Multi-sensor architectures for thermal imaging missions

Pieter du Toit
Earth Observation Scientist
FarEarth Labs GmbH
Würzburg, Germany
pieter.dutoit@farearth.com

Philip Bouwer
Tech Lead: Analytics Engineering
FarEarth Labs GmbH
Würzburg, Germany
philip.bouwer@farearth.com

Thinus Prinsloo
FarEarth Programme Lead
FarEarth Labs GmbH
Würzburg, Germany
thinus.prinsloo@farearth.com

***Abstract*—Among the innovative offerings that NewSpace companies bring to the space industry, one is a constellation of thermal instruments.**

Thermal infrared (TIR) imagers are inherently more challenging to calibrate and geolocate than visible and near-infrared (VNIR) instruments due to:

- **lower spatial resolutions**
- **higher sensor noise levels; and**
- **a scarcity of stable reference features.**

This makes radiometric and geometric modelling more complex. Multi-sensor satellite architectures that co-locate VNIR and TIR imagers are increasingly valuable for Earth observation applications. These applications require accurate, timely surface temperature data. Integrating reflective and emissive sensors on a single platform can potentially enable improved geometric precision and radiometric consistency.

Our goal with this research is to evaluate whether multi-sensor architectures can enhance the geolocation accuracy of lower-resolution TIR sensors when combined with higher-resolution VNIR imagers through cross-calibration and shared spatial models. We explore how a combination of visible and thermal data can streamline processing workflows for bottom-of-atmosphere products, enabling rapid and reliable insights. Multi-sensor solutions have their advantages, but also pose challenges, such as time-synchronisation between imagers and onboard systems. Acquiring reference data with common features across spectral bands for spatial calibration remains difficult.

Our findings will state the strengths and limitations of each method. We will present results from multiple space missions with different instruments on board. These results will help to inform the design of future thermal imaging constellations.

Keywords—Accuracy, Cross-calibration, Geolocation, Multi-sensor, NewSpace, Radiometry, TIR, VNIR

I. Introduction

Thermal instruments on NewSpace missions have seen increased popularity over the last couple of years. This is partly due to new technologies making it technically feasible, but also because thermal data has a wide and impactful range of applications in the Earth observation market. These range from wildfire monitoring to maritime surveillance, from water-stress and crop-health monitoring to urban heat mapping.

There are a handful of large government-owned missions with thermal instruments on board, but countries globally are seeking ways to commercialise access to the data and to reduce their dependence on other nations.

Missions with a thermal sensor provide both an appealing opportunity for investors and much-needed scientific and intelligence data for consumers.

By their nature, however, thermal instruments are difficult to develop, and extracting value from the data is harder than for more traditional visible and near-infrared (VNIR) spectrum instruments. Thermal sensors typically have challenging cooling requirements and are power-intensive.

Processing thermal data is more challenging. Thermal IR sensors typically operate in the Medium Wave Infrared (MWIR) range (3-5 μm) and Long Wave Infrared (LWIR) range (8-14 μm). In these spectral ranges, energy is partially absorbed by water vapour and other atmospheric compounds. Atmospheric emissions also affect the imagery.

Radiometric calibration for thermal sensors is technically challenging. Onboard thermal influences and reference sources are technically difficult to implement. Thermal sensors drift and age more rapidly than other types of sensors. The thermal cycling of the satellite as it enters day and night cycles multiple times per day also affects them negatively. Achieving a precise detector response is difficult, negatively affecting the determination of absolute temperature values.

Multi-sensor architectures seek to address some of the challenges exhibited by thermal instruments by adding an additional instrument onboard the same satellite. Examples of classic missions that follow this approach include Landsat 8 and 9, which combine the Optical Land Imager (OLI) camera for visible and near-infrared imagery with a Thermal Infrared Sensor (TIRS) that operates in the LWIR spectrum. Suomi NPP and the NOAA JPSS series of satellites also carry multiple on-board sensors, with VIIRS, CrIS, and ATMS designed to image nearly the same area simultaneously.

This paper focuses on a pairing of a visible-near-infrared (VNIR) instrument with a thermal infrared (TIR) instrument onboard one satellite.

II. Overview of image product generation

To understand how VNIR and TIR instruments can complement each other, it is beneficial to first understand how Earth observation satellite imagery is typically processed.

Basic Earth Observation imagery processing typically follows the following sequence:

1) Imagery and ancillary data are extracted from raw data. Data is stored as Level 0 products.
2) Orthorectified, map-ready imagery that is radiometrically calibrated (to represent the top-of-atmosphere what the satellite sees) is created and stored as Level 1 products.

3) Corrections for the impact that the atmosphere has on reflective and emissive imagery are applied to represent the bottom-of-atmosphere (what the surface of the Earth looks like). These are stored as Level 2 products.

During the creation of Level 1 products, processing can be divided into radiometric and geometric corrections. Radiometric correction is typically applied first, during the creation of a Level 1A product. Fig. 1 shows an overview of processing reflectance (VNIR) data from Level 0 to Level 1A, with a final conversion to a top-of-atmosphere reflectance product as Level 1B.

The process is similar for emissive data, with some variations to allow for the creation of brightness temperature products instead of reflectance.

During Level 1A to Level 1B processing, geometric corrections are also applied. Geometric processing can be divided into ortho model generation and band-alignment steps, as shown in Fig. 2 . Systematic model generation relies solely on the accuracy of ancillary data and calibration

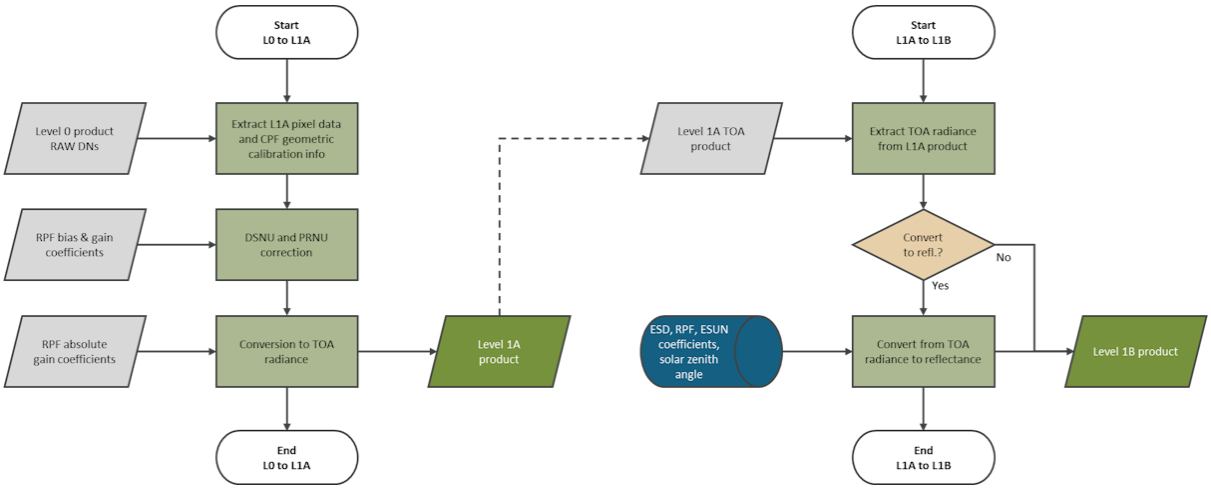

Fig. 1. Diagram showing the processing of Level 0 data to Level 1B products

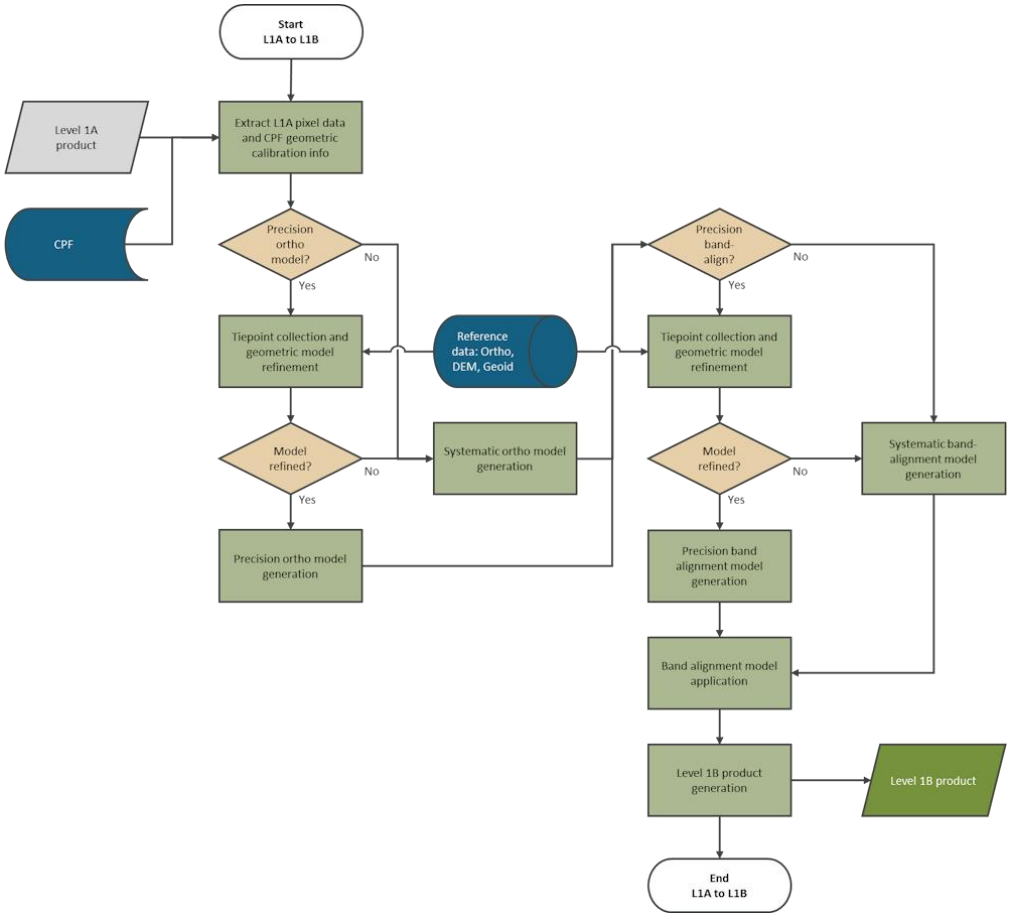

Fig. 2. Geometric processing from Level 1A data to Level 1B products

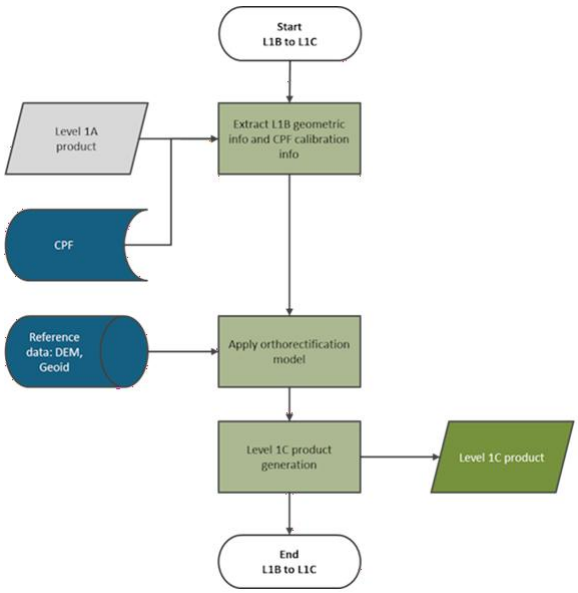

Fig. 3. Orthorectification from Level 1B data to Level 1C products

parameters. Precision refinement, in addition, uses tie-point collection between imagery to achieve greater geolocation accuracy.

The pixel values have been adjusted to achieve an accurate top-of-atmosphere radiometric calibration. The final step of Level 1 processing is to apply the orthorectification model, as shown in Fig. 3. In this stage, the pixels are resampled into a new regular grid to facilitate inter-satellite image analysis.

Next, the Level 1C product is processed into a Level 2 product, as shown in Fig. 4. Reflective data (VNIR) is processed differently from emissive data (TIR), with an additional step to estimate surface emissivity.

III. Mechanisms Negatively Impacting Image Products

Several mechanisms contribute to reducing the quality of image products, both in terms of geolocation accuracy and absolute radiometric certainty.

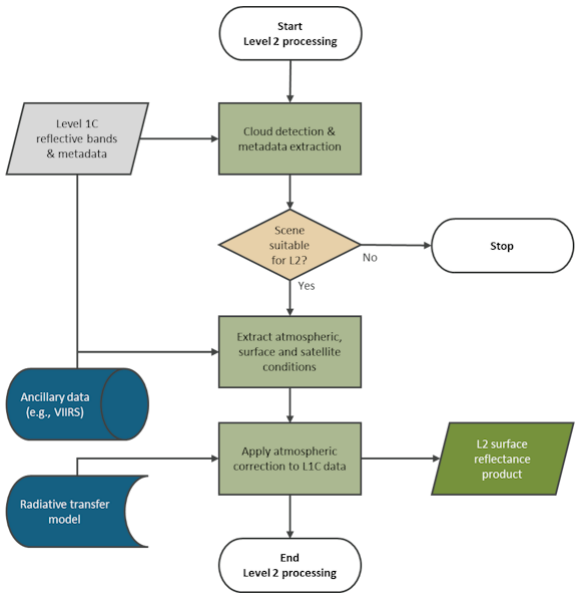

Fig. 4. Level 2 surface reflectance algorithm

Below, we highlight some of the most important mechanisms that require careful consideration. We focus on aspects that a multi-sensor architecture seeks to mitigate.

A. Mechanisms that influence image geometry

The geometric quality of an image product measures how accurately each pixel's location represents the real world. This is also referred to as the absolute geolocation accuracy.

It is common practice to process satellite imagery to a common reference instead of a very large database of high-resolution ground-truth points. This way, data can be used interchangeably across different satellites. This is referred to as relative geolocation accuracy, reflecting that the accuracy is relative to another source, which may itself have uncertainty.

A common way to assess relative geolocation accuracy is to compute residuals between a reference scene and the evaluated product. Accuracy can be expressed as a circular error of a certain percentage. For example, the CE95 metric represents a radius of a circle, centred on the reference ground location, within which 95% of the geolocation errors fall. It expresses how far off an image location can be, 95% of the time.

A key requirement for being able to measure the CE95 is to be able to identify points on both the product being evaluated and the reference data (tie-points).

The mechanisms that influence geolocation accuracy can be fixed in nature (for example, an angular offset between the ADCS and the imager's boresight) or random (for example, noise in the ADCS measurements). We will focus on random mechanisms and assume that fixed offsets can be adequately compensated for through calibration.

Geolocation errors are introduced during the processing of Level 1 products.

1) *Accuracy of position and orientation data*

Satellites use a combination of GNSS receivers (for example, GPS, Galileo) and orbit propagation models to know where they are at any given moment. The accuracy of this measurement is determined by the accuracy of the GNSS.

For the orientation or pose of a satellite, a combination of star trackers, sun sensors, gyroscopes and magnetometers is used. These sensors are fused to provide a measurement of a satellite's attitude in orbit. The submodule responsible for this is commonly referred to as the Attitude Determination and Control System (ADCS). The accuracy of these measurements relies on the accuracy of the various sensors. Since satellites can change their orientation, the frequency at which these measurements are obtained is very important.

Assuming perfect knowledge of the position of a satellite at any given time, even a very small uncertainty in the accuracy of the ADCS can have a large impact on the geolocation error of the image product. An error of just 0.01 degrees on the ADCS will introduce a geolocation error of roughly 90m on the ground. Given that most NewSpace EO satellites aim for image resolutions of 5m or better, this error is unacceptable.

2) *Accuracy of onboard clocks*

EO satellites at an altitude of 500 km fly at a speed of approximately 8 km/s. It covers 10 m on the ground in

1.25 ms. This means that a 1-millisecond timing inaccuracy will affect geolocation accuracy by multiple pixels.

The timing information between the various subsystems must be stable. It is not uncommon for GNSS systems to have a different time source than an onboard computer (OBC) and the ADCS. Any discrepancy or drift between these time sources, it can severely affect the geolocation of image data.

Fig. 5 shows the effect of incoherent timing between two sensors, resulting in an offset between the images. The blue outline represents the location of the raw, uncorrected ancillary data. The red systematic outline represents the location after the fixed offsets have been applied. The green outline shows the final precision-aligned product. A large discrepancy between the red and green outlines indicates random variances not accounted for. In this example, the TIR imager has a relatively small offset, and the VNIR imager has a large offset of approximately 6.2 km. The cause of this offset has been determined to be a 1-second discrepancy between the ADCS and the VNIR imager (not present in the TIR imagery).

3) *Satellite jitter*

Jitter refers to small, rapid and unintended variations in the satellite's attitude during imaging. It is often caused by mechanical vibrations or disturbance, by reaction wheels, thermal expansion/contraction and structural flexing.

Jitter can cause images to smear and introduce geometric distortion. This degrades geolocation accuracy and radiometric quality. This is especially true for high-resolution imagers. Fig. 6 shows an example of jitter introducing image smearing. The coloured dots represent the offset between the red and green bands, showing oscillation between small (blue) and large (red) offsets. This oscillating pattern is typical for instruments experiencing jitter.

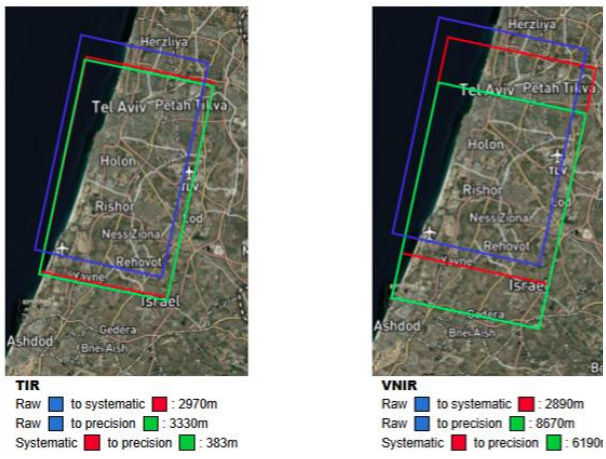

Fig. 5. Example shows the effect of timing offsets between two imagers on the same platform

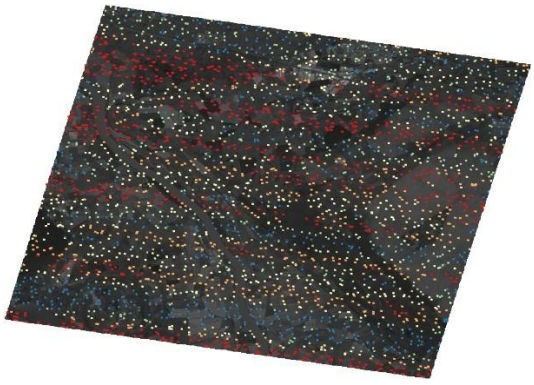

Fig. 6. Example of jitter resulting in smearing of imagery. Small offsets are indicated by blue dots, and large offsets are indicated by red dots. The offsets were measured between red and green bands

B. *Mechanisms that influence image radiometry*

Radiometric quality refers to how accurately the radiometric information within a pixel represents the truth. Radiometric values are typically determined during the calibration process as a measure relative to ground-based measurements taken at the same time, such as those from RadCalNet [1]. A satellite sits above the atmosphere, and ground-truth data is measured at ground level. Since the atmosphere affects the transmission and reflectance of light passing through it, some level of modelling is required to relate top-of-atmosphere acquisitions to bottom-of-atmosphere ground-truth data.

We focused our research on the potential benefits of VNIR and TIR data to each other when Level 2 products are generated.

1) *Low signal-to-noise ratio and readout noise*

A key performance parameter that influences geometric and radiometric modelling and processing of images is the signal-to-noise ratio (SNR). This is the ratio between the useful signal in an image and the unwanted noise. When the signal is insufficient (for example, due to low-light conditions) or the noise is increased (for example, by high sensor temperatures), the SNR decreases, and image quality degrades.

For push-broom imagers, when the useful signal is not significantly higher than the noise introduced by the readout electronics, horizontal stripes appear in the imagery, referred to as horizontal streaking.

Fig. 7 shows an example image with horizontal streaking due to readout noise and low SNR. The water vapour band suffers from a low SNR as it is used for detection in the water absorption band, where signals are typically low. The red band (right) does not suffer from low SNR and therefore has almost no horizontal streaking. The effect of readout noise on the water vapour band is that tie-point collection is difficult in areas with significant streaking.

2) *Calibration inaccuracies*

A factor influencing radiometric and geometric quality is the accuracy of the imager's radiometric non-uniformity calibration. For any imager, there will be a non-uniform radiometric response across the sensor due to spatial non-uniformities, as well as other effects such as vignetting caused

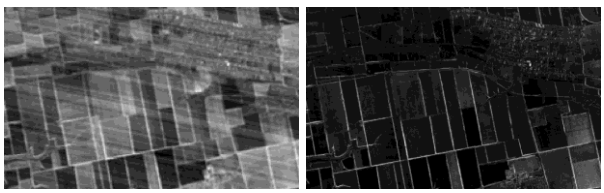

Fig. 7. Example of poor SNR and readout noise. Water vapour band (left), red band (right)

by optical elements in the light path. For push-broom sensors, this non-uniform behaviour across the sensor results in vertical streaking across the image.

The non-uniformity can be corrected by applying dark-signal and photo-response non-uniformity (PRNU) corrections. However, some residual non-uniformity will remain after correction. In the case of inadequate calibration data, not all effects will have been accounted for, resulting in incorrect calibration coefficients being used. This could happen when calibration coefficients were determined using datasets with optimal SNR, while new imagery is acquired under low-SNR conditions, resulting in over- or under-correction. Streaking can also be caused by faulty or inferior detectors.

Streaking can be corrected for using de-striping techniques, which minimise the negative effects on geometric tie-point collection down the line.

Fig. 8 shows an example of a thermal image exhibiting vertical streaking.

Fig. 9 shows examples of thermal bands with various degrees of vertical streaking for the same area.

3) *Temperature management*

Imaging instruments are designed for a specific operational temperature range. Thermal expansion affects sensor and optical geometries, potentially degrading focusing and sensor alignment. Fig. 10 shows an example of a blurred TIR image. Changes in temperature affect the sensitivity of sensor components, particularly for thermal instruments.

Obtaining and maintaining the proper operating temperature for both the sensor and the optical front end is important.

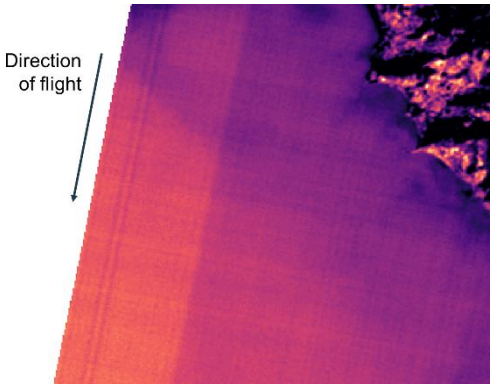

Fig. 8. Example of imagery with vertical streaking due to non-uniformity correction issues

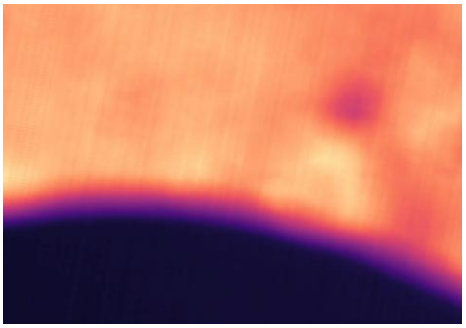

Fig. 10. Example of a blurred TIR image due to temperature outside of the ideal operating range

IV. Reliance on Reference Data

A pure systematic geometric correction of imagery, based on a geometric model derived from calibration data, is insufficient to ensure adequate geolocated imagery. Reference data is required to perform precision geolocation and orthorectification of imagery.

A. Using reference data to improve geometry

Reference data can be used to improve the geolocation and orthorectification of satellite imagery. Sufficiently cloud-free acquisitions from reference satellites, such as Landsat-8/9 and Sentinel-2, can be used to determine tie-points between the satellite and the reference images. Tie-points are salient features that are correlated between the satellite and reference image. The tie-points are then used to derive a refined model that corrects the imagery for any geometric offsets that the systematic model could not account for. Due to the effects outlined above, such as horizontal and vertical streaking, tie-point collection is not always optimal.

Various sources can be used as references for geometric correction. In the visible domain, these sources include the Sentinel-2 Global Reference Image (GRI), either as multi-layer cloud-free imagery or as a database of Ground Control Points (GCPs). Landsat-8/9 imagery, along with GCPs, is also available. Alternatively, cloud-free Sentinel-2 and Landsat-8/9 imagery acquired near the time of the satellite acquisition can be used. This comes with the benefit of the reference imagery being temporally closer to the satellite acquisition. This typically translates into better spectral and geometric correlation and improved tie-point collection.

There is a lack of reference data for thermal imagery compared to the visible domain. It is difficult to acquire thermal GCPs and imagery due to the complexity of operating thermal imagers. Significant thermal changes occur in a short

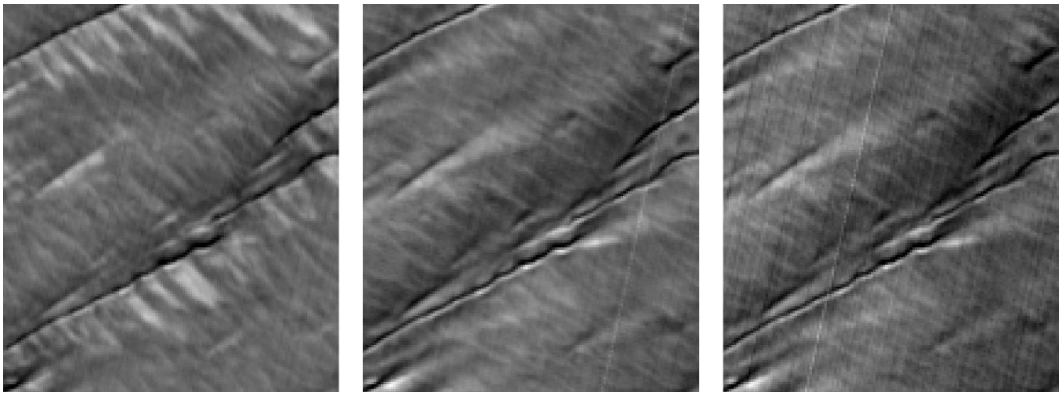

Fig. 9. Example of various degrees of vertical streaking on different thermal bands due to good non-uniformity correction (left), fair correction (centre) and poor correction (right)

period, driven by changes in surface emissivity and atmospheric conditions.

For processing of reflective imagery (VNIR), the Sentinel-2 and Landsat-8/9 OLI references contain reflective bands that match most NewSpace imagers in terms of spectral coverage. Sentinel-2 references provide imagery at 10 m resolution, while Landsat-8/9 OLI resolution is 30 m.

Fig. 11 shows the ideal scenario: matching VNIR sensor data to VNIR reference data that has a good correlation in spectral response. The overview on the left shows a uniform, dense distribution of very high-quality tie-points across the entire scene, enabling optimal geolocation of each pixel.

For processing of emissive imagery, thermal bands show good correlation with SWIR imagery under certain conditions. For this reason, the Sentinel-2 SWIR band at 20 m resolution can be used as a reference for orthorectifying thermal imagery, but this is not always suitable. Landsat-8/9 TIRS imagery can serve as a reference, with spectral characteristics closer to those of most thermal imagers. However, the coarser resolution of the TIRS thermal bands (100 m) compared to Sentinel-2 limits the achievable geometric accuracy.

Fig. 12 shows an example in which the choice of reference significantly influences tie-point collection. Red dots indicate disparities up to 100 m, and green dots indicate disparities less than 20 m. When using the Sentinel-2 SWIR band as a reference to orthorectify the thermal band, only a few tie points are available. When using the Landsat-9 TIR band as a reference, significantly more tie points are collected. This is due to the limited correlation between the SWIR spectral band in this scenario, while the Landsat-9 TIR band matches well.

When using cross-sensor approaches in which both VNIR and TIR sensors are available on the same platform, the difficulty of finding tie-points between bands with low spectral correlation is a significant problem. As was the case when using the Sentinel-2 SWIR band as a reference for thermal orthorectification, the same issue arises when trying to correlate the thermal bands with the VNIR bands, which may have significantly different features.

Fig. 13 and Fig. 14 show examples of different levels of correlation between VNIR and TIR imagery. The first example shows a high number of high-quality (green) tie-points. The second example shows only a few low-quality (orange) tie-points.

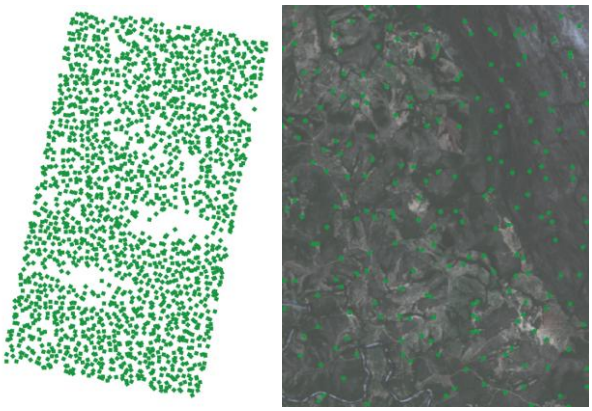

Fig. 11. Example of a VNIR scene correlating very well with VNIR reference data. Overview of scene tie-points (left), close-up showing the density of high-quality (green) tie-points (right)

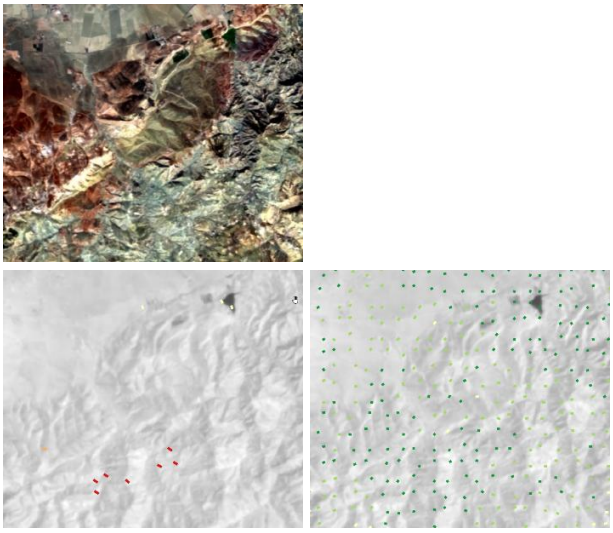

Fig. 12. Example of the use of different references for thermal image orthorectification. RGB (top-left), Sentinel-2 as reference (bottom-left), Landsat-9 as reference (bottom-right)

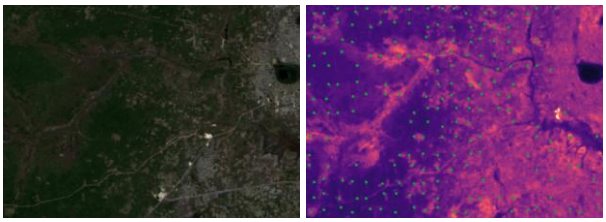

Fig. 13. Example of good correlation between VNIR (left) and TIR (right, with tie-points)

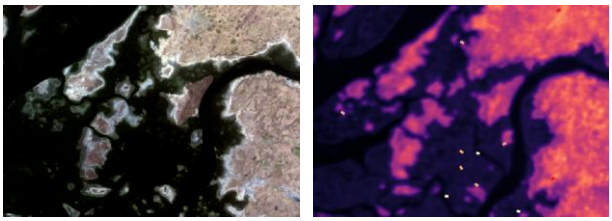

Fig. 14. Example of poor correlation between VNIR (left) and TIR (right, with tie-points)

The success of salient feature matching depends heavily on the type of surface being imaged. Scenes with snow or ice typically show a weak correlation between VNIR and TIR. For scenes with high vegetation content, such as forest canopies, grasslands and croplands, the VNIR and TIR bands tend to correlate moderately well. Urban and desert areas exhibit a strong positive correlation.

Due to the significant changes that can occur in thermal signatures over time, temporal decoherence between the satellite imagery and the reference plays a key role in the geometric processing of thermal imagery. This can be seen in Fig 15, where the closest available reference image was several days prior to the satellite acquisition. Note the significant changes in the water surface area.

B. Using reflective data to improve thermal radiometry

The availability of VNIR data enables land cover classification and the subsequent estimation of surface properties, such as surface emissivity. Various classification techniques are available that use visible-to-short-wave infrared data to determine surface types. As surface emissivity is a major contributor to surface temperature estimation,

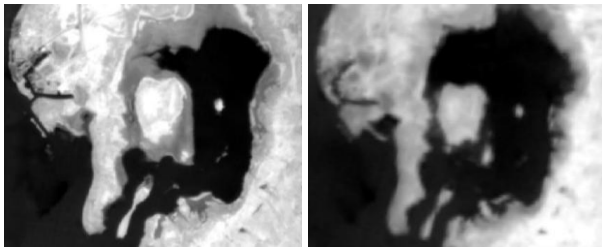

Fig. 15. Example of temporal decoherence between a satellite image (left) and a reference image (right)

accurate emissivity estimates are critical for accurate surface temperature determination.

When multiple thermal bands are available, they can be used to estimate the surface emissivity using temperature- and emissivity-separation algorithms [2]. When only a single thermal band is available, this is no longer possible. However, with the availability of reflective data, alternative methods can be used to determine surface emissivity. For example, the NDVI-based emissivity retrieval method [3] enables a rough estimate of surface emissivity, which can greatly reduce errors introduced by assuming a uniform emissivity across a scene.

Reflective data in the VNIR and SWIR ranges can also be used to perform aerosol retrieval [4]. As reference data for aerosol optical thickness (AOT) may not be readily available at the time of processing thermal data, real-time information on atmospheric properties, such as AOT, is highly beneficial. It is used to correct atmospheric effects. With no reflective data available, one is dependent on historical atmospheric data or weather forecasts, which can have significant latency.

Below is an example showing how the surface emissivity can be determined over an area with diverse surface types, from vegetation to desert, when reflective data is available. The VNIR bands were used to perform surface classification and estimate the surface emissivity using the NDVI-based method.

Fig. 16 shows the various vegetation and desert areas in the scene, with these having clear differences in emissivity. Without the VNIR bands, the emissivity would simply be estimated as a fixed value across the scene and would not account for the changes in surface types. This can lead to large errors in surface temperature, as the surface emissivity can range from high (above 0.95) for water, vegetation, and ice, to low (below 0.9) for deserts and semi-arid regions.

Fig. 17 shows an example of a Landsat-9 thermal image (Band 10). The Landsat Level 1 product was processed using the VNIR bands to retrieve or "detect" the emissivity and AOT, and using a default fixed emissivity of 0.95 and AOT of 0.125 (middle). The difference between the "detected" and "fixed" results shows that the surface temperature estimates are highly dependent on the atmospheric and surface properties used as input to the estimation. There is an approximate 3 K offset between the "detected" and "fixed" surface temperature products. This is caused by offsets between the emissivity and AOT values, with the detected emissivity in the range 0.9 to 1.0, while the AOT is closer to 0.08.

RadCalNet [1] provides spectral data of reference sites from the visible to short-wave infrared range (400 nm to 2500 nm) for both the bottom-of-atmosphere and top-of-atmosphere at regular intervals. Satellite data can be compared

Fig. 16. Landsat-9 RGB image (left) and emissivity estimate based on NDVI (right)

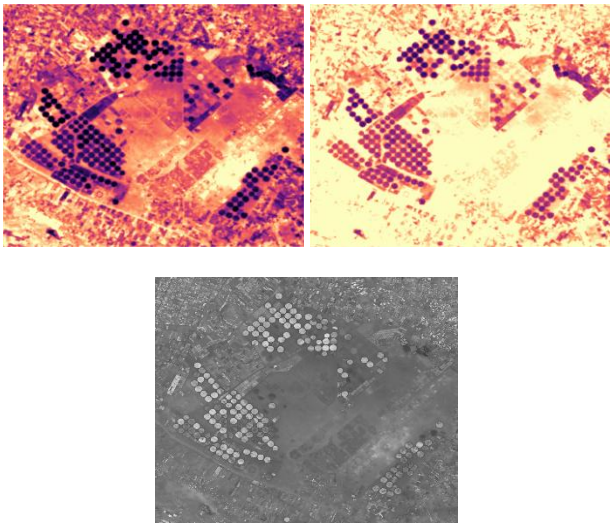

Fig. 17. Landsat Band 10 TIR image processed using AOT and emissivity retrieved from reflective data (top-left) and processed using a fixed AOT and emissivity (top-right). The difference between the two images is also shown (bottom)

with bottom-of-the-atmosphere levels to validate the radiometric performance. This can be used to validate VNIR imagery, but does not provide reference data for thermal imagery.

For thermal imagery, reference sites such as SURFRAD [5] or JPL sites [6] can be used. Alternatively, overpasses from satellites with thermal bands (for example, Landsat-8/9) can be used for validation. In the future, TIRCalNet sites [7] will provide reference data through a calibration network of SI-traceable instruments, thereby offering more opportunities for validating thermal data.

V. Conclusion

In this paper, we have described the general challenges of processing TIR imagery, both geometrically and radiometrically, for NewSpace satellites. We introduced various methodologies for addressing these challenges.

Relying on external reference data for geolocation accuracy means identifying matching tie-points between the reference and acquired imagery. Scenes with large, homogeneous areas, such as deserts, water, snow, or areas covered by clouds, will fail to provide tie-points. This is also true for multi-sensor architectures, where it will be difficult to find tie-points between the sensors if a systematic model alone does not provide sufficient accuracy. However, there is no risk for temporal decoherence, and if the useful thermal imagery is cloud-free, the VNIR imagery will be as well.

Using satellite VNIR data to orthorectify TIR data is ideal. However, due to difficulties in finding tie-points between VNIR and TIR, caused by variable timing offsets and spectral mismatches, this is not always possible.

When treating VNIR and TIR instruments as standalone, the type of reference data to use for orthorectification becomes very important. Using SWIR reference data to geolocate TIR data can yield good results if the right combination of tie-point collection strategies and correlation techniques is used. Alternatively, using TIR reference data to geolocate TIR data can work well; however, the performance is more sensitive to temporal decoherence than that of typical VNIR data. Using different references for VNIR (such as Sentinel-2 for higher resolution) and TIR (such as Landsat for better spectral match) can introduce unwanted errors, as the two reference sources have a level of geometric uncertainty between them.

When processing TIR data to a Level 2 product, access to VNIR data acquired at the same time will assist in a more accurate estimation of atmospheric properties when no real-time reference data are available. When no VNIR data is used, relying on fixed values for surface and atmospheric properties, such as emissivity and AOT@550nm, leads to significant errors in surface temperature estimation.

Our research highlights the potential benefits of multi-sensor architectures (specifically, VNIR and TIR instruments). Cost and bandwidth restrictions aside, the benefits of having a VNIR sensor seem to outweigh the downsides. Care must be taken to ensure that timestamps for ADCS and imagery are as accurate as possible. One must accept that not all areas will provide suitable tie-points for additional geolocation refinement between the two sensors.

[1] M. Bouvet et al., "RadCalNet: A Radiometric Calibration Network for Earth Observing Imagers Operating in the Visible to Shortwave Infrared Spectral Range," Remote Sensing, vol. 11, no. 20, p. 2401, Jan. 2019, doi: https://doi.org/10.3390/rs11202401.

[2] A. Gillespie, S. Rokugawa, T. Matsunaga, J. S. Cothern, S. Hook, and A. B. Kahle, "A temperature and emissivity separation algorithm for Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER) images," IEEE Transactions on Geoscience and Remote Sensing, vol. 36, no. 4, pp. 1113–1126, Jul. 1998, doi: https://doi.org/10.1109/36.700995.

[3] E. Valor, "Mapping land surface emissivity from NDVI: Application to European, African, and South American areas," Remote Sensing of Environment, vol. 57, no. 3, pp. 167–184, Sep. 1996, doi: https://doi.org/10.1016/0034-4257(96)00039-9.

[4] H. Zhang et al., "An enhanced VIIRS aerosol optical thickness (AOT) retrieval algorithm over land using a global surface reflectance ratio database," Journal of Geophysical Research. Atmospheres, vol. 121, no. 18, Sep. 2016, doi: https://doi.org/10.1002/2016jd024859.

[5] "SURFRAD (Surface Radiation Budget) Network," Noaa.gov, 2018. https://gml.noaa.gov/grad/surfrad/ (accessed Apr. 29, 2026)

[6] "JPL Radiometer Validation Sites," Calibration/Validation, 2026. https://calval.jpl.nasa.gov/sites (accessed Apr. 29, 2026).

[7] M. Chapelier et al., "TIRCalNet: Towards a Global Network for Thermal Infrared Calibration," presented at the VH-RODA 2025, Frascati (RM), Italy: European Space Agency, Nov. 2025.