

JPL D-73327

Earth Observing System



Multi-angle
Imaging
Spectro-
Radiometer

MISR Level 2 Cloud Product Algorithm Theoretical Basis

Kevin Mueller¹
Catherine Moroney¹
Veljko Jovanovic¹
Michael J. Garay¹
Jan-Peter Muller²
Larry Di Girolamo³
Roger Davies⁴

¹Jet Propulsion Laboratory, California Institute of Technology

²Mullard Space Science Laboratory, University College London

³University of Illinois at Urbana-Champaign

⁴University of Auckland



Jet Propulsion Laboratory
California Institute of Technology

March 5, 2013

JPL D-73327

Multi-angle Imaging SpectroRadiometer (MISR)

MISR Level 2 Cloud Product Algorithm Theoretical Basis

Approval:

David J. Diner
MISR Principal Investigator

The MISR web site should be consulted to determine the latest released version of this document (<http://misr.jpl.nasa.gov>).

Approval signatures are on file with the MISR Project.



**Jet Propulsion Laboratory
California Institute of Technology**

March 5, 2013

Copyright 2012 California Institute of Technology. Government sponsorship acknowledged.

The research described in this publication was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration

Document Change Log

Revision	Date	Affected Portions and Description
	5 March 2013	All, original release

This document applies to the MISR L2 Cloud data products.

TABLE OF CONTENTS

LIST OF TABLES	viii
LIST OF FIGURES.....	viii
1 INTRODUCTION	1
1.1 <i>PURPOSE</i>	1
1.2 <i>SCOPE</i>	2
1.3 <i>MISR DOCUMENTS</i>	3
1.4 <i>REVISIONS</i>	3
2 EXPERIMENT OVERVIEW	4
2.1 <i>OBJECTIVES OF MISR LEVEL 2 CLOUD PRODUCT</i>	4
2.2 <i>INSTRUMENT CHARACTERISTICS</i>	4
2.3 <i>MISR LEVEL 2 CLOUD PRODUCT STRATEGY</i>	5
2.3.1 Stereo height vector reconstruction.....	6
2.3.2 Height-resolved cloud motion vector reconstruction.....	7
3 ALGORITHM DESCRIPTION.....	11
3.1 <i>PROCESSING OUTLINE</i>	11
3.2 <i>ALGORITHM INPUT</i>	15
3.2.1 MISR data.....	15
3.2.1.1 Ellipsoid-projected TOA Radiance and Data Quality Indicator	15
3.2.1.2 Instrument radiometric uncertainties	16
3.2.1.3 Ancillary Geographic Product inputs.....	16
3.3 <i>THEORETICAL DESCRIPTION:</i>	16
3.3.1 Correspondence.....	16
3.3.1.1 HSAD Correspondence.....	17
3.3.1.1.1 Search area determination	17
3.3.1.1.2 Image normalization and resampling.....	18
3.3.1.1.3 Discrete correspondence.....	18
3.3.1.1.4 Subpixel Refinement	19
3.3.1.2 Conjugate Cluster Analysis.....	20
3.3.1.3 M23 Correspondence	21
3.3.1.3.1 Low contrast screening.....	21
3.3.1.3.2 Multipoint Matcher (M2).....	22
3.3.1.3.3 Multipoint matcher using Medians (M3).....	24
3.3.1.3.4 Elimination of Ambiguous Matches.....	24
3.3.1.3.5 Image Pyramid and Previous Match Methods for M23.....	25
3.3.2 Reconstruction	28
3.3.2.1 Determining time and look vectors from L1B2 image coordinates	28
3.3.2.2 Camera Model Corrections.....	28
3.3.2.2.1 Determining Camera Model Corrections.....	29
3.3.2.3 Reconstruction of height resolved motion vectors from conjugate triplets.....	31
3.3.2.4 Reconstruction of stereo vectors from conjugate pairs	31
3.3.2.5 Stereo height vector sensitivity to along-track motion	33
3.3.2.6 Instrument heading of height-resolved motion vectors and stereo height vectors.....	33
3.3.3 Height Resolved Motion Vector Operations	34
3.3.3.1 Merge forward and aft and assess quality of height resolved motion vectors.....	34
3.3.3.2 Assess quality and screen motion vectors per orbit.....	37

3.3.3.3	Generate motion derived cloud mask	37
3.3.3.3.1	Derive thresholds for cloud masking.....	38
3.3.4	Stereo Height Vector Operations.....	40
3.3.4.1	Merge forward and aft and assess quality of stereo vectors	40
3.3.4.1.1	Screen less reliable cross-track motion outside interior of MISR swath.....	41
3.3.4.2	Project stereo vectors from ellipsoid-referenced to feature referenced	42
3.3.4.3	Apply along-track correction to stereo height vectors where applicable	43
3.3.4.4	Generate Stereo Derived Cloud Mask.....	43
4	ASSUMPTIONS AND LIMITATIONS	45
4.1	<i>ASSUMPTIONS</i>	45
4.2	<i>LIMITATIONS</i>	45
5	REFERENCES	47
A.	GLOSSARY OF ACRONYMS.....	48
B.	APPENDIX.....	50
	<i>Constant thresholds</i>	50
	<i>SOM-y clipping boundaries for cross-track motion</i>	51

LIST OF TABLES

Table 1: Grids and parameters in the Level 2 Cloud Product.....	2
Table 2: Determinants of different cameras.....	10
Table 3: Level 2 Cloud inputs.....	15
Table 4: Threshold values.....	50
Table 5: SOM-y clipping boundaries for cross-track motion (for top of block)	51

LIST OF FIGURES

Figure 1. Schematic of correspondence and reconstruction for single conjugate.....	5
Figure 2. Simplified geometry of stereo height vector reconstruction	7
Figure 3. Geometry of MISR imaging of cloud in the along-track direction (from [7]).....	9
Figure 4. Conventions used in processing flow diagrams	12
Figure 5. Level 2 Cloud Processing Diagram.....	12
Figure 6. Motion Reconstruction Operations Diagram.....	13
Figure 7. Stereo Reconstruction Operations Diagram	13
Figure 8. Motion Vector Operations Diagram.....	14
Figure 9. Stereo Vector Operations Diagram	14
Figure 10. Two-dimensional cross-section of histogram refinement iteration	21
Figure 11. Flowchart for M2/M3 Processing.....	27
Figure 12. Bimodal retrieval altitude minus terrain altitude distribution	39
Figure 13. Bimodal retrieval cross-track motion distribution.....	40
Figure 14. Example of cross-swath coverage and clipping	42
Figure 15. Schematic of ellipsoid-referenced versus feature-referenced positions	42

1 INTRODUCTION

1.1 PURPOSE

The Multi-angle Imaging SpectroRadiometer (MISR) Level 2 Cloud product provides the climate, modeling, and general atmospheric research communities a unique set of atmospheric parameters associated with the observation of cloud frequency, height, and motion. The product contains a cloud mask and cloud top height (CTH) parameters gridded at 1.1 km resolution, height-resolved cloud motion vector (CMV) parameters at 17.6 km resolution, and information on the *cross-track* component of cloud motion at 1.1 km resolution. These parameters are recorded throughout the 380 km MISR swath and distributed for each orbit. The available parameters are divided into three categories, each represented by a distinct HDF-EOS grid as listed in Table 1. MISR CMV parameters are provided at 17.6 km resolution within the *Motion_17.6_km* grid. MISR CTH and cross-track cloud motion parameters are provided at 1.1 km resolution on two grids: *Stereo_1.1_km*, containing parameters that have taken cloud motion into account, and *Stereo_WithoutWindCorrection_1.1_km*, containing parameters that have not.

The intent of this document is to identify and describe sources of Level 2 Cloud input data required for parameter retrievals, provide the physical theory and mathematical background underlying parameter derivations, include implementation details, and describe key assumptions and limitations of the adopted approach. This document is used by the MISR Science Data System Team to establish requirements and functionality of the data processing software.

Table 1: Grids and parameters in the Level 2 Cloud Product

Motion_17.6_km		
Parameter name	Units	Description
CloudTopHeightOfMotion	m	Retrieved height of feature relative to Earth ellipsoid
CloudMotionEastward	ms ⁻¹	Retrieved eastward motion of feature
CloudMotionNorthward	ms ⁻¹	Retrieved northward motion of feature
MotionDerivedCloudMask	-	Flag denoting with low or high confidence whether retrieval is or is not a cloud
MotionQualityIndicator	-	Flag quantifying quality of retrieval
Stereo_1.1_km		
Parameter name	Units	Description
CloudTopHeight	m	Retrieved height of feature relative to Earth ellipsoid
CloudMotionCrossTrack	ms ⁻¹	Retrieved cross-track motion of feature
CloudMotionCrossTrackHeading	deg	Cross-track heading clockwise from north
StereoDerivedCloudMask	-	Flag denoting with low or high confidence whether feature is or is not a cloud
StereoQualityIndicator	-	Flag quantifying quality of retrieval
Stereo_WithoutWindCorrection_1.1_km		
Parameter name	Units	Description
CloudTopHeight_WithoutWindCorrection	m	Apparent altitude of feature relative to Earth ellipsoid
CloudMotionCrossTrack_WithoutWindCorrection	ms ⁻¹	Apparent cross-track motion of feature
CloudMotionCrossTrackHeading_WithoutWindCorrection	deg	Cross-track heading clockwise from north
StereoDerivedCloudMask_WithoutWindCorrection	-	Flag denoting with low or high confidence whether feature is or is not a cloud
StereoQualityIndicator_WithoutWindCorrection	-	Flag quantifying quality of retrieval

1.2 SCOPE

This document covers the algorithm theoretical basis for the Level 2 Cloud product that is routinely generated at the Langley Research Center (LaRC) Atmospheric Sciences Data Center

(ASDC).

Chapter 1 describes the purpose and scope of the document. Chapter 2 provides a brief overview. The processing concept and algorithm description are presented in Chapter 3. Chapter 4 summarizes assumptions and limitations. Literature references are indicated by a number in italicized square brackets (e.g., [1]).

1.3 MISR DOCUMENTS

Reference to MISR Project Documents is indicated by a number in italicized square brackets as follows (e.g., [M-1]). The MISR web site (<http://misr.jpl.nasa.gov>) should be consulted to determine the latest released version of each of these documents.

[M-1] Experiment Overview, JPL D-13407, Rev. A.

[M-2] Data Product Description for the MISR Level 2 Cloud Product, JPL D-72327.

[M-3] Level 1 Georectification and Registration Algorithm Theoretical Basis, JPL D-11532, Rev. D.

[M-4] Level 1 Ancillary Geographic Product Algorithm Theoretical Basis, JPL D-13400, Rev. A.

[M-5] Level 1 In-flight Radiometric Calibration and Characterization Algorithm Theoretical Basis, JPL D-13398.

[M-6] Level 2 Cloud Detection and Classification, JPL D-11399, Rev. D

[M-7] Level 1 In-Flight Geometric Calibration, JPL D-13399, Rev. B

1.4 REVISIONS

This is the original version of the document, although a significant portion of its content has been sourced from [M-6].

2 EXPERIMENT OVERVIEW

2.1 OBJECTIVES OF MISR LEVEL 2 CLOUD PRODUCT

The MISR Level 2 Cloud product provides a cloud mask and cloud top height (CTH) parameters gridded at 1.1 km resolution, height-resolved cloud motion vector (CMV) parameters at 17.6 km resolution, and information on the *cross-track* component of cloud motion at 1.1 km resolution. Both cloud fraction (as can be derived from a cloud mask) and CTH are crucial meteorological and climatological variables used to evaluate and develop parameterizations for both mesoscale and global climate models. The Level 2 Cloud product provides estimates of CTH that are uniquely valuable amongst passive sensors due to their fine vertical precision and independence from radiometric calibration. CMVs are a valuable proxy observation of the horizontal atmospheric wind field at the retrieved altitude of the cloud. They have long been recognized as a critical element of assimilation for numerical weather prediction (NWP) [1]. More than a decade of global sampling and near-instantaneous retrievals with fine 17.6 km horizontal resolution are unique advantages offered by MISR CMVs. The 1.1 km resolution of cross-track cloud motion field affords the opportunity to assess cloud dynamics at an even finer scale. For all parameters, the projected span of available data is noteworthy. The MISR instrument has been operational since February 2000 and is not projected to cease operation prior to 2019, ultimately providing a record spanning nearly two decades.

2.2 INSTRUMENT CHARACTERISTICS

The MISR instrument consists of nine pushbroom cameras. It obtains global coverage every nine days, and flies in a 705-km descending polar orbit on the EOS-Terra platform. The cameras are arranged with one camera pointing toward the nadir (designated An), one bank of four cameras pointing in the forward direction (designated Af, Bf, Cf, and Df in order of increasing off-nadir angle), and one bank of four cameras pointing in the aftward direction (using the same convention but designated Aa, Ba, Ca, and Da). Images are acquired with nominal view zenith angles, relative to the Earth surface reference ellipsoid, of 0°, 26.1°, 45.6°, 60.0°, and 70.5° for An, Af/Aa, Bf/Ba, Cf/Ca, and Df/Da, respectively. Each camera uses four Charge-Coupled Device (CCD) line arrays in a single focal plane. The line arrays consist of 1504 photoactive pixels plus 16 light-shielded pixels per array. Each line array is filtered to provide one of four MISR spectral bands. The spectral band shapes are approximately Gaussian and centered at 447, 558, 672, and 866 nm.

MISR contains 36 parallel signal chains corresponding to the four spectral bands in each of the nine cameras. The zonal overlap swath width of the MISR imaging data (i.e., the swath seen in common by all nine cameras along a line of constant latitude) is approximately 380 km, which provides global multi-angle coverage of the entire Earth in 9 days at the equator, and 2 days near the poles. The instantaneous field of view (IFOV) for each camera has slightly different along-track and cross-track characteristics, where along-track and cross-track refer to the orientation along and orthogonal to the subsatellite ground-track of the platform, respectively. The cross-track IFOV and sample spacing of each pixel is 275 m for all of the off-nadir cameras, and 250 m for the nadir camera. Along-track IFOVs depend on view zenith angle, ranging from 214 m in

the nadir to 707 m at the most oblique angle. Sample spacing in the along-track direction is 275 m in all cameras.

Additional background on the instrument design is provided in [M-1].

2.3 MISR LEVEL 2 CLOUD PRODUCT STRATEGY

The underlying mechanism of the MISR cloud top height and motion retrievals consists of two steps, *correspondence* and *reconstruction*. That is, (1) identifying *conjugates*, the apparent image coordinates associated with some real feature as captured by two or more of MISR's nine cameras during an overpass, and (2) subsequently inferring the true position and motion of each such feature by intersecting known camera lines of sight associated with conjugate image coordinates. Figure 1 illustrates these concepts. In the first step (a), correspondence identifies a conjugate consisting of image coordinates, $(x,y)_{An}$ and $(x,y)_{Af}$, within images captured by cameras An and Af . In the second step (b), reconstruction infers the position, (x,y,z) , of the feature associated with that conjugate. These are the same principles underlying the MISR Level 2 Top-of-Atmosphere/Cloud Stereo product software, upon which the Level 2 Cloud product software is based [M-6]. The underlying theory shared by both products is well documented [2][4][5][7][9][10].

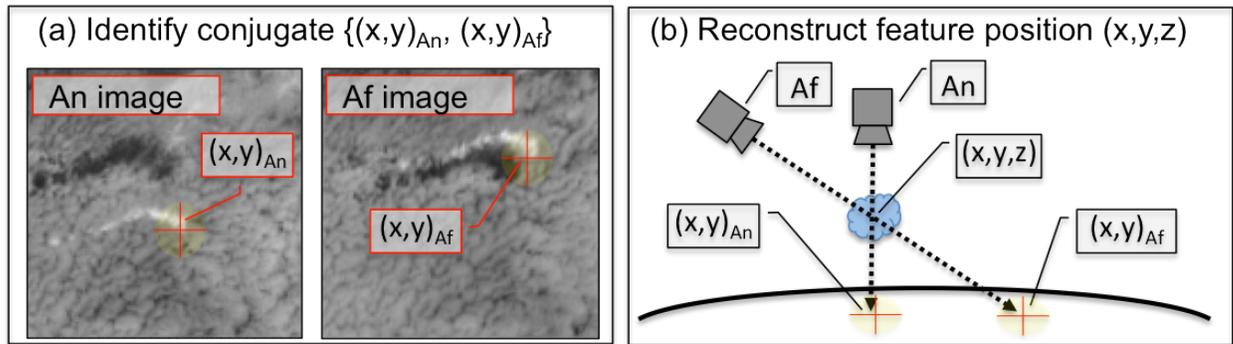


Figure 1. Schematic of correspondence and reconstruction for single conjugate

The structure of Level 2 Cloud product software is governed by practical considerations associated with correspondence and reconstruction. Correspondence algorithms rely on the premise that the pattern of radiance reflected toward one camera by a feature will be similar to the pattern reflected toward another camera. They perform optimally when the compared cameras differ minimally in observation time and viewing angle. On the other hand, accurately reconstructing *motion vectors* consisting of height and both axes of horizontal motion requires input conjugates specifying coordinates for three cameras (i.e., *conjugate triplets*) spanning a large view angle difference (as is detailed in 2.3.2). Robustly producing such conjugates further requires aggregating correspondence information over a large support region, coarsening input conjugate resolution, which is a major factor underlying the choice of 17.6 km resolution for the motion vector parameters. However, as detailed in 2.3.1, *stereo height vectors* consisting of height and only the cross-track axis of horizontal motion, can be reconstructed from conjugates specifying coordinates for two cameras (i.e., *conjugate pairs*) with minimal view angle difference, subject to the *a priori* assumption that features do not move along-track during the

interval between camera views. These conjugates can be robustly produced at 1.1 km resolution, yielding CTH and cross-track cloud motion components at that resolution. Furthermore, the error in estimated height induced by any along-track motion can, in principle, be corrected *a posteriori* using external knowledge of cloud motion. The Level 2 Cloud product provides both *a priori* heights and heights corrected for along-track motion wherever a corresponding 17.6 km resolution CMV has been retrieved.

2.3.1 Stereo height vector reconstruction

Stereo height vectors can be geometrically reconstructed from conjugates specifying a feature's coordinates in two MISR cameras. For simplicity, we illustrate the principles of the reconstruction algorithms in Figure 1 [10], by reducing the problem from three to two dimensions, neglecting the curvature of the Earth surface ellipsoid, and neglecting the curvature of the satellite path relative to Earth. As depicted in Figure 1, for a scene point along the edge of a cloud, parallax causes a disparity, $\overline{AB_1}$, proportional to cloud height, h , between different camera view zenith angles, a and b . For a cloud remaining stationary during the interval, t_1 - t_0 , between MISR camera views, the cloud height, h , is simply given by equation 1.

$$h = \frac{\overline{AB_1}}{\tan a - \tan b} \quad (1)$$

For a cloud with a non-zero along-track component velocity, u , the measured disparity ($\overline{AB_2}$) is the sum of contributions from parallax ($\overline{AB_1}$) and from along-track cloud motion ($\overline{B_1B_2}$). Assuming no vertical cloud motion, we can use the look rays \vec{a}_2 and \vec{b}_2 to yield an apparent cloud height, h' , which is biased relative to the true height as shown in equations 2 and 3. The bias is proportional to the product of the along-track motion, u , and a sensitivity parameter, h_s , governed predominantly by the view angle.

$$h' = \frac{\overline{AB_2}}{\tan a - \tan b} = h + uh_s \quad (2)$$

$$h_s = \frac{t_2 - t_0}{\tan a - \tan b} \quad (3)$$

The cross-track component of cloud motion, v , may be determined from the cross-track disparity between conjugate coordinates. Perpendicular to the cross-section shown in Figure 1, one may imagine a point B_2' on the surface ellipsoid. The cross-track disparity is then $\overline{B_2B_2'}$, and cross-track motion is simply $v = \overline{B_2B_2'} / (t_2 - t_0)$.

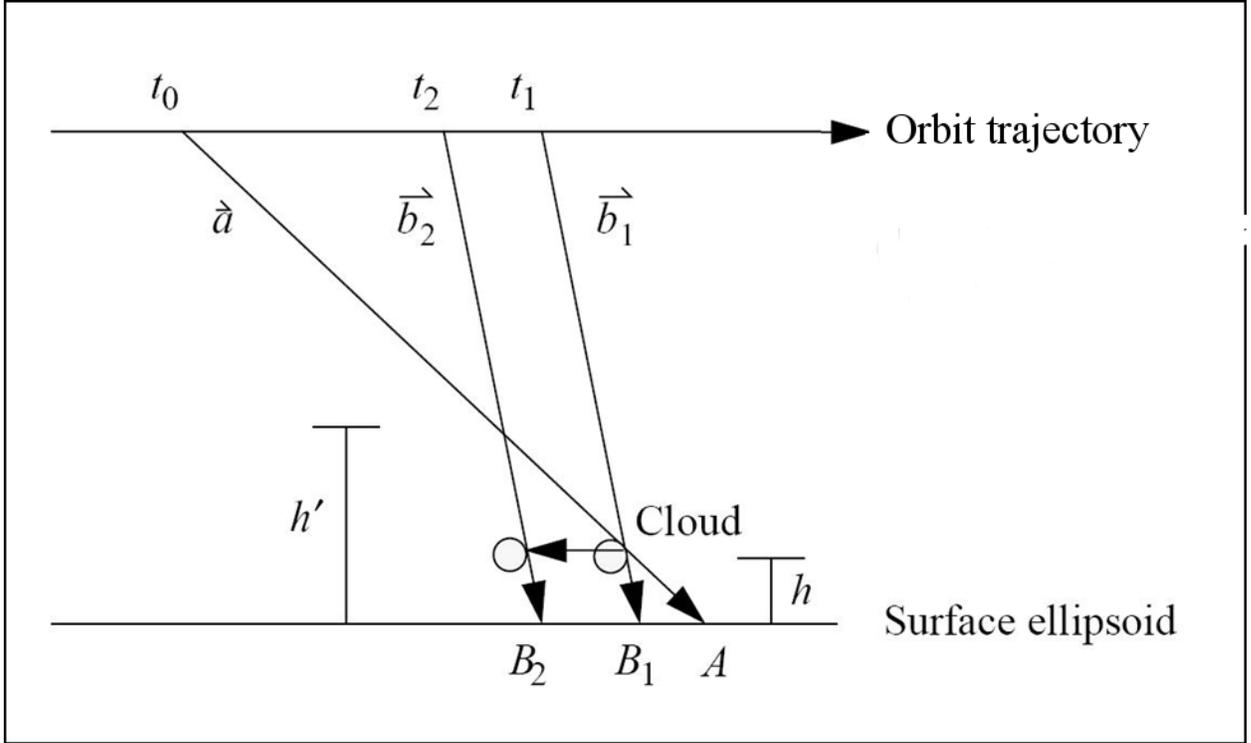


Figure 2. Simplified geometry of stereo height vector reconstruction

2.3.2 Height-resolved cloud motion vector reconstruction

The previous section outlines the reconstruction of stereo height vectors including cloud height and cross-track cloud motion information using only two MISR cameras and shows that a bias is introduced by the assumption of zero along-track cloud motion. This section outlines the more comprehensive reconstruction of height-resolved cloud motion vectors from conjugates specifying a feature's coordinates in three, rather than two, MISR cameras. In this case, the curvature of Earth's surface along MISR views plays a critical role in the separation of cloud height and the along-track cloud motion.

Figure 3 shows the geometry of MISR imaging in the along-track direction where cloud height and motion are highly correlated due to large view zenith angle differences from one camera to another [10]. The circles indicate the locations of a cloud at different times t_i , assuming there is a constant along-track horizontal cloud motion v_c and no vertical component of cloud motion. The projections of the cloud on the surface ellipsoid are at the locations x_i at these times, associated with the discrete MISR camera view zenith angles θ_i . In this discussion, the cloud motion and the camera look vectors are assumed to lie in the along-track plane for simplicity.

If a cloud edge is seen by two cameras with different view zenith angles at times t_1 and t_2 , respectively, then the traveling distance of the spacecraft and that of the cloud during this time interval can be described by the following equations:

$$v_s(t_2 - t_1) = (R + H)(\alpha_2 - \alpha_1) - (R + H)(\Gamma_2 - \Gamma_1) \quad (4)$$

$$v_c(t_2 - t_1) = (R + h)(\alpha_2 - \alpha_1) - (R + h)(\gamma_2 - \gamma_1) \quad (5)$$

where v_s and v_c are the velocities of the spacecraft and the cloud in the along-track direction, respectively, R is the radius of the Earth, H is the orbit altitude of the spacecraft above the Earth's surface, and h is the cloud height above the ellipsoid. As shown in Figure 3, α_1 and α_2 are the angles between the initial radial line at the time t_0 and the radial line passing the image locations x_1 and x_2 , respectively; Γ_1 and Γ_2 are the angles between the radial lines to the spacecraft and the corresponding image locations x_1 and x_2 ; and γ_1 and γ_2 are the angles between the radial lines to the cloud and the corresponding image locations x_1 and x_2 .

Since $h \ll R$, Eq. 5 can be rewritten with variables that we are interested in:

$$v_c(t_2 - t_1) = (x_2 - x_1) + h(\tan \theta_1 - \tan \theta_2) \quad (6)$$

For multi-angle images, Eq. 5 can be generalized into a linear system as follows:

$$v_c(t_j - t_i) - h(\tan \theta_i - \tan \theta_j) = (x_j - x_i) \quad (7)$$

$\{i, j = 1, 2, \dots, n, i \neq j, n \geq 3\}$

The linear system expressed by Eq. 7 represents a straight line in the Δx versus Δt space in which each matching pair contributes a point to this line. The slope and intercept of the line determine v_c and h . Matches from at least three images with different view zenith angles θ_i are required to solve for both v_c and h . Any two linear equations are dependent on each other if their determinant is zero, i.e., if

$$\det A = (t_i - t_{i-1})(\tan \theta_i - \tan \theta_{i+1}) - (t_{i+1} - t_i)(\tan \theta_{i-1} - \tan \theta_i) = 0 \quad (8)$$

In such a case, v_c and h are inseparable using these two equations. For a straight flight line over a planar underlying surface, this singularity will always occur.

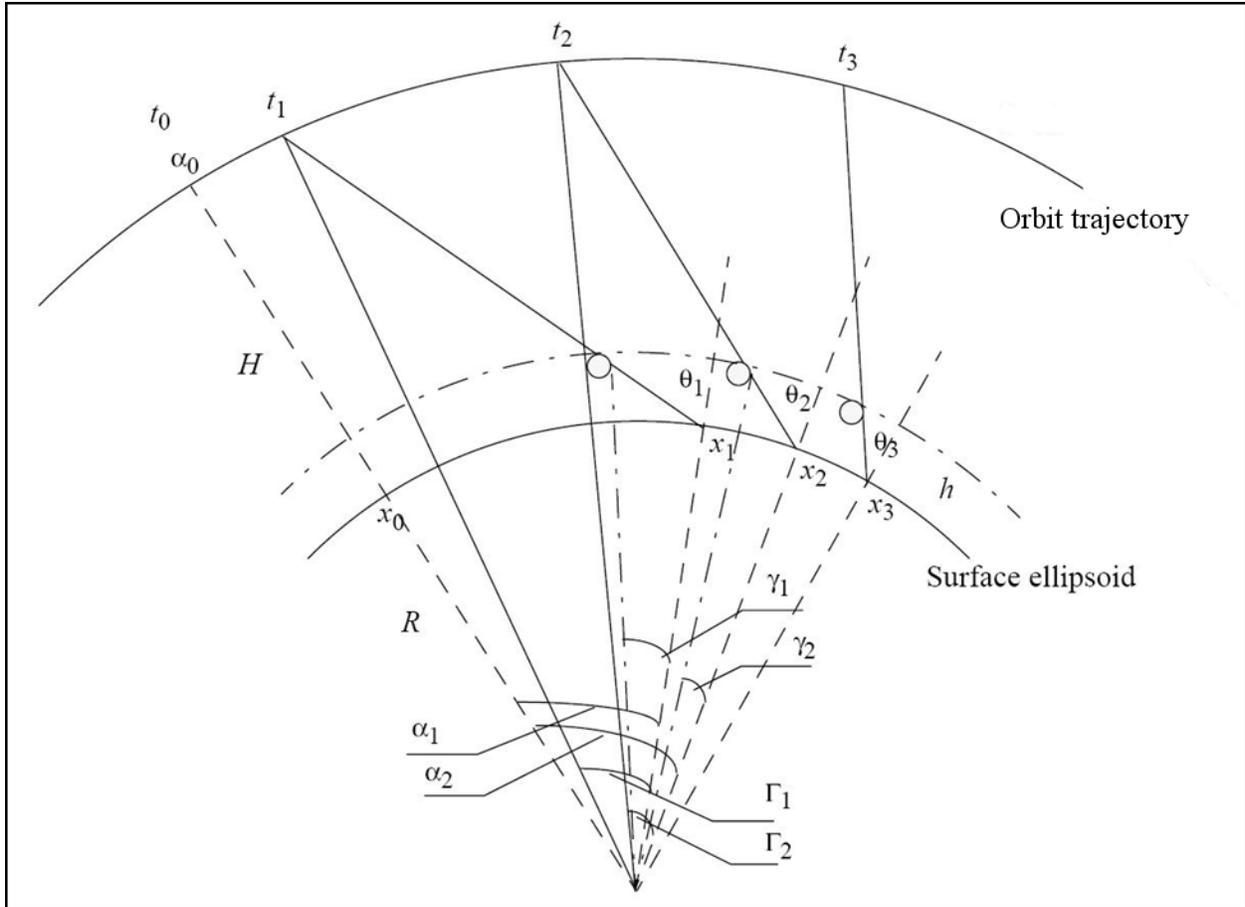


Figure 3. Geometry of MISR imaging of cloud in the along-track direction (from [10])

However, for a satellite with an elliptical orbit over an ellipsoidal surface, the time interval between a pair of camera views is linearly proportional to the tangent of the view zenith angles [10], which allows MISR to separate effects from cloud along-track motion and height [4]. Table 2 shows the relationship of MISR camera view zenith angles with respect to spacecraft traveling times with all cameras looking at the same point on the surface ellipsoid. The values shown in Table 2 for view zenith angles and timing are typical for MISR on the Terra satellite. They vary slightly with orbital location due to the ellipticity of the orbit and non-sphericity of the Earth, but not to an extent that affects this analysis. The last column contains the determinants of the three indicated cameras in square brackets with the cameras numbered consecutively in the order in which they acquire data (e.g., [234] means Cf-Bf-Af, etc.) according to Eq. 8. The bold numbers indicate cases where non-consecutive cameras are used. Cloud height and motion will cause the values in the time column to vary somewhat, but will not change the result. This table shows that the separability of height from along-track motion is most optimal for combinations of the least and most oblique viewing zenith angles.

Table 2: Determinants of different cameras [4]

Camera	θ_i (deg.)	$\tan\theta_i$	t_i (s)	$\Delta t_{i(i-1)}$ (s)	$\Delta \tan\theta_{i(i+1)}$	$\det A$
1. Df	70.5	2.82	0	–	1.09	–
2. Cf	60.0	1.73	60	60	0.71	-366 [123]
3. Bf	45.6	1.02	113	53	0.53	-115 [234]
4. Af	26.1	0.49	159	45	0.49	-38 [345]
5. An	0.0	0.0	204	45	0.49	0 [456]
6. Aa	-26.1	-0.49	249	45	0.53	38 [567]
7. Ba	-45.6	-1.02	295	46	0.71	115 [678]
8. Ca	-60.0	-1.73	348	53	1.09	273 [578]
9. Da	-70.5	-2.82	408	60	–	1217 [579]

The time between a Df camera view and an An camera view for MISR on Terra is about 3.4 minutes. A cloud feature moving along-track at 20 ms^{-1} would be displaced by 4080 m during this interval. That displacement would further lengthen the interval between views by 0.6 s. The longer the time span, the higher the accuracy of the cloud motion retrieval. For reconstruction, the selection of the An, Bx, and Dx cameras in the last box in Table 2 provides not only the best separability of cloud motion and height parameters, but also the highest accuracy in cloud motion retrieval.

3 ALGORITHM DESCRIPTION

3.1 PROCESSING OUTLINE

The standard processing at the LaRC ASDC used to generate the MISR L2 Cloud product can be divided into the six steps outlined below, with sequence and dependencies diagrammed in Figure 5 following conventions shown in Figure 4. The bullet points provided below correspond to more detailed descriptions provided in §3.3.

1. Low Contrast Screening
 - 1.1. Identify regions where the L1B2 An camera has insufficient contrast to allow robust correspondence (§3.3.1.3.1)
2. HSAD and M23 Correspondence
 - 2.1. Independently apply HSAD correspondence algorithm to four pairs of L1B2 camera images, specifically Bf-An, Bf-Df, Ba-An, and Ba-Da (§3.3.1.1)
 - 2.2. Independently apply M23 correspondence algorithm to An-Af and An-Aa pairs of L1B2 camera images (§3.3.1.3)
3. Conjugate Cluster Analysis
 - 3.1. Perform conjugate cluster analysis to extract An-Bf-Df and An-Ba-Da conjugate triplets from Bf-An, Bf-Df, Ba-An, and Ba-Da conjugate pairs (§3.3.1.2)
4. Reconstruction Operations (Figure 6)
 - 4.1. Apply camera model corrections to all conjugates (§3.3.2.2)
 - 4.2. Independently reconstruct conjugates for camera sets An-Bf-Df and An-Ba-Da (§3.3.2.3)
 - 4.3. Independently reconstruct conjugates for camera sets An-Af and An-Aa (§3.3.2.4)
5. Height-resolved Motion Vector Operations (Figure 8)
 - 5.1. Merge forward and aft and assess quality of height-resolved motion vectors (§3.3.3.1)
 - 5.2. Generate motion derived cloud mask (§3.3.3.3)
6. Stereo Height Vector Operations (Figure 9)
 - 6.1. Merge forward and aft and assess quality of stereo height vectors (§3.3.4.1)
 - 6.2. Feature-reference stereo height vectors (§3.3.4.2)
 - 6.3. Correct for along-track cloud motion (§3.3.4.3)
 - 6.4. Generate stereo derived cloud mask (§3.3.4.4)

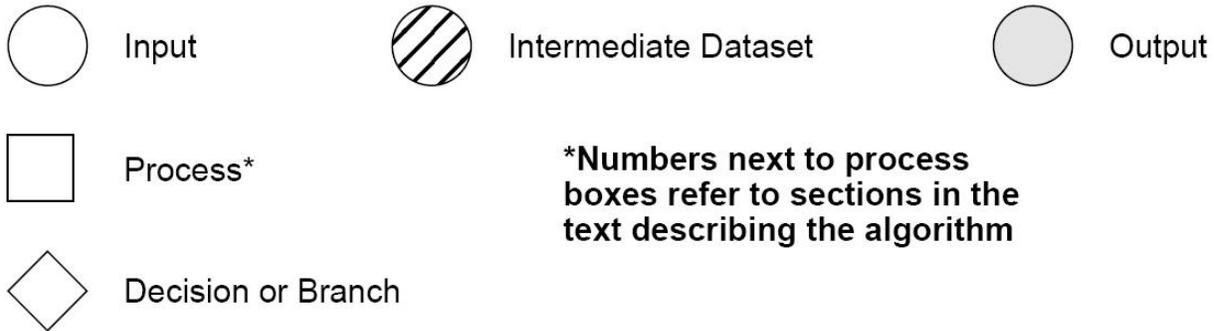


Figure 4. Conventions used in processing flow diagrams

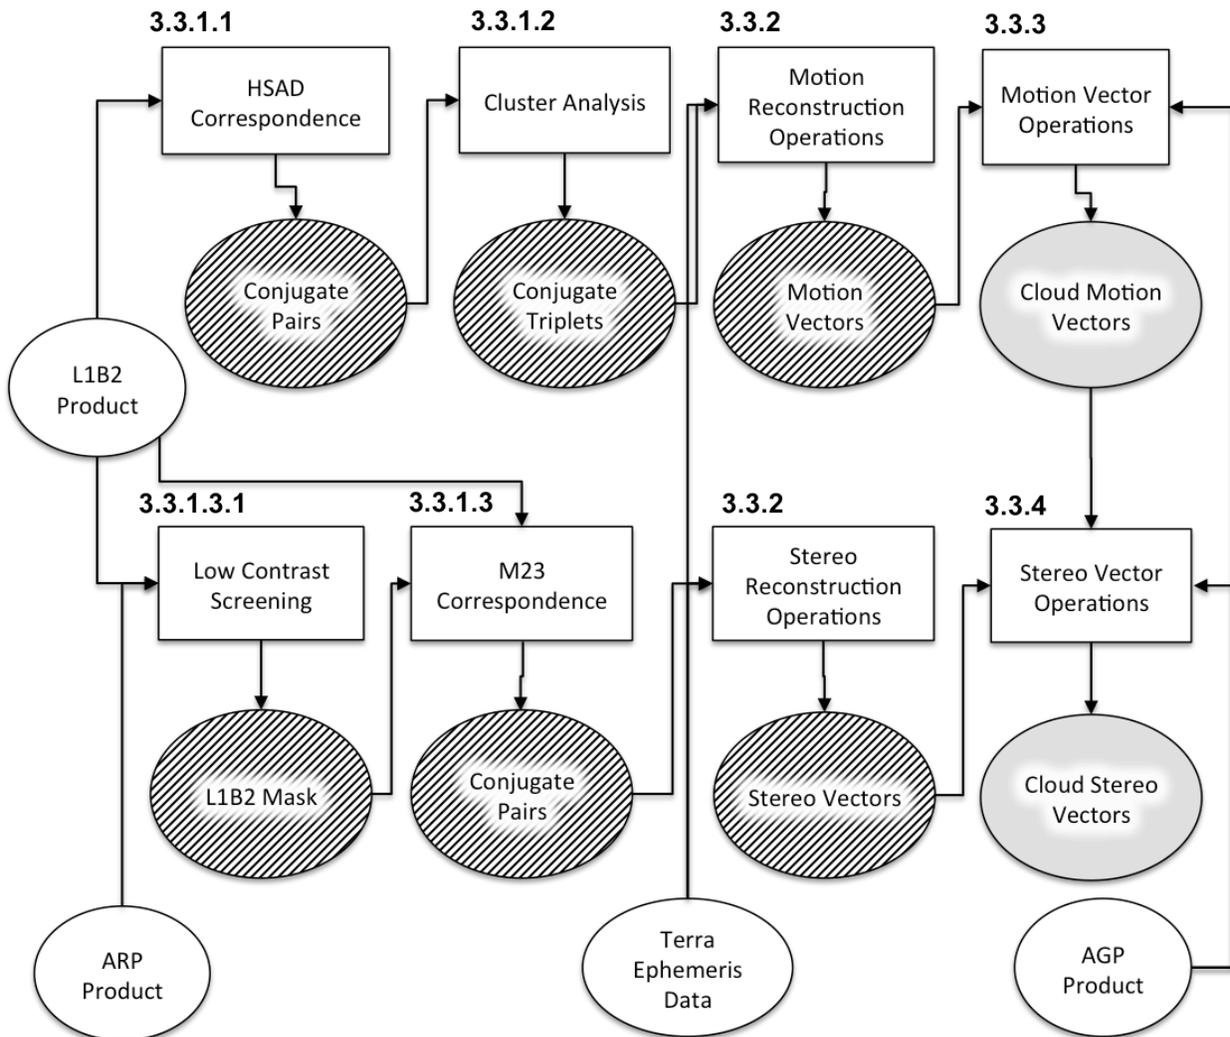


Figure 5. Level 2 Cloud Processing Diagram

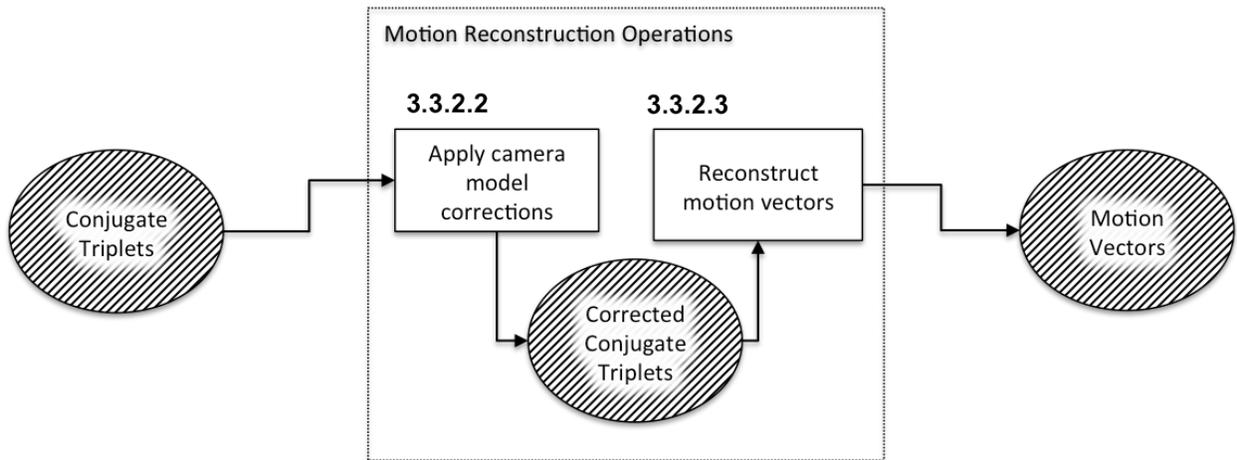


Figure 6. Motion Reconstruction Operations Diagram

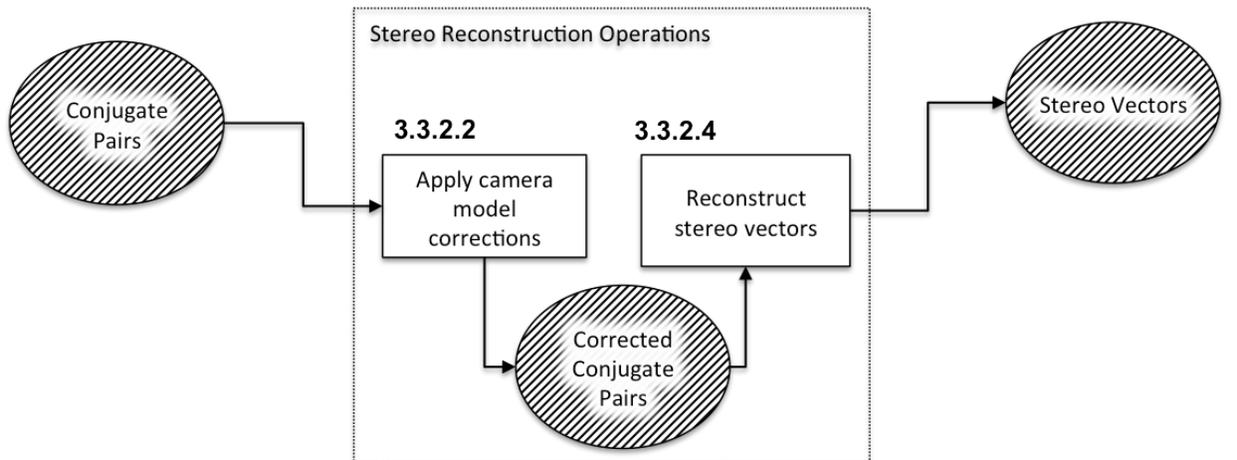


Figure 7. Stereo Reconstruction Operations Diagram

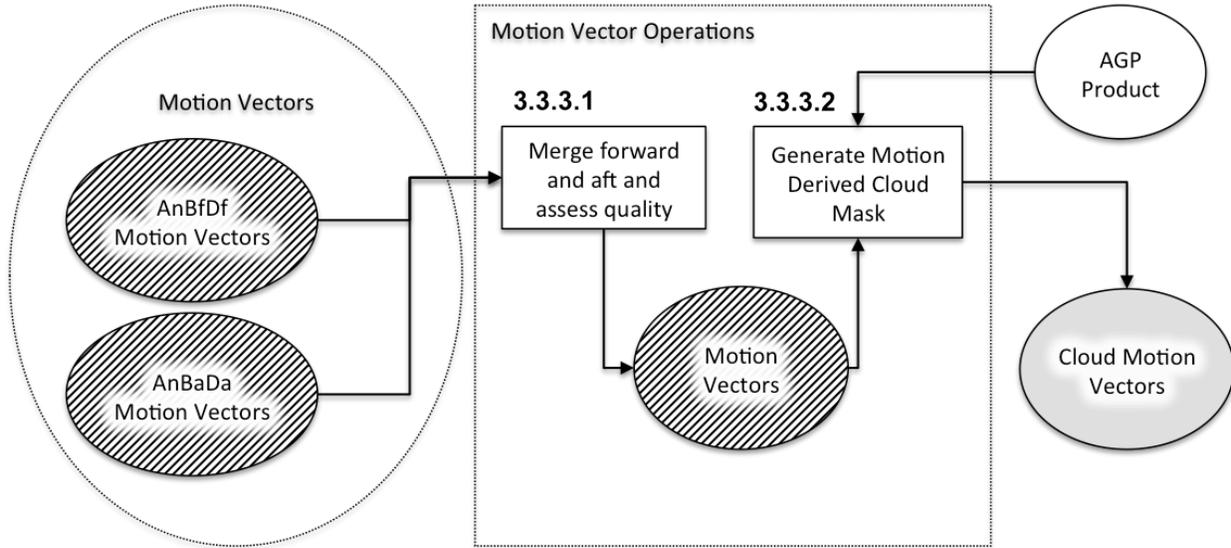


Figure 8. Motion Vector Operations Diagram

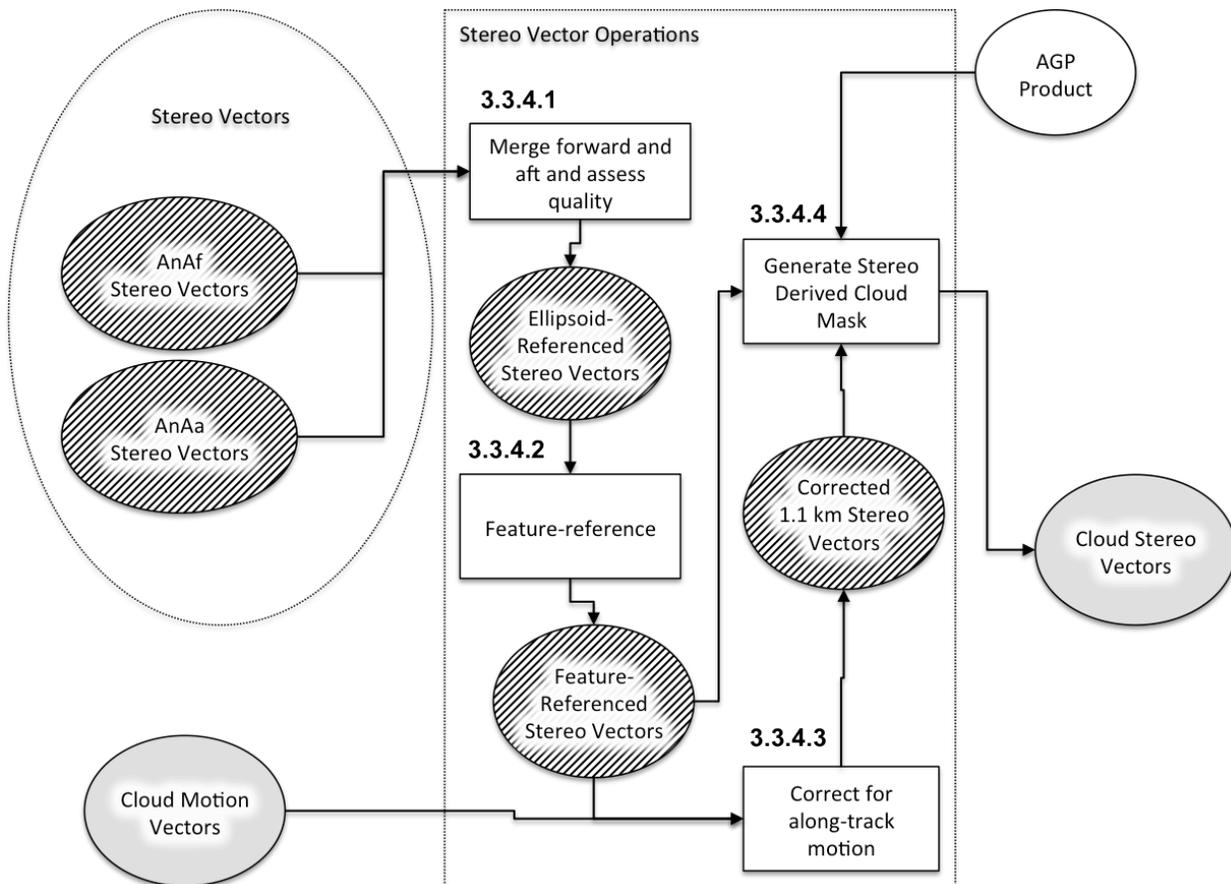


Figure 9. Stereo Vector Operations Diagram

3.2 ALGORITHM INPUT

The required inputs for the Level 2 Cloud software come from MISR and non-MISR sources and are summarized individually in the following paragraphs.

3.2.1 MISR data

Required inputs for the Level 2 Cloud software obtained from other MISR data products are summarized in Table 3. More information about each of the inputs is provided below.

Table 3: Level 2 Cloud inputs

Input data	Source of data	Reference
Ellipsoid-projected TOA Radiance	L1B2 GRP Ellipsoid product	[M-3]
Data Quality Indicators and Data Flags	L1B2 GRP Ellipsoid product	[M-3]
AGP Product Inputs	Level 1 Ancillary Geographic Product	[M-4]
Instrument radiometric uncertainties	Ancillary Radiometric Product	[M-5]
Terra Ephemeris records	Terra	
L1B2 camera model corrections	Level 2 Cloud Configuration	

3.2.1.1 Ellipsoid-projected TOA Radiance and Data Quality Indicator

The ellipsoid-projected TOA radiance parameters are derived from the MISR Level 1B2 product and consist of geolocated, registered, and time tagged radiances in all 36 channels of the instrument projected onto the World Geodetic System 1984 (WGS84) ellipsoid. A resampling process is required in order to implement this projection. In Level 1B2, these radiances are stored as 16-bit integer words. Additionally, a Radiometric Data Quality Indicator (RDQI) is assigned to each radiance value. A multiplicative scale factor, obtained from the Ancillary Radiometric Product (ARP), is applied to convert these integers to floating point radiances, with units of $W m^{-2} sr^{-1} \mu m^{-1}$. Only the 672 nm (red-band) radiances in the Dx, Bx, Ax, and An cameras produced at 275 m resolution are used in the generation of the Level 2 Cloud product.

The ellipsoid-projected radiances are used as the input to image matching in order to identify conjugate features while associated time tags are used to obtain corresponding look vectors. RDQI values are used to assess where radiance measurements are sufficiently reliable for correspondence.

3.2.1.2 Instrument radiometric uncertainties

Retrieval of stereo height vectors is not attempted where image contrast is insufficient to guarantee robust correspondence. The criteria for contrast sufficiency include instrument radiometric uncertainties that may contribute spurious contrast [M-5].

3.2.1.3 Ancillary Geographic Product inputs

Ancillary Geographic Product (AGP) inputs are static Space Oblique Mercator (SOM) gridded fields derived by projecting various map and digital elevation model parameters into SOM coordinate space. The *SurfaceFeatureID* flag categorizes 1.1 km resolution SOM grid domains into identifiers such as *Land*, *Coastline*, *Ocean*, and *Deep Inland Water*. Digital Elevation Models (DEMs) of surface elevation and elevation standard deviation relative to the WGS84 Earth ellipsoid are provided by the AGP at 17.6 km grid resolution (*RegAveSceneElev* and *StdDevRegSceneElev*) and 1.1 km grid resolution (*AveSceneElev* and *StdDevSceneElev*).

The MISR Level 2 Cloud Product uses AGP classifications and DEM information to distinguish retrievals associated with clouds from those likely associated with surface features. Specifically, the AGP *SurfaceFeatureID* is used to define a binary land mask associated with 1.1 km grid cells classified as *Land* or *Coastline*. A coarser resolution land mask is also derived, specifying as land any 17.6 km grid cells containing 1.1 km grid cells classified as *Land* or *Coastline*.

3.3 THEORETICAL DESCRIPTION:

3.3.1 Correspondence

The MISR Level 2 Cloud software employs similar but distinct approaches to produce the conjugate triplets used to reconstruct 17.6 km resolution CMVs and to produce the conjugate pairs used to reconstruct 1.1 km resolution feature height and cross-track motion. Both approaches employ local area stereo matchers. The retrieval of CMVs applies the Hierarchical Sum of Absolute Differences (HSAD) stereo matcher (§3.3.1.1) to each of the Bf-An, Bf-Df, Ba-An, and Ba-Da camera pairs, then clusters and averages these conjugate pairs into representative 17.6 km resolution conjugate triplets as described in §3.3.1.2. The retrieval of feature heights and cross-track motion applies the M23 stereo matcher (§3.3.1.3) to each of the An-Af and An-Aa camera pairs.

Both the M23 and HSAD stereo matchers identify conjugates for each reference coordinate by estimating the comparison camera coordinate that minimizes a univalued *cost function* quantifying pattern dissimilarity as a function of the respective radiance values within *local windows* surrounding a reference and comparison camera coordinate. This entails two steps: (1) calculating the cost function at regular intervals throughout a *search area* of possible matching comparison coordinates, and (2) estimating the location of the cost function minimum from the field of calculated cost function values. The search area spans the range of possible disparities between comparison and reference coordinate, and is governed by the geometry of the two cameras and the range of possible feature parallax and motion during the interval between

camera views.

3.3.1.1 HSAD Correspondence

Applied to the Bf-An, Bf-Df, Ba-An, and Ba-Da pairs of MISR L1B2 camera images, the HSAD stereo matching algorithm operates on each reference camera image coordinate throughout a regular 1100 m SOM grid, yielding either a non-retrieval value or the associated comparison camera image coordinate of a conjugate. The algorithm establishes a search area representing the range of possible comparison coordinates matches for each reference coordinate (§3.3.1.1.1). Next, the Sum-of-Absolute-Differences (SAD) cost function is calculated throughout each search area in order to find the discrete coordinate of the cost minimum representing a successful match (§3.3.1.1.3). The previously computed cost function values surrounding this minimum are then used to interpolate the cost function and thereby determine the subpixel location of minimum cost, which establishes the final conjugate coordinate value (§3.3.1.1.4).

To reduce computational burden, the SAD cost function is not directly applied to the L1B2 image data. Instead, a coarse-to-fine hierarchy of normalized image data is produced, with sampling at 1100 m, 550 m, and 275 m resolution (§3.3.1.1.2). The cost function is only calculated throughout the entire search area at the coarsest 1100 m resolution. The locations of 1100 m input cost minima are then used to prescribe 3300 m × 3300 m centered, surrounding search areas for subsequent determinations of cost function minima at 550 m resolution. The process is then repeated to refine the cost function minima from 550 m to 275 m resolution.

3.3.1.1.1 Search area determination

HSAD defines a search area representing the range of possible disparities between a reference camera coordinate and its conjugate comparison camera coordinate. This search area is determined by the nominal camera view angles and timing as presented in Table 2. It is designed to accommodate a range of possible feature heights from $h_{min} = -500$ m to $h_{max} = 20000$ m and feature motion in any horizontal direction at a velocity of up to $v_{max} = 50$ m/s. These ranges were selected to accommodate the range expected for the majority of tropospheric clouds and winds. For simplicity, the calculations approximate the ground track as being along the SOM x-axis. The inaccuracies associated with this approximation are readily overcome by selecting a sufficiently large search area. The minimum and maximum SOM x and SOM y disparities (Δx , Δy) are approximated by:

$$\Delta x_{min} = h_{min}(\tan \theta_{cmp} - \tan \theta_{ref}) - |v_{max}(t_{cmp} - t_{ref})| \quad (9)$$

$$\Delta x_{max} = h_{max}(\tan \theta_{cmp} - \tan \theta_{ref}) + |v_{max}(t_{cmp} - t_{ref})| \quad (10)$$

$$\Delta y_{min} = |v_{max}(t_{cmp} - t_{ref})| \quad (11)$$

$$\Delta y_{max} = |v_{max}(t_{cmp} - t_{ref})| \quad (12)$$

3.3.1.1.2 Image normalization and resampling

Prior to cost function computations, the 672 nm Red band DN values from MISR L1B2 for each camera are each independently resampled and normalized to produce per camera hierarchies of normalized images sampled at 275 m, 550 m, and 1100 m resolution. Resampling is performed first by averaging the values within the 2×2 and 4×4 grids of collocated 275 m pixels to obtain coarser 550 m and 1100 m resolution pixels, respectively. Normalization is next performed in order to facilitate magnitude invariant pattern comparison. For each image, the normalized values at each pixel location (i,j) are determined by taking the difference of the value at the pixel location relative to the mean value of the neighboring pixels, then scaling this difference by the standard deviation of the values of the neighboring pixels. That is:

$$I_{normalized}(i,j) = \frac{I(i,j) - I_{mean}(i,j)}{I_{deviation}(i,j)} \quad (13)$$

The mean and standard deviation of neighboring values are calculated from the 0th, 1st, and 2nd moments of a Gaussian distribution function, i.e.:

$$I_{mean}(i,j) = \frac{S_1(i,j)}{S_0(i,j)} \quad I_{deviation}(i,j) = \frac{\sqrt{S_0(i,j)S_2(i,j) - S_1(i,j)^2}}{S_0(i,j)} \quad (14)$$

where each nth moment, S_n , is a convolution of image pixel values at relative disparities $(\Delta i, \Delta j)$ governed by parameter, σ :

$$S_n(i,j) = \sum_{\Delta i, \Delta j} I(i,j)^n e^{-\frac{\Delta x^2 + \Delta y^2}{2\sigma^2}} \quad (15)$$

The magnitude of σ varies with sampling resolution. At 275 m resolution, $\sigma=4.2$; at 550 m resolution, $\sigma=2.1$; and at 1100 m resolution, $\sigma=1.05$.

The Radiometric Data Quality Indicator (RDQI) is used to screen 275 m resolution L1B2 pixels for which measurement is not within quality specifications [M-5]. Not only at swath boundaries, but also in the rare circumstance of measurement difficulties, the local neighborhood employed during the above convolutions would include pixels for which no valid data is present. Such pixels are simply excluded. Only a single input pixel of valid data is necessary to yield a resampled and normalized output pixel.

3.3.1.1.3 Discrete correspondence

Discrete correspondence determines, for each valid reference image pixel, the location of the

SAD cost function minimum for the comparison image pixels within a specified search area. If less than half of the search area has valid comparison image data, the algorithm returns a no-retrieval. The SAD cost function $Cost(i,j,m,n)$ is calculated throughout the valid search area and a one pixel margin surrounding it's bounds, quantifying the pattern difference between reference pixel $I_{ref}(i,j)$ and comparison pixel $I_{cmp}(m,n)$ as:

$$Cost(i,j;m,n) = \frac{\sum_{\Delta i, \Delta j} |I_{ref}(i + \Delta i, j + \Delta j) - I_{cmp}(m + \Delta i, n + \Delta j)|}{\sum_{\Delta i, \Delta j} 1} \quad (16)$$

The range of $(\Delta i, \Delta j)$ values excludes locations where either the reference or comparison image data is not valid and varies from a 7×7 pixel window (i.e., $|\Delta i| \leq 3$; $|\Delta j| \leq 3$) for 1100 m sampled input images, to a 13×13 pixel window for 550 m input, to a 25×25 pixel window for 275 m input. For each reference coordinate, (i,j) , the value of (m,n) that minimizes $Cost(i,j;m,n)$ is then taken as a match. If that (m,n) value is located in the one pixel margin outside the boundaries of the search window, the algorithm returns a non-retrieval. Otherwise the values $(i,j;m,n)$ represent a conjugate pair calculated at the discrete precision of the input image.

3.3.1.1.4 Subpixel Refinement

Subpixel refinement recomputes the location of the discrete cost function minimum by interpolating the cost function (and the location of its minimum) from cost values calculated at the nine pixel coordinates in a 3×3 window centered on the location of the discrete minima. A prerequisite for this approach is that the interpolated cost function has a local minimum, which requires a positive partial second derivative with respect to both the SOM-x and SOM-y pixel coordinates. For the cost function, $C_{m,n}$, at discrete comparison coordinates (m,n) , the interpolated partial second derivatives are estimated by:

$$\frac{\partial^2 C_{m,n}}{\partial m} = \frac{C_{m-1,n} - 2C_{m,n} + C_{m+1,n}}{2} + \frac{C_{m-1,n-1} - 2C_{m,n-1} + C_{m+1,n-1}}{4} + \frac{C_{m-1,n+1} - 2C_{m,n+1} + C_{m+1,n+1}}{4} \quad (17)$$

$$\frac{\partial^2 C_{m,n}}{\partial n} = \frac{C_{m,n-1} - 2C_{m,n} + C_{m,n+1}}{2} + \frac{C_{m-1,n-1} - 2C_{m-1,n} + C_{m-1,n+1}}{4} + \frac{C_{m+1,n-1} - 2C_{m+1,n} + C_{m+1,n+1}}{4} \quad (18)$$

When $\frac{\partial^2 C_{m,n}}{\partial m} \leq 0$ or $\frac{\partial^2 C_{m,n}}{\partial n} \leq 0$, no local minimum can be found from the interpolation process. In this case, the discrete location of the minimum is maintained with no refinement. Otherwise, the adjusted location (r,s) of the cost minimum can be determined from the first and second partial derivatives:

$$r = m - \frac{\partial C_{m,n}}{\partial m} \bigg/ 2 \frac{\partial^2 C_{m,n}}{\partial m^2} \quad s = n - \frac{\partial C_{m,n}}{\partial n} \bigg/ 2 \frac{\partial^2 C_{m,n}}{\partial n^2} \quad (19)$$

where:

$$\frac{\partial C_{m,n}}{\partial m} = \frac{C_{m+1,n} - C_{m-1,n}}{4} + \frac{C_{m+1,n-1} - C_{m-1,n-1}}{8} + \frac{C_{m+1,n+1} - C_{m-1,n+1}}{8} \quad (20)$$

$$\frac{\partial C_{m,n}}{\partial n} = \frac{C_{m,n+1} - C_{m,n-1}}{4} + \frac{C_{m-1,n+1} - C_{m-1,n-1}}{8} + \frac{C_{m+1,n+1} - C_{m+1,n-1}}{8} \quad (21)$$

3.3.1.2 Conjugate Cluster Analysis

A clustering algorithm is used to independently derive forward and aft 17.6 km resolution fields of modal conjugate triplets (An-Bx-Dx) from the 1.1 km resolution fields of forward and aft conjugate pairs (Bx-An and Bx-Dx) produced by HSAD, where Bx and Dx denote Bf and Df for forward and Ba and Da for aft. Clustering is accomplished by recursively shrinking a domain enclosing the most populated bins within a multi-dimensional histogram constructed from vectors associated with conjugates. The list of such vectors at each 17.6 km grid cell is prepared in two steps: (1) compose a list of Bx-An-Dx conjugates from all Bx coordinates for which both a Bx-An and a Bx-Dx conjugate are defined, and for which the An coordinate (from the Bx-An conjugate) is enclosed by the 17.6 km grid cell (2) construct list of vectors specifying the disparities between the Bx and Dx coordinates and associated An coordinates of each Bx-An-Dx conjugate. Each resultant disparity vector comprises four dimensions representing the SOM x and y coordinates of two disparities. During clustering, histograms consisting of 7^4 bins segmenting each of the vector dimensions into 7 equally spaced intervals are iteratively produced. With each iteration, the domain spanned by those intervals in each dimension is re-centered with respect to a new modal disparity vector estimate, and reduced in size to hone in on the dominant signal. One such iteration is illustrated for a two-dimensional cross-section in Figure 10.

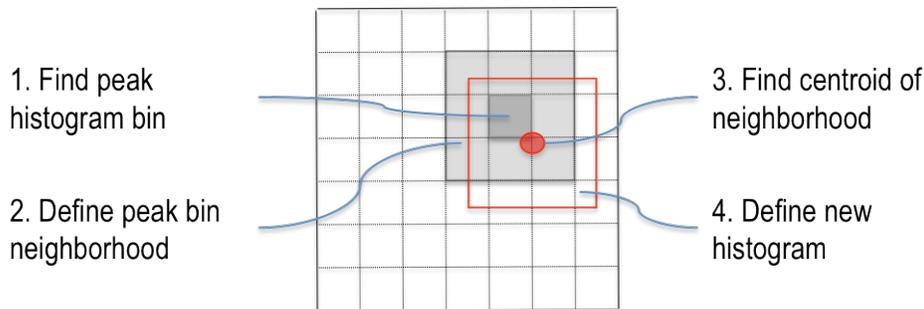


Figure 10. Two-dimensional cross-section of histogram refinement iteration

In the first iteration: the modal disparity vector is estimated as the mean of the minimum and maximum disparity vector values in each dimension; and the interval lengths are given by the greater of 275 m and 1/7 the difference between the maximum and minimum in each dimension. In subsequent iterations, the modal disparity vector is estimated as the centroid of vectors in the most populated and adjacent bins of the prior histogram; and the interval lengths are given by the greater of 275 m and 3/7 the previous interval length. The size of each histogram domain is shrunk in each dimension by 3/7 until 275 m interval lengths are obtained. Iterative clustering exits successfully if and when the interval lengths for all dimensions are 275 m, and exits unsuccessfully if and when the histogram domain ceases to enclose at least 3 vectors. Successful clustering yields a subset of conjugates associated with the modal conjugate vector. The centroid of conjugates' coordinates from this subset is taken as the final modal conjugate vector.

3.3.1.3 M23 Correspondence

The M23 stereo algorithm operates on the An-Af and An-Aa pairs of MISR L1B2 camera images, iterating over all valid 1100 m SOM grid image coordinates of the reference camera (i.e. An) image, yielding either a non-retrieval value or a matching comparison camera image coordinate.

Two cost metrics are potentially employed to find each match. First a match associated with the M2 metric is sought. If this metric does not resolve a conclusive match, then a match relative to the M3 metric is alternatively sought. This approach yields the enhanced coverage and robustness of M3 relative to M2, while mitigating the increased computational burden. Each metric is applied only where contrast is deemed sufficient (§ 3.3.1.3.1). An ambiguity test is used to determine whether a given M2 or M3 match is conclusive (§3.3.1.3.4). The conjugates returned by M2 and M3 are also subject to the same subpixel refinement as HSAD (§3.3.1.1.4).

3.3.1.3.1 Low contrast screening

M23 correspondence for stereo retrieval is only applied to image regions where contrast is determined to be sufficient to yield robust results (unlike HSAD correspondence which is applied to all image regions unconditionally). Contrast sufficiency is quantified as a factor relative to the estimated contrast resolvable by the MISR instrument. Specifically, the contrast sufficiency mask is given by:

$$q = \frac{\sigma_{M2} \cdot SNR_{1 \times 1}[\rho_{\langle R \rangle}] \cdot (R_{\max} - R_{\min})}{N_{\text{pts}} \cdot \langle R \rangle} \geq q_{\text{thresh}} \quad (22)$$

where q_{thresh} is a configurable threshold and $SNR_{1 \times 1}[\rho_{\langle R \rangle}]$ is the instrument signal-to-noise ratio for unaveraged (1×1) data, evaluated at the equivalent reflectance $\rho_{\langle R \rangle}$ that corresponds to $\langle R \rangle$,

the mean BRF. The value of σ_{M2} is given by:

$$\sigma_{M2} = \sum_{i,j} \left\| \left[\frac{R(x_i, x_j) - \langle R \rangle}{R_{\max} - R_{\min}} \right] \right\| \quad (23)$$

We note that $\rho_{\langle R \rangle} = \mu_0 \langle R \rangle$, so we obtain from the Ancillary Radiometric Product the table of signal-to-noise ratios, SNR_{am} , for the tabulated set of equivalent reflectance values for the appropriate camera and the red band and averaging mode $\text{am} = 1 \times 1$, and linearly interpolate these to the required value of $\rho_{\langle R \rangle}$. The expression in Eq. 22 implies that if the variation in BRF within the target patch is solely due to instrument radiometric random noise (i.e., $R(x_i, y_j) \cong \langle R \rangle + \text{noise}$) and not actual scene contrast, then $q \approx 1$. We require the contrast to be at least a factor q_{thresh} larger than the noise so we set $q_{\text{thresh}} = 1$. Experimentation has shown that this is a good compromise between increased coverage and runtime, as well as removing the poor quality data associated with the very-low-contrast retrievals.

3.3.1.3.2 Multipoint Matcher (M2)

For a stereo image pair, M2 takes the target patch in the reference image and a set of comparison patches within a search window in the comparison image and computes a matching metric. This metric is computed by taking all the BRF values in each patch, subtracting the mean BRF within the patch from each pixel, then normalizing by the difference in the maximum and minimum BRFs. Then, the absolute difference between these values in the target patch and the corresponding values in the comparison patch, summed over the area of the patches and normalized by an uncertainty estimate, is tested against a threshold. The M2 metric is defined as follows:

$$S_{M2} = \frac{\sum_{i,j} \left\| \left[\frac{R(x_i, y_j) - \langle R \rangle}{R_{\max} - R_{\min}} \right] \right\| - \sum_{i,j} \left\| \left[\frac{C(x_i, y_j) - \langle C \rangle}{C_{\max} - C_{\min}} \right] \right\|}{\sigma_{M2}} \quad (24)$$

where $R(x_i, y_j)$ is the BRF in the reference image pixel value at (i, j) , and $C(x_i, y_j)$ is the corresponding BRF in the comparison image. In order to apply this matcher we require that every value of $R(x_i, y_j)$ and $C(x_i, y_j)$ within their respective patches be associated with a Radiometric Data Quality Indicator $(\text{RDQI}) \leq \text{RDQI}_3$. We set $\text{RDQI}_3 = 1$. Thus, a single ‘‘bad’’ pixel in either patch will cause matching to be skipped. This is done in the interests of minimizing computational time. In the notation used in Eqs. 23 and 24, i and j are relative indices within the patches for summation, but the absolute values of the pixel coordinates within the reference and comparison images are not necessarily the same. In Eq. 24, $\langle R \rangle$ and $\langle C \rangle$ are the

average BRF values within the target and comparison patches, respectively; R_{max} and R_{min} are the maximum and minimum BRF values within the target patch, respectively; and C_{max} and C_{min} are the maximum and minimum BRF values within the comparison patch, respectively. The quantity, σ_{M2} , by which Eq. 24 is normalized, is an estimate of the uncertainty in the numerator of Eq. 24, and is defined above. For later use, we also define

$$N_{pts} = \sum_{i,j} 1 \quad (25)$$

i.e., it is the number of pixels in the patch.

We first note that the uncertainties in the terms in the numerator of Eq. 23 are independent of the absolute radiometric accuracy of the instrument, the band-to-band uncertainty (because only the red band data are used), or the camera-to-camera uncertainty (because any unknown scaling error cancels out as a result of the form of the expression). Additionally, we do not include systematic pixel-to-pixel uncertainties as the patches are localized with the cameras' fields-of-view. Thus the principal contributors to σ_{M2} are random pixel-to-pixel uncertainties and the effects of any subpixel misregistration between the patches, where the latter are expected to be more significant due to the high signal-to-noise ratio of the MISR cameras. We estimate the value of σ_{M2} by setting it equal to the value that the numerator of Eq. 23 would take if a target patch were being matched to itself, but with a misregistration in the x or y directions. We assume that the mean, maximum, and minimum BRF of the misregistered patch is the same as the target patch, and that the value of the misregistered patch BRF, $R_{misreg}(x_i, y_j)$, is estimated as

$$R_{misreg}(x_i, y_j) = \frac{R(x_{i-1}, y_j) + R(x_{i+1}, y_j) + R(x_i, y_{j-1}) + R(x_i, y_{j+1})}{4} \quad (26)$$

i.e., the average of the BRFs at the four nearest locations, since we do not know in which direction the hypothetical misregistration occurs. We then make the approximation that this average may be estimated by the average BRF in the whole patch, $\langle R \rangle$. Using these arguments, we derive

$$\sigma_{M2} = \sum_{i,j} \left| \left| \frac{R(x_i, y_j) - \langle R \rangle}{R_{max} - R_{min}} \right| \right| \quad (27)$$

Note that for a patch which is nearly uniform in BRF, the effect of misregistration is not so significant and σ_{M2} merely provides an estimate of the random noise on the data. As the contrast in the patch increases, σ_{M2} increases in magnitude, since the effect of misregistration becomes more significant for patches with spatially varying BRF. The inclusion of σ_{M2} as a normalization factor in Eq. 10 means that a value of S_{M2} on the order of unity corresponds to a match to within

the capability of the instrument to detect a difference in the patches.

For a given target, the candidate x and y values of disparity are those for which S_{M2} is smaller than or equal to a threshold T_{M2} , which is established to discard matching blunders. If $S_{M2} > T_{M2}$, the disparity is deemed an unsuccessful match and is discarded.

3.3.1.3.3 Multipoint matcher using Medians (M3)

M3 works similarly to M2, but uses medians rather than means, and applies a normalization scheme that yields better coverage than M2, with minor loss of accuracy. M3 is applied in a similar manner as M2, except that the matching metric is given by:

$$S_{M3} = \frac{\text{median}_{i,j} \left\{ \left| \frac{R(x_i, y_j)}{\text{median}\{R\}} - \frac{C(x_i, y_j)}{\text{median}\{C\}} \right| \right\}}{\sigma_{M3}} \quad (28)$$

where the inclusion of σ_{M3} follows the same rationale as with M2 but is defined differently, i.e., the BRF values in each patch are normalized by dividing by the median value of each patch. A distribution of the absolute values of the differences in the corresponding pixels between the patches is then formed (i.e., the absolute value of the difference of the top right hand pixels in each patch, etc.), and the final metric is the median of this distribution, normalized by σ_{M3} as shown in Eq. 28.

The derivation of σ_{M3} uses similar methodology as in the derivation of σ_{M2} , with the exception that instead of estimating the value in Eq. 29 using the mean BRF in the patch, we use the median. With this assumption, we obtain

$$\sigma_{M3} = \text{median}_{i,j} \left\{ \left| \frac{R(x_i, y_j)}{\text{median}\{R\}} - 1 \right| \right\} \quad (29)$$

Again, for a given target patch, the candidate x and y values of disparity associated with the M3 test are those for which S_{M3} is smaller than or equal to a threshold T_{M3} . A different threshold value may be used for M3 than is used in M2. However, as with M2, we consider a successful matching to have occurred if the application of the threshold results in only a single (i.e., unique) disparity, or if the ambiguity test (3.3.1.3.4) allows a single selection to be chosen from multiple disparities, in which case the x and y values of disparity associated with the minimum value of S_{M3} are chosen.

We note that if $\text{median}(R) = 0$ or $\text{median}(C) = 0$, the M3 matcher cannot be applied.

3.3.1.3.4 Elimination of Ambiguous Matches

Ideally, the application of the threshold test will eliminate all candidate disparities except

one. If a unique disparity (x and y value) is obtained, the match is considered successful. The match is considered to have failed if either (a) all candidate disparities are eliminated by the threshold test, or (b) multiple candidate disparities remain following application of the threshold test, and an unambiguous “best” answer cannot be chosen.

The ambiguity test is employed for both the M2 and M3 matchers. Its specifics are illustrated below for the M2 matcher, with M3 following the same paradigm. If multiple matches from M2 satisfy the threshold criterion, take the “best” match to be the one that minimizes S_{M2} . Let $S_{M2,min}$ be the corresponding value of the matcher metric. Now, set a secondary threshold A_{M2} , given by, $A_{M2}=f*S_{M2}$ where f is a value slightly larger than unity. We set $f= 1.1$. If the best match is the only one for which $S_{M2} < A_{M2}$, it is considered an unambiguous successful match. If there are still multiple matches below the threshold A_{M2} , we apply a spatial cluster test to these matches. To implement the cluster test, we compute two parameters: s_x and s_y , the maximum absolute value of the x and y distances between all conjugates that passed both the primary ($S_{M2} \leq T_{M2}$) and secondary threshold tests. If s_x is less than a threshold value p_x , and simultaneously s_y is less than a threshold value p_y , the disparities are said to be spatially clustered and the matcher result is therefore ambiguous. At present, we set $p_x=p_y=3$. If the disparities satisfy the clustering criteria, the best match is taken as the successful result. Otherwise, we deem the results as ambiguous and consider there to be no successful match. The ambiguity test is used after application of the “previous-match” method and at all but the last level (highest resolution) of the pyramid scheme.

3.3.1.3.5 Image Pyramid and Previous Match Methods for M23

In order to reduce the computational time required by application of the M23 matchers, we employ two algorithms (ImagePyramid and PreviousMatch) to allow us to only search a subset of the entire search window. See Figure 11 for a flowchart of this process.

The PreviousMatch algorithm restricts its search to a small window around the results from the pixel one row up, and one column to the left. Since matching proceeds from left to right and up to down, both of these given pixels are guaranteed to have already been matched. To form the search window for pixel (x,y) the algorithm looks for a successful match (with a tighter than normal metric value) for pixels $(x\text{-step}, y)$ and $(x, y\text{-step})$ where x and y are the row and column numbers (counting from the top-left) and step refers to the step-size of the matcher. Since we only attempt matching on each 1.1 km grid-center yet use the 275 m red-band data as input, the step-size is 4 at the bottom level of the pyramid. In other words, the algorithm steps to every fourth pixel when it is cycling through the reference points. At the top level of the pyramid, the step-size is 2 since we are still matching every 1.1 km grid center, but the input imagery is provided at 550 m resolution.

To be considered as a seed-point for the PreviousMatch algorithm, a retrieval needs to pass both the ambiguity test (3.3.1.3.4) and also have a metric value less than or equal to half the usual threshold (T_{M2} or T_{M3} depending on the identity of the matcher used). If a retrieval meets these criteria, the search window is set to a small 5x5 square centered on the disparity of the seed point. This process is repeated for both the $(x\text{-step}, y)$ and $(x, y\text{-step})$ locations and the final search window is the union of the two squares.

If at least one valid seed-point exists, matching proceeds as usual on the smaller search window and if a result that meets both the threshold and ambiguity criteria is found it is accepted without need to check the rest of the full window.

The ImagePyramid method is an implementation of a standard hierarchical matcher which again decreases computational time by reducing the size of the search window. M2 and M3 use a two-level pyramid in contrast to HSAD which employs 3 levels. The first step in the process is down-sampling the reference and comparison images down to half the original resolution (275 m down to 550 m) in both the x and y directions. The search window is also reduced by a factor of 2 in each direction so the area covered by the search window remains the same, but the number of pixels needed to cover that area is a factor of 4 smaller.

After the degraded images have been generated, matching proceeds as previously described in with the same threshold values and ambiguity test. The PreviousMatch method is attempted first, only proceeding to searching the full (at 550 m resolution) search window in case PreviousMatch fails.

After matching has completed for the top-level of the pyramid, we fall down to the full resolution data, but only with a reduced search window that is calculated using the following algorithm. The disparity field is upsampled to 275 m resolution by first multiplying all the disparities by 2 (to convert back to 275 m sampling) and then copying each data point to a 2×2 area in order to form a complete set of results at the desired resolution. The original disparities were expressed in terms of 550 m pixels, so the multiplication-by-2 is needed to ensure that the disparity (as measured in meters) remains constant. Then, for each 1.1 km grid center we look at the value of the retrieved disparity retrieved from the top level of the pyramid. If no valid retrieval from the previous matching invocation exists, this pixel is skipped entirely. Therefore, we never search the full search window at 275 m resolution. If a valid match is found, a small 5×5 square is drawn around that result and then we run M2 followed by M3. The ambiguity test is not applied at the bottom level of the pyramid.

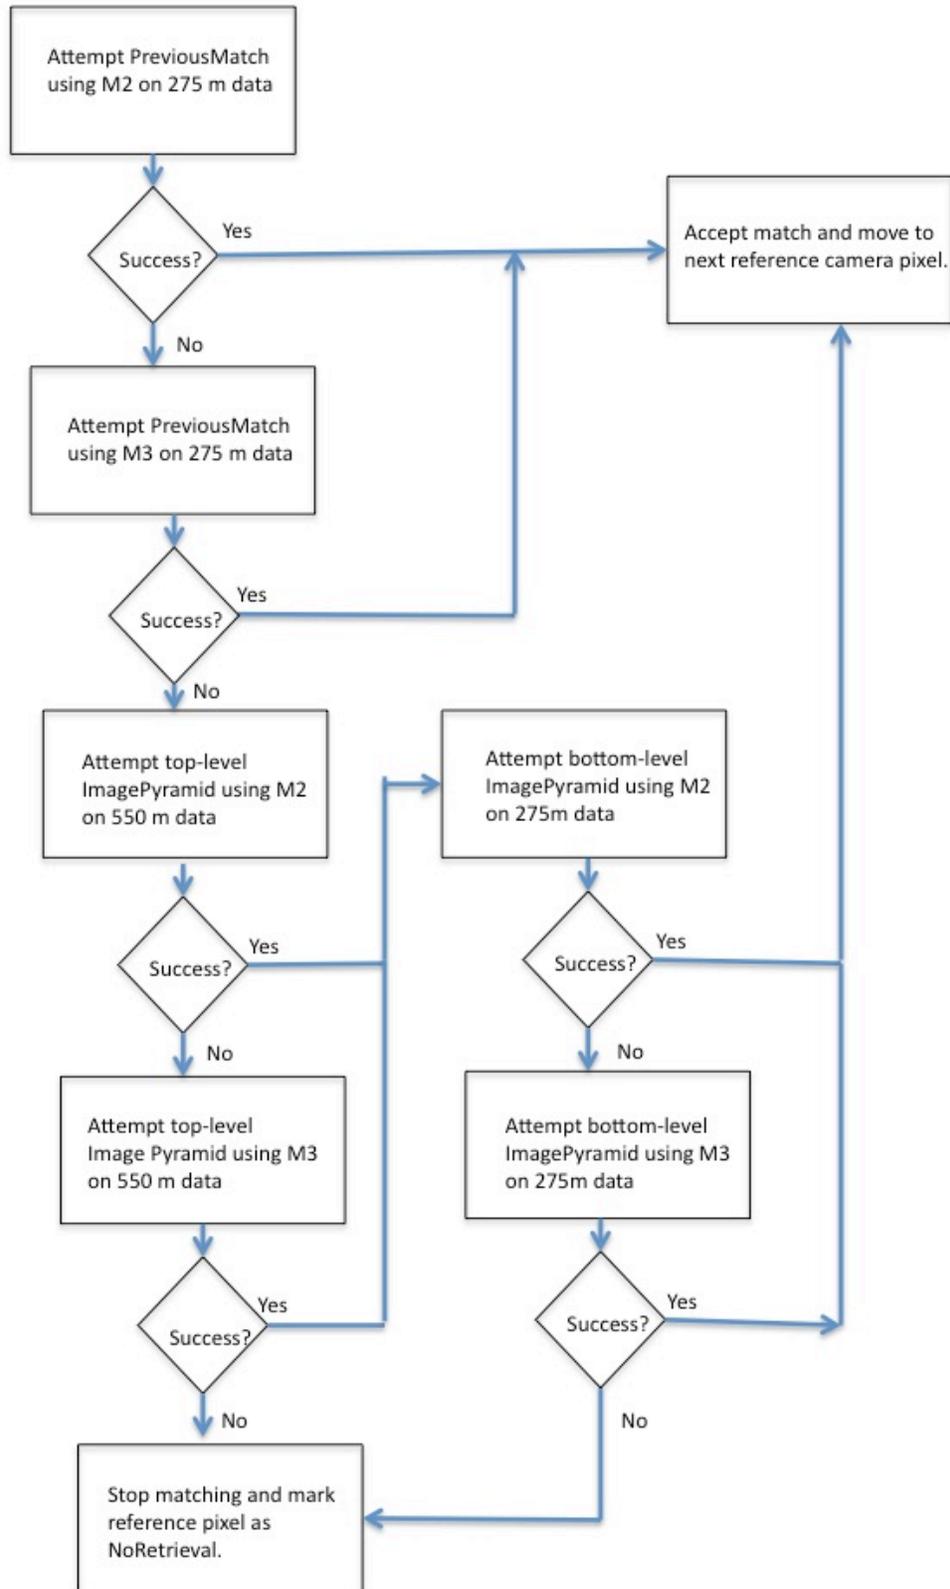


Figure 11. Flowchart for M2/M3 Processing

3.3.2 Reconstruction

Sections 2.3.1 and 2.3.2 review the principles by which feature height and cross-track motion are reconstructed from conjugate pairs and by which height resolved motion vectors are reconstructed from conjugate triplets. Both of these approaches to reconstruction (further detailed in §3.3.2.4 and §3.3.2.3, respectively) first require determining the precise time, camera position, and look vector in Earth Centered (Conventional Terrestrial) Reference (ECR) coordinates associated with each conjugate camera's Ellipsoid-projected SOM image coordinates (§3.3.2.1). Bias associated with the determination of imaging times associated with image coordinates, and stemming from lens aberrations, is mitigated by first applying a fixed set of position dependent corrections to conjugate image coordinates (§3.3.2.2).

3.3.2.1 Determining time and look vectors from L1B2 image coordinates

The imaging time and look vectors corresponding to the L1B2 image coordinates for each camera of a conjugate pair or triplet are required to reconstruct feature position and movement. Imaging times, are used in conjunction with spacecraft navigation data, solar ephemeris information, and the MISR camera model to calculate the vectors in Earth Centered Reference (ECR) coordinates.

Imaging times for L1B2 image grid points are derived from transform domain information specified per half-block of radiance data. Each domain specifies: a reference imaging time, t_0 ; a reference SOM-projected grid point $(i_0, j_0)_{som}$; a reference line and sample (l_0, s_0) in MISR red band imagery; and transform coefficients. The transform coefficients facilitate the mapping of any SOM-projected grid point, $(i, j)_{som}$, within the domain to a corresponding line and sample (l, s) [M-3]. The imaging time of that SOM grid point is then given by:

$$t_{i,j} = t_0 + l_{\Delta t}(l - l_0) \quad (30)$$

where $l_{\Delta t} = 0.0408$ seconds is the image line repeat time.

Given the imaging time of each SOM grid point, the spacecraft position $P_{spacecraft}(x, y, z)$ in ECR coordinates is determined for that time using the spacecraft ephemeris provided by Product Generation System (PGS) Toolkit functions. The look vector corresponding to the (ellipsoid referenced) SOM grid point is then determined as the normalized vector pointing from the spacecraft location $P_{spacecraft}(x, y, z)$ to the SOM grid location in ECR coordinates, $P_{som}(x, y, z)$.

3.3.2.2 Camera Model Corrections

The transform coefficients from which imaging times and then look vectors are derived

(as described in §3.3.2.1) exhibit biases that are mitigated by applying corrections to the conjugate image coordinates themselves prior to this derivation. These biases vary as a function of cross-swath position and stem from lens aberrations that are not accounted for during L1B2 processing. Bias magnitudes have been estimated by measuring the magnitude of disparities between the Terrain referenced coordinates of terrain features observed by all nine MISR cameras in scenes with no clouds. As detailed in §3.3.2.2.1, these biases are estimated relative to the An camera for each camera other than An and for each 17.6 km interval within the SOM-y coordinate space to which L1B2 data is projected. The per camera biases associated with each conjugate set are found by mapping the An camera SOM-y coordinate to an associated 17.6 km bias domain. An equal and opposite correction to each conjugate coordinate is applied for each of the biases.

3.3.2.2.1 Determining Camera Model Corrections

The Camera Geometric Model parameters implemented in support of L1B2 Georectification and Registration processing do not fully account for optical distortions of the system [M-3]. So called pixel-theta coefficients are applied to mitigate distortions in the direction across field angle (i.e. across-swath). However, the remaining distortion in the along-track direction has not been modeled.

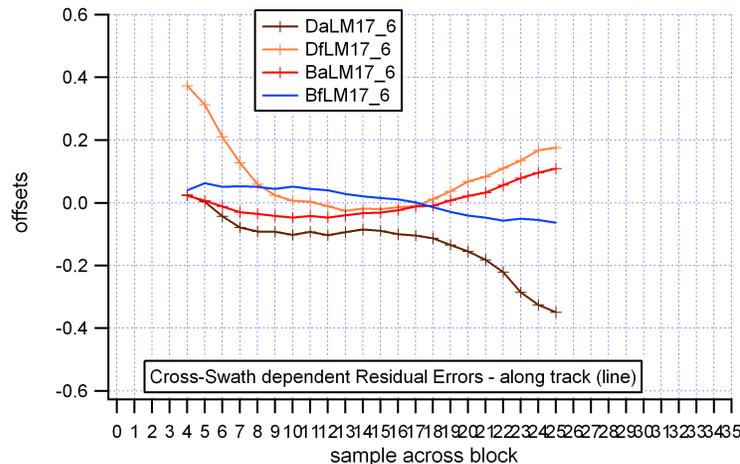


Figure 10 – Outstanding along track co-registration errors between An (nadir) and Da, Df, Ba, and Bf cameras. It has been measured using a validation tool based on the image matching over cloud free regions. Reported on 17.6 km grid sampling across field of view (i.e. across-swath) in pixel units.

The outstanding error is also field angle dependent with a magnitude of between +/- 0.4 of pixel for most oblique Da and Df camera (see Figure 10). Given that the accuracy of reconstructed CMVs are particularly sensitive to the co-registration error in this direction, an ancillary input containing corrections to mitigate the errors has been created and used as the part of processing.

PGE21, the software currently used as part of standard processing and validation, has been modified so that required corrections can be generated. Nominally, PGE 21 is used to measure co-registration between An and other cameras such that pairs of numbers are produced representing co-registration error (along-track and across-swath) not for every 70.4 km of orbit extent (in the flight direction) but for the entire field of view i.e. whole orbit swath. As such, PGE21 is used to measure and monitor dynamic camera pointing changes resulting in the co-registration changes. Co-registration offsets are determined by image matching of the points between An and other cameras in the terrain projected imagery. Prominent features are identified, then subjected to cross-correlation matching to determine initial conjugates. These initial conjugates are refined by a higher accuracy, sub-pixel, least-square matching approach. The values of the measured offset for each conjugate reflect pointing error, except when the matching is spurious or the feature matched is not associated with terrain. In order to deal with spurious matches, features that have not been successfully matched in all eight cameras are screened. Matches over cloudy surfaces are identified and screened using the distribution of offsets across all eight cameras other than An. If after filtering a significant number of 13 or more points per 70.4 km orbit extent are still present, mean corrections along with standard deviations for that orbit extent are computed.

Modifications of the PGE21 are focused on evaluating static pointing errors assumed to be due to the un-modeled camera optical distortions as opposed to the measurement of the dynamic pointing changes in the nominal version. Two main modifications are made. First, to pull out static error, statistical summaries are made over an extended time period of several months combined from the different time of the year. Second, these statistical summaries are broken down into 17.6 km across-swath regions to fully represent field of view behavior of the optical distortions. So, as a result ancillary information contains, for each of the Da, Df, Bf, Bf, Aa, and Af a string of numbers i.e. “image line corrections” each associated with a particular 17.6 km grid index in the SOM map projection space. There are 23 corrections for each of the cameras corresponding to the 23 central grid centers out of possible 32. The remaining 9 are on the edges of the potential ground footprint are unlikely to be observed by all nine MISR cameras.

In the particular implementation used for the first version of these ancillary data two sets of the 1000 orbits representing periods of Jan/Feb and Jun/July 2007 have been used to determine corrections. It should be noted that prior to accepting these data, tests using representative sample time periods throughout the mission time line show no time dependency on generated corrections. Also, tests in which certain regions of the globe have been excluded have been conducted to evaluate possible spatial dependency due to variable quality of the underlying projection surface, i.e. global DEM. These tests have not shown measurable influence of the DEM upon estimated bias. They did identify coastal regions where conditions favor low clouds moving across track. Such clouds are prone to being misidentified as terrain, influencing the estimation of pointing bias. These regions (see table) have therefore been removed from the input used to produce corrections.

3.3.2.3 Reconstruction of height resolved motion vectors from conjugate triplets

The reconstruction of feature position, height, and horizontal motion for conjugates of three MISR cameras is undertaken by solving a system of equations as described below.

- Knowns:
 - Ellipsoid projected position vectors, for each of 3 cameras: $\vec{X}_1, \vec{X}_2, \vec{X}_3$
 - Unit look vectors from ellipsoid to camera, for each of 3 cameras: $\hat{P}_1, \hat{P}_2, \hat{P}_3$
 - Precise times of observation for each of n cameras: t_1, t_2, t_3
 - Reference camera: r
 - WGS84 ellipsoid flattening factor: $f=0.00335281$
- Unknowns:
 - Distance along look vector from ellipsoid to feature position at time of observation by each of n cameras: d_1, d_2, d_3
 - Motion vector of feature: \vec{V}

The position of a feature when observed by camera, c , relative to reference camera, r , is then given by:

$$\vec{X}_c + d_c \hat{P}_c = \vec{X}_r + d_r \hat{P}_r + (t_c - t_r) \vec{V} \quad (31)$$

The second and fourth terms in the above equation represent the vector from the Ellipsoid projected position to the true position of the feature as observed by each camera. The fifth term represent the displacement of the feature during the interval between camera views. Limiting \vec{V} to horizontal motion yields the following additional constraint:

$$\left[\begin{array}{ccc} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & (1-f)^{-2} \end{array} \right] \cdot (\vec{X}_r + d_r \hat{P}_r) \cdot \vec{V} = 0 \quad (32)$$

Equations **31** and **32** comprise an overconstrained non-linear system of ten equations in six unknowns, whose solution can be iteratively refined from an initial estimate of d_r by prescribing that estimate as a constant in equation **32** and then calculating an improved estimate of d_r from the least squares solution to equation **31**. In practice, the algorithm has been found to require only two calculations of d_r following an initial estimate of $d_r = 0$. (Convergence such that further iterative estimates of d_r differ by less than 0.01 meters is immediate because $\vec{X}_r \gg d_r$.)

3.3.2.4 Reconstruction of stereo vectors from conjugate pairs

The reconstruction of feature position, height, and cross-track motion for conjugates of two MISR cameras is undertaken by solving a system of equations as described below.

- Knowns:
 - Ellipsoid-projected position vectors, for a reference and comparison camera: \vec{X}_r, \vec{X}_c

- Unit look vectors from ellipsoid to camera, for a reference and comparison camera: \hat{P}_r, \hat{P}_c
- Precise times of observation reference and comparison camera: t_r, t_c
- WGS84 ellipsoid flattening factor: $f=0.00335281$
- Unknowns:
 - Distance along look vector from ellipsoid to feature position at time of observation by reference and comparison camera: d_r, d_c
 - Unit cross-track motion vector motion, \hat{V}_{CT}
 - Magnitude of cross-track motion vector, m

Solving the above system requires defining and calculating \hat{V}_{CT} for each conjugate. \hat{V}_{CT} is determined by two constraints. The first constraint is achieved by defining a *pseudo-epipolar* normal associated with apparent feature displacement that cannot be due to parallax. This normal is defined by the two camera look vectors:

$$\hat{N}_{EPIPOLAR} = \frac{\hat{P}_r \times \hat{P}_c}{\|\hat{P}_r \times \hat{P}_c\|} \quad (33)$$

The second constraint is that \hat{V}_{CT} comprise only horizontal motion. Horizontal motion is defined as motion orthogonal to the Ellipsoid normal at the feature position, where the ellipsoid normal is given by:

$$\hat{N}_{ELLIPSOID} = \frac{\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & (1-f)^{-2} \end{bmatrix} \cdot (\vec{X}_r + d_r \hat{P}_r)}{\left\| \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & (1-f)^{-2} \end{bmatrix} \cdot (\vec{X}_r + d_r \hat{P}_r) \right\|} \quad (34)$$

\hat{V}_{CT} is then given by the component of $\hat{N}_{EPIPOLAR}$ directed along the ellipsoid:

$$\hat{V}_{CT} = \frac{\hat{N}_{ELLIPSOID} \times (\hat{N}_{EPIPOLAR} \times \hat{N}_{ELLIPSOID})}{\|\hat{N}_{ELLIPSOID} \times (\hat{N}_{EPIPOLAR} \times \hat{N}_{ELLIPSOID})\|} \quad (35)$$

This yields the following equation for feature position observed at t_0 relative to that at t_1 :

$$\vec{X}_r + d_r \hat{P}_r + m(t_c - t_r) \hat{V}_{CT} = \vec{X}_c + d_c \hat{P}_c \quad (36)$$

Equations 34, 35, and 36 express a non-linear system of three equations in three unknowns,

whose solution can be iteratively refined from an initial estimate of d_r by prescribing that estimate as a constant in equation 34 and then calculating an improved estimate of d_r from the solution to equation 36. In practice, the algorithm requires only two calculations of d_r following an initial estimate of $d_r = 0$. (Convergence such that further iterative estimates of d_r differ by less than 0.01 m is immediate because $\vec{X}_r \gg d_r$.)

3.3.2.5 Stereo height vector sensitivity to along-track motion

During stereo reconstruction, the height sensitivity to along-track motion (introduced in §2.3.1) is calculated precisely as a function of the look vectors and observation times.

The precise view zenith angle, θ , of each look vector is given by:

$$\theta = \arccos(\hat{P} \cdot \hat{N}_{ELLIPSOID}) \quad (37)$$

As first introduced in equation 2, the height sensitivity is then trivially determined by the reference and comparison viewing zenith angles, θ_r and θ_c , and times. That is:

$$h_s = \frac{t_r - t_c}{\tan \theta_r - \tan \theta_c} \quad (38)$$

3.3.2.6 Instrument heading of height-resolved motion vectors and stereo height vectors

During reconstruction, an instrument heading is calculated for each height-resolved motion vector and for each stereo height vector. For the height-resolved motion vector, this calculation facilitates subsequent quality control procedures by allowing the motion to be decomposed into an along-track and cross-track component. For the stereo height vector, the related cross-track heading output parameter is equal to the instrument heading minus 90 degrees.

The instrument heading is calculated from the geometry of the look vectors used for reconstruction. The heading is specifically derived from an orientation vector associated with the intersection of a plane tangent to the ellipsoid at the location of retrieval (normal given by $\hat{N}_{ELLIPSOID}$) and a pseudo-epipolar plane defined by the look vectors (normal given by $\hat{N}_{EPIPOLAR}$). A pseudo-epipolar normal is given by the cross-product of any pair of look vectors (see equation 33 of §3.3.2.4). For stereo height vectors, there is only one pair of look vectors, from whose cross-product $\hat{N}_{EPIPOLAR}$ is derived. For height-resolved motion vectors, there are two pairs of look vectors, so $\hat{N}_{EPIPOLAR}$ is derived from the average of these cross-products:

$$\hat{N}_{EPIPOLAR} = \frac{\sum \hat{P}_r \times \hat{P}_c}{\|\sum \hat{P}_r \times \hat{P}_c\|} \quad (39)$$

The intersection of this pseudo-epipolar plane with a tangential plane to the ellipsoid,

$\hat{N}_{ELLIPSOID}$ (defined by equation 35), at the retrieval location defines an orientation vector, \vec{V}_{orient} :

$$\vec{V}_{orient} = \frac{\hat{N}_{EPIPOLAR} \times \hat{N}_{ELLIPSOID}}{\|\hat{N}_{EPIPOLAR} \times \hat{N}_{ELLIPSOID}\|} \quad (40)$$

Once translated from ECR coordinates into Local North coordinates, this orientation vector is readily decomposed into an eastward component, u_{orient} and a northward component, v_{orient} . The instrument heading, λ , is then given by:

$$\lambda = 90 - \frac{360}{\pi} \arctan \left(\frac{v_{orient}}{(\sqrt{u_{orient}^2 + v_{orient}^2}) + u_{orient}} \right) \quad (41)$$

Note that the along-track component, at , and cross-track component, ct , of a given motion vector are derived from the instrument heading, λ , eastward component, u , and northward component, v , as follows:

$$\begin{aligned} at &= u \sin \lambda + v \cos \lambda \\ ct &= -u \cos \lambda + v \sin \lambda \end{aligned} \quad (42)$$

3.3.3 Height Resolved Motion Vector Operations

The raw height resolved feature motion vectors reconstructed from conjugate triplets undergo flagging and quality control operations that ultimately yield the height resolved cloud motion vectors provided to end users. First, the forward and aft motion vector fields are merged into one representative vector field. During this process, each such vector is assigned a quality indicator and screened on the basis of agreement between collocated forward and aft vectors and between adjacent vectors (§3.3.3.1). Next, the statistics of forward and aft agreement throughout the orbit are assessed, with poor agreement resulting in all motion vectors for the orbit being screened from the product (§3.3.3.2). Last, each feature motion vector is assigned a Motion Derived Cloud Mask (MDCM) value labeling it as labeled as either cloud or near surface (i.e. likely not to be cloud) with either low or high confidence (§3.3.3.3)

3.3.3.1 Merge forward and aft and assess quality of height resolved motion vectors

A single height resolved motion vector with associated Quality Indicator (QI) value is synthesized at each grid cell from collocated preliminary vectors independently extracted from the forward and aft sets of camera images. One or both preliminary vectors may be unavailable due to unsuccessful retrieval or due to being discarded during quality assessment. The final vector ultimately consists of one of three possibilities: (1) the mean of the forward and aft vectors; (2) the forward or the aft vector; or (3) a non-retrieval. Choosing between these requires assessing the absolute value of the height difference between forward and aft motions, $\Delta H_{CMVFWD AFT}$; the vector magnitude difference between the forward and aft horizontal motion

components, $\Delta V_{CMVFWD AFT}$; and the minimum vector magnitude difference between motion in a grid cell relative to neighboring vectors of similar height, $\Delta V_{CMVNEIGHBOR}$. These differences also govern the assigned QI value, which consists of the mean of predefined transform functions applied to each difference.

The height and vector motion differences between forward and aft preliminary motion vectors are ultimately used to calculate QI, but are also used to identify *low confidence* and *high confidence* grid cells. Grid cells where both the forward and aft preliminary vectors have been retrieved are categorized as *low confidence* or *high confidence* on the basis of these difference magnitudes. *High confidence* grid cells are defined as those where the height and vector horizontal motion differences are within thresholds widely bracketing expected precision, specifically $H_{CMVFWD AFT} = 1000$ m and $V_{CMVFWD AFT}$ is 12 ms^{-1} . That is:

$$\Delta H_{FWD AFT} = |h_{fwd} - h_{aft}| < H_{CMVFWD AFT} \quad (43)$$

$$\Delta V_{FWD AFT} = \sqrt{(u_{fwd} - u_{aft})^2 + (v_{fwd} - v_{aft})^2} < V_{CMVFWD AFT} \quad (44)$$

Where forward and aft preliminary vectors are not both available, $\Delta H_{FWD AFT}$ is assigned a the value of $H_{CMVFWD AFT}$ and $\Delta V_{FWD AFT}$ is assigned the value of $V_{CMVFWD AFT}$.

The *neighbor motion difference* is governed by the minimum difference between each preliminary vector at a grid cell and all preliminary vectors of similar height retrieved from *high confidence* neighboring grid cells. Neighbors are defined as the other grid cells within a 3x3 surrounding window. Neighbor height similarity is governed by the threshold, $H_{CMVNEIGHBOR} = 500$ m. That is:

$$\Delta H_{NEIGHBOR} = |h - h_{neighbor}| < H_{CMVNEIGHBOR} \quad (45)$$

Where there is at least one surrounding preliminary vector of similar height, the preliminary neighbor difference, $\Delta V_{PRENEIGHBOR}$, is then given by:

$$\Delta V_{PRENEIGHBOR} = \min \left\{ \sqrt{(u_{pre} - u_{neighbor})^2 + (v_{pre} - v_{neighbor})^2} \right\} \quad (46)$$

Prior to computing the *neighbor motion difference*, the preliminary neighbor differences are used to identify and mask likely erroneous preliminary vectors associated with two error conditions. Masked preliminary vectors are treated as unavailable (as if retrieval had been unsuccessful) when computing $\Delta H_{FWD AFT}$, $\Delta V_{FWD AFT}$, and $\Delta V_{NEIGHBOR}$ and when synthesizing a final motion vector from the preliminary forward and aft vectors. The first error condition consists of preliminary vectors that agree with their neighbors in height, but not horizontal

motion. These are masked on the basis of $\Delta V_{PRENEIGHBOR}$ exceeding the threshold, $V_{CMVNEIGHBOR} = 12 \text{ ms}^{-1}$. The second error condition consists of previously characterized *low confidence* grid cells where preliminary forward and aft vectors have both been retrieved, but do not agree. For such grid cells, one or both preliminary vectors will be masked. To not be masked, a preliminary vector must agree with at least one neighboring vector in height, and must also not have a value of $\Delta V_{PRENEIGHBOR}$ exceeding $V_{CMVNEIGHBOR} = 12 \text{ ms}^{-1}$. If both preliminary vectors meet this requirement, the one with the greater value of $\Delta V_{PRENEIGHBOR}$ will be masked.

The final neighbor motion difference, $\Delta V_{NEIGHBORDIFF}$, is synthesized from the preliminary forward and aft motion neighbor differences. Its value is assigned the mean of the forward and aft neighbor differences, where both are available. If only one preliminary difference is available, that is assigned. Otherwise, a default value of $V_{CMVNEIGHBOR} = 12 \text{ ms}^{-1}$ is assigned.

The final QI value is calculated from the mean of three metrics, $M_{HFWDAFTDIFF}$, $M_{VFWDAFTDIFF}$, and $M_{NEIGHBORDIFF}$ that are respectively generated by applying transform functions to $\Delta H_{CMVFWDAFT}$, $\Delta V_{CMVFWDAFT}$, and $\Delta V_{CMVNEIGHBOR}$. The purpose of these transforms is twofold. First, they yield values that increase with expected retrieval quality over a range from 1 to 100. Second, the functions are normalized using coefficients whose values can be configured so as to obtain comparable expected distributions of QI values for each metric. That is, the probability that a given value is assigned to $M_{HFWDAFTDIFF}$ is intended to be the same as the probability that that value is assigned to $M_{VFWDAFTDIFF}$ or $M_{NEIGHBORDIFF}$. This was found to be true of test data spanning the year 2000 produced using the prescribed coefficients. The three transform functions are given by:

$$M_{HFWDAFT} = 100 - 100 \left(\tanh \left(\frac{\Delta H_{CMVFWDAFT}}{H_{CMVFWDAFT}} \right) \right)^{0.9} \quad (47)$$

$$M_{VFWDAFT} = 100 - 100 \left(\tanh \left(\frac{\Delta V_{CMVFWDAFT}}{V_{CMVFWDAFT}} \right) \right)^{1.0} \quad (48)$$

$$M_{NEIGHBOR} = 100 - 100 \left(\tanh \left(\frac{\Delta V_{CMVNEIGHBOR}}{V_{CMVNEIGHBOR}} \right) \right)^{0.9} \quad (49)$$

The final QI value is then:

$$QI = \text{mean}\{M_{NEIGHBOR}, M_{VFWDAFT}, M_{HFWDAFT}\} \quad (50)$$

After the final QI value has been computed, retrievals with a QI value less than $QI_{CMVTHRESH} = 25$ are set to fill values. Setting this QI threshold roughly corresponds to requiring that $\Delta H_{CMVFWDAFT} < H_{CMVFWDAFT}$, $\Delta V_{CMVFWDAFT} < V_{CMVFWDAFT}$, and $\Delta V_{CMVNEIGHBOR} < V_{CMVNEIGHBOR}$.

3.3.3.2 Assess quality and screen motion vectors per orbit

Since the derivation of the motion vectors is highly sensitive to the georegistration quality of the entire orbit, we calculate an “orbit_qa” flag based on the fwd-aft differences of the wind retrievals. If the orbit is deemed to be so poorly registered as to invalidate any motion vectors, we set all the winds and wind-corrected cloud-top heights and their associated parameters to NoRetrieval. There are two such orbit-qa flags: Orbit_QA and orbit_qa_winds with the former being calculated in the Level 1B2 processing and the latter being derived following the methodology described below. If either of these flags are set to -1.0 (their default value is 0.0), then we reset the data to NoRetrieval.

The flag is calculated by looking at the mean of the absolute values of the fwd-aft differences of the three components of the motion vectors (along-track, cross-track, and height). For each of these components, all the NoRetrieval and outlier values are screened out and the mean of all the remaining data is calculated and stored. If any of these mean values is larger than a given threshold then the orbit is called “bad” and all the motion vectors and wind-corrected heights are reset to NoRetrieval.

Because the fwd-aft differences sometimes contain very large numbers that affect the calculation of the mean, a prior filtering step is necessary to remove these values from the algorithm. Calculation of mean fwd-aft differences includes only the differences whose magnitude is less than ten times the relevant threshold.

Component	Orbit is poorly registered if mean of absolute values exceeding this threshold.	Screen out all individual fwd-aft differences whose absolute value exceeds this threshold.
Along-track	12.0 m/s	120.0 m/s
Cross-track	3.0 m/s	30.0 m/s
Height	990.0 m	9900.0 m

3.3.3.3 Generate motion derived cloud mask

The purpose of this stage is to produce the Motion Derived Cloud Mask (MDCM) that classifies motion vector retrievals as either *Cloud* or *Near Surface* with *Low Confidence* or *High Confidence*. *Cloud* retrievals are associated with the advection of cloud features and should be considered proxy observations of wind, while *Near Surface* retrievals are associated with static terrain or orographic clouds. The distinction between these two modes of retrieval is evident from bimodal distribution statistics of retrieved height and motion components (see §3.3.3.1). *Near Surface* retrievals are characterized by low speeds and heights at or near the terrain surface, whereas *Cloud* encompass a broad range of feature heights and speeds. The height and motion distributions are bimodal, but not entirely distinct, so discrimination based on height and motion cannot be perfect, hence the *High Confidence* and *Low Confidence* designations.

Categorization as either *Cloud* or *Near Surface* is applied independently to any successful forward or aft camera retrieved motion vector at each grid cell. Where both forward and aft are

present, and agree in their discrimination, the retrieval is categorized as *High Confidence*. It is otherwise categorized as *Low Confidence*. Each forward or aft motion vector is categorized as *Cloud* if and only if at least one of the following conditions is true: the retrieval altitude, h , is well above the terrain (Eq. 51), the retrieval has significant cross-track cloud motion, ct (Eq. 52), or the retrieval is over ocean and has significant along-track motion, at (Eq. 53). All other motion vectors are classified as *Near Surface*. The third criterion is applied only over ocean, due to the greater frequency of near surface CMVs associated with cloud advection, as opposed to over land, where the overwhelming majority of near surface CMVs appear to be associated with static terrain (see Figure 3).

$$h > H_{MDCM} + h_{terrain} + 2\sigma_{h,terrain} \quad (51)$$

$$|ct| > CT_{MDCM} \quad (52)$$

$$|at| > AT_{MDCM} \quad (53)$$

The terrain altitude, $h_{terrain}$, and terrain altitude variance, $\sigma_{h,terrain}$, and whether the retrieval is over land are all derived from the AGP DEM. The three constant thresholds ($H_{MDCM}= 330$ m, $CT_{MDCM}= 1.2$ ms⁻¹, and $AT_{MDCM}= 4.0$ ms⁻¹) are determined through an analysis of Level 2 TC_STEREO cloud motion vectors described in §3.3.3.3.1.

3.3.3.3.1 Derive thresholds for cloud masking

To facilitate use of TC_CLOUD motion vectors as wind observations, we seek a mask that successfully identifies at least 95% of Level 2 Cloud height-resolved motion vectors that are not associated with cloud advection. To determine appropriate thresholds for accomplishing this, a proxy set of TC_STEREO motion vectors retrievals collected from all valid retrievals in the time span from December 2000 to November 2009. In order to focus only on the retrievals of at least moderate quality, this set was further screened to include only retrievals with a vector magnitude of forward-aft difference less than 12 ms⁻¹.

To determine an appropriate value of threshold H_{MDCM} used in equation 42, we review the distribution of retrieval altitude minus terrain altitude for retrievals over land and water shown in Figure 12. Over land there is a distinct mode of terrain surface retrievals that is well characterized as a normal distribution centered at 0 m, with standard deviation of ~200 m. The width of this *Near Surface* distribution is governed by the mean standard deviation of terrain altitude throughout the 70.4 km square domain, ~100 m, and by the standard deviation of retrieval altitude error. Over water, there is a less pronounced mode of surface retrievals. Because surface altitude and variance are negligible over water, the standard deviation of this mode is governed only by altitude retrieval error, which appears to be ~165 m. We therefore set $H_{MDCM}= 330$ m, or twice the standard deviation of the distribution. By accounting for two

standard deviations of uncertainty, this threshold should successfully identify 95% of non-advection retrievals.

To choose an appropriate value for CT_{MDCM} used in equation 52, we review the distribution of retrieved cross-track motion for retrievals over land shown in Figure 13. There is a distinct mode of static retrievals that is well characterized as a normal distribution centered at 0 ms^{-1} with a standard deviation of 0.6 ms^{-1} .

We set $AT_{MDCM}=4.0 \text{ ms}^{-1}$, i.e. the standard deviation of the difference between forward and aft along-track motion. The along track cloud motion distribution is not distinctly bimodal, requiring a threshold determined by independent analysis of the expected error as predicted by statistical differences between forward and aft camera retrieved motion vectors.

Figure 12 and Figure 13 also show the distributions of retrievals labeled *Cloud* and those labeled *Near Surface* by the methodology of §3.3.3.3. The *Near Surface* retrieval distribution has a Gaussian shape for retrieval height above terrain (over land) and retrieval cross-track motion. The *Near Surface* retrieval altitude (over water) has a skewed distribution, suggesting a tendency to mislabel a small fraction of low altitude, low speed retrievals of cloud advection over water.

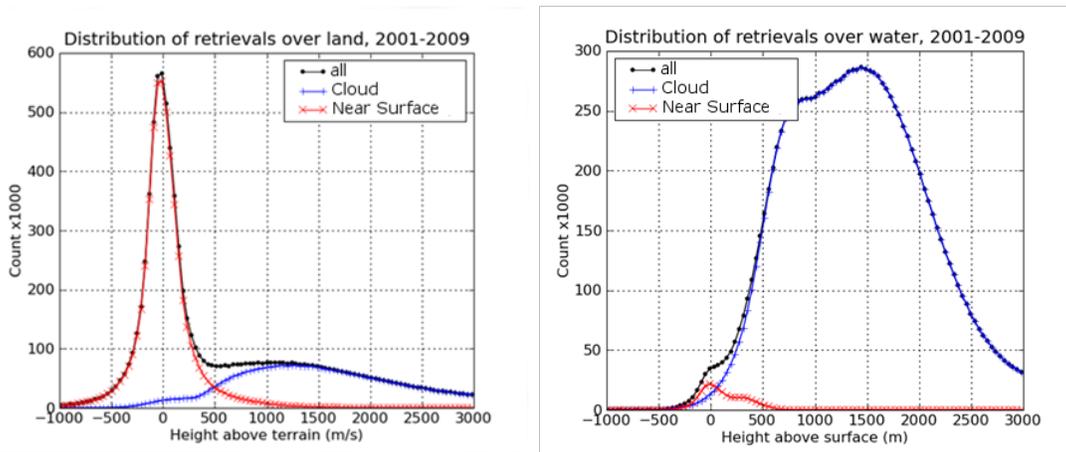


Figure 12. Bimodal retrieval altitude minus terrain altitude distribution

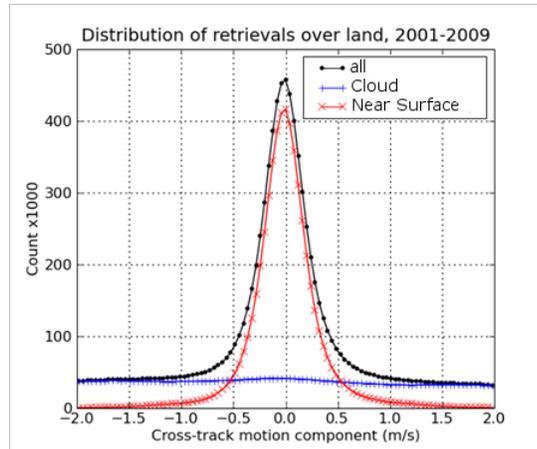


Figure 13. Bimodal retrieval cross-track motion distribution

3.3.4 Stereo Height Vector Operations

After being reconstructed, stereo height vectors undergo various operations before yielding the parameters provided to end users. First, the independently determined, but collocated forward (i.e. An-Af) and aft (i.e. An-Aa) camera derived grids of stereo height vectors are merged onto a single grid with at most one vector per grid cell. During this process, each vector is assigned a quality indicator and screened on the basis of agreement between forward and aft height and cross-track motion (§3.3.4.1). Screening additionally excludes cross-track motion magnitude and heading parameters retrieved at grid cells outside of a predefined interior of the MISR swath (§3.3.4.1.1). Outside this interior, L1B2 registration uncertainties reduce the reliability of cross-track motion retrieval. The field of stereo retrieval parameters is then reprojected from its original ellipsoid-referenced coordinate grid to a feature-referenced coordinate grid (§3.3.4.1.1). At this stage, processing diverges from a single raw stereo retrieval field into one field that is corrected for along-track feature motion and another that is not. Corrections are applied using collocated and applicable height resolved motion vectors (§3.3.4.3). The corrected stereo retrieval field includes only retrievals for which a correction is available. After the correction is applied, the corrected and uncorrected stereo retrieval fields differ in height and in sampling. Lastly, each and every stereo retrieval from both fields is assigned a Stereo Derived Cloud Mask (SDCM) value labeling the retrieved feature as either *Cloud* or *Near Surface* with *Low Confidence* or *High Confidence* (§3.3.4.4)

3.3.4.1 Merge forward and aft and assess quality of stereo vectors

A single stereo vector with associated Quality Indicator (QI) value is synthesized at each grid cell from collocated preliminary stereo vectors independently extracted from the forward and aft sets of camera images. One or both preliminary stereo vectors may be unavailable due to unsuccessful retrieval or due to having been identified as likely erroneous. The final stereo vector ultimately consists of one of four possibilities: (1) the mean of the forward and aft vectors; (2) the forward vector; (3) the aft vector; or (4) a non-retrieval. The first of these choices is only possible where the forward and aft camera stereo vectors do not differ in height cross-track

motion as assessed by $M_{FWD AFT}$ determined by equation 54. Where $M_{FWD AFT} \leq 1.0$, indicating forward and aft height differences less than $H_{HFWD AFT}=840$ m and cross-track differences less than $CT_{HFWD AFT}=9 \text{ ms}^{-1}$, the mean of the two is assigned at that grid cell.

$$M_{FWD AFT} = \max \left\{ \frac{|h_{fwd} - h_{aft}|}{H_{HFWD AFT}}, \frac{|ct_{fwd} - ct_{aft}|}{CT_{HFWD AFT}} \right\} \quad (54)$$

A similar quality metric, $M_{NBRFWD AFT}$, is calculated at grid cells where both preliminary forward and aft vectors are available, but $M_{FWD AFT} > 1.0$, and at grid cells where only one vector is available. This metric is obtained by computing $M^{i,j}_{FWD AFT}$, a corollary to $M_{FWD AFT}$ obtained by using the opposite camera stereo vector from a nearby grid cell at relative coordinate (i, j) . $M_{NBRFWD AFT}$ is taken to be the minimum value of $M^{i,j}_{FWD AFT}$ within a 5×5 neighborhood (equation 55). If no opposite camera stereo vectors are available for comparison, $M_{NBRFWD AFT}$ is not estimated; the reference stereo vector is discarded. At grid cells where $M_{NBRFWD AFT}$ is obtained from both a forward and aft stereo vector, the vector with a greater $M_{NBRFWD AFT}$ is discarded.

$$M_{NBRFWD AFT} = \min \{ M^{i,j}_{FWD AFT} \mid -2 \leq i \leq 2, -2 \leq j \leq 2 \} \quad (55)$$

QI values are assigned to grid cells on the basis of $M_{FWD AFT}$ where available, and $M_{NBRFWD AFT}$ otherwise. The QI is merely a transformation of this metric, given by:

$$QI = 100 - 100 \tanh M_{FWD AFT} \quad (56)$$

Grid cells assigned QI values less than 23 are discarded. This threshold is roughly equivalent to enforcing that respective differences not exceed $H_{HFWD AFT}$ or $CT_{HFWD AFT}$.

3.3.4.1.1 Screen less reliable cross-track motion outside interior of MISR swath

Cross-track motion magnitude and heading parameters retrieved at grid cells outside of a predefined interior of the MISR swath are screened from the product. Outside this interior, lens aberrations not accounted for in L1B2 registration procedures reduce the reliability of cross-track motion retrieval, though height retrieval is not adversely affected. The unclipped cross-track motion distribution derived from retrievals labeled as *Near Surface High Confidence* (see §3.3.4.4) in January and July of 2003, 2005, 2007, and 2009 show motion magnitudes appropriately constrained to within 2 m/s for all but the leftmost 18.75% (6/32) and rightmost 18.75% (6/32) of the sensor domain. Additionally, test cases have exhibited artifacts related to registration in the SOM-projected swath edges corresponding to these less well constrained portions of the sensor domain. An example of these artifacts is shown in Figure 14.

Conservative SOM-y axis boundaries associated with nominal mappings of the leftmost 23.4375% (15/64) and rightmost 23.4375% were defined for the top of each MISR block (i.e. minimum SOM x-axis coordinate). The top boundaries at the x-minimum of each block, together with the top boundaries of the next block, define a bounding parallelogram, outside of which the

cross-track motion and cross-track instrument heading parameters are set to fill. The top per-block clipping boundaries are given in Table 1, while Figure 14 shows an example of unclipped cross-track motion coverage (colored by retrieved motion value) and clipping boundaries. Figure 14 also shows clipped artifacts in the top of the blocks 31-35 on the left, where abrupt color shifts mark retrieval discontinuities near the swath edge.

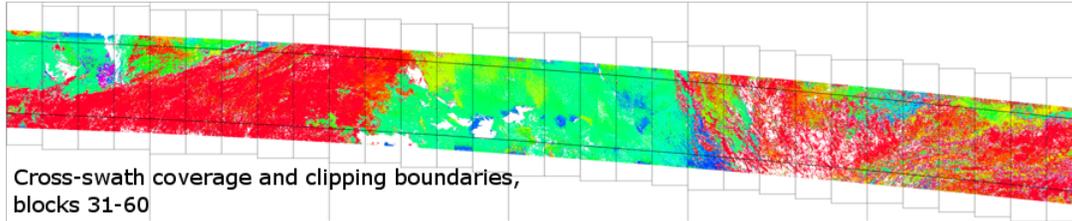


Figure 14. Example of cross-swath coverage and clipping

3.3.4.2 Project stereo vectors from ellipsoid-referenced to feature referenced

Stereo vectors are initially gridded with respect to their observed coordinates within the ellipsoid-referenced (ER) camera imagery, requiring a projection procedure to instead grid them with respect to their feature-referenced (FR) coordinate. Feature referenced (coordinates differ from ellipsoid reference coordinates by an offset proportional to the feature height and dependent upon the viewing zenith angle, z , as depicted in Figure 15.

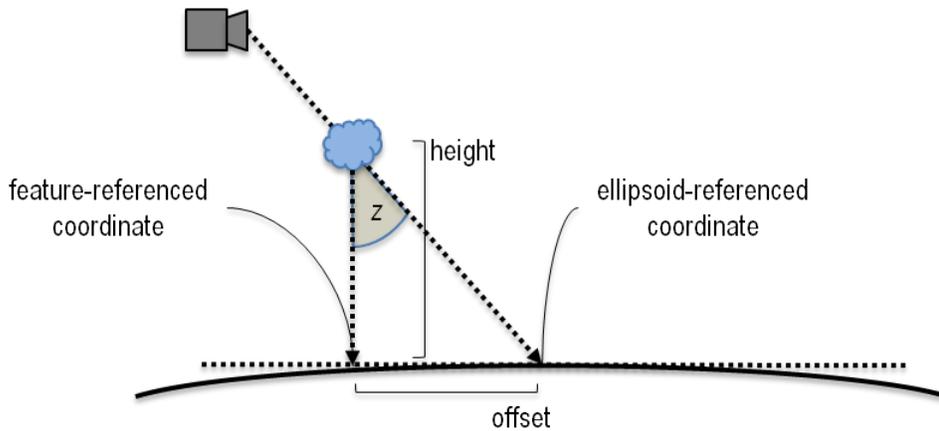


Figure 15. Schematic of ellipsoid-referenced versus feature-referenced positions

Because the scale of the ER to FR offset (1-5 km) is negligible relative to Earth's radius (~3600 km), the scene geometry can be well approximated by treating the ellipsoid as a plane. The offset magnitude is then given by $h \cdot \tan z$, with orientation given by viewing azimuth angle. Equivalently, the horizontal translation vector, \vec{T} , from ER to FR is computed from the feature height and the unit look vector from observed pixel in Local North coordinates, $\hat{P} = \langle P_0, P_1, P_2 \rangle$:

$$\vec{T} = -h \cdot \left\langle \frac{P_0}{P_2}, \frac{P_1}{P_2} \right\rangle \quad (57)$$

The FR grid of stereo vectors is populated iteratively. For each ER grid cell containing a stereo vector, the above translation vector is calculated and applied to determine an associated FR grid cell. Multiple ER stereo vectors can map to the same FR grid cell, in which case the stereo vector with the largest QI value will be assigned.

3.3.4.3 Apply along-track correction to stereo height vectors where applicable

Height corrections are applied to stereo retrievals where applicable using height resolved motion vectors retrieved in collocated or neighboring 17.6 km resolution grid cells. The corrections are designed to negate the bias proportional to along-track motion, u , with sensitivity, h_s , (introduced in §2.3.1). The correction is approximated by equation 58, which is merely a rearrangement of equation 2. The planar approximation of Earth's ellipsoid was deemed sufficient for the 26.1° view angle difference of the An-Af and An-Aa camera pairs, which constrain the horizontal scale (specifically \overline{AB}_1 in Figure 2) to 1:1000 that of Earth's radius for cloud heights up to 20 km.

$$h = h' - uh_s \quad (58)$$

For each stereo retrieval, motion vectors applicable for correction include only those whose height does not differ by more than $H_{CORRECTIONDIFF}$ (840 m) from the height of the stereo retrieval *after correction*. If the 17.6 km grid cell collocated with the stereo retrieval grid cell has an applicable motion vector, its along-track component is used for the correction. Otherwise, if neighboring grid cells (i.e. within a 3x3 window) contain applicable motion vectors, the along-track component of the one whose height differs the least from the stereo height after correction is used. Where applicable motion vectors are not available, the stereo retrieval is masked out.

3.3.4.4 Generate Stereo Derived Cloud Mask

The purpose of this stage is to produce the Stereo Derived Cloud Mask (SDCM) that classifies stereo retrievals as either *Cloud* or *Near Surface* with *Low Confidence* or *High Confidence*. *Near Surface* retrievals are discriminated from *Cloud* retrievals solely on the basis of their height with respect to underlying terrain. The specific test is expressed in equation 59, where the threshold, H_{SDCM} is 560 m, and the terrain height, $h_{terrain}$, and terrain height variance, $\sigma_{h,terrain}$ are read from the AGP.

$$|h| > H_{SDCM} + h_{terrain} + 2\sigma_{h,terrain} \quad (59)$$

Categorization as either *Cloud* or *Near Surface* is applied independently to any successful forward or aft camera retrieved stereo vector at each grid cell. Where both forward and aft vectors are present, and agree in their discrimination, the retrieval is categorized as *High Confidence*. It is otherwise categorized as *Low Confidence*. Where the forward and aft vectors disagree in their discrimination, the average of the two heights is used in *Low Confidence*

retrievals, the average height is assessed by equation **59** to distinguish between *Cloud* and *Near Surface*.

4 ASSUMPTIONS AND LIMITATIONS

4.1 ASSUMPTIONS

A number of assumptions are inherent to the retrieval of Level 2 Cloud parameters. The most noteworthy of these relate to the local area correspondence algorithms used to identify conjugate features in pairs of MISR images. Theoretically, the success of these algorithms is contingent upon conditions that are never strictly met, and often blatantly violated [8]. Well known pitfalls that cause undefined behavior include:

- Rapidly evolving scene: If a feature is present in the scene observed by one camera, and absent when observed by another, correspondence will fail. Intense convection has been known to cause this problem.
- Occlusion: A feature observed by a reference camera is obscured from the view of another, preventing successful correspondence.
- Transparency: A partially transparent object (e.g. cirrus cloud) is present, resulting in the observed pattern of radiance being associated with the convolution of signals from multiple scene points associated with the object and underlying objects. The algorithm is solving for only one.
- Discontinuity: A discontinuity in the scene structure, especially when unpredictable (e.g. a small isolated cloud), also results in multiple scene points contributing to the observed pattern of radiance. The algorithm is solving for only one.
- Insufficient contrast: A low-texture surface can contribute a patterns of radiance not sufficiently distinct to be uniquely associated with any specific image coordinate.

The principle method of addressing the above issues employed by Level 2 Cloud is the required agreement of parameters independently derived by applying correspondence to distinct forward and aft sets of cameras. Instead of producing incorrect retrievals, the algorithm thereby produces no retrieval. However, this method assumes that all correspondence error is random in nature, unlikely to produce the same error from two independent views.

Additional assumption include:

- Heights of cloud motion and stereo vector heights are assumed to be associated with the top of the cloud. This has been consistent with ongoing evaluations of cloud top height accuracy, but may not be strictly true, especially for optically thin clouds.

4.2 LIMITATIONS

The following limitations apply to the Level 2 Cloud retrievals:

- Stereo and motion retrievals do not exclusively contain values representing meteorological clouds. The underlying retrieval mechanism operates on identifiable features within MISR camera images that may be associated with cloud, aerosol, or terrain, and may also be spurious.
- The search area employed during correspondence (§3.3.1.1.1) is configured to obtain cloud motion vectors at heights between 0 and 20 km, with speeds up to 50

m/s. Valid heights and speeds outside this range can be retrieved and are included in the product. However, as the algorithm is configured, the possibility of the obtaining such retrievals is not independent of the instrument orientation, cloud height, and wind speed.

- Correspondence can produce spurious results that are nevertheless consistent between forward and aft camera retrieval. Some fraction of spurious motion vectors possess obviously unrealistic values. Unrealistic values among reported motion vectors in test data spanning 2004-2011 occur with a frequency ranging from ~0.02% in the mid-latitudes to ~0.10% at the poles. Of these unrealistic motion vectors, some 80% are assigned Quality Indicator (QI) values less than 50, indicating low confidence in the retrieval.
- Cloud motion vector accuracy is sensitive to the georegistration accuracy of the MISR L1B2 input. This sensitivity is estimated to range from 6 m/s/pixel to 16 m/s/pixel depending on the camera view angle [Davies et al., 2007; Zong et al., 2002]. Georegistration accuracy can vary significantly by camera and is notably a function of the quantity of terrain observed by MISR during a particular orbit. A small fraction of orbits (e.g., 7% in data from 2007) are identified by quality control procedures as having insufficient georegistration accuracy for cloud motion vector retrieval. Such orbits are flagged as poorly registered and provide only cloud top heights without wind correction.
- Cloud motion vectors are subject to bias related to L1B2 georegistration that is weakly correlated with cross-track position within the MISR swath.
- Cloud motion vectors are not obtained independently from their associated cloud top heights. Notably, error in the height of a cloud motion vector is not independent of error in the along-track component of motion [Zong et al., 2002].

5 REFERENCES

- [1] Atlas, R. (1997), Atmospheric observations and experiments to assess their usefulness in data assimilation, *Journal of the Meteorological Society of Japan*, 75, 111–130.
- [2] Davies, R., Á. Horváth, C. Moroney, B. Zhang, and Y. Zhu (2007), Cloud motion vectors from MISR using sub-pixel enhancements, *Remote Sensing of Environment*, 107, 194–199.
- [3] Hinkelman, L. M., R. T. Marchand, and T. P. Ackerman (2009), Evaluation of Multiangle Imaging Spectroradiometer cloud motion vectors using NOAA radar wind profiler data, *Journal of Geophysical Research*, 114, D21207, doi:10.1029/2008JD011107.
- [4] Horvath, Á. and R. Davies (2001), Feasibility and error analysis of cloud motion wind extraction from near-simultaneous multiangle MISR measurements, *Journal of Atmospheric and Oceanic Technology*, 18-4, 591-608
- [5] Horvath, Á. and R. Davies (2001), Simultaneous retrieval of cloud motion and height from polar-orbiter multiangle measurements, *Geophysical Research Letters*, 28-15, 2915-2918
- [6] Marchand, R. T., T. P. Ackerman, and C. Moroney (2007), An assessment of Multiangle Imaging Spectroradiometer (MISR) stereo-derived cloud top heights and cloud top winds using ground-based radar, lidar, and microwave radiometers, *Journal of Geophysical Research*, 112, D06204, doi:10.1029/2006JD007091.
- [7] Moroney, C., R. Davies, and J.-P. Muller (2002), Operational retrieval of cloud-top heights using MISR data, *IEEE Transactions on Geoscience and Remote Sensing*, 40-7, 1532-1540
- [8] Mueller, K.J. (2008), Stereo observations of polar stratospheric clouds, *University of Illinois*, M.S. Thesis
- [9] Muller, J.-P., A. Mandanayake, C. Moroney, R. Davies, D.J. Diner, and S. Paradise (2002), MISR stereoscopic image matchers: Techniques and results, *IEEE Transactions of Geoscience and Remote Sensing*, 40-7, 1547-1559
- [10] Zong, J., R. Davies, J.P. Muller, D.J. Diner (2002), Photogrammetric retrieval of cloud advection and top height from the Multi-angle Imaging Spectroradiometer (MISR), *Photogrammetric Engineering and Remote Sensing*, 68, Number 8, 821-829, 2002.

A. GLOSSARY OF ACRONYMS

A

AGP (Ancillary Geographic Product)
ASDC (Atmospheric Sciences Data Center)
ATB (Algorithm Theoretical Basis)

C

CCD (Charge-Coupled Device)
CF (Climate and Forecast)
CMV (Cloud Motion Vector)
CMVP (Cloud Motion Vector Product)

E

ECI (Earth Centered Inertial)
ECR (Earth Centered Rotational)
EOS (Earth Observing System)
ER (Ellipsoid Referenced)

F

FR (Feature Referenced)

G

GDQI (Geometric Data Quality Indicator)
GOES (Geostationary Operational Environmental Satellite)

I

IFOV (Instantaneous Field of View)

L

L2C (Level 2 Cloud Product)
LaRC (Langley Research Center)

M

MISR (Multi-angle Imaging SpectroRadiometer)
MODIS (Moderate Resolution Imaging Spectroradiometer)

N

NetCDF (Network Common Data Form)

P

PGS (Product Generator System)

S

SOM (Space Oblique Mercator)

W

WGS84 (World Geodetic System of 1984)

B. APPENDIX

Constant thresholds

Table 4: Threshold values

Threshold name	Value	Description
q_{thresh}	1	Lowest acceptable ratio of L1B2 image contrast for correspondence relative to expected contrast due to MISR instrument noise
$H_{CMVFWD AFT}$	1000 m	Largest acceptable difference between collocated forward and aft camera motion vector heights
$V_{CMVFWD AFT}$	12 ms ⁻¹	Largest acceptable motion vector magnitude difference between collocated forward and aft camera motion vectors
$H_{CMVNEIGHBOR}$	500 m	Largest acceptable height difference between motion vectors in neighboring grid cells to allow for comparison of horizontal motion components
$V_{CMVNEIGHBOR}$	12 ms ⁻¹	Largest acceptable horizontal vector motion difference between
$QI_{CMVTHRESH}$	25	Smallest acceptable Quality Indicator value for inclusion of motion vector in product
H_{MDCM}	330 m	Largest height above terrain possible for motion vector are classified as Near Surface
CT_{MDCM}	1.2 ms ⁻¹	Largest cross-track component magnitude of motion vector classified as Near Surface
AT_{MDCM}	4.0 ms ⁻¹	Largest along-track component magnitude of motion vector (over ocean) classified as Near Surface
H_{SDCM}	560 m	Largest height above terrain possible for stereo vector classified as Near Surface
$H_{CORRECTIONDIFF}$	840 m	Maximum difference between corrected stereo height and height of motion vector used for correction
$H_{FWD AFT}$	840 m	Largest acceptable height difference between collocated (or nearly collocated) forward and aft stereo vectors. Also used to determine QI.
$CT_{FWD AFT}$	9 ms ⁻¹	Largest acceptable cross-track motion component difference between collocated (or nearly collocated) forward and aft stereo vectors. Also used to determine QI.
$QI_{HTHRESH}$	23	Smallest acceptable Quality Indicator value for inclusion of stereo vector in product

SOM-y clipping boundaries for cross-track motion

Table 5: SOM-y clipping boundaries for cross-track motion (for top of block)

Block	SOM y [min, max]						
1	[661650, 947650]	46	[-878350, -587950]	91	[-161150, 122650]	136	[568150, 853050]
2	[682550, 968550]	47	[-909150, -617650]	92	[-219450, 64350]	137	[532950, 817850]
3	[700150, 986150]	48	[-935550, -644050]	93	[-276650, 7150]	138	[494450, 779350]
4	[713350, 1000450]	49	[-959750, -667150]	94	[-333850, -48950]	139	[453750, 738650]
5	[724350, 1010350]	50	[-979550, -686950]	95	[-391050, -105050]	140	[410850, 694650]
6	[729850, 1016950]	51	[-994950, -702350]	96	[-446050, -160050]	141	[364650, 648450]
7	[733150, 1019150]	52	[-1008150, -715550]	97	[-499950, -213950]	142	[316250, 600050]
8	[730950, 1018050]	53	[-1016950, -723250]	98	[-551650, -265650]	143	[266750, 549450]
9	[726550, 1012550]	54	[-1021350, -728750]	99	[-602250, -315150]	144	[213950, 497750]
10	[717750, 1003750]	55	[-1022450, -729850]	90	[-650650, -363550]	145	[161150, 443850]
11	[705650, 991650]	56	[-1020250, -726550]	101	[-697950, -408650]	146	[106150, 388850]
12	[689150, 975150]	57	[-1013650, -721050]	102	[-741950, -452650]	147	[50050, 332750]
13	[669350, 955350]	58	[-1002650, -710050]	103	[-782650, -493350]	148	[-7150, 276650]
14	[646250, 932250]	59	[-988350, -696850]	104	[-821150, -530750]	149	[-64350, 218350]
15	[619850, 905850]	60	[-970750, -679250]	105	[-856350, -565950]	150	[-122650, 161150]
16	[590150, 875050]	61	[669350, 955350]	106	[-889350, -597850]	151	[-180950, 102850]
17	[557150, 842050]	62	[689150, 975150]	107	[-917950, -626450]	152	[-238150, 45650]
18	[520850, 805750]	63	[704550, 991650]	108	[-944350, -651750]	153	[-296450, -11550]
19	[481250, 766150]	64	[717750, 1003750]	109	[-966350, -673750]	154	[-352550, -68750]
20	[439450, 724350]	65	[726550, 1012550]	110	[-985050, -692450]	155	[-408650, -123750]
21	[395450, 679250]	66	[730950, 1018050]	111	[-999350, -707850]	156	[-463650, -177650]
22	[349250, 631950]	67	[732050, 1019150]	112	[-1011450, -718850]	157	[-517550, -231550]
23	[299750, 583550]	68	[729850, 1016950]	113	[-1018050, -725450]	158	[-569250, -282150]
24	[249150, 532950]	69	[724350, 1010350]	114	[-1022450, -729850]	159	[-618750, -331650]
25	[196350, 480150]	70	[713350, 1000450]	115	[-1022450, -728750]	160	[-667150, -378950]
26	[142450, 426250]	71	[700150, 986150]	116	[-1018050, -725450]	161	[-712250, -424050]
27	[87450, 370150]	72	[682550, 969650]	117	[-1010350, -717750]	162	[-755150, -465850]
28	[31350, 314050]	73	[661650, 947650]	118	[-998250, -706750]	163	[-795850, -506550]
29	[-25850, 256850]	74	[637450, 923450]	119	[-983950, -691350]	164	[-833250, -542850]
30	[-84150, 199650]	75	[609950, 895950]	120	[-964150, -672650]	165	[-868450, -576950]
31	[-142450, 141350]	76	[579150, 864050]	121	[675950, 961950]	166	[-899250, -607750]
32	[-199650, 84150]	77	[545050, 829950]	122	[694650, 980650]	167	[-927850, -635250]
33	[-257950, 25850]	78	[507650, 792550]	123	[710050, 996050]	168	[-952050, -659450]
34	[-315150, -30250]	79	[468050, 751850]	124	[721050, 1007050]	169	[-972950, -680350]
35	[-372350, -86350]	80	[425150, 708950]	125	[728750, 1014750]	170	[-990550, -697950]
36	[-427350, -142450]	81	[380050, 663850]	126	[732050, 1019150]	171	[-1003750, -711150]
37	[-482350, -196350]	82	[332750, 616550]	127	[732050, 1019150]	172	[-1013650, -721050]
38	[-535150, -248050]	83	[283250, 567050]	128	[728750, 1014750]	173	[-1020250, -727650]
39	[-585750, -298650]	84	[231550, 515350]	129	[721050, 1007050]	174	[-1022450, -729850]
40	[-635250, -347050]	85	[178750, 462550]	130	[710050, 996050]	175	[-1021350, -728750]
41	[-682550, -394350]	86	[123750, 407550]	131	[694650, 981750]	176	[-1015850, -723250]
42	[-726550, -438350]	87	[68750, 351450]	132	[675950, 963050]	177	[-1007050, -714450]
43	[-769450, -480150]	88	[11550, 295350]	133	[653950, 939950]	178	[-993850, -701250]
44	[-809050, -518650]	89	[-45650, 238150]	134	[628650, 914650]	179	[-977350, -685850]
45	[-845350, -554950]	90	[-102850, 179850]	135	[600050, 886050]	180	[-957550, -664950]
						181	[-949850, -658350]