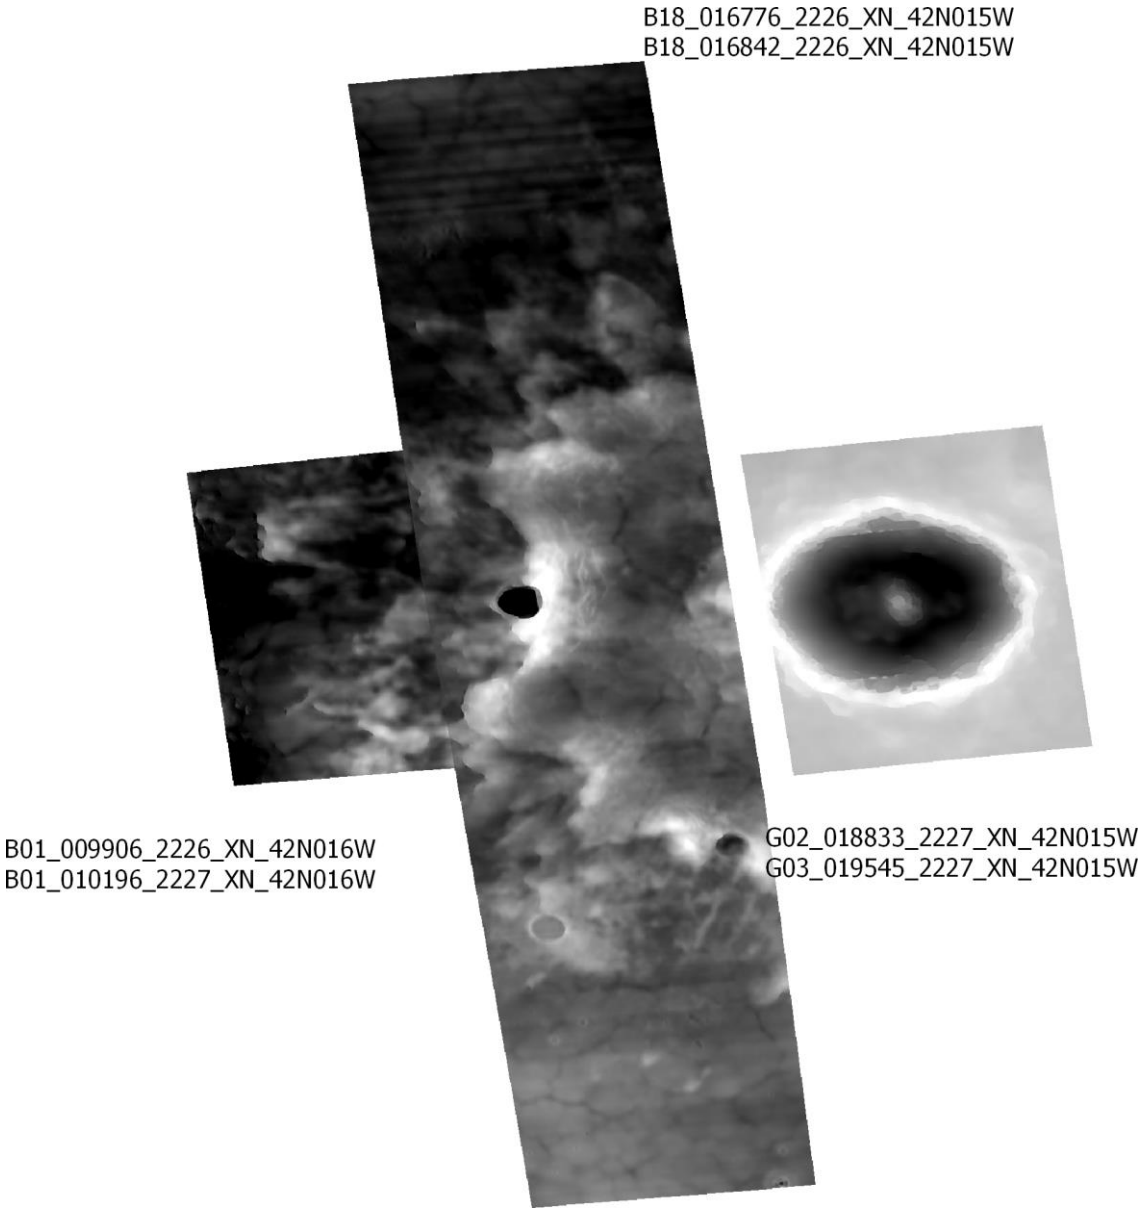


Figure 13e. Example of DTMs rated as 4 or 5 with very good qualities including DTMs for polar area.

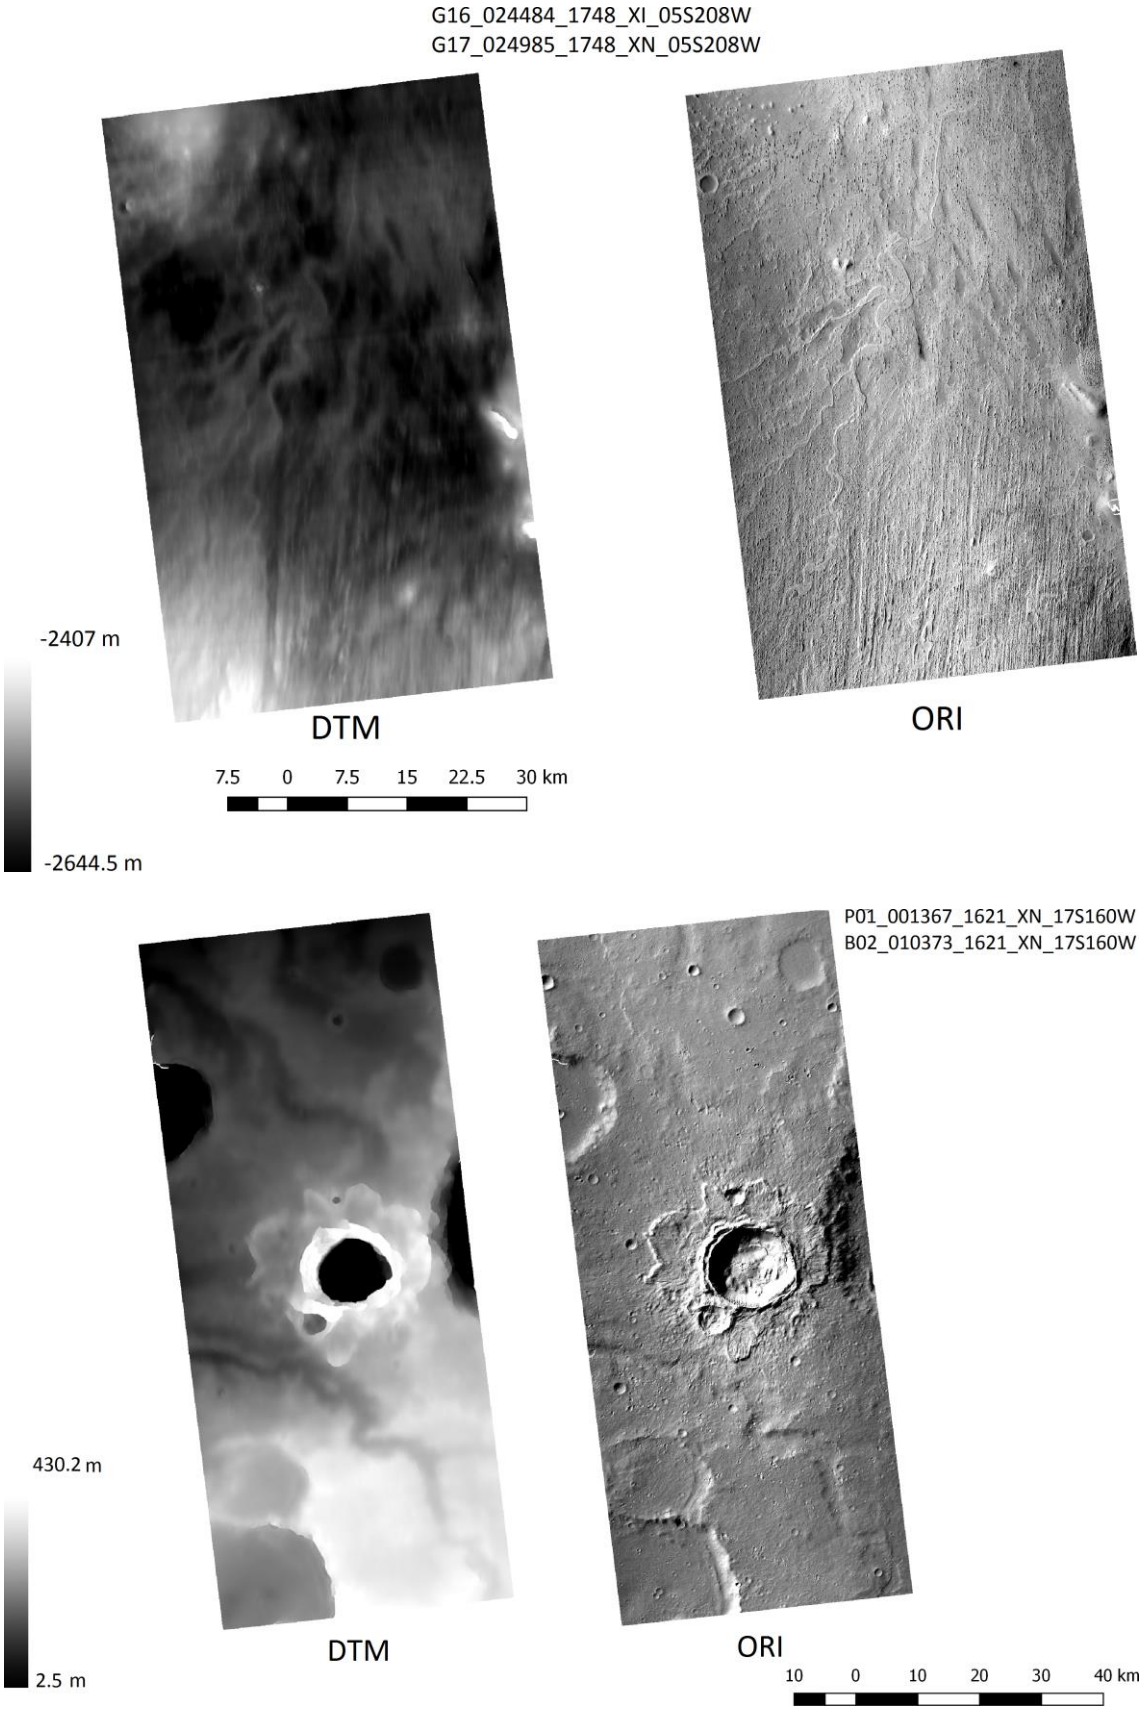


Figure 14. Final DTM and ORIs for 2 sample areas.

# 5. References

Broxton, M. J. and L. J. Edwards. 2008. The Ames Stereo Pipeline: Automated 3D Surface Reconstruction from Orbital Imagery. Lunar and Planetary Science Conference 39, abstract #2419.

Moratto, Z. M., M. J. Broxton, R. A. Beyer, M. Lundy, and K. Husmann. 2010. Ames Stereo Pipeline, NASA's Open Source Automated Stereogrammetry Software. Lunar and Planetary Science Conference 41, abstract #2364.

Olson, C., Maximum-Likelihood Image Matching, 2002, IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 24, no. 6.

Shin, D. and J.-P. Muller, Progressively weighted affine adaptive correlation matching for quasi-dense 3D reconstruction. Pattern Recognition, 2012. 45(10): p. 3795 -3809.

Stein, M. L. (1999), Statistical Interpolation of Spatial Data: Some Theory for Kriging, Springer, New York.

Sidiropoulos, P., and J.-P. Muller (2015a) "Matching of Large Images Through Coupled Decomposition," IEEE Trans. Image Processing, vol. 24 (7) pp. 2124-2139. DOI: 10.1109/TIP.2015.2409978

Sidiropoulos, P. and J. P. Muller (2015b) "On the status of orbital high-resolution repeat imaging of Mars for the observation of dynamic surface processes," *Planetary and Space Science*, vol. 117, pp. 207–222. DOI: 10.1016/j.pss.2015.06.017

Tao, Y.,J.-P. Muller and W. Poole (2016) "Automated localisation of Mars rovers using co-registered HiRISECTX-HRSC orthorectified images and wide baseline Navcam orthorectified mosaics", *Icarus*, vol. 280, pp. 139–157, DOI: 10.1016/j.icarus.2016.06.017

Tao, Y. and J.-P. Muller (2017) "Automated planet-wide DTM generation from NASA MRO data – a status report," *Lunar and Planetary Science*, vol. XLVIII, p2965 http://www.hou.usra.edu/meetings/lpsc2017/pdf/2965.pdf