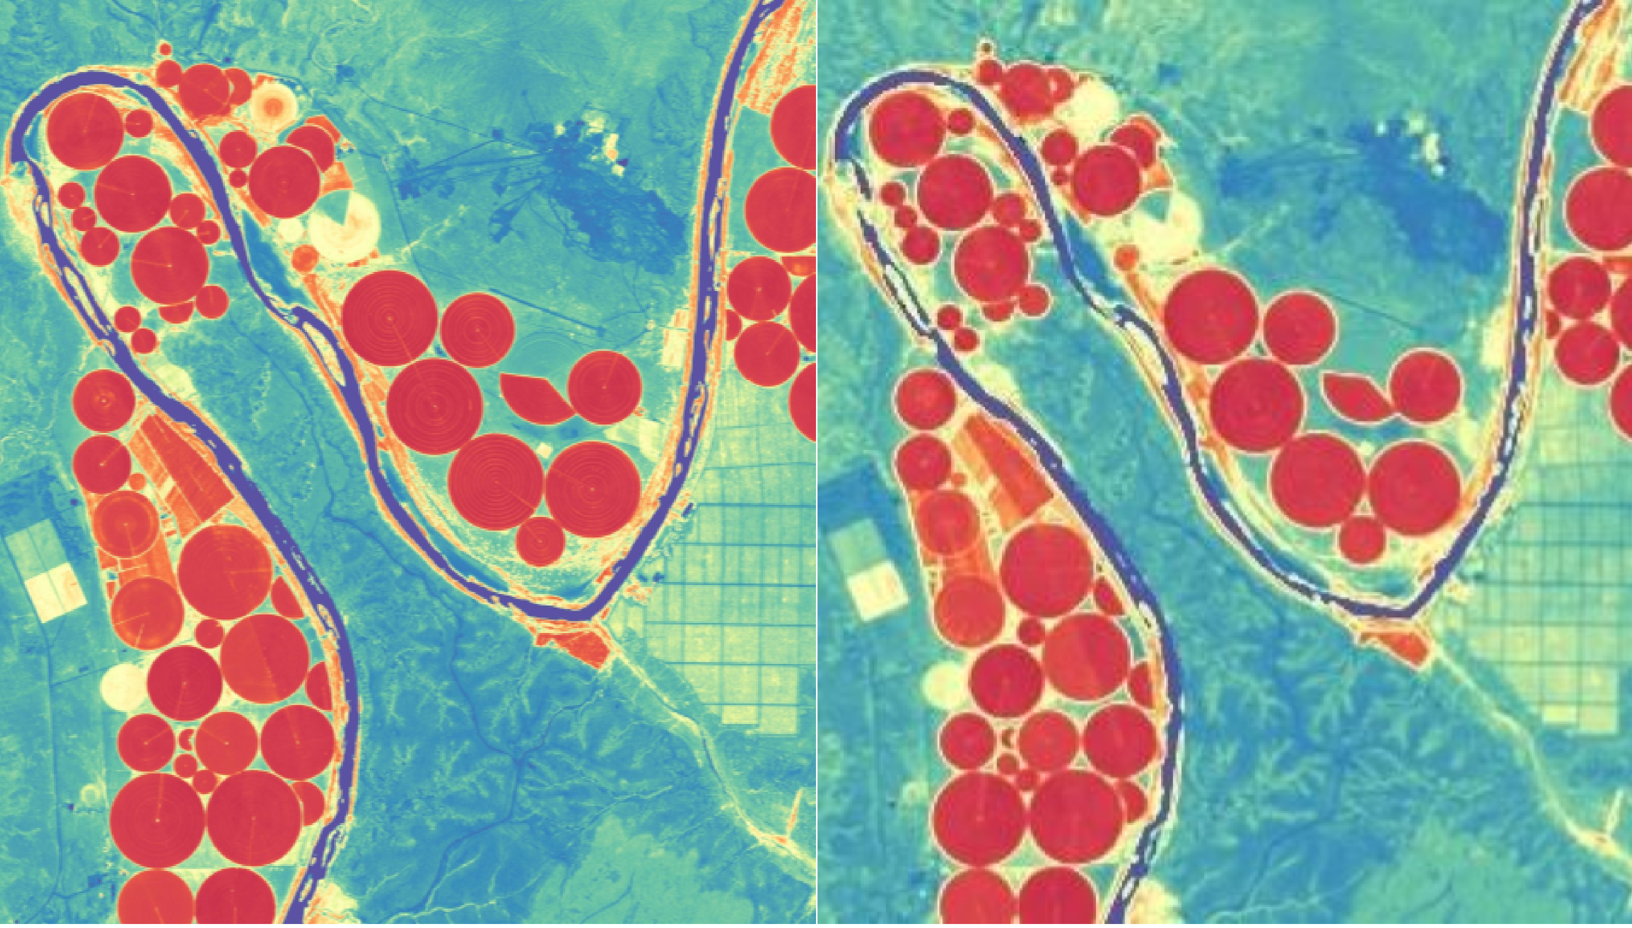

Shown LEFT is Planet's Fusion Monitoring Product (PlanetScope data)
Shown RIGHT is LandSat 8 data of the same area

PLANET FUSION MONITORING
TECHNICAL SPECIFICATION

Version 1.0.0-beta.3, March 2021

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

Planet
sales@planet.com
www.planet.com
645 Harrison St, Floor 4
San Francisco CA, 94107
Tel. +1 (415) 419-2194

TABLE OF CONTENTS

1. PLANET FUSION OVERVIEW

The PlanetScope constellation of 100+ 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., 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).

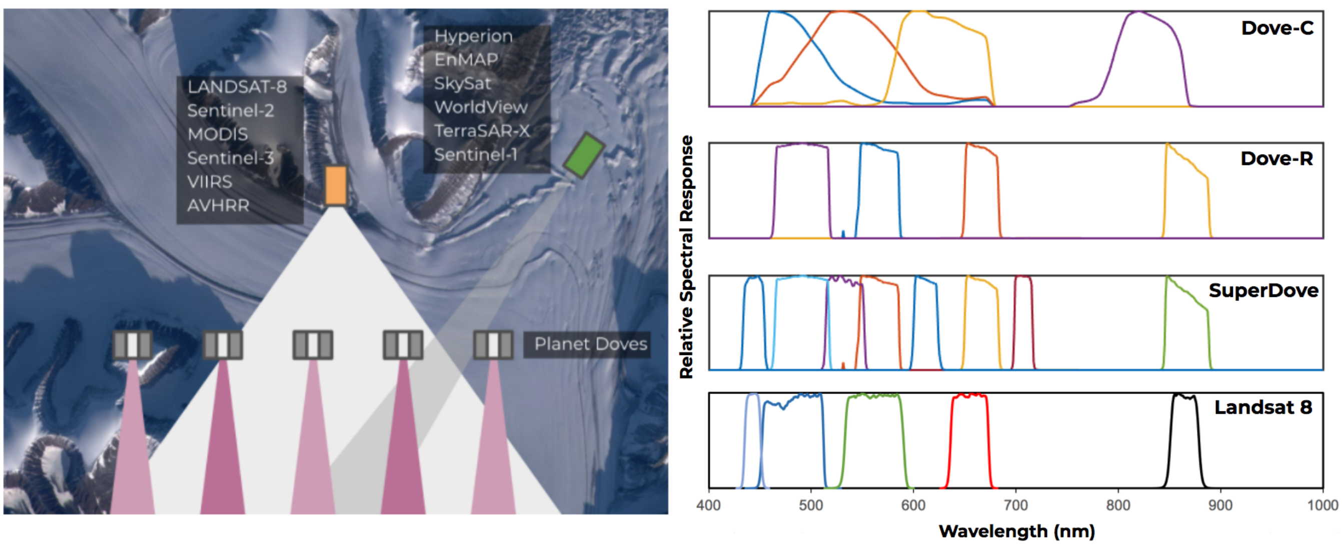

At Planet, we have implemented and improved a rigorous methodology to enhance, harmonize, inter-calibrate, and fuse cross-sensor data streams. The CubeSat-Enabled Spatio-Temporal Enhancement Method (CESTEM) (Houborg and McCabe 2018a,b) leverages rigorously calibrated publicly accessible multispectral satellites (i.e., Sentinel, Landsat, MODIS) 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 surface reflectance (SR) data product. Planet Fusion integrates all of the best features from both public and private satellite sensor resources.

Planet Fusion processing ingests data from multiple sensors with differing radiometry, quality, and resolution characteristics 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. It automatically generates a uniform, consistent spatial (i.e., 3 m) and spectral (4-band) resolution time-series dataset. 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, time-series analysis, phenological monitoring, vegetation growth and disturbance monitoring, physically-based model retrieval, machine learning, and other applications relying on 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 2018. The Planet Fusion processing translates original PS Top Of Atmosphere Radiance (TOAR) inputs into Surface Reflectances (SR) generally consistent with Landsat 8 (L8) and Sentinel-2 (S2) 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.

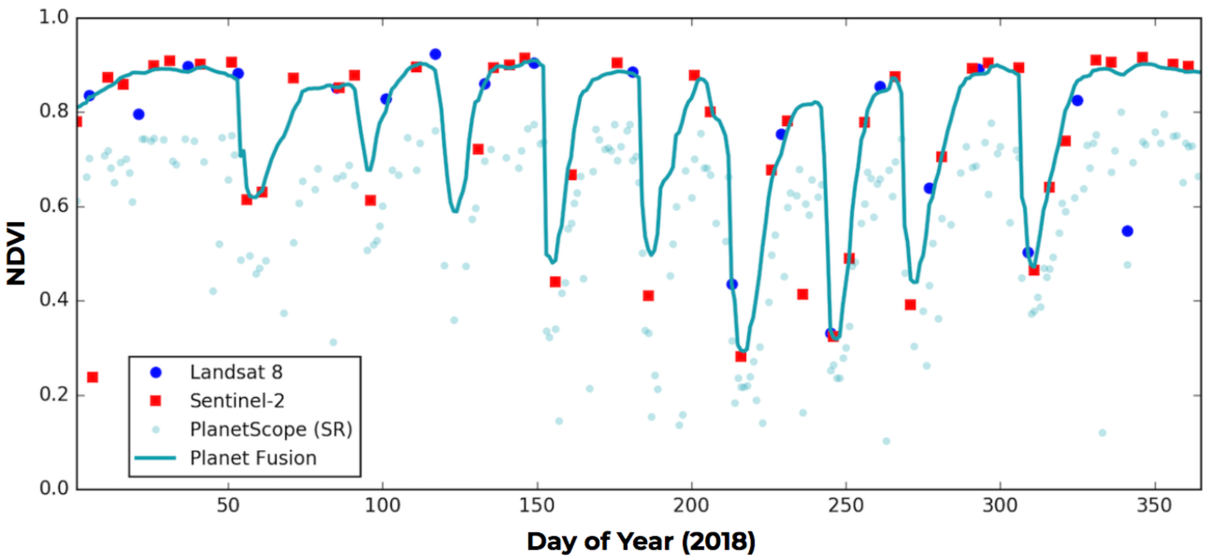

The unique features of Planet Fusion data can be summarized as:

- **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 Surface Reflectance values delivered as regularly gridded UTM raster tiles (24 km x 24 km) on a daily basis.**
- **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 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 100+ 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. 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 (TOAR) Radiance inputs contribute the bulk of the observations used to make Fusion, 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-TOAR	Scene-based PlanetScope L1B - Top Of Atmosphere Radiance (TOAR) (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/mcd43a4v006/)
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 S2 data have been spectrally adjusted to match L8 spectral band passes. Based on the Framework for Operational Radiometric Correction for Environmental Monitoring (FORCE) (https://www.uni-trier.de/index.php?id=63673)
S2-L1C	Sentinel-2 L1C TOA Reflectance (10 - 20 m)
L8-TOA	Landsat 8 TOA Reflectance (30 m)

We use a scalable implementation of the Framework for Operational Radiometric Correction for Environmental Monitoring (FORCE; 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). 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 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 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 combined (i.e., Terra and Aqua) MCD43A4 product that provides daily 500 m SR in 7 bands corrected for reflectance anisotropy (MODIS has a ~110° 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., ~7 days latency). The VIIRS products (VNP43IA4, 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.

Native resolution Sentinel-2 (10 - 20m) and Landsat 8 (30 m) 4-band (VNIR) TOA reflectance data (Table 1) are needed as input to the data fusion module (see Section 4.3). These products will go through the Planet Fusion calibration and harmonization process to produce cross-sensor consistent SR data matching the Sentinel 2 radiometry. The spatial resolution of the Sentinel-2 and Landsat 8 data is sharpened to 3 m during this process for consistency with the PlanetScope constellation. It follows that a Planet Fusion data product may feature inputs from both PlanetScope and Sentinel-2/Landsat 8 with prioritization of PlanetScope data where clear-sky observations are available from both sources.

3. PLANET FUSION PRODUCTS

The current 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 3.3).

Table 2: The current Planet Fusion (PF) product line

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 Sentinel-2/Landsat 8 data fusion (gap-free, daily, 3 m)
PF-QA	Planet Fusion Quality Assurance product

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 TOAR (converted to reflectances) 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). Due to the nature of the Planet Fusion cross-calibration approach (see Section 4.3), the significant uncertainties related to the 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. QA 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 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) (100) versus actual (PlanetScope and/or S2/L8) data (1) used to generate pixel value	16-bit signed integer	1 - 100	1	0
Layer 2	Days to closest scene with actual observation data (- is before, + is after prediction day) *See embedded metadata for all dates [yyyymmdd] used to gap-fill* 900 (-900): Sentinel 2 (Landsat 8) data **acquired on the prediction day** used to gap-fill *See embedded metadata for the scene IDs of the S2/L8 data (if any) used to gap-fill*	16-bit signed integer	-900 to 900	1	0
Layer 3	PlanetScope cloud and cloud shadow mask 1 = Clear 2 = Bright Clouds 3 = Shadows 4 = Other cloud elements 5 = Adjacent clouds 6 = Additional cloud/shadow/haze elements based on a cross-scene correlation detection approach 7 = Pixel masked out due to band-to-band mis-alignment issue -999 = PlanetScope scene data not available	16-bit signed integer	1 - 7, and -999	1	0
Layer 4	PlanetScope pixel traceability mask (*see embedded metadata for scene IDs*) -999 = PlanetScope scene data not available	16-bit signed integer	1 - 200, and -999	1	0
Layer 5	Total number of FORCE L8/S2 reference scenes used during calibration -999 = PlanetScope scene data not available	16-bit signed integer	0 - 500, and -999	1	0
Layer 6	Blue band uncertainty estimate (absolute percentage)	16-bit signed integer	0 - 200	1	0
Layer 7	Green band uncertainty estimate (absolute percentage)	16-bit signed integer	0 - 200	1	0

Layer 8	Red band uncertainty estimate (absolute percentage)	16-bit signed integer	0 - 200	1	0
Layer 9	NIR band uncertainty estimate (absolute percentage)	16-bit signed integer	0 - 200	1	0

QA layer 1 represents a raster mask indicating the percentage of "synthetic" (i.e., gap-filled) versus observation data used to produce each pixel value across the tile domain. Note that the observation data can be a combination of PlanetScope and Sentinel 2/Landsat 8 (see QA layer 2). A value of 1 indicates no gap-filling (i.e., real PS/S2/L8 data) whereas a value of 100 indicates an entirely gap-filled pixel value. The pixel-specific weighting is done to ensure smooth and gradual transitions across interfaces between gaps and real 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).

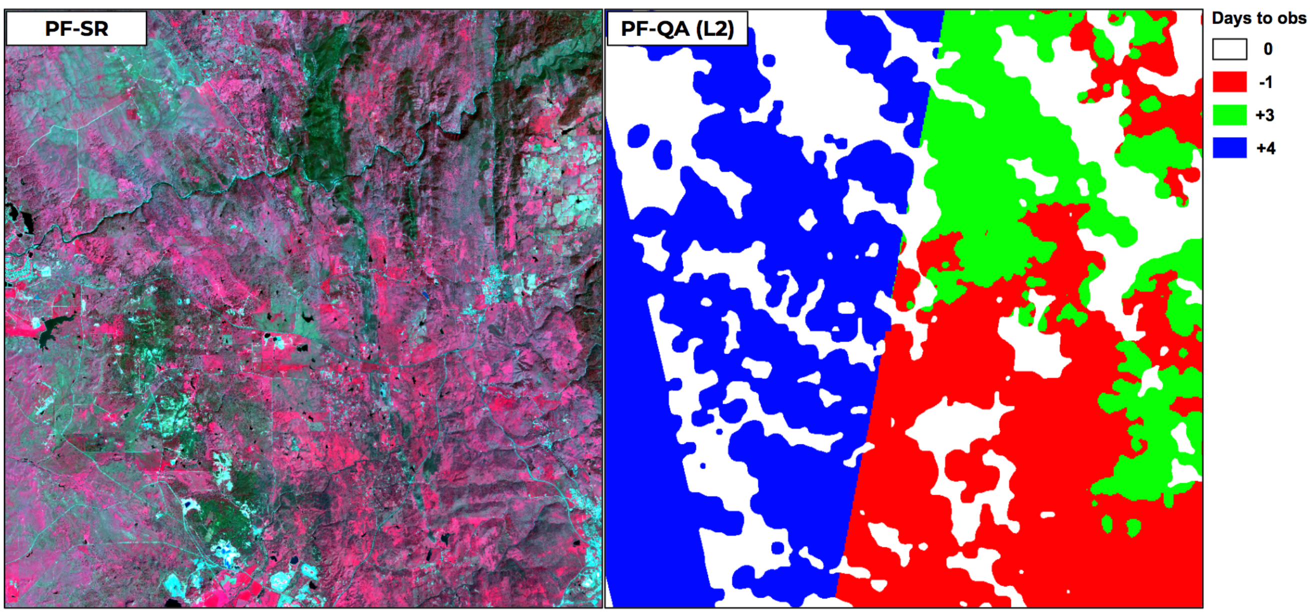

QA layer 2 is a gap-filling confidence metric, providing pixel-specific information on the number of days to a day with actual observation data. The longer the day gap, the higher the uncertainty is likely to be in the gap-filled estimate although several factors will play a role in the actual retrieval uncertainty (see Section 4.4). This QA layer will also identify pixels where Planet Fusion processed Sentinel-2/Landsat 8 data on the prediction date have been used to fill gaps resulting from clouds or gaps in PlanetScope coverage (Table 5). The actual dates of all the scenes 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 and 5 exemplify this metric and the associated cloud and cloud shadow mask for an AOI close to Sacramento (CA).

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

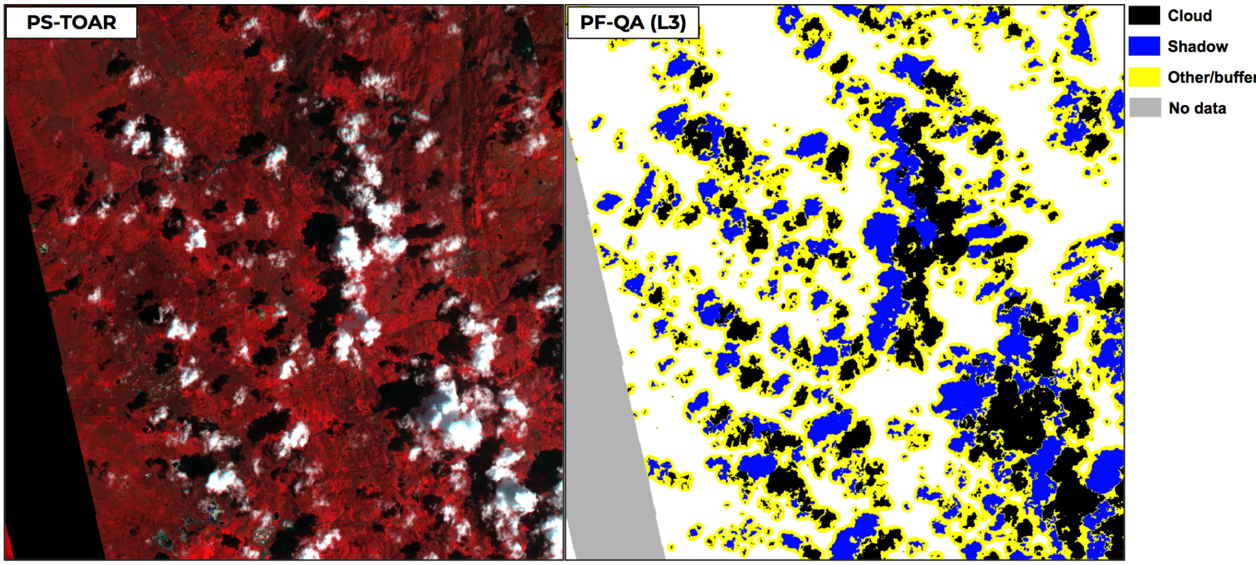

A pixel traceability mask is also provided (Figure 6). This raster layer identifies the footprints of the PlanetScope scenes used to produce any given tile image (cloudy or clear). Each domain is associated with a unique integer value that is linked to a scene identifier (i.e., Itemtype/sceneID) embedded 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.

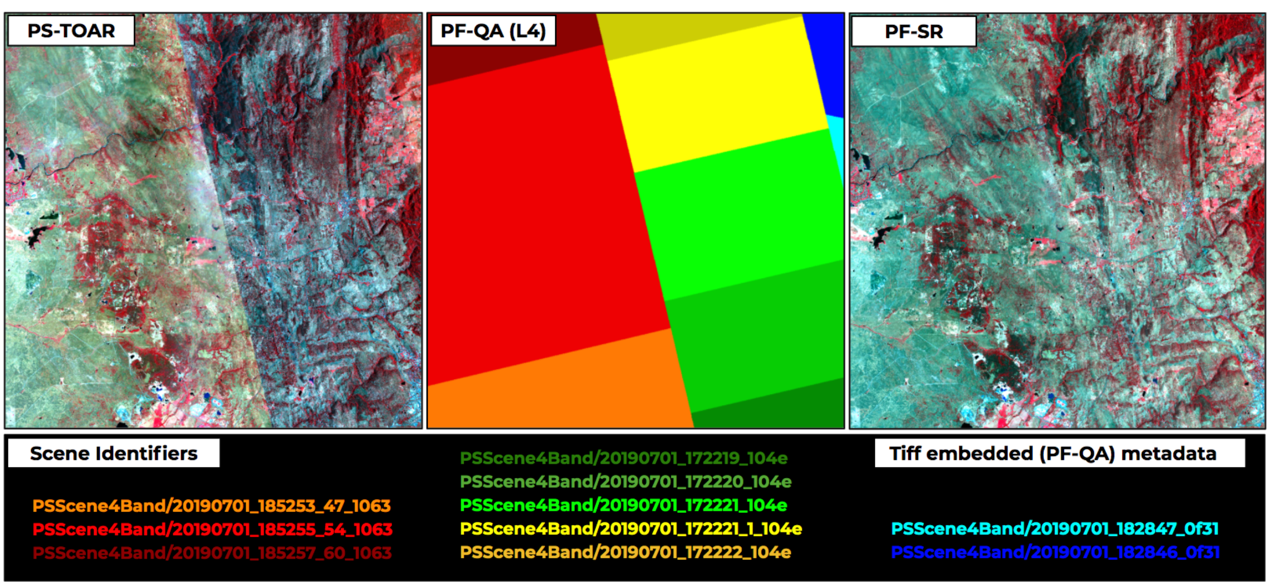

QA layer 5 (Table 5) identifies the number of FLS-SR scenes (i.e., Sentinel-2 or Landsat 8) used for each pixel during reference sampling and CESTEM calibration. In general the larger the number of reference scenes, the more robust the calibration is expected to be (see Section 4.3).

QA layers 6 - 9 provide band-specific confidence information for the "synthetic" (i.e., gap-filled) pixel values (real data will have a value of 0). The process for deriving these are described in Section 4.6. The confidence information is reported as an absolute percentage to indicate the deviation of the synthetic value from a real 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. 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), 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 below for a description of each of the product keys.

Product Key	Description
PF-SR	Planet Fusion Surface Reflectance product
PF-QA	Planet Fusion Quality Assurance product

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).

Currently, the Planet Fusion products can be delivered in "near real-time" with a 48 hour latency (i.e., data requested for Monday will be delivered Wednesday).

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: Flow diagram of Planet Fusion processing elements for any given tile and TOI.

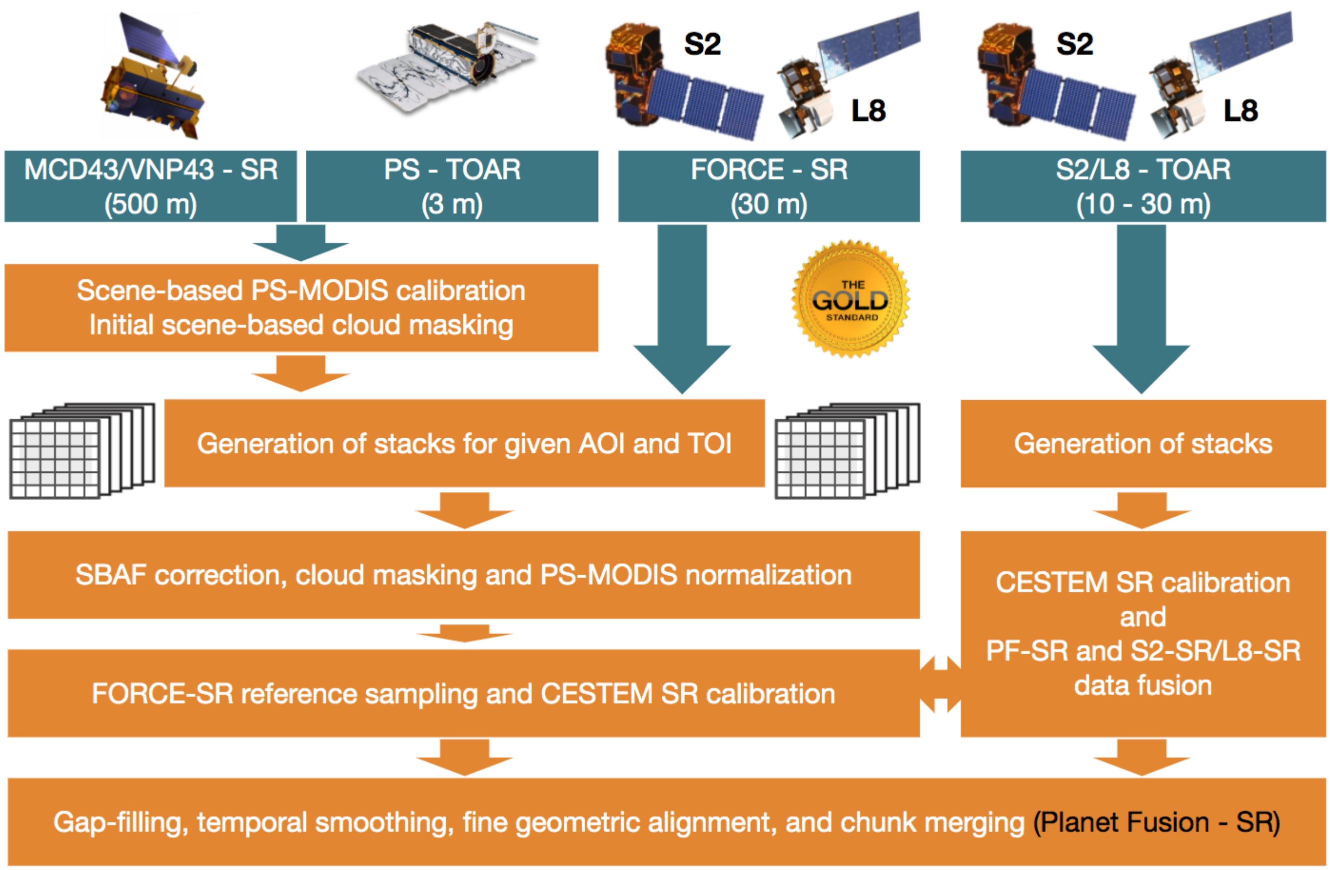

4.1. STACK GENERATION

After identifying the scenes (PS-TOAR) and tiles (FORCE, MCD43, S2, L8) 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). In the case of PS-TOAR, several scenes from multiple sensors may be overlapping with the tile domain (e.g., Figure 6). In order to retain the best data for stack generation, priority is determined as a function of Dove sensor type (prioritizing Dove-R and SuperDove over Dove-Classic), initial cloud fraction (prioritizing scenes with the lowest cloud fraction), 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. Planet Fusion processing is done on slightly overlapping (i.e., a buffer of 750 m) chunks with a ~13 by 13 km extent. The tile products will be generated by merging data from 4 chunks, utilizing a gradual weighting approach across chunk overlap zones to avoid visible boundaries in the final outputs.

4.2. CLOUD MASKING

Planet Fusion cloud and cloud shadow detection is performed at 30 m resolution based on MODIS/VIIRS calibrated PlanetScope SR stacks (i.e., PS-MODIS). The MODIS/VIIRS calibration is performed at the PS scene level using the MCD43/VNP43 NBAR products (Table 1) as the radiometric reference, translating PS-TOAR into MODIS/VIIRS-consistent SR data. During this initial calibration stage, an initial scene-based cloud mask is produced exploiting in part cloud 2.0 (UDM2) detections, when available. A rigorous cloud verification approach is also implemented at this step in an attempt to reduce commission/omissions errors.

After stack generation (Figure 7) an iterative temporal cloud masking approach is invoked to refine the initial scene-based mask. This approach utilizes clear-sky (as identified by the initial mask during the first iteration) time series information over 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, or cloud shadow. 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. 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.

4.3. FLS-SR REFERENCE SAMPLING AND CESTEM 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 for cross-calibration/harmonization
- 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-TOAR)

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

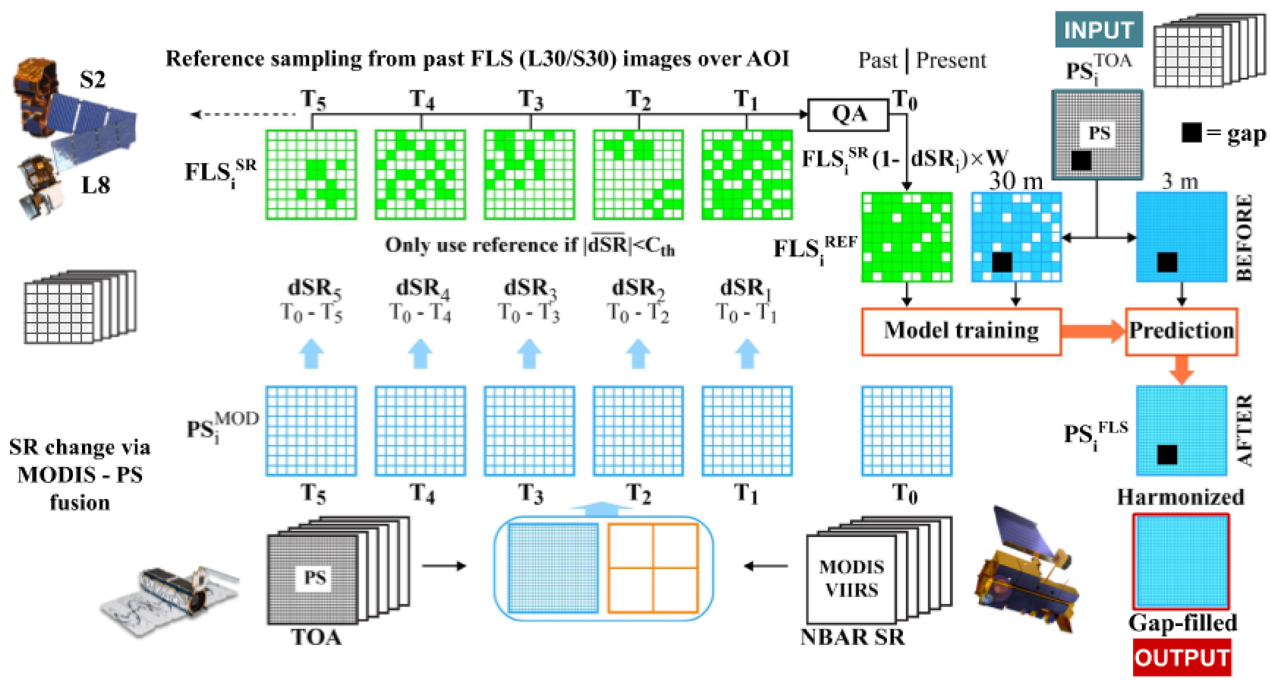

The CESTEM-based harmonization is applied to all images acquired over a defined Time Of Interest (TOI) drawing spectral (i.e., blue, green, red, and NIR) reference data from a pool of S2/L8 SR images acquired over a predefined "calibration window" (typically one year centered around the TOI). In order 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 spectral changes over given PS - S2/L8 acquisition time spans (Figure 8). 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 - S2/L8 acquisition time spans and the magnitude of surface reflectance change relative to the prediction day (Figure 8). 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 layer (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 TOAR (PS^{TOA}) 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 9, 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 9: Illustrating the robustness of the CESTEM harmonization to noise in the input data streams (FLS and PS)

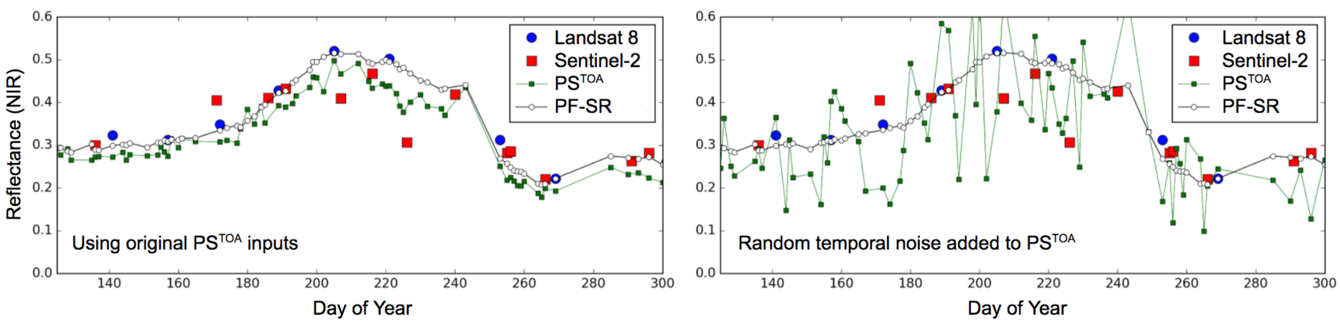

4.4. GAP-FILLING AND DATA FUSION

The Planet Fusion gap-filling ensures a spatially complete and temporally continuous (daily) 3 m product irrespective of the actual acquisition coverage and cloud environment. 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 is informed by daily MODIS/VIIRS information (Table 1) in combination with cloud masked (but not yet gap-filled) Planet Fusion SR (PF-SR) data from both before and after (if available) the prediction date. Up to 6 Planet Fusion tiles are typically used for this purpose, which may feature CESTEM-based S2-SR and L8-SR data, if available. During PF-SR and S2-SR/L8-SR data fusion, the S2/L8 data are 1) harmonized to be consistent with FLS-SR, 2) sharpened to the resolution of the PS data (3 m), and 3) sub-pixel shifted (as needed) to align with the PS data. Cloud masking of the S2/L8 data is conducted concurrently using the framework established for PS data (Section 4.2). The associated PF-QA product (Table 5) will provide full traceability of the source data used to fill the gaps.

Figure 10: A generalized overview of the Planet Fusion gap-filling framework

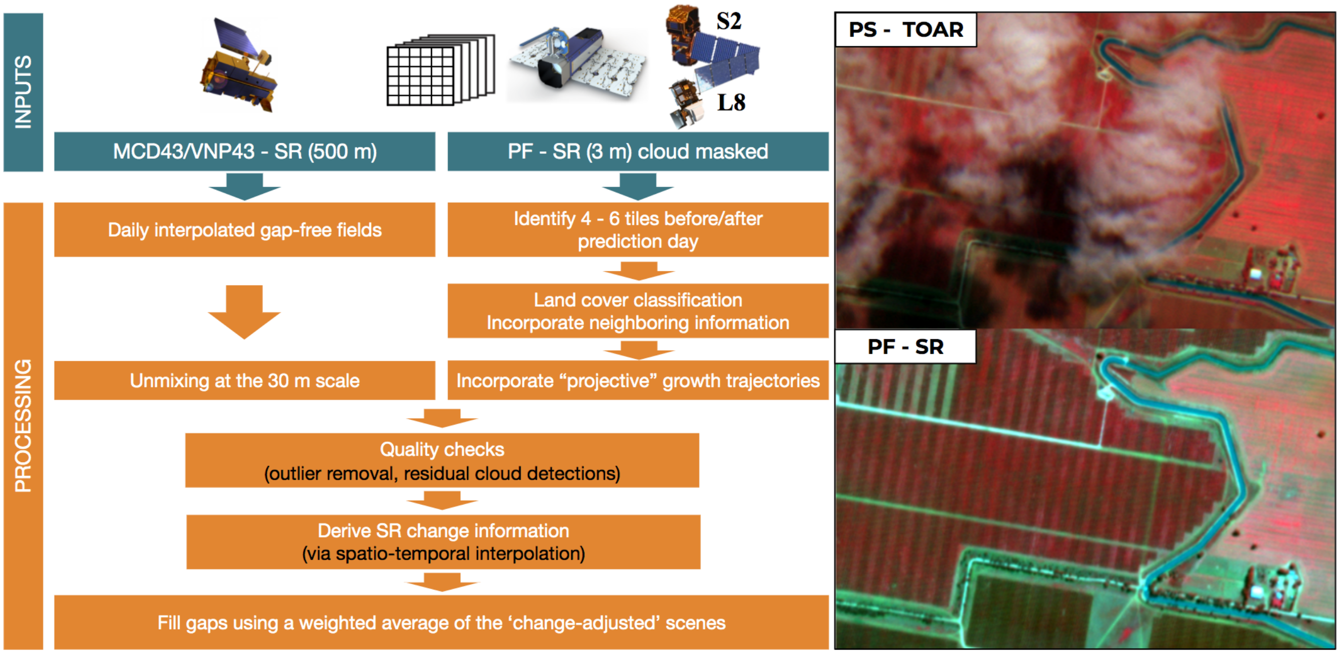

The PF-SR tile selection (for filling the gaps) prioritizes relatively clear images acquired with the shortest time lag relative to the prediction date. Combined with the daily MODIS/VIIRS information, SR change information (i.e., relative to the prediction date) specific to each of the selected PF-SR tiles is derived using a carefully weighted interpolation technique. The gap-filling uses both spatial (i.e., neighboring and class-specific pixel information) and temporal interpolation techniques to provide an informed and weighted guess of the pixel values. In addition, clustered full-season spectral soil and vegetation growth trajectories computed based on past Planet Fusion observations (e.g., past growing season) and matched with current trajectories/phenologies, are used to guide the predictions when data availability after the prediction date is limited. This projective feature is critical to ensuring realistic predictions when producing Planet Fusion data close to present time (e.g., near real-time). A series of quality checks are in place for residual cloud detection and outlier removal before using a weighted average of the "change-adjusted" scenes to fill the gaps. While the gap-fill confidence metric (Table 5) only identifies the tile with the highest weight, tiles from 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 dates of all the tiles used to inform the gap-filling can be found in the PF-QA tiff file as embedded metadata.

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 tiles (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 11 showcases the Planet Fusion gap-free NDVI time series plot for two fields (alfalfa and corn) in South Dakota during the 2018 growing season. The Planet Fusion processing is seen to be effective at translating the relatively noisy PlanetScope TOAR data into FLS-consistent SR by effectively detecting, leaving out, and gap-filling the frequent occurrences of cloud contaminated pixels. The frequency of the FORCE-based L8 and S2 clear-sky observations (as identified by the associated FORCE cloud mask) is substantially reduced relative to the Planet Fusion clear-sky observations, which can severely limit the utility of the FLS-SR products for tracking dynamic crop phenology in environments such as this. In contrast, the daily Planet Fusion data capture the various crop development stages and phenological transitions with high fidelity. Importantly, the enhanced temporal consistency does not affect the ability of the Planet Fusion data to reproduce and detect rapid transition events.

Figure 11: Planet Fusion-based NDVI relative to L8/S2 and the original PlanetScope data exemplified for two fields in South Dakota during the 2018 growing season

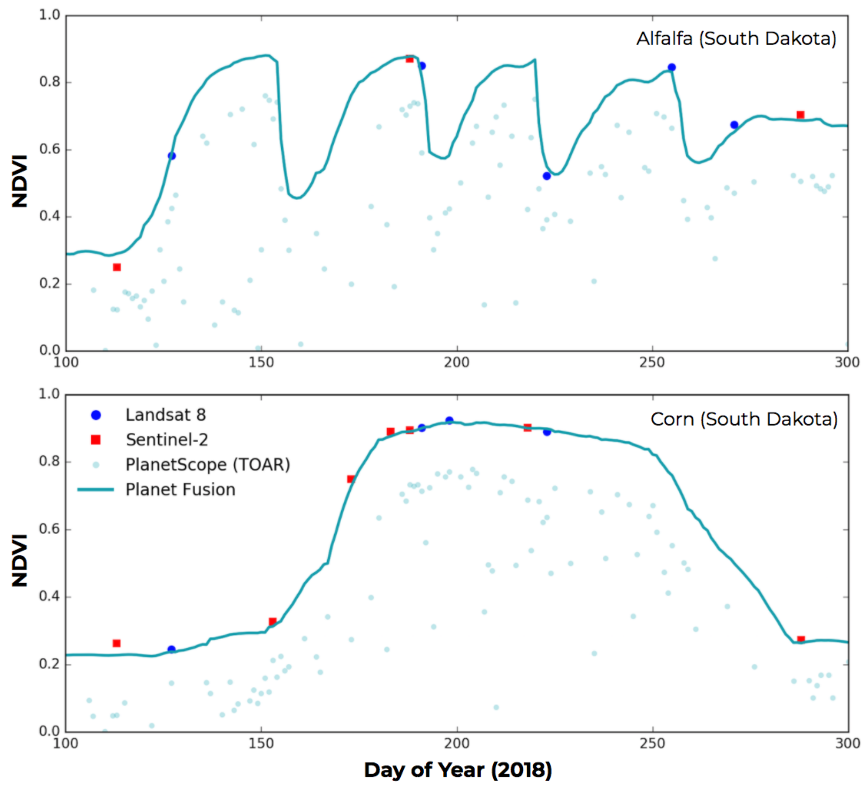

The enhanced ability to track high frequency land surface dynamics is further supported by Planet Fusion NDVI time-series over a frequently harvested crop field in Imperial Valley (CA) (Figure 12). Figure 12 shows individual sensor contributions to the time series and the tile-based clear-sky coverage and revisit time over the TOI (DOY 1 to 288). The Dove-Classic constellation is seen to have the highest revisit frequency (2.58). Adding Dove-R imagery increases the clear-sky coverage from 39% to 57%. The addition of SuperDove imagery bumps it up to 66% (Figure 12). Finally, fusing in Landsat 8 and Sentinel-2 sources adds 17 new actual observation points (i.e., the yellow circles) and increases clear-sky coverage from 66% to 72%. As previously mentioned, the fusion module will use CESTEM-harmonized and super-resolved (i.e., 3 m) S2 or L8 data to fill gaps in PlanetScope clear-sky coverage when possible. The combination and harmonization of multiple sensor inputs helps resolve and monitor dynamic land surface processes. By balancing the strengths of the different sensors, Planet Fusion attempts to produce consistent, comprehensive, and sensor-agnostic data to help unlock the full potential of rich multi-sensor datasets for a wide range of applications.

Figure 12: Multi-source NDVI time series for a crop field in Imperial Valley (CA), exemplifying the use of Planet Fusion-integrated L8/S2 data (yellow circles) to fill gaps in PlanetScope coverage. The table to the right provides a clear-sky coverage breakdown for the associated Planet Fusion tile (TOI: DOY 1 - 288) highlighting individual sensor contributions

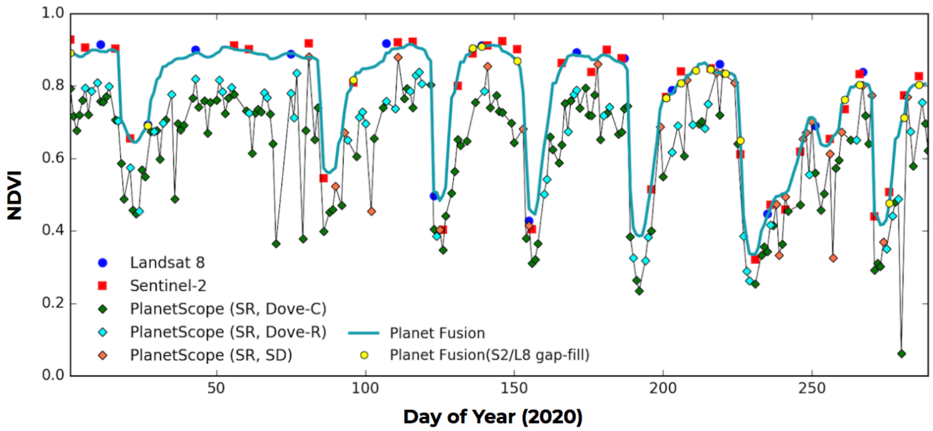

Source Imagery	Clear-Sky Coverage	Revisit time
Dove C Dove R SuperDove Sentinel-2 Landsat 8	72%	1.38
Dove C Dove R SuperDove	66%	1.51
Dove C Dove R	57%	1.75
Dove C	39%	2.58

4.5. FINE GEOMETRIC ALIGNMENT

The input imagery that feeds into Planet Fusion (e.g., L8, S-2, PS) have 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 km x 13 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 fine aligned 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.

Geometric fine alignment in temporal space is performed during and after gap-filling. As the gap-filling relies on surface reflectance inputs from multiple images acquired at different dates, it can be important to align the inputs to the same reference image first. The reference image is generated for each prediction date using a weighted average of the four closest predominantly clear-sky images. Each of the input images needed for gap-filling are then sub-pixel aligned to that reference using the associated cloud masks to identify an optimal clear-sky image subset to use for assessing band-specific shifts. A final band-specific alignment step is performed on the deep stack of gap-filled imagery using a seven day rolling surface reflectance average as the geometric reference.

While the implemented fine alignment techniques will also help address band to band mis-alignments, artifacts will still occur when the shifts are not global (i.e., not uniform across the image domain). This can occur due to registration issues during scene composition or as a result of parallax effects (typically in the vicinity of clouds). Resulting “rainbow” artifacts are fairly common in 4-stripe (red, green, red, NIR) Dove-R and SuperDove imagery. A “rainbow” detector has been developed as part of Planet Fusion in an attempt to capture these residual image artifacts. These detections will appear with a mask value of 7 in the PlanetScope cloud and cloud shadow mask (Table 5).

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 to assess the uncertainty of the estimate relative to a real surface reflectance observation.

The model used to derive the confidence information is based on comparisons between real and gap-filled data over a selection of representative AOIs and TOIs covering a diversity of surface characteristics and cloud environments. The absolute relative difference between the real and gap-filled surface reflectance (for the same pixels) serves as the target variable (i.e., the uncertainty), which is related to a suite of explanatory variables using a data mining technique. As the uncertainty of the synthetic pixel retrievals is expected to depend in part on the number of days to scenes with real observation data, the set of explanatory variables include “distance” metrics to the closest, second closest, and third closest scene used to inform the gap-filling of any given scene. The 4-band spectral data is also included as metrics of the surface characteristics and land cover. A few additional variables are included to help assess the robustness of the predictions across different climates and biomes. 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 13: The upper panel shows the agreement between real and synthetic (gap-filled) data for the entire training dataset. The lower panel shows the performance of the trained model for predicting the relative absolute differences (rMAD) between the real and synthetic data when the observational gap is greater than 2 days.

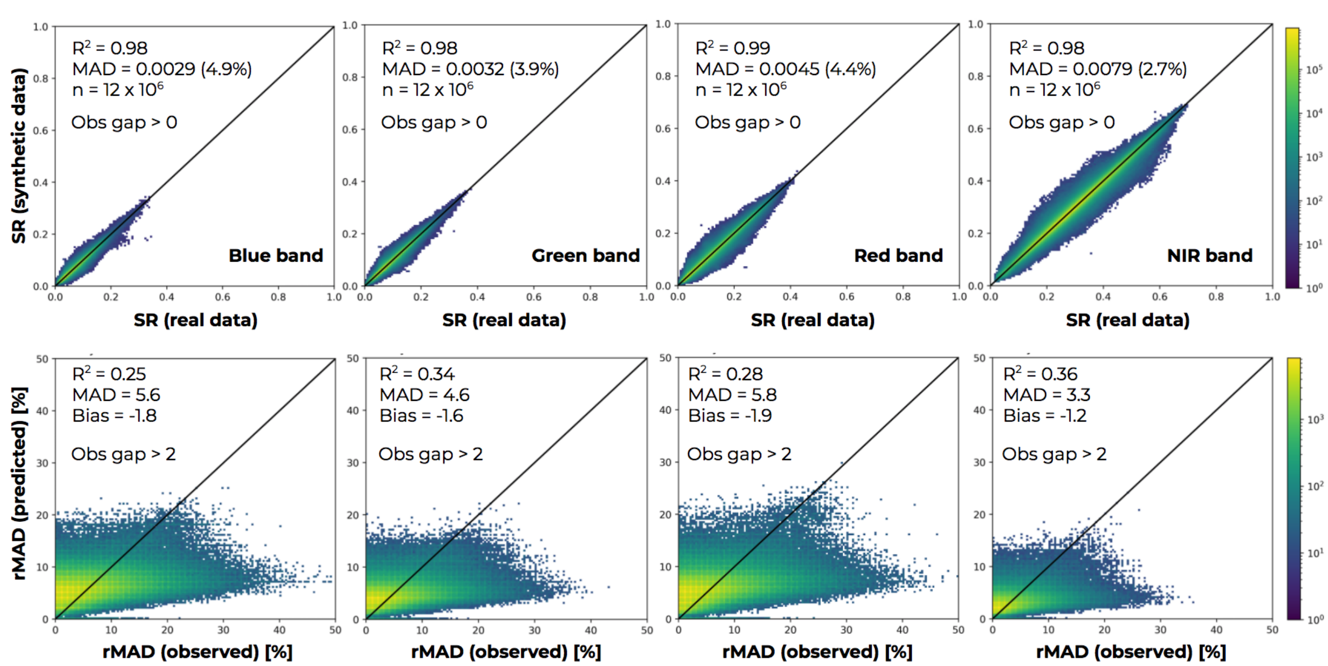

Figure 13 (upper panel) displays the agreement between the real and gap-filled data for the entire training dataset that includes a wide range (1 to 60) in the observational gaps (i.e., days to the closest scene with clear-sky observation data). In general the synthetic data is in good agreement with the real data with the overall relative mean absolute differences (rMAD) ranging between 2.7 and 4.9 %. The lower panel demonstrates the ability of the trained model to predict the observed uncertainties (i.e., relative absolute differences) from the suite of selected explanatory variables when the observational gap is greater than 2 days. The predictive ability (e.g., R^2 on the order of 0.3 and MAD ranging from 3.3 to 5.8) was assessed based on an independent dataset and does demonstrate the challenge of deriving the uncertainty of synthetic pixel retrievals, which is controlled by a complex interaction of several factors. While the uncertainty of the confidence predictions is reduced when the observational gap <= 2 (MAD ranging from 1.9 to 3.5) it is clear that the provided estimates should be considered as a rough approximation with a fair degree of uncertainty at this time.

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 1 - 2 year period is always initially required in order to establish deep temporal image stacks required 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 look 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 14 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 NDVI (with a mean absolute difference 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) 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 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) conservative temporal smoothing (Section 4.4), which will effectively minimize 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 14: Quality (relative to L8/S2) of the Planet Fusion NDVI harmonization process showcased for two regions with day-coincident PlanetScope and L8/S2 acquisitions

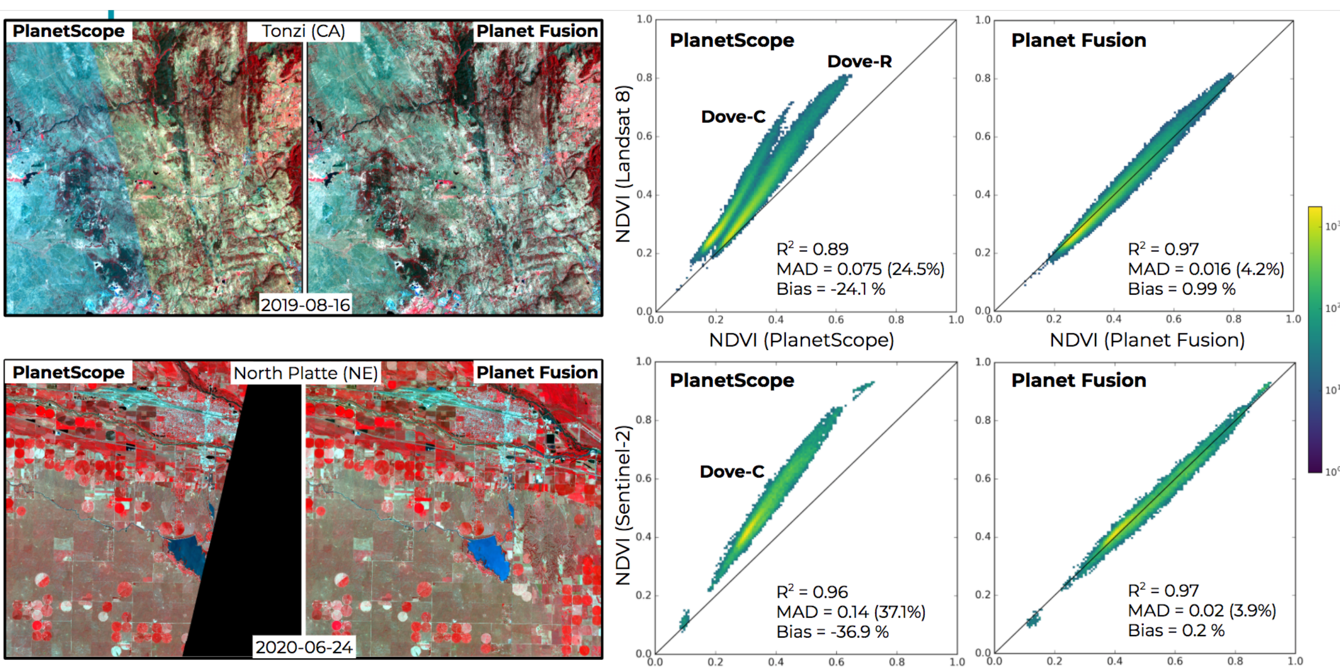

Figure 15: Real versus synthetic (gap-filled) Planet Fusion data for different observation gap intervals (false color: NIR, Red, Green)

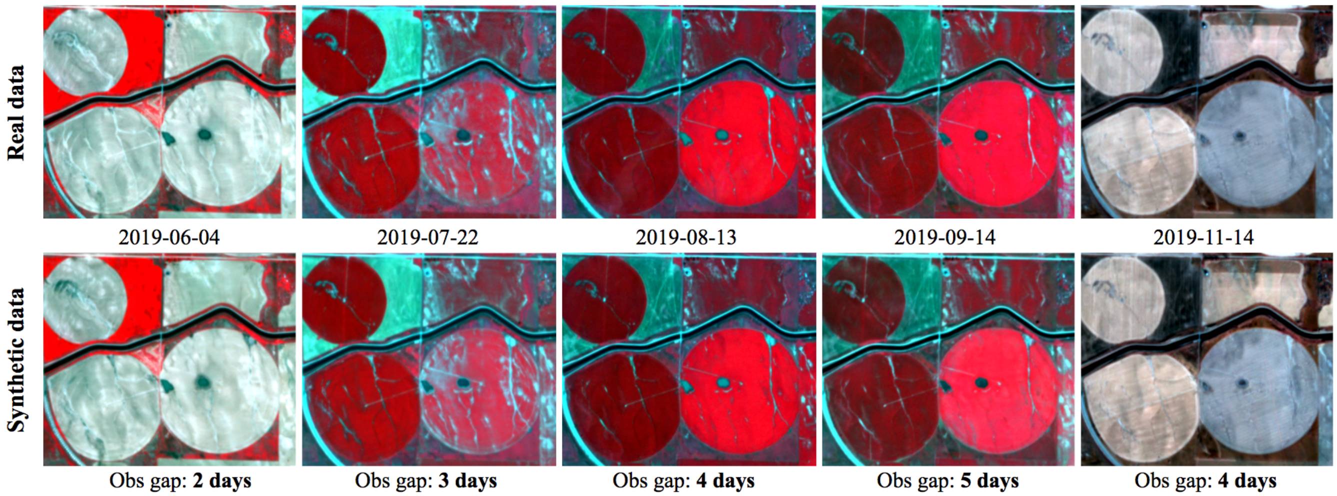

Regarding the synthetic (gap-filled) component of Planet Fusion data, users are advised to consult with the included QA metrics (Table 5) to determine the synthetic data retrieval confidence. Figure 15 above demonstrates the fidelity of the gap-fill data values for different observation gaps in an agricultural 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 likely be within reasonable limits over the majority of landscapes and environments.

Figure 16 showcases real versus synthetic data comparisons for a region in Nebraska over a 3 year period (2017 - 2019), during which the day gaps to actual observation data ranged from 1 - 6 days. The full-tile density scatter plots on select days (top panel) illustrate an excellent agreement between real and synthetic NDVI with the relative Mean Absolute Difference (MAD) ranging from 1.3% to 3.4%. The bottom panel shows a box-and-whisker representation of the band-specific relative MADs between a total of 242 real and synthetic images. These plots display the minimum, lower quartile, median, upper quartile, and maximum MAD as a function of the observation gap (1 - 6 days). In general, the median MAD ranges from ~2% to 6% with an increasing trend as the number of days to real observation data increases.

Figure 16: Top panel: full-scene (agricultural region in Nebraska) statistical evaluations and density scatter plots of synthetic (gap-filled) versus real NDVI. Bottom panel: Box-and-whisker plots of the band-specific relative mean absolute differences [%] between synthetic and real observations for the Nebraska domain during 2017 - 2019 (a total of 242 scenes).

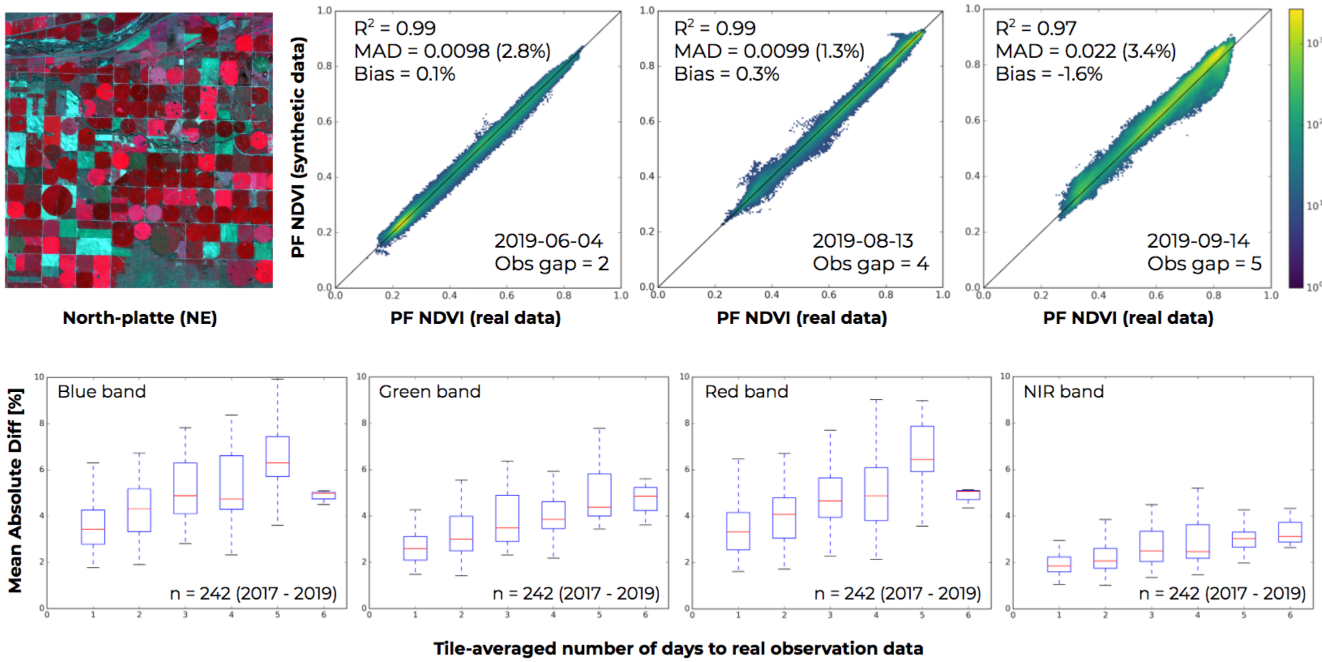

6. KNOWN LIMITATIONS AND CAVEATS

- **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 significant terrain and shadowing. Significant effort has gone into developing automated techniques to differentiate between actual change and atmospheric contamination, but in some cases commision 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. This does mean that 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.

- 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. This capability is expected to become available within the foreseeable future.

REFERENCES

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.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. Kovalskyy. 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

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

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

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