

Shown LEFT is Planet's Fusion Monitoring Product (PlanetScope data)
Shown RIGHT is LandSat 8 data of the same area



PLANET FUSION MONITORING TECHNICAL SPECIFICATION

Version 1.0.0, March 2022

Calibration, Analysis Ready Data, and InterOperability (CARDIO) operations



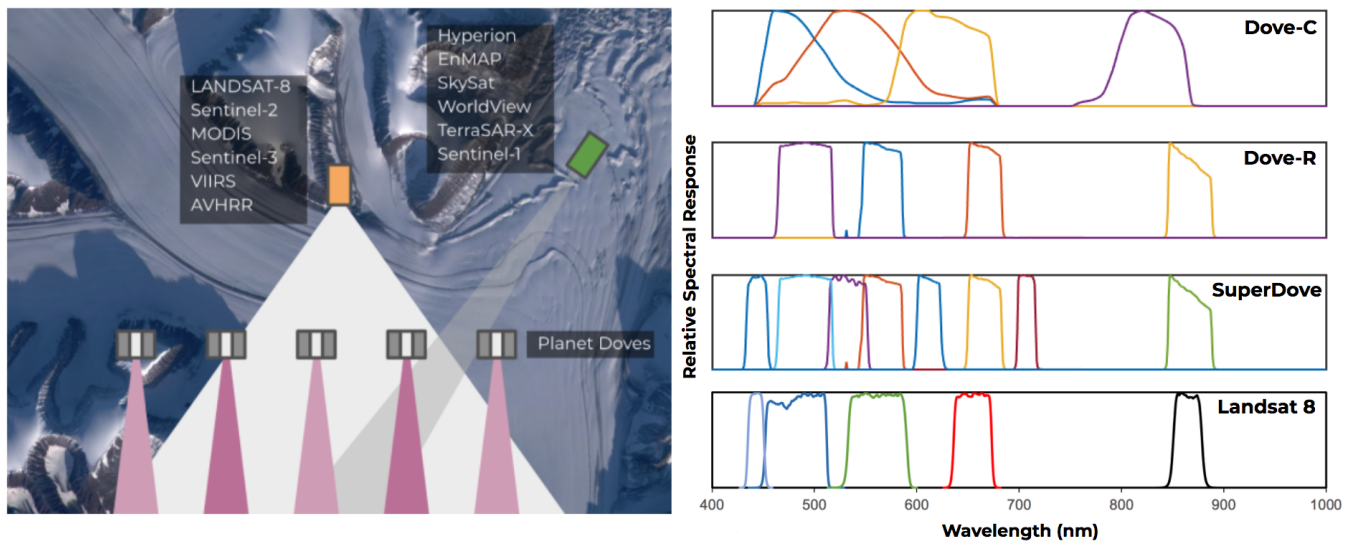
TABLE OF CONTENTS

TABLE OF CONTENTS	2
PLANET FUSION OVERVIEW	3
PLANET FUSION INPUTS	5
PLANET FUSION PRODUCTS	6
SURFACE REFLECTANCE PRODUCT (PF-SR)	7
QUALITY ASSURANCE PRODUCT (PF-QA)	7
SPATIOTEMPORAL ASSET CATALOG (STAC)	10
PROJECTION, GRIDDING, FILE NAMING, AND DELIVERY	11
PLANET FUSION METHODOLOGY	12
TILE (CHUNK)-LEVEL STACKING	13
CLOUD AND CLOUD SHADOW MASKING	13
REFERENCE SAMPLING AND CROSS-CALIBRATION	14
GAP-FILLING AND TEMPORAL FILTERING	16
GEOMETRIC HARMONIZATION	19
CONFIDENCE INFORMATION	22
BACKFILL VERSUS FORWARD-FILL OPERATION	24
UNCERTAINTY ESTIMATES	24
KNOWN LIMITATIONS AND CAVEATS	27
REFERENCES	28

1. PLANET FUSION OVERVIEW

The PlanetScope constellation of 200+ CubeSats in low earth orbits represents a novel observational resource, which when combined with advances in conventional spaceborne sensing has resulted in a proliferation of satellite sensor data with unprecedented spatial, temporal, and spectral resolution. This constitutes a revolution in the ability to derive time-critical, location-specific insights about dynamic land surface processes. However, the potential for these systems to support decision making is often limited by sensor interoperability issues (Figure 1), cross-calibration challenges, and atmospheric contamination. These obstacles can stand in the way of realizing the full potential of these rich datasets.

Figure 1: Sources of interoperability issues. Left: The reflectance field of the same observation target can appear very different at the point of the satellite sensor due to differences in satellite viewing and sun illumination angles augmented by shadow effects and non-lambertian surface characteristics (i.e., as described by the Bidirectional Reflectance Distribution Function; BRDF). Right: Differences in spectral bands and spectral response functions can result in poor sensor interoperability. This is particularly pronounced when comparing Dove-Classics (i.e., first generation PlanetScope) with public sensor sources (e.g., Landsat 8).

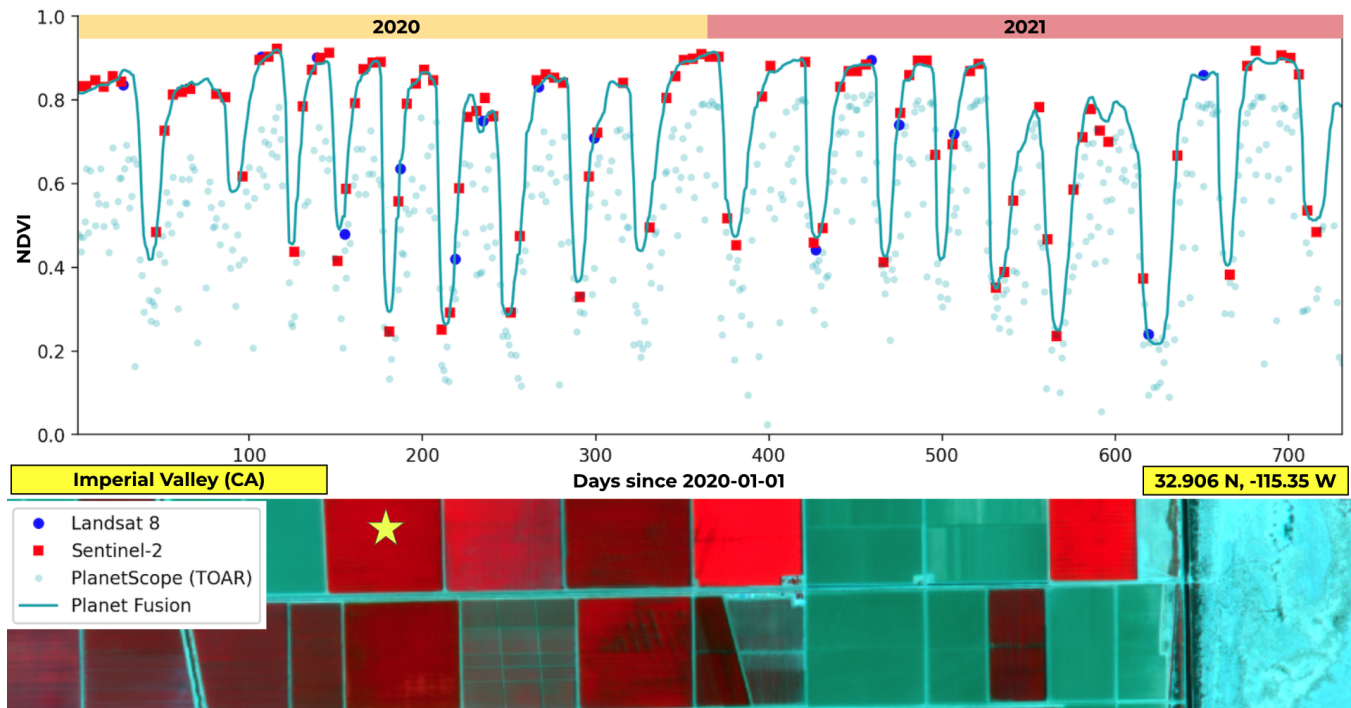


At Planet, we have implemented and improved a rigorous methodology to enhance, harmonize, inter-calibrate, and fuse cross-sensor data streams. The CubeSat-Enabled Spatio-Temporal Enhancement Method (CESTEM) (Houborg and McCabe 2018a,b) leverages rigorously calibrated publicly accessible multispectral satellites (i.e., Sentinel, Landsat, MODIS, VIIRS) to work in concert with the higher spatial and temporal resolution data provided by Planet’s Dove CubeSats. The result is a next generation, analysis ready, harmonized Level-3 data product, which delivers a clean (i.e. free from clouds and shadows), gap-filled (i.e., daily, 3 m), temporally consistent, and radiometrically accurate 4-band surface reflectance (SR) data product.

Planet Fusion integrates all of the best features from both public and private satellite sensor resources. Data from multiple sensors with differing radiometry, quality, and spatiotemporal resolution characteristics are fused in order to produce an entirely new, sensor-agnostic dataset that inherits the best traits from each sensor while ensuring radiometric consistency with a suite of widely used “reference” satellite platforms. This next generation Analysis Ready Data (ARD) product is suitable for analytic and data science purposes and is particularly beneficial for inter-day change detection and disturbance monitoring, dynamic land use and land cover classification, phenology and vegetation growth monitoring, and biophysical properties and process retrievals (e.g., Nieto et al., 2022; Ziliani et al., 2022; Aragon et al., 2021; Kong et al., 2022), in addition to other applications

relying on high frequency, gap-free and radiometrically, geometrically and temporally consistent information (Figure 2).

Figure 2: Normalized Difference Vegetation Index (NDVI) time series over a dynamic multi-cut alfalfa field in Imperial Valley (CA) over the course of two years (2020 - 2021). The Planet Fusion processing translates original PS Top Of Atmosphere (TOA) reflectance inputs into Surface Reflectances (SR) consistent with Landsat 8 and Sentinel-2 clear-sky observations. Gaps in the Planet Fusion time series (due to clouds, missing acquisition data) have been filled using an elaborate gap-filling and data fusion technique ([Section 4.4](#)) to provide a spatially complete daily product.



The unique features of Planet Fusion data can be summarized as:

- **Advanced radiometric harmonization which leverages rigorously calibrated third-party sensors (MODIS/VIIRS, Landsat 8, and Sentinel-2) for full fleet interoperability**
- **Rigorous, temporally driven, cloud and cloud shadow detection**
- **Cleaned and gap-filled (daily) Surface Reflectance values delivered as regularly gridded UTM raster tiles (24 km x 24 km) with a 48 hour latency**
- **Fusion of Sentinel-2 and Landsat 8 data to help fill gaps in PlanetScope coverage**
- **CubeSats with near-nadir field of view result in minimal BRDF variation effects**
- **Designed to provide high radiometric accuracy and spatio-temporal consistency**
- **Includes precise co-registration and sub-pixel fine alignment of disparate image sources**
- **Includes pixel traceability information to easily identify source imagery and assess the confidence of gap-filled data for every data point**

2. PLANET FUSION INPUTS

Table 1 lists the data sources currently used in Planet Fusion production.

Planet’s collection of 200+ CubeSats operate in sun synchronous orbits (altitude ~475 km) with a midmorning equatorial overpass time (9:30–11:30 a.m., local solar time) providing global near-nadir (~5° field of view) imaging on a near-daily basis (Roy et al., 2021). Three generations of “Doves” are used as input to Planet Fusion products; the Dove-Classic constellation (2016-) is characterized by broad and partly overlapping spectral bands in the visible and near-infrared (NIR) spectrum (Figure 1), whereas the Dove-R (2019-) and SuperDove (2020-) constellations are improved to be directly interoperable with the visible and narrow NIR bands of Sentinel-2. The nominal ortho scene size (at 475 km altitude) is also larger for Dove-R (~25 km x 23 km) and SuperDove (~32.5 km x 19.6 km) relative to Dove-Classic (~25 km x 11.5 km). PlanetScope Top Of Atmosphere (TOA) Radiance inputs contribute the bulk of the observations used to make Fusion, and they are derived from 4-band Orthorectified Scene Products that have a pixel size of 3 m. While the SuperDove is an 8-band sensor (i.e., adds bands in the visible and red-edge domain), currently only the blue, green, red, and NIR bands are used in Planet Fusion production.

Table 1: List of inputs currently used in Planet Fusion production.

Product	Description
PS-TOA	Scene-based PlanetScope L1B - Top Of Atmosphere (TOA) Radiance (4-band, 3 m)
MCD43A4, MCD43A4N	Tile-based MODIS Surface Reflectance (SR) normalized to a nadir view direction and local solar noon (daily, 500 m) (https://lpdaac.usgs.gov/products/mcd43a4v061/)
VNP43IA4	Tile-based SR (VIIRS imagery bands) normalized to a nadir view direction and local solar noon (daily, 500 m) (https://lpdaac.usgs.gov/products/vnp43ia4v001/)
VNP43MA4	Tile-based SR (VIIRS moderate bands) normalized to a nadir view direction and local solar noon (daily, 1000 m) (https://lpdaac.usgs.gov/products/vnp43ma4v001/)
FLS-SR	In-house implementation for tile-based generation of Nadir BRDF Adjusted Reflectances (NBAR) from Landsat 8 and Sentinel-2 data (4-band, 30 m). The L8 data have been spectrally adjusted to match S2 spectral band passes. Based on the Framework for Operational Radiometric Correction for Environmental Monitoring (FORCE) (https://github.com/davidfrantz/force)

We use a scalable implementation of the Framework for Operational Radiometric Correction for Environmental Monitoring (FORCE version 3.7.7; Frantz 2019a) for generating a combined Landsat 8 and Sentinel-2 surface reflectance product (FLS-SR) to be used as the “gold reference” during the radiometric calibration and normalization of Planet Fusion products. FORCE includes state-of-the-art atmospheric correction, terrain correction, cloud and cloud shadow detection, spatial co-registration, and view angle normalization (Frantz 2019a). FORCE infers surface reflectance from Landsat 8 and Sentinel-2 imagery using an implementation of the 5S (Simulation of the Satellite Signal in the Solar Spectrum) code (Tanre et al., 1990). The aerosol optical depth is estimated from the imagery using a dark object based approach whereas the water vapor content is either estimated on a pixel-specific basis (Sentinel-2) or derived from a global MODIS-based database (Landsat 8) (Frantz et al., 2019b). Clouds and shadows are detected using a modified version of Fmask (Zhu and Woodcock, 2012) that exploits parallax effects to improve detections for Sentinel-2 images (Frantz et al., 2018). However the FORCE derived cloud masks may be further refined as part of Fusion processing ([Section 4.2](#)). A global

assessment of the FORCE atmospheric correction approach was conducted as part of the Atmospheric Correction Inter-comparison Exercise (ACIX) (Doxani et al., 2018).

Our FORCE implementation maps the Landsat 8 and Sentinel-2 data onto a common grid (i.e., the UTM-based Military Grid Reference System) to produce 30 m resolution L8/S2 data with a 2 - 3 day frequency. A spectral bandpass adjustment (Claverie et al. 2018) is applied to Landsat 8 to align with Sentinel-2 radiometry. Only the blue (0.45 - 0.51 μm), green (0.53 - 0.59 μm), red (0.64 - 0.67 μm), and narrow NIR (0.85 - 0.88 μm) bands are currently used for the Planet Fusion production.

MODIS or VIIRS surface reflectance (SR) data normalized to nadir view and local solar noon is a required input to the Planet Fusion reference sampling and calibration process (Section 4.3). Planet Fusion uses the version 6.1 combined (i.e., Terra and Aqua) MCD43A4 product that provides daily 500 m SR in 7 bands corrected for reflectance anisotropy (MODIS has a $\sim 110^\circ$ field of view) using a semiempirical bidirectional reflectance distribution function (BRDF) (Schaaf et al. 2002). The BRDF utilizes the best observations from both Terra and Aqua sensors collected over a 16-day period centered on the day of interest where observations at the day of interest are emphasized in the daily retrieval. Only the blue (0.459 - 0.479 μm), green (0.545 - 0.565 μm), red (0.62 - 0.67 μm), and NIR (0.841 - 0.876 μm) bands are ingested for Planet Fusion processing. The near real-time product version (MCD43A4N) is used when the standard product isn't available (i.e., as dictated by the ~ 7 days latency). The VIIRS products (VNP431A4, VNP43MA4) have been designed to ensure continuity of MCD43 and are used as a backup should MCD43A4/MCD43A4N become unavailable. The VIIRS-based processing ingests the red (0.60 - 0.68 μm) and NIR (0.85 - 0.88 μm) Imagery bands (500 m) in addition to the blue (0.478 - 0.488 μm) and green (0.545 - 0.565 μm) Moderate bands (1000 m). Since the two latter bands are provided at a coarser spatial resolution (1000 m), the finer resolution (500 m) red band is used to super-resolve the data for consistency.

3. PLANET FUSION PRODUCTS

The Planet Fusion (PF) product line is outlined in Table 2. The Planet Fusion products are spatially complete (i.e., gap-free), cloud-free, provided at a daily cadence, and orthorectified onto a fixed grid with a 3 m resolution (see Section 3.4). The products can be produced starting from January 1, 2018 and up till present, deliverable with a 48 hour latency (i.e., data requested for Monday will be delivered Wednesday).

Table 2: Overview descriptions of the Planet Fusion (PF) products.

Product key	Description
PF-SR	Planet Fusion Surface Reflectance (SR) product. PS TOA Reflectance radiometrically harmonized to 4-band FLS-SR using the CESTEM methodology . Cloud masked and gap-filled via PS, MODIS/VIIRS, and Landsat 8/Sentinel-2 data fusion (gap-free, daily, 3 m)
PF-QA	Planet Fusion Quality Assurance product
PF-STAC	Planet Fusion SpatioTemporal Asset Catalog items and catalog

3.1. SURFACE REFLECTANCE PRODUCT (PF-SR)

The Planet Fusion Surface Reflectance product (PF-SR) records gridded (3 m) and gap-free orthorectified data in four spectral bands (blue, green, red, NIR) at a daily interval. The data is stored in 16-bit integer format (with a multiplication factor of 10,000) as cloud optimized geotiffs compressed using LZW compression. During Planet Fusion processing and harmonization, 4-band PS TOA Reflectance (PS-TOAR) (converted from the radiances) are transformed into surface reflectances ensuring radiometric consistency with Sentinel-2. As a result the spectral bands and spectral response functions of Planet Fusion data (Table 3) will be equivalent to the blue (B2), green (B3), red (B4), and narrow NIR (B8a) bands of Sentinel-2 (ESA 2021).

The Planet Fusion SR data represent Normalized BRDF Adjusted Reflectances (NBAR) as the Landsat 8 and Sentinel-2 data used for cross-calibration have been normalized to nadir view (Roy et al. 2016, 2017). As the Planet Fusion cross-calibration adopts a multi-temporal reference sampling approach (see [Section 4.3](#)), the significant uncertainties related to the L8/S2-based BRDF normalization (Roy et al. 2017) are likely to cancel out. In addition, in contrast to Landsat 8 (~15° field of view) and particularly Sentinel-2 (~21° field of view), the PlanetScope sensors are nadir viewing natively (~5° field of view), which will act to further minimize view angle BRDF effects.

Table 3: PF-SR data format specifications

Layer	Description	Date Type	Valid range	Scale factor
Band 1	Blue band (0.45 - 0.51 µm) SR (NBAR)	16-bit signed integer	1 - 10,000	0.0001
Band 2	Green band (0.53 - 0.59 µm) SR (NBAR)	16-bit signed integer	1 - 10,000	0.0001
Band 3	Red band (0.64 - 0.67 µm) SR (NBAR)	16-bit signed integer	1 - 10,000	0.0001
Band 4	NIR band (0.85 - 0.88 µm) SR (NBAR)	16-bit signed integer	1 - 10,000	0.0001

3.2. QUALITY ASSURANCE PRODUCT (PF-QA)

The Planet Fusion Quality Assurance (QA) product is a 9 layer thematic raster using the same spatial grid as the corresponding Planet Fusion spectral data. It contains information denoting gap-filling (layers 1 and 2), cloud and cloud shadow detection (layer 3), pixel traceability (layer 4), number of L8/S2 reference scenes used during calibration (layer 5), and spectral confidence estimates for synthetic (gap-filled) data (layer 6 - 9) (Table 5).


Table 5: PF-QA data format specifications. Note that additional metadata have been embedded in the tiff file, as described in the table

Layer	Description	Date Type	Valid range	Scale factor	Offset
Layer 1	Percentage of synthetic (gap-filled) versus actual PlanetScope data used to generate pixel value. A 100% synthetic pixel will have a value of 100.	16-bit signed integer	0 - 100	1	0

Layer 2	<p>Number of days to closest pixel-level observation used to gap-fill (- is before, + is after prediction day)</p> <p>See <i>embedded metadata for the dates [yyyymmdd] of the closest imagery used to gap-fill</i></p> <p>900 (-900): Sentinel-2 (Landsat 8) data acquired on the prediction day used to gap-fill</p> <p>See <i>embedded metadata for the scene IDs of the L8/S2 data (if any) used to gap-fill</i></p>	16-bit signed integer	-900 to 900	1	0
Layer 3	<p>PlanetScope Cloud and cloud shadow mask</p> <p>1 = Clear</p> <p>2 = Bright Clouds</p> <p>3 = Shadows</p> <p>4 = Haze</p> <p>5 = Adjacent clouds</p> <p>6 = Additional cloud/shadow/haze elements based on a cross-scene correlation detection approach</p> <p>-999 = PlanetScope scene data not available</p>	16-bit signed integer	1 - 6, and -999	1	0
Layer 4	<p>PlanetScope pixel traceability mask (see <i>embedded metadata for scene IDs</i>)</p> <p>-999 = PlanetScope scene data not available</p>	16-bit signed integer	1 - 200, and -999	1	0
Layer 5	<p>Total number of FORCE L8/S2 reference scenes used during calibration</p> <p>-999 = PlanetScope scene data not available</p>	16-bit signed integer	0 - 500, and -999	1	0
Layer 6	Blue band uncertainty estimate (absolute percentage)	16-bit signed integer	3 - 200	1	0
Layer 7	Green band uncertainty estimate (absolute percentage)	16-bit signed integer	3 - 200	1	0
Layer 8	Red band uncertainty estimate (absolute percentage)	16-bit signed integer	3 - 200	1	0
Layer 9	NIR band uncertainty estimate (absolute percentage)	16-bit signed integer	3 - 200	1	0


QA layer 1 represents a raster mask indicating the percentage of “synthetic” (i.e., gap-filled) versus observed PlanetScope data used to produce each pixel value across the tile domain. A value of 0 indicates no gap-filling (i.e., 100% observed PS data) whereas a value of 100 indicates an entirely (100%) gap-filled pixel value. The pixel-specific weighting (0 - 100) in the neighborhood of gaps is done to ensure smooth and gradual transitions across interfaces between synthetic and observed data points.

Figure 4: QA layer 2 (gap-filling confidence metric) and PF-SR product for a tile east of Sacramento (CA) on a day with significant cloud and cloud shadow contamination. In QA layer 2, pixels containing real observation data will have a value of 0, whereas synthetic (gap-filled) pixels will be represented by the number of days since a valid observation (a negative/positive day value indicates an observation acquired before/after the given day).



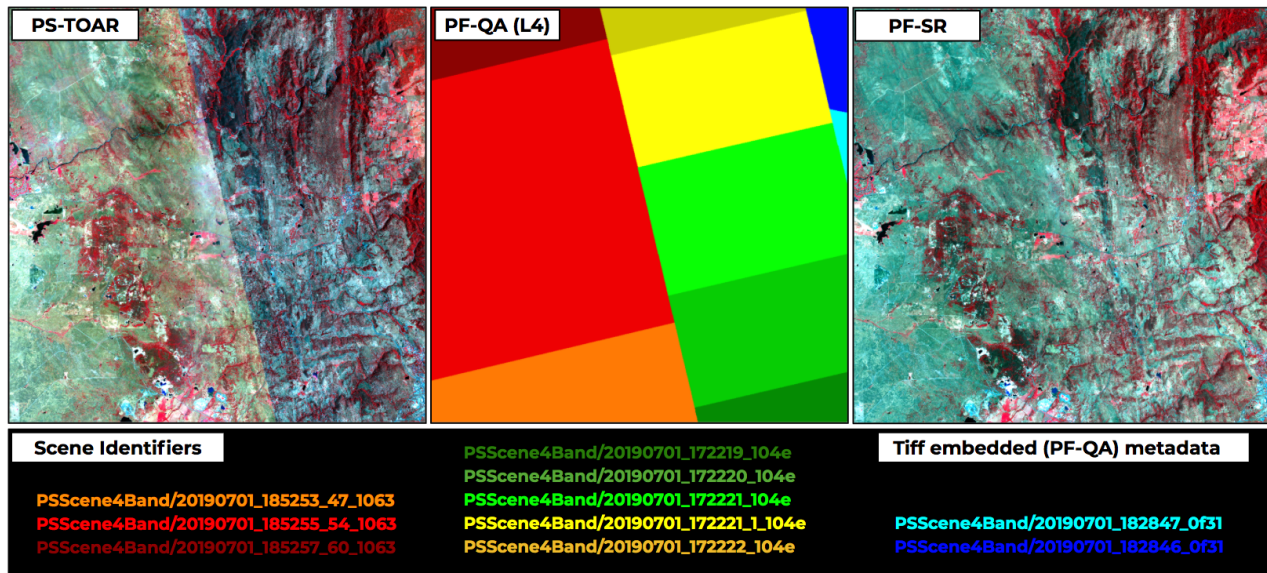
QA layer 2 is a gap-filling confidence metric, providing pixel-specific information on the number of days to a day with observed satellite data. The longer the day gap, the higher the uncertainty is likely to be in the gap-filled estimate although several factors will play a role in the actual retrieval uncertainty (see [Section 4.6](#)). This QA layer will also identify pixels where FORCE-based L8/S2 SR data (FLS-SR) on the prediction date have been used to help fill gaps resulting from clouds or gaps in PlanetScope coverage (Table 5). The actual dates of the closest images used to gap-fill have been included as tiff embedded metadata. The embedded metadata will also store the scene identifiers of the Sentinel or Landsat data (if any) used to fill gaps in PS coverage. Figure 4 exemplifies this metric for an AOI close to Sacramento (CA). Figure 5 depicts the associated cloud and cloud shadow mask.

Figure 5: QA layer 3 (cloud and cloud shadow mask) and input PS-TOAR product exemplified for the AOI used in Figure 4.



A pixel traceability mask is also provided (Figure 6). This raster layer identifies the footprints of the PlanetScope scenes used to produce any given tile image (cloudy or clear). Each domain is associated with a unique integer value that is linked to a scene identifier (i.e., Itemtype/scenelD) embedded as metadata in the QA geotiff. The scene identifier (e.g., PSScene4Band/20190701_172222_104e) provides the information needed to locate and access the source data through Planet’s API. The embedded metadata may be displayed using GDAL (e.g. gdalinfo name_of_file.tif).

Figure 6: QA layer 4 (pixel traceability mask) for the tile in Figure 4 (8000 x 8000 pixels) on July 1, 2019. In this case, a total of 10 PS ortho scenes from three separate strips were used to construct the tile image. The associated scene identifiers were extracted from the PF-QA embedded metadata. The visible seam lines in the PS-TOAR product result from merging Dove-R (reddish strip) and Dove-C (greenish strip) scene data with quite different spectral bands and Relative Spectral Response (RSR). Note that these transitions are not visible in the final PF-SR output.



QA layer 5 (Table 5) identifies the number of FLS-SR scenes (i.e., Sentinel-2 or Landsat 8) used for each pixel during reference sampling and calibration. In general the larger the number of reference scenes, the more robust the calibration is expected to be (see [Section 4.3](#)).

QA layers 6 - 9 provide band-specific confidence information for the “synthetic” (i.e., gap-filled) pixel values (a fixed 3% uncertainty is used for observed data). The process for deriving these is described in [Section 4.6](#). The confidence information is reported as an absolute percentage to indicate the deviation of the synthetic value from an actual observation. Accordingly, the absolute uncertainty in reflectance units can be calculated by multiplying the provided confidence information (x 0.01) with the corresponding surface reflectances recorded in the PF-SR product (Table 3).

3.3. SPATIOTEMPORAL ASSET CATALOG (STAC)

All Planet Fusion products are delivered with a SpatioTemporal Asset Catalog (STAC) file that conforms to the [STAC specifications](#). This STAC item contains information that summarizes the properties of the SR and QA products. This includes the version of Planet Fusion that the products were generated with, the dates that the products were created, and support for several STAC extensions.

STAC Extension	Description
Electro-Optical	Specification of the spectral bands contained with the SR product and descriptions of the information bands contained within the QA product
Projection	Coordinates representing the bounding geometry of the tile's footprint
Raster	Properties of the tile pixels

A STAC catalog is produced alongside all STAC items that contains an organizational structure to browse the STAC items. These STAC items and the static STAC catalog are delivered alongside the SR and QA products. A new STAC catalog with links to all new STAC items is delivered every time Planet Fusion products are delivered and overwrites any previously delivered STAC catalogs.

3.4. PROJECTION, GRIDDING, FILE NAMING, AND DELIVERY

Planet Fusion products are generated as regularly gridded raster tiles. Tiles have a 3 m pixel size, a 24 by 24 km extent (8000 pixel width and height), and are projected in the UTM zone intersected by their extent using the WGS-84 horizontal datum. Tile identifiers are based on a "{i}E-{j}N" template, where "i" is the zero-based easting index and "j" is the zero-based northing index using the origin of the UTM zone's coordinate reference system.

Tile Property	Description
Tile Size	24 km (8000 lines) by 24 km (8000 columns). Storage size varies with LZW compression. A typical Surface Reflectance tile consumes 300 MB. A typical Quality Assurance product consumes 1 to 5 MB.
Pixel Size	3 m
Spatial Reference	WGS-84 UTM zones based on tile intersection
Identifier	{i}E-{j}N Using coordinate system origin, i is the zero-based easting index, and j is the zero-based northing index.

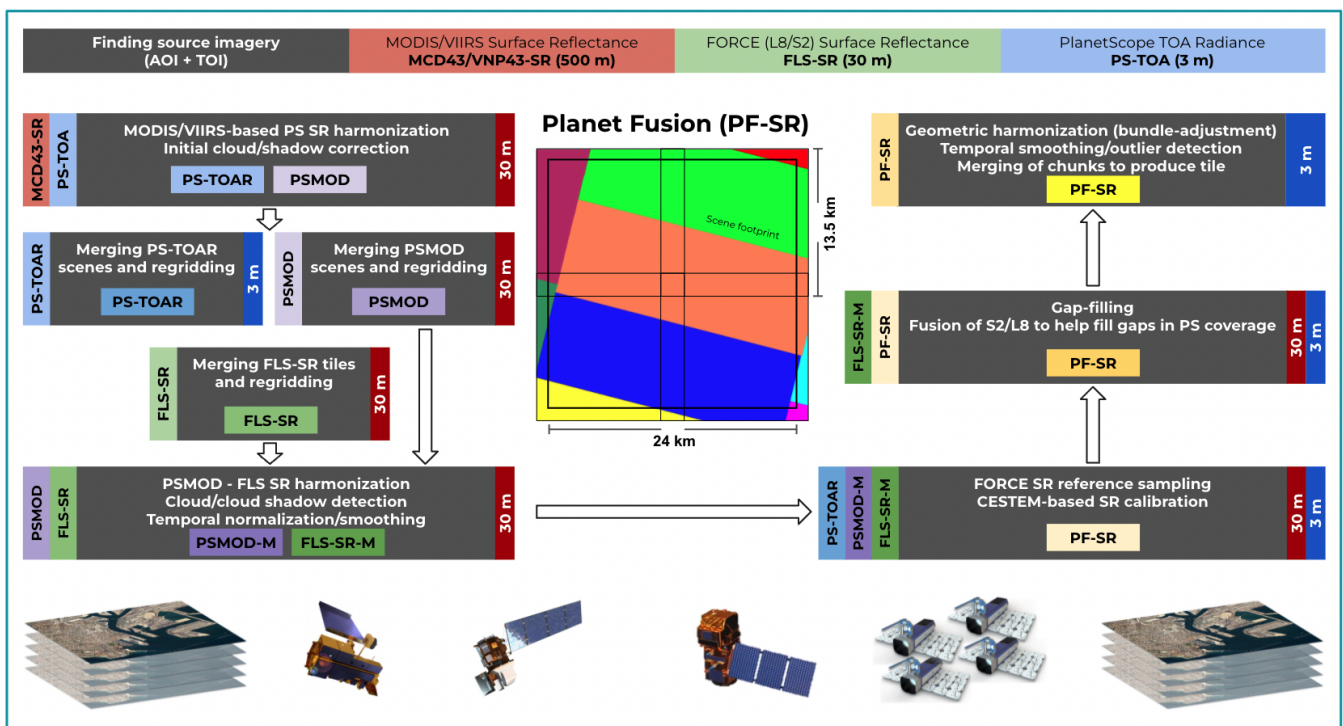
The naming of the generated output files depends on the delivery mechanism. In general, file names are composed of the tile identifier, the product key, and the date. See Table 2 for a description of each of the product keys.

For Google Cloud Storage delivery, products will be in a folder structure like "{scheme}/{zone}/{tile-id}/{product-key}/{date}.tif". For example, a Surface Reflectance tile for January 2, 2006 would be named like UTM-2400/15N/17E-192N/PF-SR/2006-01-02.tif (where UTM-2400 refers to the 24 km tiling scheme, 15N is the zone intersected by the tile, and 17E-192N is the specific tile id).

4. PLANET FUSION METHODOLOGY

The overall methodological elements of Planet Fusion processing are diagrammed in Figure 7. Planet Fusion products are based on an implementation of the CubeSat-Enabled Spatio-Temporal Enhancement Method (CESTEM), which has been described in detail in Houborg and McCabe (2018a, 2018b). Planet Fusion processing also includes significant refinements and additional functionality related to cloud masking, gap-filling, and sensor data fusion. Key elements of the approach and processing specifics are briefly outlined below.

Figure 7: A generalized overview of Planet Fusion processing elements for any given Fusion tile and TOI. Processing is done on slightly overlapping chunks (13.5 x 13.5 km) that are merged to produce the full tile (24 x 24 km). The diagram highlights the key intermediate and final product artifacts, the associated pixel resolution (3 or 30 m) and processing features (e.g., cloud masking, radiometric harmonization, geometric harmonization, gap-filling, time-series analyses, temporal filtering).



4.1. TILE (CHUNK)-LEVEL STACKING

After identifying the source imagery (PS-TOAR, FLS-SR, PSMOD) that overlap with a given Planet Fusion tile over a specified Time Of Interest (TOI), the respective input streams are stacked and re-gridded to the Planet Fusion tiling system ([Section 3.3](#)) with either a 3 m or 30 m pixel resolution (Figure 7).

In the case of the scene-based PlanetScope TOA reflectance (PS-TOAR) data, several scenes from multiple sensors may be overlapping with the tile domain (e.g., Figure 6). In order to retain the best data for the tile domain, priority is determined as a function of image quality category (prioritizing “standard” over “test” quality) (for more details on the image quality categorization see: Roy et al., 2021), initial cloud percentage (prioritizing scenes with the lowest cloud percentages), sun elevation (prioritizing scenes with the highest sun elevation), and scene overlap with the tile domain. The scene to tile conversion will also prioritize merging of scenes acquired from a single satellite in a single pass (i.e., strips) to reduce seam lines and spatial discontinuities introduced by cross-sensor inconsistencies. A phase correlation technique (see [Section 4.5](#)) is used to ensure that the PlanetScope scenes are geometrically aligned/co-registered (with sub-pixel precision) before compositing the tile. In addition, the re-aligned PS-TOAR scenes are brightness harmonized across the tile domain to facilitate a seamless Surface Reflectance calibration ([Section 4.3](#)). The brightness harmonization utilizes the clear-sky overlap between scenes to derive band-specific regression coefficients that are used to normalize the reflectance magnitudes across the scenes.

A 30 m resolution stack of MODIS/VIIRS Surface Reflectance-calibrated PlanetScope imagery (PSMOD) is also produced. The MODIS/VIIRS calibration is performed at the PlanetScope scene level using day-coincident MCD43/VNP43 NBAR products (Table 1) as the radiometric reference, translating PS-TOAR into MODIS/VIIRS-consistent SR data (Houborg and McCabe, 2018ab). The scenes are prioritized as described above for PS-TOAR when producing the tile. This also involves co-registration of the 30 m PSMOD scenes prior to merging and tile composition.


The FORCE-based Surface Reflectance tiles (FLS-SR) are mapped onto the Fusion tiling grid in a similar way.

Planet Fusion processing is done on slightly overlapping (i.e., a buffer of 750 m) chunks with a 13.5 by 13.5 km extent. The tile products are generated by merging data from 4 chunks, utilizing a gradual weighting and harmonization approach across chunk overlap zones to avoid visible boundaries in the final outputs.

4.2. CLOUD AND CLOUD SHADOW MASKING

Planet Fusion cloud and cloud shadow detection (Figure 8) is performed at 30 m resolution using a temporally-driven approach that takes advantage of PlanetScope and FORCE (L8, S2) SR information in a synergistic way. The approach is iterative and initialized with the original scene-based cloud and cloud shadow masks associated with both the PlanetScope and FORCE data. For PlanetScope, the initial cloud and cloud shadow mask is produced during the MODIS/VIIRS-based calibration stage (described above) informed by cloud 2.0 (UDM2) detections, when available. A rigorous cloud verification approach is implemented at this step in an attempt to reduce commission/omissions errors.

Figure 8: Example of PF cloud (white) and cloud shadow (black) detections over a region in Bolivia. Yellow represents the applied buffer around the detections.



The temporally-driven detection approach can accommodate PS and FORCE data acquired at different times on the same day and includes geometric harmonization of the multi-source input data in order to improve detection results and reduce commission issues resulting from pixels not being properly aligned. It utilizes clear-sky observations (as identified by the initial cloud masks during the first iteration) over a flexible temporal window (up to ± 1 year) for any given 30 m pixel to flag spectral outliers potentially resulting from cloud or cloud shadow contamination. During this process, clear-sky background images are created for each acquisition day over the defined TOI which are then used in combination with the actual acquisition images to identify spectrally distinct classes using an unsupervised K-means clustering approach. A suite of spectral difference metrics, such as the difference in red and NIR reflectance between the background and actual imagery, combined with a set of carefully defined spectral thresholds and a number of other constraints are then used to classify each cluster as clear, cloud, cloud shadow, or haze. Additional cross-correlations tests between the background and actual imagery serve to further resolve and label any cloud, cloud shadow, and haze contaminations. Furthermore, a series of automated techniques are implemented to verify these classifications and avoid (to the extent possible) masking out actual change. This includes the integration of hillshade information to help reduce commission issues in landscapes with significant terrain shadowing. In periods with prolonged cloudiness, the final cloud mask will reflect the initial scene-based mask to a greater extent in order to avoid excessive cloud commission issues. The cloud detected areas are expanded using a buffer zone (i.e., adjacent cloud domain) in order to ensure that most of the contaminated pixels are removed from the final outputs. The outlined processing is repeated once in order to take advantage of the updated cloud mask in the construction of the clear-sky background images, which serve as critical inputs to reliably classifying cloud and cloud shadow clusters in the acquired imagery on any given day. The scheme will refine and update the scene-based cloud masks (30 m) associated with both the PS (PSMOD-M) and FORCE (FLS-SR-M) data (Figure 7).

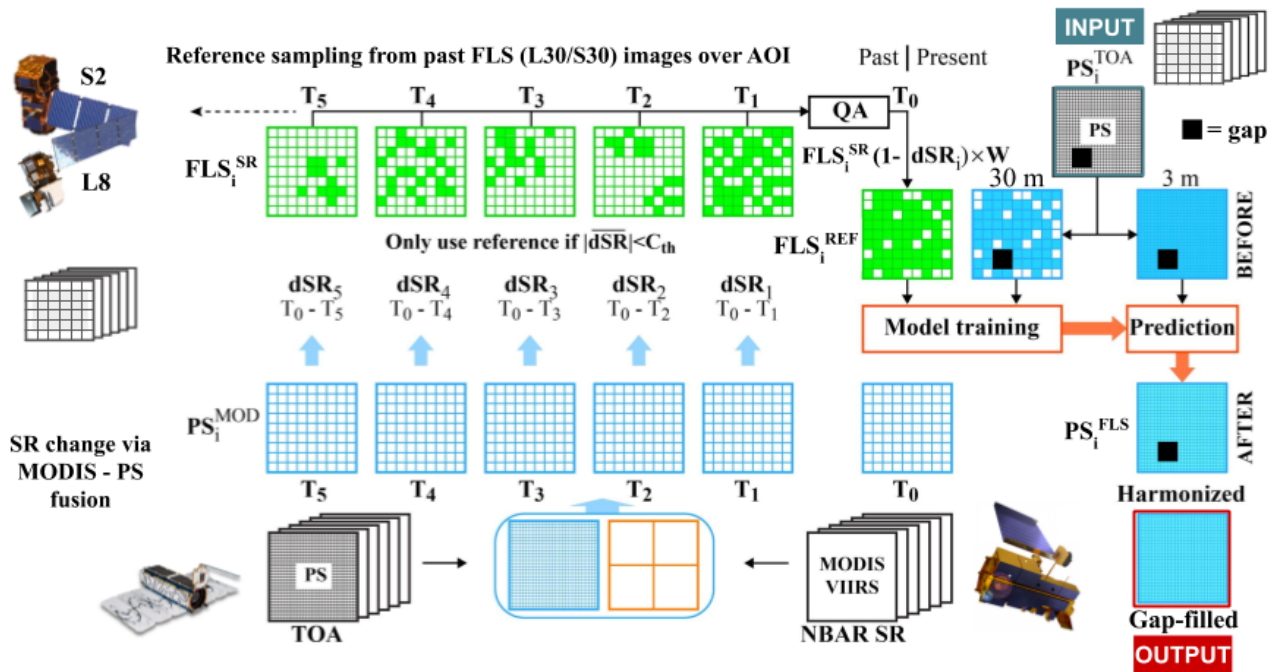
4.3. REFERENCE SAMPLING AND CROSS-CALIBRATION

As its core, the Cubesat-Enabled Spatio-Temporal Enhancement Method (CESTEM) serves as a flexible mechanism to harmonize multi-sensor spectral data into a consistent radiometric surface reflectance standard (i.e., the “gold standard”). The FORCE-based surface reflectance product (30 m) (Table 1; FLS-SR) is currently adopted as the gold standard. CESTEM is characterized by a number of unique features:

- Does not require day co-incident acquisitions (PS versus L8/S2) for cross-calibration/harmonization

- Implements a secondary MODIS/VIIRS-based harmonization to correct for surface reflectance changes occurring over given PS and L8/S2 acquisition time spans
- Effectively minimizes uncertainties associated with both PS and L8/S2 data via temporally-driven outlier detection
- Can harmonize data from sensors with contrasting spectral bands and Relative Spectral Responses
- Is largely insensitive to noise (e.g., calibration uncertainties) in the input data (e.g., PS-TOA)

Figure 9: Diagram of the CESTEM radiometric harmonization framework, translating image stacks of PlanetScope TOA radiance (PS^{TOA}) into FORCE-consistent Landsat 8/Sentinel-2 Surface Reflectance (PS^{FLS}).



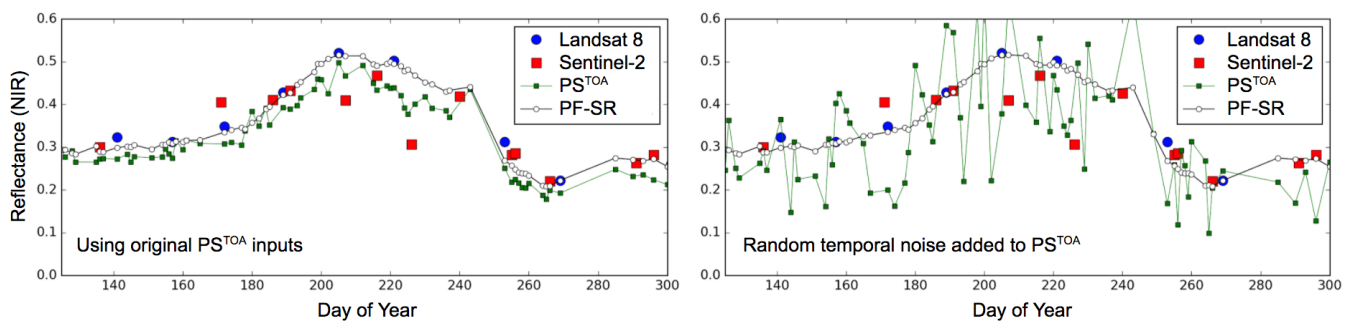
The CESTEM-based harmonization is applied to all images acquired over a defined Time Of Interest (TOI) drawing spectral (i.e., blue, green, red, and NIR) reference data from a pool of L8/S2 SR images acquired over a predefined “calibration window” (typically two years centered around the TOI). In order to reliably use past or “future” FLS images for calibration purposes they must be associated with a day-coincident PS acquisition. Critical to this process is the use of MODIS/VIIRS-consistent (i.e., MCD43/VNP43) PlanetScope data (PS^{MOD}) to quantify relative surface reflectance changes over given PS - L8/S2 acquisition time spans (Figure 9). It follows that data from multiple L8/S2 acquisitions will be sampled to generate a given PS coincident calibration reference image (FLS^{REF}), using weights derived as a function of PS - L8/S2 acquisition time spans and the magnitude of surface reflectance change relative to the prediction day (Figure 9). In addition, the multi-temporal FLS inputs are quality assured (QA) during the sampling step and outliers removed. Importantly, the scheme can handle significant lags between the PS imagery to be harmonized and suitable FLS images as the calibration references can be sampled from images in the past or future (relative to the prediction date). As a result, the harmonization approach will continue to perform well over extended periods of cloudiness and FLS unavailability as long as a sufficient number of good L8/S2 scenes can be identified within the “calibration window.” The associated layer 5 metadata (Table 5) will keep a pixel-specific record of the number of reference scenes available on any given day.

The multi-sensor and multi-time sampling approach effectively minimizes potential issues and uncertainties (e.g., atmospheric contamination, cloud masking, BRDF effects, calibration inaccuracies) associated with both the PS and L8/S2 data to create a very robust and temporally consistent radiometric reference. With this in

place, a multivariate linear regression and decision tree approach (Houborg and McCabe 2018a) is employed to learn non-linear scene, sensor, and band-specific translational associations. The resulting models are then used to convert PS TOA (PS^{TOA}) reflectances into FLS-consistent SR (PS^{FLS}). During this process, a number of techniques are implemented to avoid overfitting and to preserve band-specific textural features and spatial gradients present in the original 3 m resolution PS imagery.

The CESTEM-based radiometric harmonization has been designed to be largely insensitive to temporal inconsistencies associated with the input data (PS^{TOA}), which may result from calibration uncertainties and cross-sensor spectral differences. In fact, as showcased in Figure 10, adding significant random temporal noise to the PS input data has a largely indistinguishable impact on the harmonized results. Another noteworthy feature is the low sensitivity to issues with the reference data, in this case resulting from cloud omission errors associated with the S2 data (Figure 9).

Figure 10: Illustrating the robustness of the CESTEM harmonization to noise in the input data streams (FLS and PS)



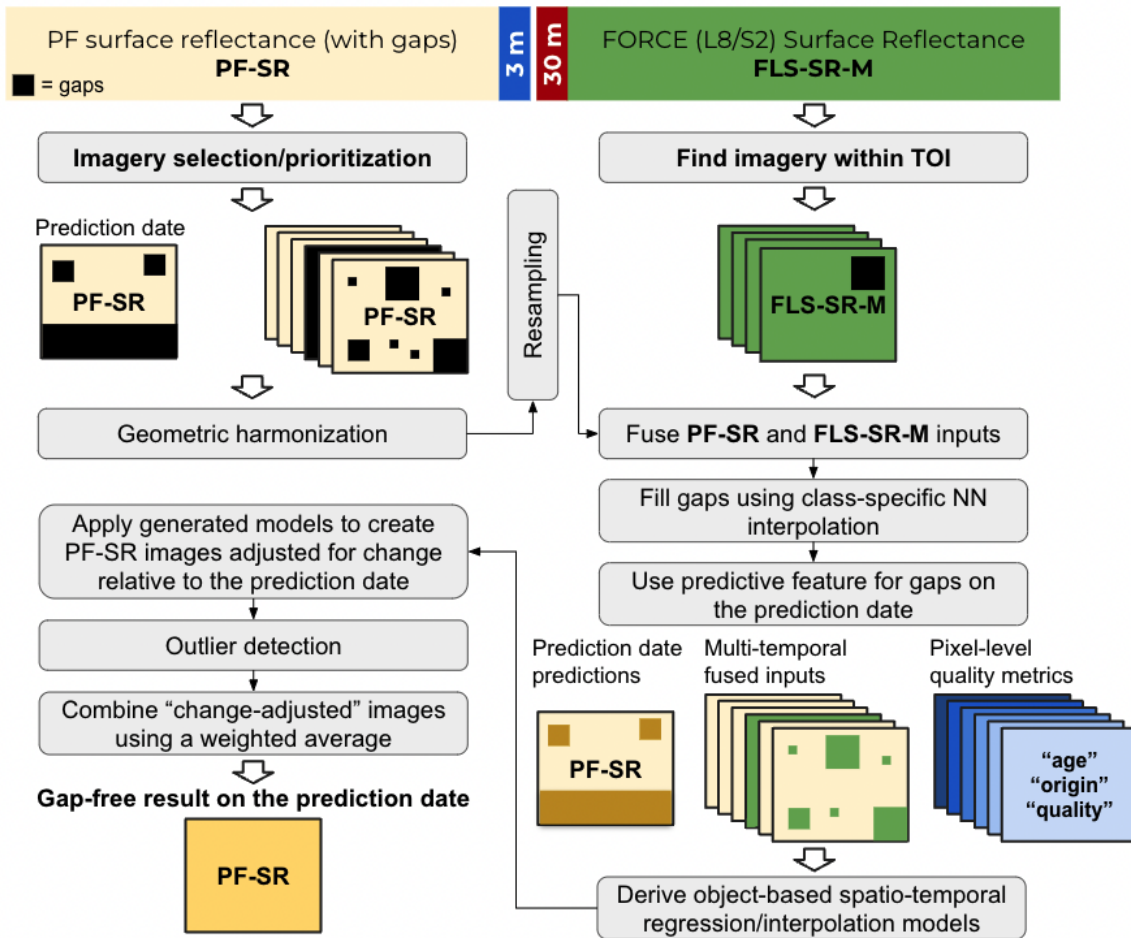
4.4. GAP-FILLING AND TEMPORAL FILTERING

The Planet Fusion gap-filling ensures a spatially complete and temporally continuous (daily) 3 m product irrespective of the actual acquisition coverage and cloud conditions. As this involves generation of estimated (synthetic) radiometric data, users are advised to utilize the associated QA metrics (Table 5) as measures of the synthetic data retrieval confidence. In general, the longer the daily interval gap is to the actual observation data (QA layer 2) the larger the overall uncertainty is in the estimated gap fill values. Several other factors, such as surface characteristics, crop type, and vegetation dynamics are also thought to play an important role in shaping the uncertainty associated with estimated Planet Fusion data values. These interacting factors will be reflected by the derived confidence estimates to some extent (Table 5 and [Section 4.6](#)).

Currently, the gap-filling subprocess (Figure 11) is informed by cloud masked (but not yet gap-filled) Planet Fusion SR (PF-SR) data in combination with complementary FORCE (30 m) L8/S2 data (FLS-SR-M) available within a temporal buffer centered on the prediction date. Data from up to 6 Planet Fusion tiles are typically used for this purpose, which may originate from both before and after (if available) the prediction date.

The PF-SR chunk-level images to be used for filling the gaps are selected as a function of image quality category (prioritizing “standard” over “test” quality), cloud percentage (prioritizing scenes with the lowest cloud percentages), and time lag relative to the prediction date (prioritizing images acquired with the shortest lag). The selected input imagery is geometrically harmonized (i.e., to ensure pixels align across the stack) before undergoing a series of anomaly-based quality checks to flag (and mask out) pixel-level outliers resulting from e.g. residual cloud contamination.

Figure 11: Flowchart of 3 m (left side) and 30 m (right side) processing elements in the PF gap-filling module.



These images will then be resampled to 30 m in order to perform an object-based spatio-temporal interpolation. FORCE (L8/S2) data (30 m) acquired within a short temporal buffer of the prediction date will be used to inform the interpolation, if available. In addition, near-coincident FORCE data will be queried to help fill gaps in the resampled PF-SR images. Any remaining gaps will be filled using a class-specific nearest neighbor interpolation technique with confidences assigned as a function of the spatial proximity to an observed pixel value. The resulting multi-temporal "fused" inputs are used to generate a detailed (i.e., ~200 - 300 classes) classification (via k-means clustering), which is used to derive class-level surface reflectance change information relative to the prediction date for each of the selected PF-SR images. Class and band-specific regression/interpolation models (first or second order) are derived using class aggregated multi-temporal SR information with weights assigned as a function of a number of metrics including the pixel "age" relative to the prediction date, the pixel "origin" (e.g., observed or nearest neighbor spatially interpolated), and the pixel "quality" (Figure 11).

A predictive component is invoked when data availability after the prediction date is limited, which can occur during extended periods of cloudiness or when running Fusion in forward-fill mode (Section 4.7). This feature predicts surface reflectances on the prediction date from past imagery via pre-established class-level spectral trajectories. It works by first computing annual land surface spectral trajectories (e.g., phenological time series) based on past Planet Fusion SR observations (e.g., past growing season) for a suite of spectrally distinct land cover classes (derived via k-means clustering using normalized multi-temporal inputs for the year in question).

These are then matched and fused with the spectral trajectories representative of the most current observations using dynamic time warping to find the optimal non-linear alignment between time series, and subsequently used to extrapolate/project the spectral information content from the time of the last acquired imagery to a given prediction date. This predictive information feeds into the object-based spatio-temporal interpolation scheme described above in order to better guide the interpolation when there is no data available beyond the prediction date (Figure 11). This projective feature is critical to ensuring realistic predictions when producing Planet Fusion data close to present time (i.e., when running in forward-fill mode).

The relative SR change information described above is then applied to the 3 m resolution PF-SR images and the “change-adjusted” images are combined using a weighted average to fill the gaps on the prediction date (Figure 11). The gap-fill confidence metric (QA layer 2, Table 5) will only record the number of days to the closest pixel-level observation even though data observed on other dates may also feed into the synthetic data estimates albeit with a rapidly decreasing weight as the number of days passed since the prediction date increases. The PF-QA product will provide full traceability of the source data used to fill the gaps, which includes tiff embedded metadata on the dates of all the images used to inform the gap-filling.

A conservative three-point (i.e., 3-day) temporal smoothing algorithm is applied to the resulting gap-free daily image stack in an attempt to correct for any remaining anomalous spectral behavior. This algorithm will only invoke smoothing on pixels flagged as “anomalous” based on spectral comparisons with corresponding pixel retrievals from the two bounding images (i.e., ± 1 day). The smoothing weights are determined as a function of the magnitude of the spectral differences and take into account neighborhood information and spatial texture in order to avoid introducing spatial discontinuities/artifacts. Figure 12 exemplifies the process from PS-TOA reflectance to gap-free surface reflectance imagery using the outlined methodology.

Figure 12: PF cloud screening and gap-filling exemplified over the Orange River region in South Africa. It shows the original PS TOA reflectance data on the left and the resulting gap-free PF Surface Reflectance data on the right.


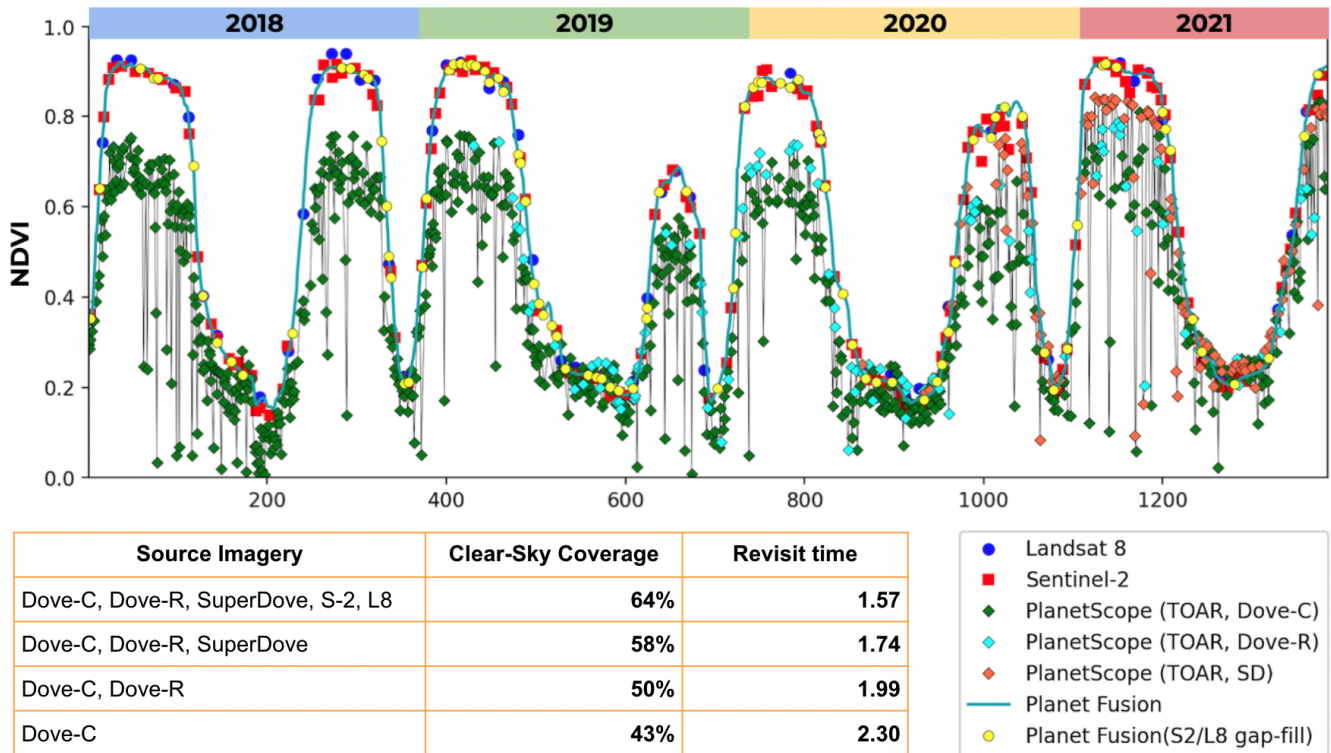


Figure 13: Individual sensor (Dove-C, Dove-R, SD, S2, L8) contributions to a gap-filled (daily) PF time series of NDVI over the Orange River region in South Africa. A breakdown of the corresponding clear-sky contributions (100% = daily coverage) and revisit times (1 = daily coverage) is also provided.



The integration of data from multiple sensor sources enhances the ability to track high frequency land surface dynamics. Figure 13 shows contributions from individual sensor sources (Dove-C, Dove-R, SuperDove, S2, and L8) over a 4-year period for the Orange River region in South Africa. Combining all sources gives a clear-sky coverage (100% = daily coverage) of 64% which translates into a revisit time of 1.57 (1 = daily coverage). Using only PlanetScope, the clear-sky coverage amounts to 58% and contributions from the different PlanetScope generations are shown to be highly time dependent (e.g., only Dove-C contributes in 2018 whereas the SuperDoves become dominant in 2021). When L8/S2 is used to gap-fill (i.e., the yellow circles) the clear-sky coverage increases to 64%, which translates into an additional 82 observation points over the TOI.

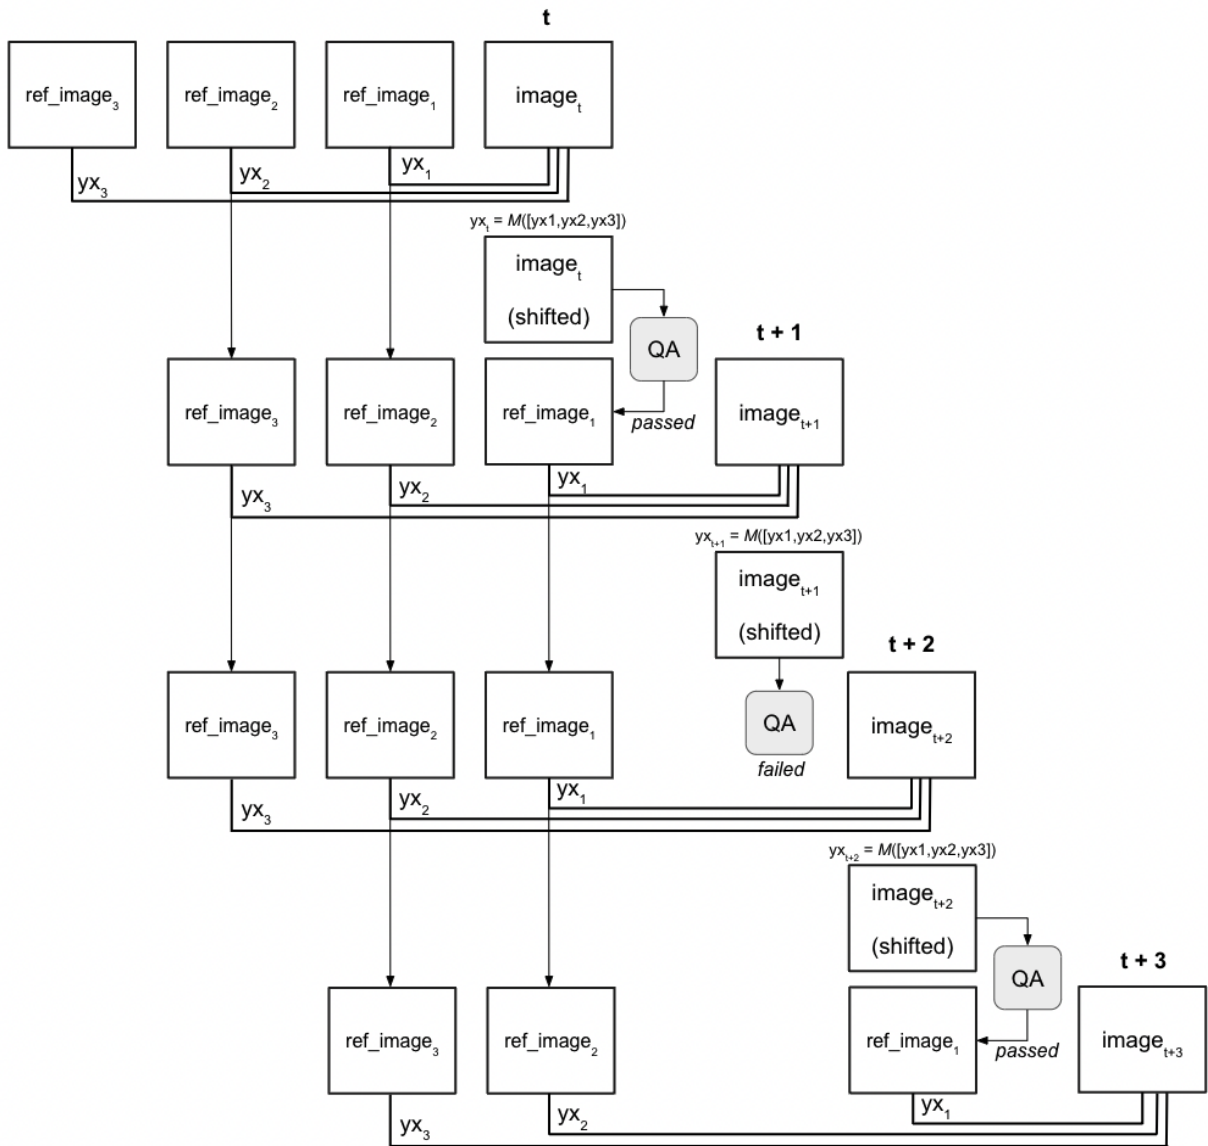
4.5. GEOMETRIC HARMONIZATION

The input imagery that feeds into Planet Fusion (e.g., L8, S2, PS) has been orthorectified using rigorous preprocessing protocols with a positional accuracy typically better than 10 m RMSE. Nevertheless, perfect image to image alignment is difficult to achieve, particularly when combining data from disparate sensor sources. As Planet Fusion relies heavily on taking advantage of temporal information content for calibration, gap-filling, and smoothing, precise co-registration and sub-pixel fine alignment of stacked imagery becomes a critical component of sensor data fusion and harmonization.

We use a phase correlation technique (Guizar-Sicairos et al., 2008) to detect the global shift between two images with sub-pixel precision at various stages of Planet Fusion processing. A PlanetScope chunk (13.5 km x 13.5 km) typically combines scenes from multiple PlanetScope sensors (Figure 6), which may sometimes be slightly mis-aligned. The clear-sky overlap (if any) within a strip or between strips (i.e., a strip signifies the set of

scenes acquired from a single satellite in a single pass) is used to assess sub-pixel shifts on a band-specific basis. If valid shifts are encountered they are applied using a Fourier transformation approach. This approach should ensure that the PlanetScope scenes within the chunk are geometrically harmonized before merging. However, it will not correct for mis-alignments within the scene due to registration issues when creating the scene composite from raw frames (done upstream of Planet Fusion processing). MODIS/VIIRS and FLS data are sub-pixel aligned to the coincident (or near-coincident) PlanetScope scene using a similar approach except that shifts are only evaluated based on the red reflectance band.

Figure 14: Schematic diagram of the bundle adjustment-based approach exemplified for a series of 4 sequential (daily) images. A triple reference image approach is used to reduce the risk of introducing drifts in the derived pixel shifts (y, x). In addition, the images undergo stringent quality control (QA) before being used to update the reference imagery.

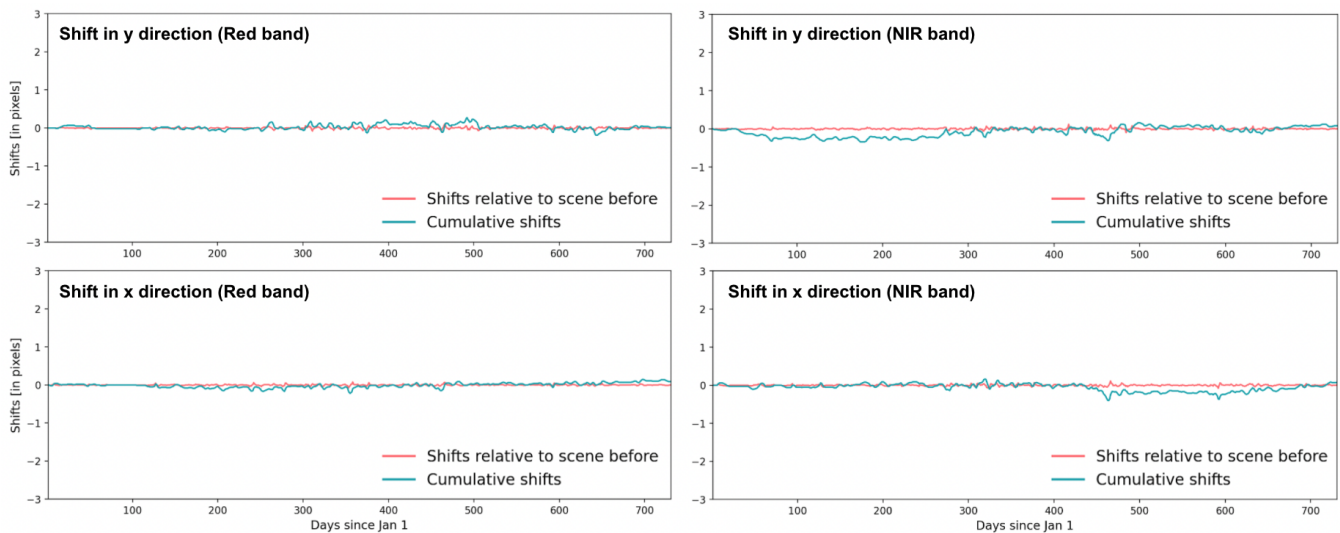


Geometric harmonization is also performed during cloud and cloud shadow masking and gap-filling to ensure proper pixel alignment before performing the temporally-driven pixel-level operations as previously described. After gap-filling a rigorous bundle adjustment-based approach is used to align the images in the chunk-level deep stacks (Figure 14). This approach utilizes the same phase correlation technique to sequentially align

images in the stack using a set of moving reference images. It works by 1) going through the images in the temporal stack starting from the second image ($t=2$), 2) calculating pixel shifts (y, x) of image_{*t*} relative to three reference images in the past (ref_image₁, ref_image₂, ref_image₃), 3) applying the median shifts to image_{*t*}, 4) updating the reference images sequentially (i.e., replace ref_image₁ with image_{*t*}, ref_image₂ with ref_image₁, and ref_image₃ with ref_image₂), and 5) calculating the pixel shifts for the next image (image_{*t+1*}) in the stack using the updated reference images (Figure 13). Importantly, the image to be shifted will undergo stringent quality checks (QA field in Figure 13) (including checks for within chunk registration issues) before being used to update the reference images, in order to not introduce drifting in the pixel shifts. Once the pixel shifts have been computed across the entire temporal stack, the shifts will be adjusted to align to the median state rather than using the first reference images in the stack as the reference state. As this processing can affect the absolute registration accuracy of the realigned images, it is followed by co-registration against original PS TOA reflectance images (PS-TOAR). For this purpose, shifts are evaluated against a series of PS-TOAR images acquired on dates with predominantly clear-sky observations. The resulting median shifts (y, x) are then added to the “relative” shifts derived during the bundle adjustment in order to inherit the absolute geometric accuracy of the native PlanetScope data.

Figure 15 demonstrates the results of the outlined geometric harmonization over Eden Landing in the San Francisco Bay Area showing the shifts (y, x) relative to the scene before as well as the cumulative shifts (drifting indicator). As seen, the approach effectively corrects for any alignment issues without introducing significant drifts in the shifts over the 2-year study period.

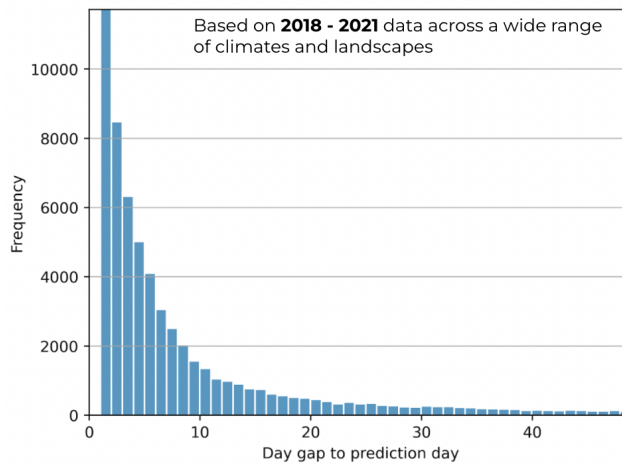
Figure 15: A temporal analysis of shifts (pixels) in y and x direction for a two-year period over Eden Landing (CA) after applying the outlined geometric harmonization steps.



4.6. CONFIDENCE INFORMATION

The QA file provides band-specific confidence information for synthetic surface reflectance retrievals (Table 5). The band-specific confidence information is reported for each 3 m pixel as an absolute percentage, which can be used as an approximate assessment of the uncertainty of the estimate relative to a real surface reflectance observation.

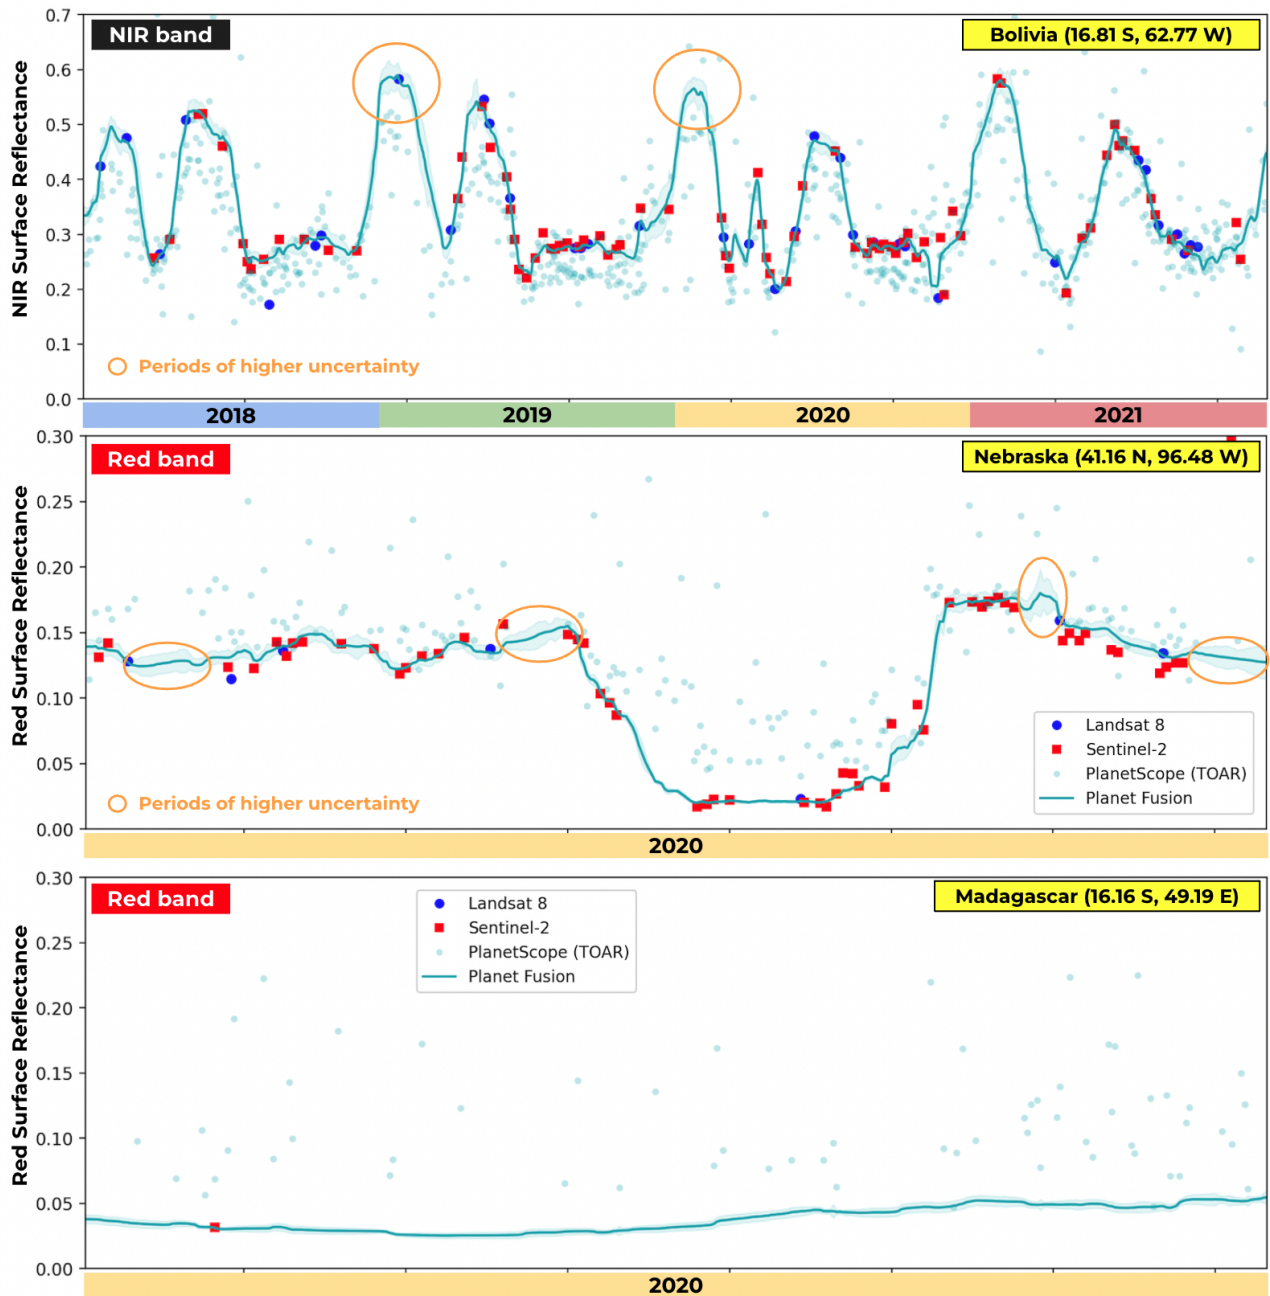
Figure 16: Histogram showing the frequency of absolute day gaps (relative to the prediction date) in the multi-AOI training dataset used to generate the confidence models.



The models used to derive the band-specific confidence information are based on comparisons between observed and gap-filled SR data over a selection of representative AOIs and TOIs covering a diversity of surface characteristics and cloud conditions (Figure 16). The training data were generated by running the Fusion gap-filling module with any valid observation data on the prediction date removed. The absolute relative difference between the observed and gap-filled surface reflectance (for the same pixels) serves as the target variable (i.e., the uncertainty), which is related to a set of explanatory variables using a data mining technique. The explanatory variables include 1) the 4-band spectral data, 2) NDVI, 3) the number of days to the closest observation used to gap-fill (Figure 16), and 4) a surface reflectance change metric. Models trained for each spectral band are then used to derive the confidence information for each Planet Fusion tile as a function of the day and tile-specific explanatory input dataset.

Figure 17 showcases the time series of derived confidence estimates over AOIs characterized by significantly different cloud conditions and cover types. Over agricultural fields in Bolivia (top panel) and Nebraska (middle panel) periods of higher predicted uncertainty generally coincide with more extended gaps in the PS observation record resulting from clouds or snow due to an expectedly tight coupling between SR uncertainties and the “age” (relative to the prediction date) of the pixels. However the red reflectance time series over a forested pixel in Madagascar has a relatively low associated uncertainty (bottom panel), despite the extreme cloud environment and significant pixel “ages” (e.g., only one clear-sky S2 observation was available over the one year period). This seemingly contradictory information can be explained by cover type differences (from agriculture to forest). The spectral signature of an evergreen forest canopy is relatively stationary in time relative to the spectral signature of a crop canopy going through various development stages (e.g., planting, green-up, maturity, senescence, harvest), which will tend to make gap-filling over such forested pixels less uncertain even during extended periods of clouds. The explanatory surface reflectance change variable is effective at capturing such subtleties during confidence model training. Still, synthetic pixel uncertainties are controlled by a complex interaction of several factors, and the provided confidence estimates should only be considered as a rough approximation with a fair degree of uncertainty at this time.

Figure 17: NIR and red band time series results over agricultural and forested AOIs in Bolivia (top), Nebraska (middle), and Madagascar (bottom) with associated confidence intervals (the colored area around the Planet Fusion solid line).



4.7. BACKFILL VERSUS FORWARD-FILL OPERATION

Planet Fusion can be run in either backfill or forward-fill mode. Backfill signifies a run over a time of interest in the past. A backfill run over a 2 year period is always initially required in order to establish deep temporal image stacks to inform the calibration, cloud masking, and gap-filling processes. In forward-fill, we run as close to present time as possible on a daily basis, processing any new imagery that has become available since the last job execution utilizing information from the deep temporal stacks generated during the backfill for cloud masking and gap-filling purposes. We currently target a 48 hour latency of delivering Planet Fusion tiles during forward-fill operation, which is a function of the latency of the input sources (i.e., primarily PlanetScope and MODIS/VIIRS) and processing time.

Gap-filling is notably different during backfill and forward-fill operation. While data is typically available both before and after a given prediction date in backfill operation, a forward-fill run is intrinsically constrained in the forward looking direction, which tends to increase the uncertainty of synthetic pixel retrievals (as you don't know what is "around the corner"). As mentioned in [Section 4.4](#), Planet Fusion integrates a projective element that exploits spectral trajectories/phenologies from past observations to make informed predictions in forward-fill mode. This projective feature is critical to ensuring realistic predictions when producing Planet Fusion data close to present time.

5. UNCERTAINTY ESTIMATES

This section provides a preliminary assessment of Planet Fusion uncertainties related to both radiometric harmonization and synthetic data generation (i.e., gap-filling). A more elaborate analysis involving ground-based spectrometer data will be documented in a future report.

Figure 18 illustrates the CESTEM-based radiometric harmonization for two AOIs in California (top) and Nebraska (bottom). The harmonization is able to robustly re-calibrate the PlanetScope input stream, producing Planet Fusion NDVI values that align well with day-coincident L8/S2 NDVI with a mean absolute difference (MAD) of ~4%. Noteworthy, is the bimodal distribution of the PlanetScope based NDVI relative to L8, which results from having sensors with different spectral bands and relative spectral responses featured in the tile. This showcases the promise of the CESTEM approach for accounting for such non-linearities in spectral associations between the input and reference (i.e., L8 or S2) stream. While CESTEM can ensure radiometric consistency with the chosen "gold reference" (i.e., FLS currently), a perfect 1:1 agreement will not always be guaranteed. This consistency results from the advanced Planet Fusion "clean up" process involving 1) enhanced cloud and cloud shadow detection ([Section 4.2](#)), 2) multi-sensor and multi-time reference sampling ([Section 4.3](#)), and 3) temporal filtering ([Section 4.4](#)), which effectively minimizes uncertainties (e.g., due to atmospheric contamination, BRDF effects, calibration inaccuracies) impacting both the PlanetScope and L8/S2 data streams. As a result, Planet Fusion data will in general be characterized by enhanced radiometric robustness and spatio-temporal consistency.

Figure 18: Quality (relative to L8/S2) of the Planet Fusion NDVI harmonization process showcased for two regions with day-coincident PlanetScope and L8/S2 acquisitions

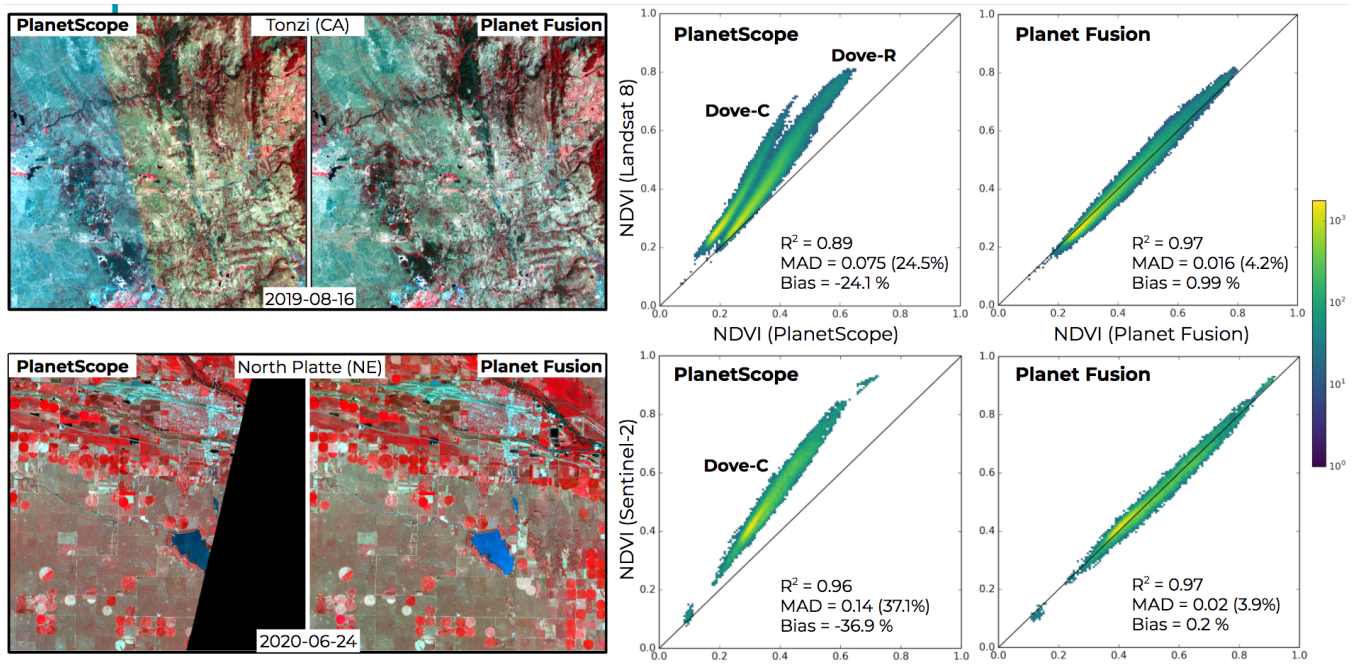



Figure 19: Real versus synthetic (gap-filled) Planet Fusion data for different observation gap intervals (false color: NIR, Red, Green)



Regarding the synthetic (gap-filled) component of Planet Fusion data, users are advised to consult with the included QA metrics (Table 5) to determine the synthetic data retrieval confidence. Figure 19 above demonstrates the fidelity of the gap-fill data values for different observation gaps in an agricultural region in Nebraska. In general, the fine-scale features and reflectance magnitudes in the actual data (top panel) are accurately reproduced in the synthetic data (bottom panel). However, as the observation gap increases further you are likely to see greater reflectance discrepancies, particularly if surface conditions change in a non-linear

fashion. However, given the availability of near-daily CubeSat imaging, the acquisition/observation gaps between clear-sky images will often be within reasonable limits over the majority of landscapes and environments.

Figure 20: Box-and-whisker plots of red (left) and NIR (right) band relative mean absolute differences [%] between synthetic and real observations as a function of the day gap to the observed data. Results were compiled from Fusion runs over a wide variety of geographies with day gaps ranging from 1 to 200 days.

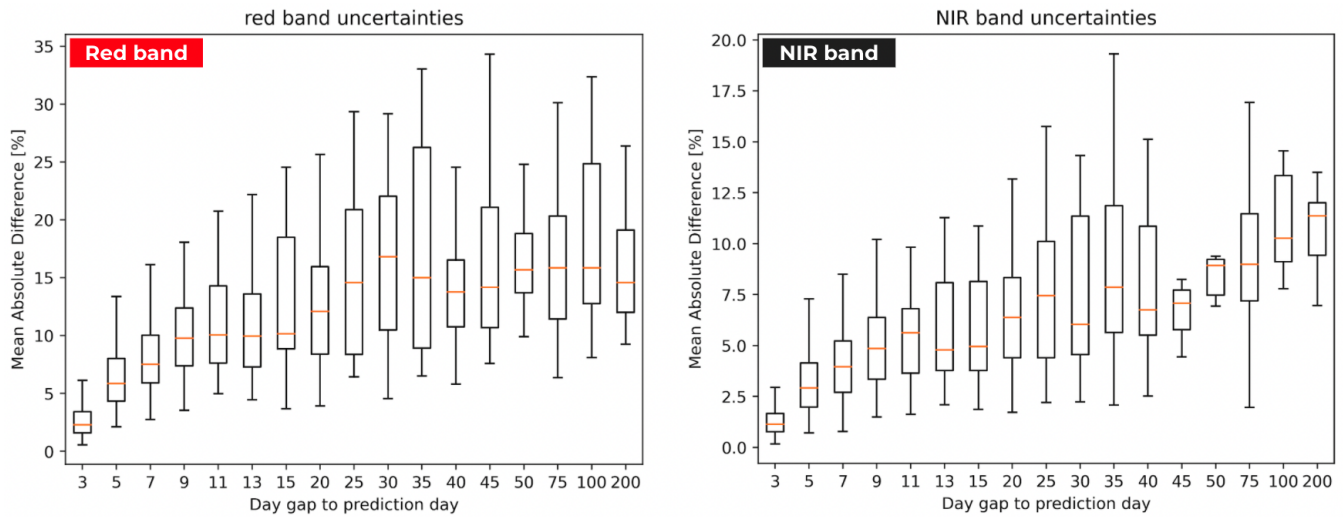


Figure 20 shows a box-and-whisker representation of band-specific relative MADs between synthetic and observed data as a function of the observation gap, which was compiled from the geographically diverse dataset used to train the confidence models (Section 4.6). These plots display the minimum, lower quartile, median, upper quartile, and maximum MAD across a wide range in the observation gaps (1 - 200 days). The median MAD shows an increasing trend from around 2.5% (1.5%) to 15% (7.5%) for the red (NIR) band as the observation gap increases from 1 to ~25 days. Beyond that the median MAD more or less plateaus although the range and quartiles of MAD continue to show significant variations.

6. KNOWN LIMITATIONS AND CAVEATS

- **False cloud/shadow detections** may occur in certain cases: 1) if surface conditions change very rapidly, 2) during prolonged cloudiness, or 3) over AOIs with complex terrain and shadowing. Significant effort has gone into developing automated techniques to differentiate between actual change and atmospheric contamination, but in some cases commission errors can still be an issue.
- **Gap-filling artifacts** may occur during periods of 1) prolonged cloudiness and 2) snow.
- **Planet Fusion is not suited for studies over snow covered surfaces.** As it is virtually impossible to robustly distinguish between snow and clouds based on 4-band VNIR data, periodic snow cover will in most cases be masked out as clouds. Snow covered scenes will be down-prioritized during gap-filling as the mixing of spectral signals from snow and no snow conditions can introduce severe artifacts. As a result, Planet Fusion reproduced surface reflectance signals may be associated with large uncertainties during times when snow is expected. It also means that occasionally you may see very large observational lags (QA layer 2, Table 5) as the gap-filling algorithm has difficulty identifying a snow free scene.
- As the Fusion processing is done independently for each tile, **tile boundary artifacts** can sometimes occur when doing cross-tile analyses. This may be manifested in the form of visible tile boundaries, small brightness differences, and small (typically sub-pixel) pixel mis-alignments. These artifacts will tend to be more prevalent for gap-filled data.
- The current Planet Fusion product suite is **4-bands only**. However, the general framework is extendable to 8-band SuperDove data (our Next Generation Doves) with some adjustments. The capability for additional spectral bands is expected to become implemented within the foreseeable future.

REFERENCES

- B. Aragon, M.G. Ziliani, R. Houborg, T.E. Frantz. 2021. CubeSats deliver new insights into agricultural water use at daily and 3 m resolutions. *Sci Rep* 11, 12131. <https://doi.org/10.1038/s41598-021-91646-w>
- M. Claverie, J. Ju, J.G. Masek, J.L. Dungan, E.F. Vermote, J.-C. Roger, S.V. Skakun, C. Justice. 2018. The Harmonized Landsat and Sentinel-2 surface reflectance data set. *Remote Sensing of Environment*, 219, 145-161: <https://doi.org/10.1016/j.rse.2018.09.002>
- G. Doxani, E. Vermote, J.-C. Roger, F. Gascon, S. Adriaensen, D. Frantz, O. Hagolle, A. Hollstein, G. Kirches, F. Li, J. Louis, A. Mangin, N. Pahleva, B. Pflug, Q. Vanhellemont. 2018. Atmospheric Correction Inter-Comparison Exercise. *Remote Sens.*, 10(2), 352: <https://doi.org/10.3390/rs10020352>
- ESA. 2021. The European Space Agency. Available online: <https://sentinels.copernicus.eu/web/sentinel/user-guides/sentinel-2-msi/resolutions/spatial> (accessed on 19 January 2021)
- D. Frantz. 2019a. FORCE—Landsat + Sentinel-2 Analysis Ready Data and Beyond. *Remote Sensing*, 11, 1124: <https://doi.org/10.3390/rs11091124>
- D. Frantz, M. Stellmes, P.A. Hostert. 2019b. Global MODIS Water Vapor Database for the Operational Atmospheric Correction of Historic and Recent Landsat Imagery. *Remote Sens.*, 11, 257. <https://doi.org/10.3390/rs11030257>
- D. Frantz, E. Haß, A. Uhl, J. Stoffels, J. Hill. 2018. Improvement of the Fmask algorithm for Sentinel-2 images: Separating clouds from bright surfaces based on parallax effects. *Remote Sens. Environ.*, 215, 471–481: <https://doi.org/10.1016/j.rse.2018.04.046>
- R. Houborg, M.F. McCabe. 2018a. Daily Retrieval of NDVI and LAI at 3 m Resolution via the Fusion of CubeSat, Landsat, and MODIS data. *Remote Sensing*, 10(6), 890: <https://doi.org/10.3390/rs10060890>
- R. Houborg, M.F. McCabe. 2018b. A Cubesat Enabled Spatio-Temporal Enhancement Method (CESTEM) utilizing Planet, Landsat and MODIS data. *Remote Sensing of Environment*, 209, 211-226: <https://doi.org/10.1016/j.rse.2018.02.067>
- M. Guizar-Sicairos, S.T. Thurman, J.R. Fienup. 2008. Efficient subpixel image registration algorithms. *Optics Letters* 33, 156-158: <https://doi.org/10.1364/OL.33.000156>
- J. Kong, Y. Ryu, J. Liu, B. Dechant, C. Rey-Sanchez, R. Shortt, D. Szutu, J. Verfaillie, R. Houborg, D.D. Baldocchi. 2022. Matching high resolution satellite data and flux tower footprints improves their agreement in photosynthesis estimates. *Agricultural and Forest Meteorology* Vol. 316, 108878, <https://doi.org/10.1016/j.agrformet.2022.108878>
- L. Nieto, R. Houborg, A. Zajdband, A. Jumpasut, P.V.V. Prasad, B.J.S.C. Olson, I.A. Ciampitti. 2022. Impact of High-Cadence Earth Observation in Maize Crop Phenology Classification. *Remote Sensing* 14, no. 3: 469. <https://doi.org/10.3390/rs14030469>
- J.W. Rouse, R.H. Hass, J.A. Shell, D. Deering. 1973. Monitoring vegetation systems in the Great Plains with ERTS-1. In: *Third Earth Resources Technology Satellite Symposium*. Washington DC, pp. 309–317
- D.P. Roy, H.K Zhang, J. Ju, J.L. Gomez-Dans, P.E. Lewis, C.B. Schaaf, Q. Sun, J. Li, H. Huang, V. Kovalsky. 2016. A general method to normalize Landsat reflectance data to Nadir BRDF Adjusted Reflectance. *Remote Sensing of Environment*, Vol. 176, pp 255–271: <https://doi.org/10.1016/j.rse.2016.01.023>
- D.P. Roy, J. Li, H.K. Zhang, L. Yan, H. Huang, Z. Li. 2017. Examination of Sentinel-2A multi-spectral instrument (MSI) reflectance anisotropy and the suitability of a general method to normalize MSI reflectance to nadir BRDF adjusted reflectance. *Remote Sensing of Environment*, Vol. 199, pp 25-38: <https://doi.org/10.1016/j.rse.2017.06.019>
- D.P. Roy, H. Huan, R. Houborg, V.S. Martins. 2021. A global analysis of the temporal availability of PlanetScope high spatial resolution multi-spectral imagery. *Remote Sensing of Environment*, Vol. 264, 112586: <https://doi.org/10.1016/j.rse.2021.112586>
- C.B. Schaaf, F. Gao, A.H. Strahler, W. Lucht, X. Li, T. Tsang, N.C. Strugnell, X. Zhang, Y. Jin, J.-P. Muller et al. 2002. First operational BRDF, albedo nadir reflectance products from MODIS. *Remote Sens. Environ.*, 83, 135–148: [https://doi.org/10.1016/S0034-4257\(02\)00091-3](https://doi.org/10.1016/S0034-4257(02)00091-3)
- D. Tanré, C. Deroo, P. Duhaut, M. Herman, J.J. Morcrette, J. Perbos, P.Y. Deschamps. 1990. Description of a Computer Code to Simulate the Satellite Signal in the Solar Spectrum: The 5S Code. *Int. J. Remote Sens.*, 11, 659–668: <https://doi.org/10.1080/01431169008955048>
- M.G. Ziliani, M.U. Altaf, B. Aragon, R. Houborg, T.E. Franz, Y. Lu, J. Sheffield, I. Hoteit, M.F. McCabe. 2022. Early season prediction of within-field crop yield variability by assimilating CubeSat data into a crop model. *Agricultural and Forest Meteorology* Vol.

313, 108736, <https://doi.org/10.1016/j.agrformet.2021.108736>

Z. Zhu, C.E. Woodcock. 2012. Object-Based Cloud and Cloud Shadow Detection in Landsat Imagery. *Remote Sens. Environ.*, 118, 83–94: <https://doi.org/10.1016/j.rse.2011.10.028>