

Polar perspectives

a deep dive into geo-referencing historical Antarctic photos

Dahle, Felix; Lindenbergh, Roderik; Wouters, Bert

DOI

[10.1080/17538947.2024.2406384](https://doi.org/10.1080/17538947.2024.2406384)

Publication date

2024

Document Version

Final published version

Published in

International Journal of Digital Earth

Citation (APA)

Dahle, F., Lindenbergh, R., & Wouters, B. (2024). Polar perspectives: a deep dive into geo-referencing historical Antarctic photos. *International Journal of Digital Earth*, 17(1), Article 2406384. <https://doi.org/10.1080/17538947.2024.2406384>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.



Polar perspectives: a deep dive into geo-referencing historical Antarctic photos

Felix Dahle, Roderik Lindenberg & Bert Wouters

To cite this article: Felix Dahle, Roderik Lindenberg & Bert Wouters (2024) Polar perspectives: a deep dive into geo-referencing historical Antarctic photos, International Journal of Digital Earth, 17:1, 2406384, DOI: [10.1080/17538947.2024.2406384](https://doi.org/10.1080/17538947.2024.2406384)

To link to this article: <https://doi.org/10.1080/17538947.2024.2406384>



© 2024 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group



[View supplementary material](#)



Published online: 22 Sep 2024.



[Submit your article to this journal](#)



Article views: 165



[View related articles](#)



[View Crossmark data](#)



Polar perspectives: a deep dive into geo-referencing historical Antarctic photos

Felix Dahle, Roderik Lindenbergh and Bert Wouters

Department of Geoscience & Remote Sensing, Delft University of Technology, Delft, The Netherlands

ABSTRACT

The utility of historical image repositories is often limited due to the lack of geo-referencing. A good example is the TriMetron Aerial (TMA) archive, a collection of historical aerial images of Antarctica between 1940 and 2000. These images are difficult to use, as their geolocation is only approximately, with location errors in the order of tens of km. This study addresses this challenge by developing an automated geo-referencing workflow that leverages recent advancements in machine-learning-based tie-point matching. We use the algorithm LightGlue, to establish tie-points between geo-referenced Sentinel-2 satellite imagery and historical non-geo-referenced aerial images. To aid the process, we use already known approximate positions of the historical images. Due to the sub-optimal and inhomogeneous quality of the aerial images, only a portion of the images can be geo-referenced directly by matching. For the remaining images, we employed alternative means of geo-referencing, again based on tie-point matching. Out of a subset of 4,459 images located inside the research area, 3,393 images could be geo-referenced, a percentage of 76%. Reasons for the geo-referencing failing are insufficient contrast or an approximate position too far away from the real position. The workflow can easily be applied to historical images from other archives, to enhance the usability of historical image repositories for scientific research.

ARTICLE HISTORY

Received 31 March 2024
Accepted 15 September 2024

KEYWORDS

Antarctica; geo-referencing; historical imagery; machine learning; tie-points

1. Introduction

Historical aerial imagery offers a window into the past, granting access to valuable information from before the satellite era. Many of these images, initially produced for military purposes on a large scale, were consigned to archives and their contents were inaccessible. However, in recent years these archives have been digitized and made available to the public. Consequently, the potential of these historical image repositories has been recognized across a spectrum of scientific disciplines, with applications ranging from archaeological (Cowley and Stichelbaut 2012; Usmanov et al. 2018), environmental (Heisig and Simmen 2021; Nicu and Stoleriu 2019) to cryospheric sciences (Geyman et al. 2022; Pope et al. 2014).

In cryospheric research, historical aerial images can provide information on the state of the ice sheet before the satellite era for a wide range of subjects: they can be used for evaluating glacier mass balance, shedding light on both historical states and recent changes (Denzinger et al. 2021; Korsgaard et al. 2016), aid in understanding the structure, dynamics, and flow of ice sheets over time

CONTACT Felix Dahle ✉ f.dahle@tudelft.nl

© 2024 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group
This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

(Yu et al. 2022), and provide insight into past ice extents in specific regions (Cook and Vaughan 2010). Additionally, these images are valuable for monitoring non-ice-related phenomena, such as identifying areas of coastal erosion (Wang et al. 2022) or observing thermo-erosion gully processes (Morgenstern et al. 2021; Perreault et al. 2017).

The use of historical aerial images is not confined to 2D analyzes. Advanced techniques like Structure-from-Motion (SfM) enable the generation of historical 3D models from these images (Child et al. 2020; Knuth et al. 2023; Zhang, Rupnik, and Pierrot-Deseilligny 2021).

However, processing archives of historical images often face a significant challenge that limits their use: many images are not, or incorrectly, geo-referenced. In this context, geo-referencing refers to the process of associating specific points within an image with their corresponding real-world locations. This typically involves identifying tie-points, which are points on two distinct images representing the same geographical location. Working with un-referenced images requires researchers to manually identify similar structures and pin-point tie-points. This is arguably the primary reason these valuable historical images are underutilized in research.

Efforts have already been made to automate this process (Chen and Tseng 2016; Cléri, Pierrot-Deseilligny, and Vallet 2014; Karel et al. 2013) by using tie-point detection and feature matching. Furthermore, geo-referencing has expanded beyond mere image comparison; photogrammetric methods like Structure from Motion (SfM) have also been employed (Frankl et al. 2015; Giordano, Le Bris, and Mallet 2018; Heisig and Simmen 2021). However, these methods typically rely on having scenes with rich textures to provide distinct features for tie-point matching. Furthermore, in general, images of the same type, such as aerial-to-aerial, are compared. Geo-referencing and SfM are particularly challenging when geo-referencing images of the cryosphere, where many images lack contrast and only have few distinct landmarks. Matching is further hindered, as corresponding modern aerial imagery for these regions is generally not available which implies that historical areal images can only be matched to contemporary satellite imagery.

To address these challenges, we have developed a workflow that applies recent advancements in neural-network-based tie-point matching techniques to historical aerial images. Our approach aims to identify tie-points between historical aerial imagery and contemporary satellite imagery, streamlining the geo-referencing process as much as possible. Recognizing the inherent difficulties in matching aerial and satellite images, we have also devised several algorithms for indirect geo-referencing.

We applied our workflow to images from the TriMetrogon Aerial (TMA) archive, a collection of historical photographs of Antarctica, acquired between 1940 and 2000, with the majority of the images taken in the 1960s. The archive has been exploited before to gain insights into the condition of Antarctica prior to the satellite era, in studies focusing on individual or a subset of glaciers, for which manual geo-referencing suffices (Child et al. 2020; Cook et al. 2005; Kunz et al. 2012; Zhao et al. 2017). For the total of around 330,000 images in the archive, such a manual approach is unfeasible. In this study, we present a workflow that can be executed with minimal user interaction and in a reasonable amount of time. We specifically target the Antarctic Peninsula due to its extensive coverage by overlapping images, which facilitated the evaluation of our geo-referencing accuracy.

While our code was tailored for this specific archive, its design allows for adaptation to other non-referenced datasets. The geo-referencing code and the resultant geo-referenced images are accessible at GitHub (Dahle 2024a). Additionally, we will provide the footprints and transformation matrices for each geo-referenced photograph of the TMA archive. Over time, we aim to geo-reference and publish the entire archive.

2. Data

2.1. TMA-archive

The TMA archive is a collection of photographs retrieved by the U.S. Navy between 1946 and 2000 along the margins of Antarctica at approximately 125,000 locations, consisting of about 330,000

images. This archive consists of black and white trimetrogon (a camera system mounted on airplanes pointing left, right and nadir) aerial photos. Some examples are shown in Figure 1. While acquiring TMA data, airplanes flew along so-called flight paths: approximate straight lines with more or less constant height. About 15,000 images were collected on the Antarctic Peninsula, most of them collected in the mid-1960 s. The imagery has been digitally scanned at high resolution (25 micron/1000 dpi) by the US Geological Survey (USGS) and the Antarctic Geospatial Information Center (University of Minnesota) and is publicly available at (Polar Geospatial Center 2023).

To date, the aerial photographs within the collection have not been geo-referenced on a large scale. Nonetheless, as can be seen in Figure 2, digitized flight lines exist, which by visual inspection provide approximate locations of the actual flight paths and camera positions. These flight lines are distributed with various attributes such as the altitude of the airplane and the acquisition date. However, these attributes are not consistently available for every camera position, and their accuracy cannot be assured.

In this work, we focus on the images with nadir view, as these are most straightforward to geo-reference. The camera position is consistent for each triple of left-right-nadir images. Therefore, although the initial goal is to geo-reference the 4,459 available nadir images located in the research area on the Antarctic Peninsula, a successful procedure will enable a relatively easy consecutive geo-referencing of corresponding left-right images as well.

2.2. Sentinel-2 satellite imagery

For our geo-referencing method, the historical images must be compared with already geo-referenced data. In our study, we employ imagery from Sentinel-2, an Earth observation mission under the Copernicus program of the European Commission. This mission offers freely available optical imagery characterized by a medium spatial resolution. Although our research is based on Sentinel-2 imagery, it is important to mention that other geo-referenced datasets can also be




Figure 1. Example images from the TMA archive.




Figure 2. Study area with the digitized flight lines for the TMA-archive with their respective recording date. The green dots highlight the position of the use cases discussed in Section 4: (a) Shoosmith glacier (b) McMorris Glacier (c) Fleming Glacier (d) Somigliana Glacier (e) Unnamed Glacier near Mt. Wilcox (f) Sibieli Glacier (g) Rothera Research station. The bold red line highlights a non-straight flight path discussed in Section 5.

suitable. The Sentinel-2 data used for geo-referencing spans from 2016 to 2023. Using Google Earth Engine, these images are compiled into a cloud-free composite by masking cloud-cover pixels and computing the median value for each pixel. Given our objective of geo-referencing this data with the historical aerial imagery, we extract bands 2, 3, and 4 to produce an RGB representation. The spatial resolution for these bands is 10 meters. To ensure radiometric consistency with historical images, all images are standardized to an 8-bit format, with RGB values ranging from 0 to 255.

3. Methodology

In our geo-referencing approach, we employ tie-point matching techniques between images, making use of recent technological advancements to design a workflow that can automatically geo-reference historical images.

Figure 3 illustrates the workflow of our methodology. The workflow is specified in detail below: first, we prepare the TMA images for matching and geo-referencing by creating approximate footprints (Section 3.2). Second, we geo-reference part of the TMA images using Sentinel-2 imagery and subsequently perform a quality assessment (Section 3.3). Due to the mixed quality of the historical images, many of them cannot be geo-referenced directly on the first attempt. Therefore, we apply alternative in-direct strategies to also ensure successful geo-referencing for these challenging images (Section 3.4). Due to the importance of geo-referencing in almost every step, our tie-point matching workflow is explained first in Section 3.1.

3.1. Tie-point matching

Various tie-point detection and matching techniques can be incorporated into our geo-referencing workflow. In this research, we use ‘LightGlue’ (Lindenberg, Sarlin, and Pollefeys 2023), a

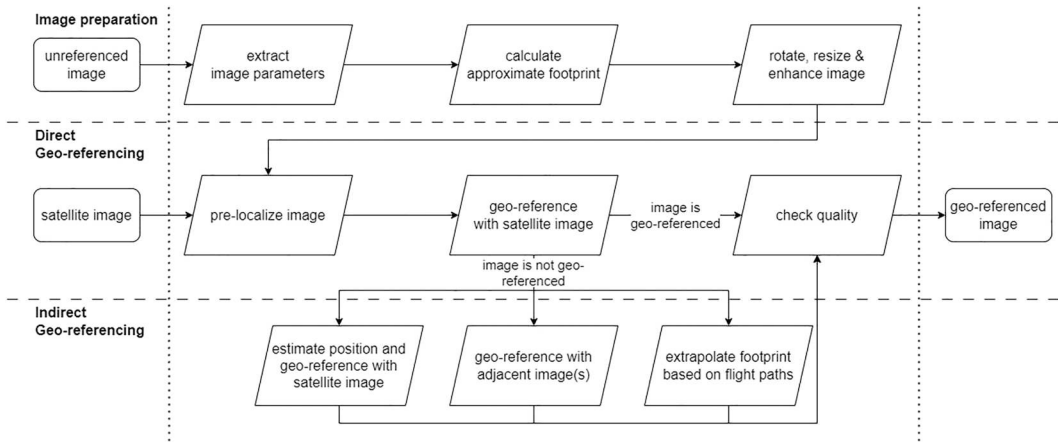


Figure 3. Geo-referencing workflow. If direct geo-referencing by matching a prepared historical image to recent Sentinel-2 data fails, indirect strategies will be used.

successor to the popular tie-point matching algorithm ‘SuperGlue’ (Sarlin et al. 2020). Both algorithms leverage the tie-point detection method ‘SuperPoint’ (DeTone, Malisiewicz, and Rabinovich 2018). SuperPoint is a fully convolutional deep neural network, trained to detect interest points (unique points in an image) and compute their corresponding descriptors between pairs of images. Given a collection of these points and descriptors, LightGlue can be used to identify matches between the images.

It is noteworthy, that both LightGlue and SuperGlue were primarily trained on datasets without historical or aerial images. Moreover, the specifics of their training parameters remain undisclosed, rendering retraining for historical images impractical. Nevertheless, they are among the most proficient tie-point matching techniques currently available that can match images of varying sizes and are already used successfully in many geo-science related applications (Ghuffar et al. 2023; Maiwald, Feurer, and Eltner 2023).

However, certain constraints exist: the algorithms are not rotation-invariant and the maximum rotation between two images should not surpass more than 45 degrees, as otherwise fewer matches are found (Tyszkiewicz, Fua, and Trulls 2020). Furthermore, image pairs must fit within the GPU memory, while scanned historical images typically have high resolutions, often surpassing 10,000 pixels in width or height. While downsizing the images is a potential solution to fit the images into the memory, it reduces the information content and consequently the number of detectable tie-points. To maintain a dense distribution of tie-points, we devised a workflow that optimizes tie-point detection using LightGlue, specifically for large images. Our tie-point matching workflow is divided in four different steps, as depicted in Figure 4. Only the first step is mandatory, all other steps can be applied optionally.

3.1.1. Initial tie-point matching

In the initial matching, the selected matching algorithm (in this case ‘LightGlue’) is applied to resized images. Both images are resized iteratively by 10% until the matching process fits into GPU memory. Parts of the images can be masked to ignore certain areas for tie-point matching. To account for the rotational dependency of the matching algorithm, the image on the right is matched four times, for four different rotations of 0, 90, 180 and 270 degrees, respectively. The rotation with most matches is fixed for the further workflow. Finally, the tie-point coordinates are converted to original images size. Figure 5 shows an initial matching between the two images, which are resized from 10000×10000 pixels to $\sim 2000 \times 2000$ pixels (a reduction of 96 % in image size). The red lines indicate matches that are masked, while the green lines represent successful matches that will be kept for the further workflow.

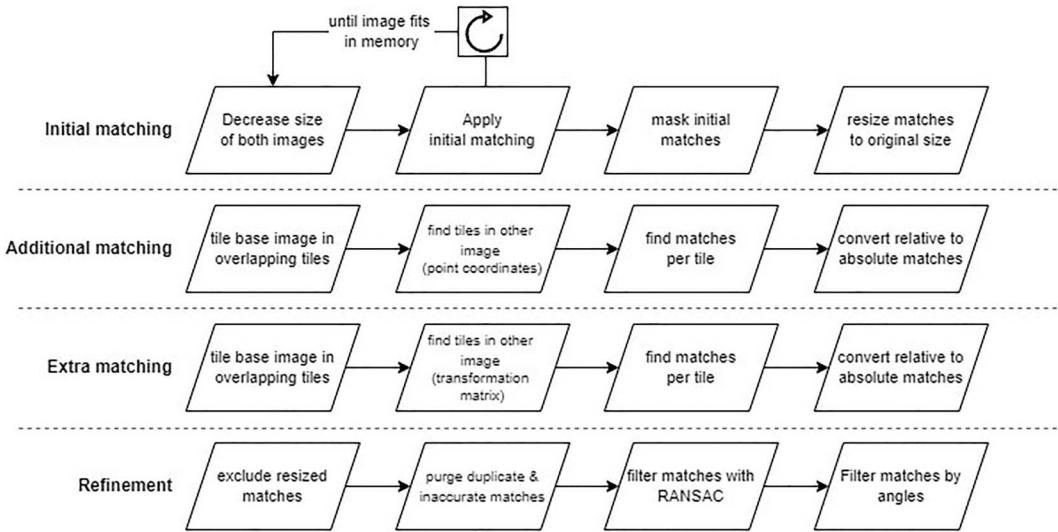


Figure 4. Tie-point matching workflow. The first row is mandatory, all other rows can be applied optionally.

3.1.2. Additional tie-point matching

In the additional matching process, the larger image (or a randomly selected image if both are of the same size) is divided into multiple overlapping tiles, similar to the approach in (Zhang, Rupnik, and Pierrot-Deseilligny 2021). In Figure 6, the left image is tiled into 30 tiles labeled A to AD. For each tile, all initial tie-points within that tile are selected. Corresponding tie-point coordinates are used to create for each tile a subset of the right image. To account for outliers in the initial matching process, this subset is adjusted in size so that a tile and a subset in both images have identical size. Tie-point matching is then applied between the tile and the subset in the original resolution. This is done iteratively for all tiles containing initial tie-points, while tiles without tie-points (e.g. the red tiles in Figure 6) are excluded.

3.1.3. Extra tie-point matching

For now, the matching only focused on parts of the image where in the initial matching step tie-points were found. To also use information in other parts of the image, a transformation matrix




Figure 5. Initial matching between two images (IDs CA216632V0282 and CA216632V0283), that are resized to 2000×2000 pixels. 881 valid matches (green lines) and 407 masked matches (red lines) are found.

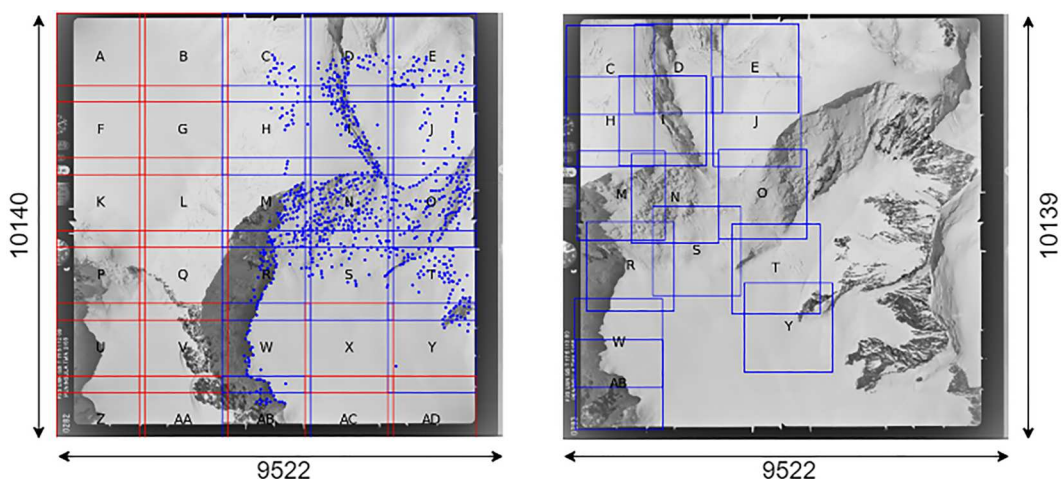


Figure 6. Additional matching in original resolution: Only tiles with initial tie-points (blue) are matched in original resolution and have a corresponding subset in the right image. Red tiles do not contain any tie-points.

is derived from the already established tie-points to determine the spatial correspondence between the two images. Again, the left image is tiled in multiple overlapping tiles. For each tile, a corresponding subset on the right image is created by applying the derived transformation matrix on the coordinates of each tile. Like in the previous additional matching step, the matching algorithm is applied at the original resolution between each tile and its correspondent subset. Subsets predicted to fall outside the image boundaries based on the transformation matrix are not included in this matching step.

3.1.4. Refinement

If additional or extra matching is applied, the tie-points detected in the initial matching are removed due to their lower accuracy, as these tie-point were obtained in down-scaled imagery of lower resolution. Subsequently, all remaining matches are combined, and two criteria are applied as detailed in Table 1: first, duplicate matches (with the same coordinates in both images) are removed. Second, based on a five-pixel distance, tie-points within 5px distance are averaged, while tie-points exceeding this 5px distance are discarded.

As a final step, tie-points are filtered with two methods. The first method employs the classical RANSAC algorithm (Fischler and Bolles 1981) for outlier removal. The second method focuses on the angles formed by lines connecting tie-points. If a predominant angle is observed among these

Table 1. Example for the refinement of matches: old tie-points T1 and T2 have identical coordinates and are combined to new tie-point T1'.

Old tie-points	x1	y1	x2	y2
T1	100	100	120	120
T2	100	100	120	120
T3	150	150	170	170
T4	150	150	172	172
T5	200	200	210	210
T6	200	200	220	220
New tie-points	x1	y1	x2	y2
T1'	100	100	120	120
T2'	150	150	171	171

Note: Old tie-points T3 and T4 have different x2/y2 coordinates, which are closer than 5px. Therefore, they are combined to new tie-point T2'. Old tie-points T5 and T6 are more than 5px apart and are therefore discarded.

lines, any tie-point line deviating more than 45 degrees from this angle is deemed an outlier and discarded. It is crucial, that the images being matched share the same resolution to ensure parallelism of these lines. To achieve this, the resolution of the smaller image is temporarily adjusted to align with that of the larger image.

3.2. Image preparation

To prepare the images for the matching, approximate footprints are calculated and the images are enhanced for higher contrast. Both steps are optional, but can speed up and enhance the quality of geo-referencing.

3.2.1. Approximate footprint estimation

Initial positions (x and y in the Antarctic Polar Stereographic coordinate system) and TMA image orientation are available from the Polar Geospatial Center (Polar Geospatial Center 2023). For each image an estimated position is available, manually determined based on historical flight documents. To compute the footprints, the x, y, z position, rotation and the camera's focal length are required. These approximate footprints are neither at the correct position nor have the correct size, but still give a good first indication.

For some images, the focal length and camera height are written on the images themselves, as illustrated in Figure 8. To extract this information, we use the Character Recognition (OCR) tool PaddleOCR (Du et al. 2020), as it offers the best performance for recognizing such handwritten text. This tool serves two purposes: extracting metadata such as focal length or height (box A) and identifying regions in the images that are excluded from tie-point searching because they contain text (box B). For images where parameter extraction is not possible, we use average values derived from all other images on the same flight path.

The boundaries of this footprint are determined using the so-called CameraCalculator, adapted from Java to Python (Engel 2020). The determination is initialized with the camera's horizontal and vertical field of view, altitude, roll, pitch, and heading. Four rays are defined to represent the camera's field of view and are rotated based on the roll, pitch, and yaw parameters. The intersection points between these rays and the ground, represented as a flat plane, are computed. These intersection coordinates are then aligned with the camera's actual position, providing the approximate footprint's boundaries.

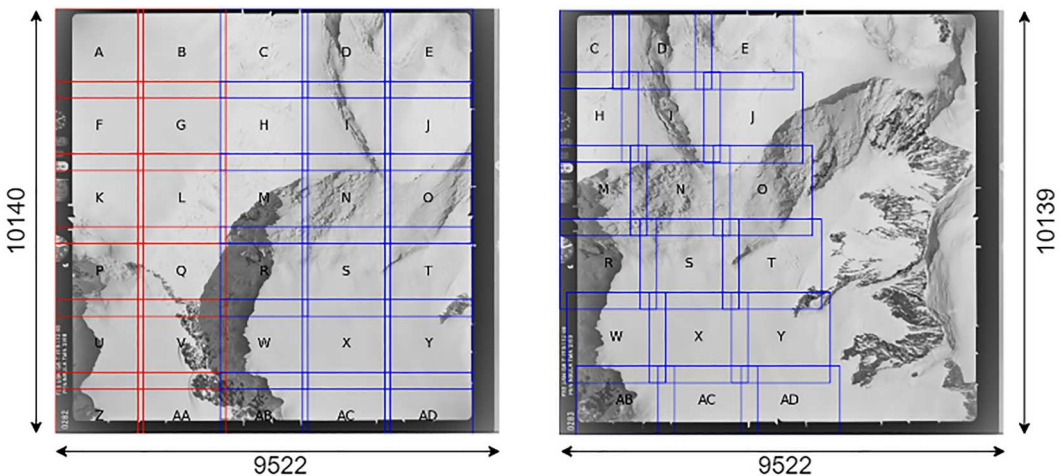


Figure 7. Extra matching in original resolution: For each tile on the left image the equivalent subset on the right image is calculated and tie-point matching is applied in full resolution. Tiles that do not have a subset are highlighted in red.




Figure 8. Historical image (CA182632V0127, located on Foyen Coast) with its metadata (A: image context information, B: height information) included in the image.

Subsequently, the initial footprint is used to get the average elevation of the encompassed area. The mean elevation is derived from raster elevation data (in this instance, REMA elevation data (Howat et al. 2019) specific to the footprint's bounding box). The camera's altitude is then adjusted to recalculate its height in relation to the surface's average height instead of sea level. This updated altitude is then used to recalculate the footprint, which is usually smaller than the initial estimate due to the altitude adjustment.

3.2.2. Image enhancement

Before the geo-referencing with the satellite imagery starts, several enhancement steps are applied. These are divided into four primary stages: border masking, resolution adjustment, image rotation, and contrast enhancement.

(1) *Border masking.* Many historical photographs were captured using analogue cameras, often featuring borders with fiducial marks. The high contrast between these borders and the main image content can introduce numerous false errors during tie-point matching. For the images in the TMA archive, all borders adhere to a consistent spatial pattern, allowing for straightforward masking during the geo-referencing process.

(2) *Resolution adjustment.* Historical images are usually scanned at 600dpi, resulting in a high resolution, typically measuring $\sim 10,000 \times 10,000$ pixels. Aligning with the approximate footprint shows that the ground sampling distance is in the order of 60 cm. In contrast, the Sentinel-2 data consists of $\sim 1,000 \times 1,000$ pixels, corresponding to a pixel spanning 10 m. Conducting tie-point matching across such disparate resolutions is challenging. To address this, the larger image (the TMA image) is resized based on its approximate footprint to match the spatial resolution of the satellite image.

(3) *Image rotation.* The images are generally not oriented facing North and require a rotation based on the known flight path angle. The necessity of this step depends on the tie-point matching method in use. For instance, LightGlue is not rotation invariant, and without this adjustment, fewer or no matches may be identified. When the rotation angle is not available, an alternative strategy involves rotating the images in increments of 90 degrees.

(4) *Contrast enhancement.* A substantial portion of the images have sub-optimal contrast or inconsistent brightness levels, further complicating the tie-point matching process. To account for this, a median filter is first applied to reduce image noise. This is followed by a linear contrast

stretch and concluded with a gamma adjustment. The enhancement approach is inspired by the SpyMicMac repository (McNabb et al. 2020).

3.3. Direct geo-referencing with satellite images

Geo-referencing historical images is based on tie-point matching between these images and contemporary satellite data. For this method to be effective, certain consistent structures in both the historical and modern imagery need to be available. In most applications, these are human-made structures such as houses or streets. While such structures are absent in Antarctica, rock formations have proven consistent over time, in this instance, spanning nearly 80 years.

3.3.1. Information content

However, not all images from the TMA archive are suitable for geo-referencing due to a lack of distinct structures. To address this, we compute an ‘information content’ value for each image prior to geo-referencing. The information content is determined based on the count of distinct points in an image, as detected by the aforementioned method SuperPoint. A count of 2500 tie-points or more is deemed the highest score with a value of 1. Counts below this are calculated as a percentage of 2500. This benchmark can be adjusted based on specific requirements and the number of tie-points deemed necessary. The footprints of images with too low information content or where the geo-referencing with satellite images failed, will be estimated indirectly, as outlined in Subsection 3.4.

3.3.2. TMA footprint relocation

In the case of the TMA archive, the starting boundary of the satellite image (that is used for matching) is defined by the approximate footprint, which however can deviate by several kilometres from the real position. To increase the possibility of a successful tie-point matching, two refining steps are employed. These adapt the the boundary of the satellite to the most likely location for a more successful geo-referencing.

(A) Segment based footprint relocation. Figure 9 illustrates the first refinement step of relocating the footprint of the satellite image: the satellite image around the initial approximate footprint is segmented into smaller sections, each characterized by the dimensions of the approximate




Figure 9. Relocating the TMA image (ID CA174632V0181) 20 km from approximate position (grey bounding box) to refined approximate position (red bounding box).

footprint. An initial tie-point matching is performed at the position of the approximate footprint (level 1). If the number of recognized tie-points falls below a predetermined threshold, a second tier (level 2) is initiated. This phase involves tie-point matching at eight surrounding positions (cell 2B – cell 2I) encircling the central location of the approximate footprint. Although the positions depicted in Figure 9 appear uniformly aligned, this is a simplified representation. In practice, these positions exhibit slight overlaps, ensuring tie-points near boundaries are incorporated. In this demonstration, a match is found at cell 2D of the second tier. Without a match in the second tier, another tier of 16 cells can be added resulting in a 25 km search radius. The search concludes if no match is found by this stage. Theoretically, one can endlessly iterate this process, albeit at increasing computational costs.

(B) Footprint relocation by tie-point maximization. This new approximate footprint is further refined in a second step: the bounds of the footprint are iteratively adjusted to maximize the tie-point count. This process is shown in Figure 10. The historical image is segmented into four quadrants by bisecting its height and width. The tie-points within each quadrant are counted and expressed as a fraction of the total number of tie-points. For each cardinal direction (N, S, E, W), the quadrant percentages for that direction (e. g. the percentages of quadrant NE and NW are averaged to a percentage for N) are calculated and then multiplied by a manually predefined step size of 2500 pixels. This yields a measure indicating the image's shift in a given direction. These values are consolidated (with east and south being negative) to determine steps in x and y direction. These adjustments redefine the footprint bounds and a new tie-point matching is executed. If the number of tie-points decreases, the refinement is stopped instantly.

Using the refined footprints, a tie-point matching between the historical image and the Sentinel-2 image is executed. With a sufficient number of tie-points between the satellite and historical images, accurate geo-referencing becomes possible. With the precise bounds of the satellite image known, the relative tie-point coordinates in the satellite image are converted to absolute coordinates. These serve as the foundation for Ground Control Points (GCPs) that are used for the geo-referencing of images. The Python library Rasterio is used next to derive a transformation matrix from these tie-points. This matrix is subsequently be applied to geo-reference the TMA image. The geo-referenced TMA images are then transformed into polygons, enabling the computation of the photo-center of the TMAImage.

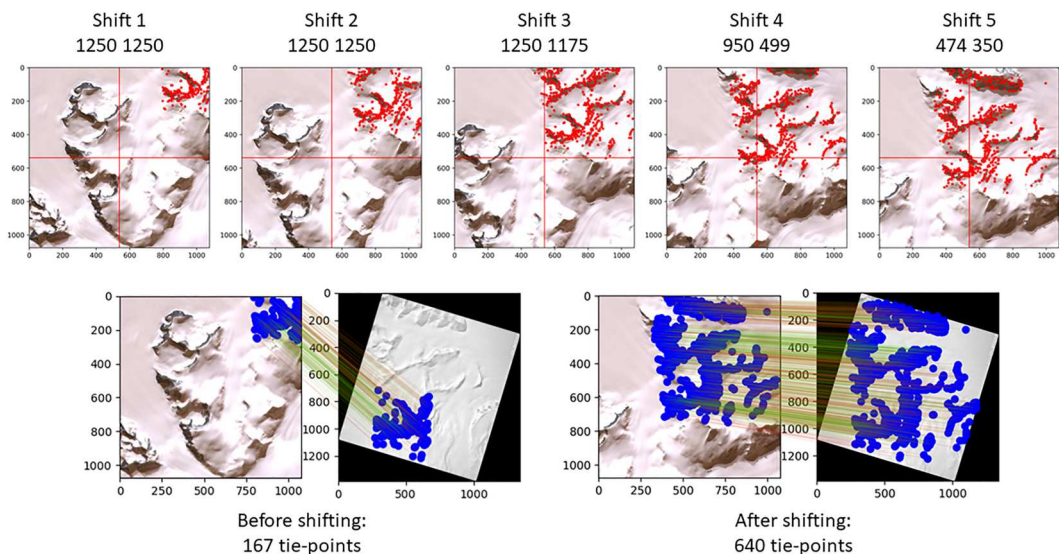


Figure 10. Fine shifting of TMA image (ID CA181232V0070) to improve the number of tie-points by 473. The images in the top row show the process with the steps in x- and y-direction. The bottom images show the tie-points before and after fine shifting.

3.4. Indirect geo-referencing

While geo-referencing with satellite images proves effective for a significant number of images, challenges are encountered in specific scenarios: some images do not yield meaningful tie-points, often attributed to factors like cloud cover or an absence of texture. In other situations, the detection of too many incorrect tie-points undermines the efficiency of the geo-referencing process. Additionally, there are instances where tie-points are primarily concentrated in a corner of an image. Even though geo-referencing might technically succeed under such conditions, the resultant footprints are unfortunately rendered invalid. For such images, we resort to one of three alternative indirect matching methods, proceeding from the first to the last as described below.

3.4.1. First option: flight line based footprint estimation

When for a TMA image direct geo-referencing is not possible, according to above criteria, a new approximate TMA footprint is estimated from already geo-referenced images from the same flight line. Subsequently, these problematic TMA images undergo a second round of geo-referencing with satellite images.

Figure 11 illustrates a segment of a flight path, highlighting the ground-projected camera positions of images from this trajectory and their corresponding footprints. The green footprints represent TMA images that were successfully geo-referenced using the direct approach. The coordinates of these TMA images are used to fit the flight path by regression. The singular red footprint highlights a TMA image that, while geo-referenced, is categorized as an outlier due to its position relative to the flight path. For the camera positions associated with the grey footprints, geo-referencing has yet to be successful. The grey footprints are estimated from the established flight path and the intervals between already geo-referenced images. They will subsequently serve as the basis for a renewed attempt at geo-referencing with Sentinel-2 images, since in this scenario the position of the footprint and the actual position are more similar, rendering the geo-referencing usually more successful.

3.4.2. Second option: matching with neighboring TMA images

Geo-referencing of the TMA images with Sentinel-2 data can sometimes be challenging due to low information content, leading to an insufficient number or unfavourable distribution of tie-points across an image. However, the aerial TMA imagery from a single flight path has a substantial overlap of up to 60%. Since these overlapping images are captured within mere seconds of each other,

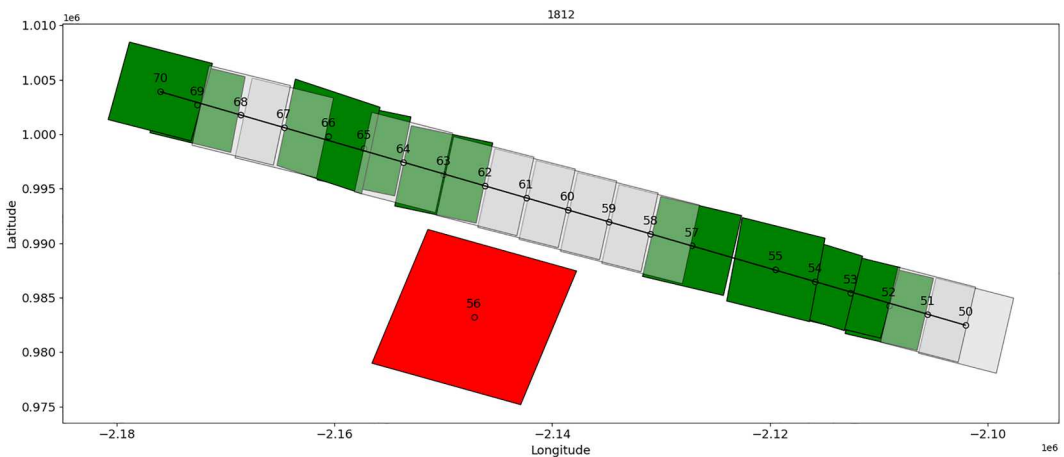


Figure 11. Estimation of footprints (grey) using successfully geo-referenced images (green) and ignoring the outlier image with id 56 (red). Ground-projected camera positions are indicated by numbered circles.

they essentially depict the same scene from slightly different angles. If one image from a sequence is already successfully geo-referenced, it is plausible to geo-reference it by tie-point matching with its immediate neighbors.

Note that this method is most accurate when limited to TMA images directly geo-referenced with Sentinel-2 data. Extending this approach beyond too many iterations may lead to the accumulation of small geo-referencing errors. As a consequence, TMA images successively geo-referenced using other TMA images can become increasingly distorted with each iteration.

3.4.3. Third option: extrapolating the footprint

Analogous to the footprint estimation in the first option, another polygon estimation is performed here. Between two geo-referenced polygons (from already geo-referenced TMA images from the same flight-path), an averaged polygon, is estimated from the weighted (higher weight for closer polygons) size and shape of adjacent polygons, as the new footprint for the un-referenced image. The four vertices of the estimated polygon serve as Ground Control Points (GCPs) for geo-referencing the image. This modest number of GCPs suffices, given their precise positioning at the four corners.

This method serves as the final solution when all other methods have proven ineffective and is also the least accurate among them. It is only operational when the images follow a discernible order, adhere to a logical flight path, and possess similar footprints. In the case of the TMA archive, this method is viable as the images are sequentially numbered and aligned within a flight path. For images aligned in straight lines, a minimum of two images along a flight path must have been successfully geo-referenced; however, the accuracy of geo-referencing improves with an increase in the number of images available.

3.5. Quality check and metrics

After applying the transformation matrix to a TMA image, images where geo-referencing went wrong must be removed.

Image size. A first quality check involves assessing the image size: images that are either too large or too small (exceeding a difference of 15% compared to other images in the flight path) are excluded, as illustrated by 'A' in Figure 12. Rather than solely focusing on the overall size of the image, it is also checked whether the pixel size of the image is consistent by deviating no more than 10% from other images in the flight path.

Footprint shape. A second quality criterion is related to the shape of the footprints. Given that the TMA cameras captured images in a square format, the transformed images should ideally be squares as well. However, there are instances where images appear as rectangles, with one side longer than the other. In certain cases, the image corners deviate from the expected 90-degree angles, resembling a parallelogram instead. Such images, depicted by 'B' in Figure 12 (with the red image representing both a parallelogram and a rectangle), are also excluded if their corner angles fall outside a minimal threshold margin of 10 degrees of a right angle.




Figure 12. Different scenarios for invalid (red) and valid (green) images. Red footprints are rejected because of (A) footprint size, (B) footprint shape, or, (C), footprint position.

Location w.r.t. flight path. In an approach similar to the first step of indirect geo-referencing, when analyzing multiple images from a flight path, it is possible to identify TMA images positioned inaccurately, as shown by ‘C’ in Figure 12. The calculated photo-centers of these TMA images are used to fit a regression line among all points, enabling the identification of outliers (points further away then two standard deviations). This approach assumes that the flight path locally follows a straight trajectory.

Error vector consistency. Recognizing outliers by position is further enhanced by using the original (inaccurate) approximate position of the images. We are calculating a vector between the camera center of the real position and the approximate position. If this so-called error vector deviates from the other error vectors in the same flight path, we also disregard this image as an outlier.

4. Results

In total, 4,459 images were used in the workflow. For the geo-referencing of the images, a desktop computer with a Quadro RTX 5000 (16GB video memory) was used. Approximately 42 h of calculation time was required, with an average referencing time per image of 30 s. First, overall results and results summarized per flight path are presented in Section 4.1. In Section 4.2, we present several case studies illustrating the application of geo-referenced historical images. For several locations on the Antarctic Peninsula, changes observed between historical photographs and modern Sentinel-2 imagery are discussed.

4.1. Overall results

Out of all images, 3,393 (76%) of the images could be geo-referenced with an average length of the error vectors being 28km. Here, an error vector connects the initial, provided location of a TMA image with the new position of the TMA image, obtained as a result of the geo-referencing procedure. Table 2 shows in detail how many images could be geo-referenced by which method. In total, there are 77 flight paths with more than 10 images located in the research area. Table A1 in the appendix give an overview of the flight paths and some averaged attributes, like information content or the length of the error vector. For 46 flight paths, all images were geo-referenced (over 99% of the images) while 15 flight paths could only be partially geo-referenced. For 13 flights, no images could be successfully geo-referenced.

Figure 13 shows the geo-referenced positions of the footprints, color-coded by the methods they were retrieved by. Each footprint has an error vector, that connects the the center of the resulting geo-referenced footprint to the center of the initially provided footprint.

Upon closer examination of various flight paths, it is evident that many images are accurately positioned, as confirmed by the alignment with geo-referenced images from different flight paths. This is exemplified in the top two rows of Figure 14. However, a detailed visual inspection reveals discrepancies for parts of some flight paths. Although the algorithms successfully compute footprints without distortions or outliers for these images, these images are misplaced and only close to their real positions. Examples of this issue are depicted in the bottom row of Figure 14. Currently, this problem impacts parts of 6 out of 77 flight paths.

Table 2. Number of geo-referenced images by method and non geo-referenced images together with their averaged parameters.

Method	Count	Information content	Tie-points	Residuals	Calculation time (s)
Direct	412	0.32	254	117.26	147
Indirect: option 1	319	0.28	80	93.94	33
Indirect: option 2	181	0.26	1112	58.18	16
Indirect: option 3	2861	0.13	4	3.32	4
not geo-referenced	665	0.12	n.a	n.a	n.a




Figure 13. Real positions of footprints and error vectors on a Sentinel-2 composite of the Antarctic Peninsula. Footprints are colored according to the method used for their estimation: direct geo-referencing and indirect option 1 (blue), indirect option 2 (turquoise) and indirect option 3 (yellow).

Unfortunately, no direct numerical estimation of the geo-referencing quality can be given, as the images have never been geo-referenced before. To have at least some numerical information, we manually placed 20 ground control points distributed over the research area covering different flight paths and calculated the distance to their equivalent geo-referenced position. The results of this comparison can be found in [Table 3](#).

4.2. Case studies: surface changes revealed by comparing historical and modern imagery

In this section, we demonstrate some practical application of accurately geo-referenced historical imagery, by highlighting its value in environmental change analysis. Our analysis is by no means meant to be comprehensive, but showcases the potential of the accurately geo-referenced TMA imagery for research purposes.

By employing the Structural Similarity Index (SSIM) (Wang et al. 2004), we quantified the similarity between historical images and their modern counterparts from Sentinel-2, considering their geo-referenced extents. The SSIM is a metric used to measure the similarity between two images, assessing changes in structural information, luminance, and contrast rather than just pixel differences. Figure 15 provides an overview of the similarity scores associated with these geo-referenced images. A blue color indicates high similarity between images, whereas a red color indicates that the historical images and their modern counterparts are not similar.

In our analysis, we observed that for most images, particularly those away from coastal regions, there is a high similarity score. This indicates that these areas have not undergone substantial changes over time. Conversely, images near the coast exhibited greater changes, primarily attributed to factors such as changes in sea ice cover. While it is important to note that these variations could also be influenced by

Table 3. Distance in meter between the geo-referenced position and the real position of sample points for different TMA images, indicated by their TMA Id.

Nr	Id	Distance (m)	Method
1	CA164432V0071	106	Indirect: option 3
2	CA174632V0171	155	Direct
3	CA182532V0011	108	Indirect: option 2
4	CA183332V0068	33	Indirect: option 1
5	CA184332V0074	99	Direct
6	CA184832V0102	119	Direct
7	CA184832V0120	47	Direct
8	CA184832V0156	14	Direct
9	CA196632V0028	28	Direct
10	CA196632V0056	19	Indirect: option 1
11	CA212332V0016	50	Direct
12	CA212332V0040	83	Direct
13	CA212432V0106	135	Indirect: option 2
14	CA214232V0162	185	Direct
15	CA214232V0193	117	Direct
16	CA216432V0215	12	Indirect: option 1
17	CA216632V0291	161	Indirect: option 1
18	CA216632V0305	67	Direct
	avg	85.44	

other elements such as cloud cover or deviations in image quality, this initial assessment provides a useful baseline for understanding environmental changes in these regions.

With the help of these similarity scores, we identified some locations of interest, i.e. locations in Antarctica, exhibiting significant environmental change over a 60-year period. Figure 16 present




Figure 14. Geo-referenced historical images from flight paths 1801/1856 (a), 2142 (b), 2146 (c), 2141 (d), 1824 (e), 1812/1824 (f) overlaid over Sentinel-2 imagery. The dashed line shows the footprint of the images.




Figure 15. Similarity between historical images and co-located modern Sentinel-2 imagery. Red indicates low similarity, while blue indicates high similarity.

these comparative examples: on the left, the geo-referenced historical imagery, on the right modern co-located Sentinel-2 imagery with enhanced natural color visualization. The dates of the historical images are given in the figure, the dates of Sentinel-2 satellite images are from the corresponding months in 2023. As for subfigure (d) no month is known, the month of January was chosen. The positions of these examples are indicated in Figure 2.

In Figure 16, a series of subfigures (a to e) illustrate glacier retreats in various highlighted areas. Focusing on subfigure (a), the Shoemith Glacier on Horseshoe Island is featured, characterized by two distinct branches: one leads into Lystad Bay to the west and the other into Gaul Cove to the north-east. From a direct image comparison, we estimate a retreat of 700 m at the western terminus, contrasted with a 300 m retreat at the northeastern terminus. Rodrigo et al. (2021) previously reported a retreat of approximately 200 m in the western part from 1999 to 2012, yet no data were available for the northeastern part. Coincidentally, a British gravity survey was conducted in this region in 1969, aligning with the time of the photograph (Smith 1973). The McMorrin Glacier, showcased in subfigure (b), is located on the Fallières Coast, extending westwards from Mount Metcalfe to Marguerite Bay. Based on the partial visibility in the historical imagery, a retreat of 1500 m is estimated, leading to the isolation of a small northern sub-glacier. In subfigure (c), Fleming Glacier is depicted, located along the Fallières Coast and terminating at the Forster Ice Piedmont, to the east of the Wordie Ice Shelf. Zhao et al. (2017) observed a significant decrease in elevation at the glacier's front, noting a reduction of 60 m from 2002 to 2014. Similarly, Friedl et al. (2018) reported a retreat of the grounding line by 6–9 km from 1996 to 2011. These findings are consistent with the historical imagery covering the northern part of the glacier, which reveals a clear change in the glacier front's position. In the modern image, the glacier front is distinctly visible, whereas in the historical image, the glacier front cannot be discerned at the marked position, nor further seawards. subfigure (d) presents the Somigliana Glacier, located on the northern Arrowsmith Peninsula and extending towards Langmuir Cove. Analysis reveals a retreat of 2,000 m from 1956 to 2023. Similarly, subfigure (e) examines a glacier near Mt. Wilcox, displaying a 600 m retreat from 1966 to 2023. In Cook et al. (2016, 2014), a surface area

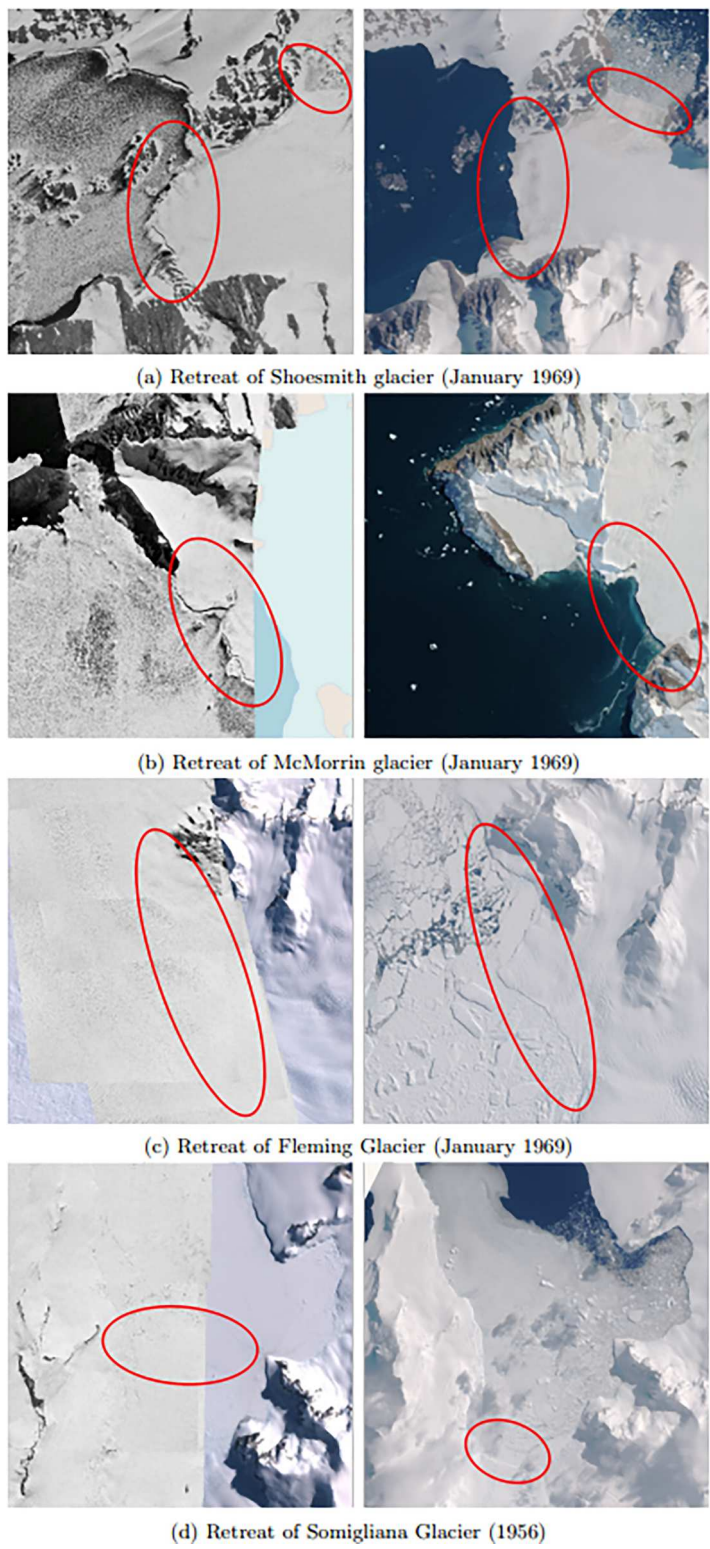


Figure 16. Examples for changes between historical images (left) and modern Sentinel-2 images (right). Notable changes are marked with red.

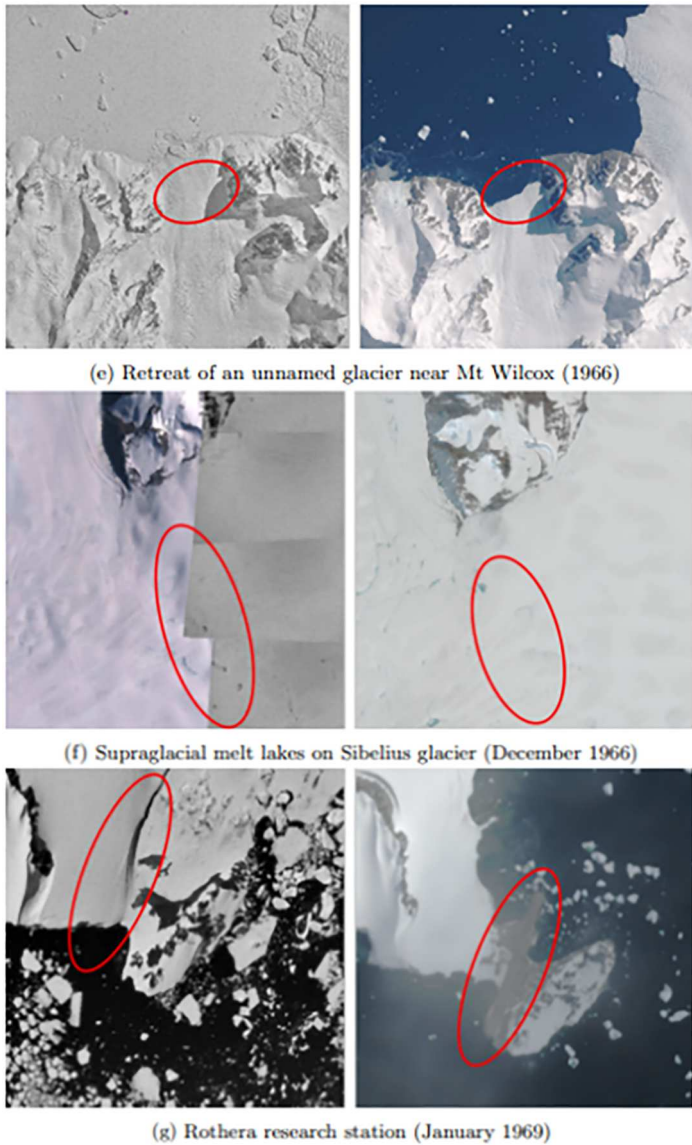


Figure 16. *Continued.*

reduction of 15% (1957 to 2009) and 5% (1940 to 2001) was reported for these glaciers, respectively. In subfigure (f), we observe supraglacial melt lakes on the Sibelius Glacier in central Alexander Island. Sentinel-2 imagery shows annual lake formations from 2014 to 2023, a phenomenon dating back to at least 1966 as per TMA images (Corr et al. 2022; Moussavi et al. 2020). Lastly, subfigure (g) captures Rothera Point on Adelaide Island, where the British Rothera Research Station is situated. These historical images predate the construction of the research base. In the satellite imagery, the station's landing strip is discernible (marked with a red square).

With the images now correctly oriented vertically, we can automatically obtain the accurate positioning of the oblique images, too, which facilitates additional comparisons. As a first example, Figure 17 presents a comparative oblique view of the Fleming Glacier (at position c) in Figure 2), juxtaposing an image from January 1969 with a Sentinel-2 image from January 2023. The latter is artificially rendered in 3D while approximating the same viewpoint.




Figure 17. Oblique view on Fleming Glacier from TMA imagery in January 1969 (left) and a Sentinel-2 3D rendering from January 2023 (right).

5. Discussion

A large fraction of the images from the TMA archive in our study area have been successfully geo-referenced. While in the past only an approximate position was available, which deviated by several kilometres and lacked exact image footprints, these images are now situated at their actual position, scaled to their true size, and aligned with their genuine orientation. As the cameras for the vertical and both oblique directions are located at the same position, the vertical geo-referencing can easily be extended to the oblique images.

Given the inherent limitations of the historical data and the lower geometric resolution of the Sentinel-2 satellite imagery, geo-referencing accuracy is often confined to several meters. Moreover, the quality of some image scans is compromised due to the poor image quality (degradation of the images & scanning errors), posing challenges in accurate geo-referencing. This level of precision, while not comparable to that of modern images, represents a substantial advancement compared to the approximate camera positions.

Even though only 20% of the historical images can be geo-referenced directly with satellite images, these images are the key to further successful geo-referencing. They play a pivotal role in the broader geo-referencing process and serve as references for geo-referencing adjacent images. Furthermore, they prove instrumental as reference points when computing the footprints of the remaining images. These direct geo-referenced key images are typically characterized by distinct rock structures, which have proven to be remarkably stable for many decades. For most of these key images successful geo-referencing is possible; however, there are three limitations: first, for some images, the changes between historical and recent images are too substantial to establish meaningful tie-points (top part of Figure 18). This issue is particularly prevalent in coastal areas, where over time the extent of glaciers and ice shelves has drastically changed. Furthermore, long-term and interannual variations in sea ice cover further frustrate the geo-referencing process when images contain a large fraction of the ocean. Second, tie-points must be dispersed across various parts of the image. When distinct structures are concentrated in just one corner of the image (middle part of Figure 18), geo-referencing tends to be unsuccessful, as in such cases, while the image might be correctly positioned, its footprint becomes distorted due to the imbalanced distribution of reference points. Lastly, even though the information content computed in Subsection 3.3 can serve as a valuable metric to identify these key images, it is not without limitations. In some cases, images may lack sufficient meaningful content (right part of Figure 18) or the tie-points are randomly distributed, often due to cloud coverage, thus hindering the identification of meaningful tie-points. On the opposite, we found that prevalent shadows in both the Sentinel-2 imagery

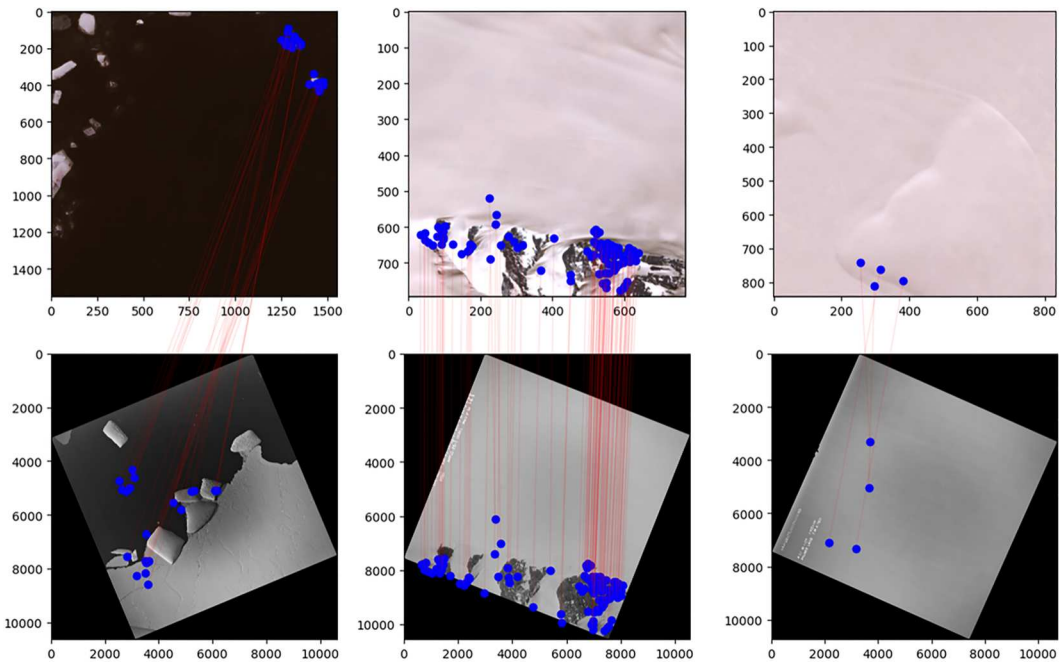


Figure 18. Examples for images where the geo-referencing failed due to changed landscape (left), unequal distribution of tie-points (middle) and too few tie-points (right). Modern Sentinel-2 imagery is on the top, historical images are on the bottom.

and the historical images are only marginally influencing the tie-point matching and geo-referencing.

The main reason that for some flight paths no images can be geo-referenced, is the low information content of all their images. The average information for these flight paths is significantly lower (average 0.07) compared to those where geo-referencing is successful (average: 0.18). For instance, the images of some flight paths predominantly feature snowfields, leading to monochromatic and uniformly white frames. These conditions render even manual inspection methods for geo-referencing ineffective, underscoring again the importance of diverse landscape features in images for accurate geo-referencing.

Conversely, a high information content can also hinder the geo-referencing. This often indicates dynamic landscapes, such as coastal regions, where significant changes in the landscape can occur. Images of this flight path feature coastal areas that have undergone substantial alterations, complicating the geo-referencing process. Additionally, images with extensive cloud coverage present a high information content as well. However, the clouds obscure essential landscape features, rendering the further geo-referencing task impossible.

Despite these challenges, when at least 2 images of a flight path can be geo-referenced successfully, the workflow can usually successfully extrapolate the position of the other images. During the geo-referencing, it was noticed that some images were scanned with an incorrect rotation, leading to accurate footprint extents but misaligned orientations. Furthermore, for some images, there were gross errors in the position of the provided approximate footprint: initially, we assumed that positional deviations in images would be limited to 30 km. However, our recent findings indicate discrepancies as large as 70 km. Increasing the search radius of the TMA footprint relocation can detect these deviations, but they impose a significant computational burden, increasing processing times to approximately 15 min per image (from the majority of images having a calculation time of 46 s).

As described in the results, for some flight paths, the images are technically successfully geo-referenced but misaligned from their true positions. This issue predominantly arises for images

geo-referenced with the third option (calculating the position of images), especially those situated further from the directly geo-referenced images within the same flight path. Minor discrepancies in angle and distance between the directly geo-referenced images accumulate, leading to the misplacement of the images with calculated positions. However, they are generally located in the vicinity to their real positions, allowing for the possibility of manual geo-referencing. Additionally, contrary to the initial hypotheses that all flight paths are straight lines, it has been observed that some flight paths deviate significantly from linear trajectories, incorporating sharp turns (see bold red flight path in Figure 2). This leads to greater deviations of the images from their actual positions, challenging the traditional assumption of uniformly straight and predictable flight paths. To address this, it is necessary to manually segment the flight path into different sections to guarantee accurate geo-referencing.

In conclusion, our observations underscore the multifaceted challenges in geo-referencing aerial images. From issues related to low information content and dynamic landscapes to errors in image rotation and underestimations of deviations, each factor significantly impacts the accuracy of geo-referenced data. As we continue to refine our methodologies and deepen our understanding of these complexities, we aim to enhance the precision and reliability of geo-referencing in aerial imaging.

6. Conclusions

In this study, we outlined the methodology for geo-referencing historical aerial photographs. We employed advanced neural network-based tie-point matching techniques, creating a unique workflow specifically optimized for high-resolution, large historical imagery. Geo-referencing is achieved by matching tie-points between the historical photos and contemporary satellite imagery, guided by preliminary footprints. If initial attempts are unsuccessful, we employ a tiered approach: first, re-attempting with refined footprints from satellite data; second, utilizing geo-referenced neighboring images; and finally, computing the image's footprint directly. Rigorous quality control measures are then implemented to guarantee the precision of the geo-referenced outputs. The efficacy of our method was showcased in a practical application, in which we successfully geo-referenced 76% of the images of an extensive collection of aerial photographs from the Antarctic Peninsula and highlighted differences between the historical and modern images for 7 scenarios.

While our geo-referencing technique proves successful for a majority of the images in the subset and these images are geo-referenced on a large scale for the first time, certain challenges persist. Many images, due to their sparse informational content, cannot be geo-referenced through tie-point matching alone, necessitating the deployment of our tiered strategy. Due to the geometric resolution of our modern data and the limited quality of the historical data, our method remains constrained to meter-level precision. For the majority of the flight paths in our research area, we were able to geo-reference at least some images from the flight path using satellite data, ensuring the following geo-referencing via extrapolations. However, should a flight path entirely lack satellite-based geo-referencing, estimating the location of its images becomes unattainable.

For the future, our primary focus is on increasing the number of images that can be directly geo-referenced with modern data. This is crucial as it significantly influences the quality of our method. Enhancements can be achieved either by refining the tie-point matching process (e.g. LoFTR Yuan et al. 2022 or DISK Tyszkiewicz, Fua, and Trulls 2020), retraining the neural network for tie-point matching with historical images, or by further optimizing the historical imagery to ease the identification of tie-points. In the current approach, an initial approximation of the image's location is required, which should be relatively close to its actual location. The step of refining the approximate position is notably time-consuming. We plan to employ machine learning, specifically U-nets, for this initial approximation refining, allowing for the faster identification of a historical image's location within a larger area. To enhance the functionality of the geo-referencing workflow, non-linear flight paths should be automatically recognized in the future and be taken into account during the line-fitting. Finally, it would be beneficial to evaluate this workflow's adaptability to

other regions and contexts, including historical aerial imagery from places like Greenland or Svalbard, or entirely different settings such as urbanized regions.

Geo-referencing the images from these historical archives offers many opportunities for subsequent research. Potential applications include monitoring changes in glaciers, ice shelves, and other cryospheric features over time, thereby furnishing tangible proof of Antarctica's response to global warming. Given that these images overlap, they can be used to construct 3D models using Structure from Motion (SfM). These models can then be compared with contemporary elevation data, like satellite stereo-photogrammetry and altimetry (e.g. ICESat-2), to procure an understanding of elevation and mass variations.

Acknowledgments

During the preparation of this work, the authors used generative AI to enhance the readability of specific sections. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This work was funded by NWO-grant ALWGO.2019.044.

Author contribution

Felix Dahle: Conceptualization, Methodology, Visualization, Writing – original draft, Data curation. **Roderik Lindenbergh:** Conceptualization, Methodology, Supervision, Writing – review & editing. **Bert Wouters:** Conceptualization, Methodology, Supervision, Writing – review & editing.

Data availability statement

The data for this submission is available at Dahle (2024b). An online visualization of the footprints of the historical data & more recent geo-referenced images can be found at <https://polar-archive.citg.tudelft.nl/>. The code for this submission is available on https://github.com/fdahle/hist_georef.

References

- Chen, H., and Y. Tseng. 2016. "Study of Automatic Image Rectification and Registration of Scanned Historical Aerial Photographs." *ISPRS – International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* XLI-B8: 1229–1236. <https://doi.org/10.5194/isprs-archives-XLI-B8-1229-2016>.
- Child, S. F., L. A. Stearns, L. Girod, and H. H. Brecher. 2020. "Structure-From-Motion Photogrammetry of Antarctic Historical Aerial Photographs in Conjunction with Ground Control Derived From Satellite Data." *Remote Sensing* 13:21. <https://doi.org/10.3390/rs13010021>.
- Cléri, I., M. Pierrot-Deseilligny, and B. Vallet. 2014. "Automatic Georeferencing of a Heritage of Old Analog Aerial Photographs." *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences* II–3: 33–40. <https://doi.org/10.5194/isprsannals-II-3-33-2014>.
- Cook, A. J., A. J. Fox, D. G. Vaughan, and J. G. Ferrigno. 2005. "Retreating Glacier Fronts on the Antarctic Peninsula Over the Past Half-Century." *Science (New York, N.Y.)* 308 (5721): 541–544. <https://doi.org/10.1126/science.1104235>.
- Cook, A. J., P. R. Holland, M. P. Meredith, T. Murray, A. Luckman, and D. G. Vaughan. 2016. "Ocean Forcing of Glacier Retreat in the Western Antarctic Peninsula." *Science (New York, N.Y.)* 353 (6296): 283–286. <https://doi.org/10.1126/science.aae0017>.
- Cook, A. J., and D. G. Vaughan. 2010. "Overview of Areal Changes of the Ice Shelves on the Antarctic Peninsula Over the Past 50 Years." *The Cryosphere* 4:77–98. <https://doi.org/10.5194/tc-4-77-2010>.

- Cook, A., D. Vaughan, A. Luckman, and T. Murray. 2014. "A New Antarctic Peninsula Glacier Basin Inventory and Observed Area Changes Since the 1940s." *Antarctic Science* 26 (6): 614–624. <https://doi.org/10.1017/S0954102014000200>.
- Corr, D., A. Leeson, M. McMillan, C. Zhang, and T. Barnes. 2022. "An Inventory of Supraglacial Lakes and Channels Across the West Antarctic Ice Sheet." *Earth System Science Data* 14:209–228. <https://doi.org/10.5194/essd-14-209-2022>.
- Cowley, D. C., and B. B. Stichelbaut. 2012. "Historic Aerial Photographic Archives for European Archaeology." *European Journal of Archaeology* 15 (2): 217–236. <https://doi.org/10.1179/1461957112Y.0000000010>.
- Dahle, F. 2024a. "Historical Geo-Referencing." https://github.com/fdahle/hist_georef.
- Dahle, F. 2024b. "Geo-Referenced Image Positions of The Historical Images of the TMA Archive." Version 1.4 4TU.ResearchData <https://doi.org/10.4121/cd746a00-15bb-479b-a669-37b5aeec44b5.v1>.
- Denzinger, F., H. Machguth, M. Barandun, E. Berthier, L. Girod, M. Kronenberg, R. Usabaliyev, and M. Hoelzle. 2021. "Geodetic Mass Balance of Abramov Glacier From 1975 to 2015." *Journal of Glaciology* 67 (262): 331–342. <https://doi.org/10.1017/jog.2020.108>.
- DeTone, D., T. Malisiewicz, and A. Rabinovich. 2018. "SuperPoint: Self-Supervised Interest Point Detection and Description." In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 337–33712. Salt Lake City, UT: IEEE. <https://doi.org/10.1109/CVPRW.2018.00060>.
- Du, Y., C. Li, R. Guo, X. Yin, W. Liu, J. Zhou, Y. Bai, et al. 2020. "PP-OCR: A Practical Ultra Lightweight OCR System." Pre-print arXiv. <https://doi.org/10.48550/ARXIV.2009.09941>.
- Engel, F. 2020. "Camera-Footprint-Calculator." Accessed September 5, 2023. <https://github.com/frank-engel-usgs/camera-footprint-calculator>.
- Fischler, M. A., and R. C. Bolles. 1981. "Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography." *Communications of the ACM* 24 (6): 381–395. <https://doi.org/10.1145/358669.358692>.
- Frankl, A., V. Seghers, C. Stal, P. De Maeyer, G. Petrie, and J. Nyssen. 2015. "Using Image-Based Modelling (SfM-MVS) to Produce a 1935 Ortho-Mosaic of the Ethiopian Highlands." *International Journal of Digital Earth* 8 (5): 421–430. <https://doi.org/10.1080/17538947.2014.942715>.
- Friedl, P., T. C. Seehaus, A. Wendt, M. H. Braun, and K. Höppner. 2018. "Recent Dynamic Changes on Fleming Glacier After the Disintegration of Wordie Ice Shelf, Antarctic Peninsula." *The Cryosphere* 12:1347–1365. <https://doi.org/10.5194/tc-12-1347-2018>.
- Geyman, E., W. Pelt, A. Maloof, H. Aas, and J. Kohler. 2022. "Historical Glacier Change on Svalbard Predicts Doubling of Mass Loss by 2100." *Nature* 601:374–379. <https://doi.org/10.1038/s41586-021-04314-4>.
- Ghuffar, S., O. King, G. Guillet, E. Rupnik, and T. Bolch. 2023. "Brief Communication: Glacier Mapping and Change Estimation Using Very High-Resolution Declassified Hexagon KH-9 Panoramic Stereo Imagery (1971–1984)." *The Cryosphere* 17:1299–1306. <https://doi.org/10.5194/tc-17-1299-2023>.
- Giordano, S., A. Le Bris, and C. Mallet. 2018. "Toward Automatic Georeferencing of Archival Aerial Photogrammetric Surveys." *ISPRS Annals of Photogrammetry, Remote Sensing and Spatial Information Sciences* IV-2: 105–112. <https://doi.org/10.5194/isprs-annals-IV-2-105-2018>.
- Heisig, H., and J.-L. Simmen. 2021. "Re-Engineering the Past: Countrywide Geo-Referencing of Archival Aerial Imagery." *PFG – Journal of Photogrammetry, Remote Sensing and Geoinformation Science* 89 (6): 487–503. <https://doi.org/10.1007/s41064-021-00162-z>.
- Howat, I. M., C. Porter, B. E. Smith, M. -J. Noh, and P. Morin. 2019. "The Reference Elevation Model of Antarctica." *The Cryosphere* 13:665–674. <https://doi.org/10.5194/tc-13-665-2019>.
- Karel, W., M. Doneus, G. Verhoeven, C. Briese, C. Ressler, and N. Pfeifer. 2013. "Oriental: Automatic Geo-Referencing and Ortho-Rectification of Archaeological Aerial Photographs." *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences* II-5/W1:175–180. <https://doi.org/10.5194/isprsannals-II-5-W1-175-2013>.
- Knuth, F., D. Shean, S. Bhushan, E. Schwat, O. Alexandrov, C. McNeil, A. Dehecq, C. Florentine, and S. O'Neel. 2023. "Historical Structure From Motion (HSfM): Automated Processing of Historical Aerial Photographs for Long-Term Topographic Change Analysis." *Remote Sensing of Environment* 285:113379. <https://doi.org/10.1016/j.rse.2022.113379>.
- Korsgaard, N. J., C. Nuth, S. A. Khan, K. K. Kjeldsen, A. A. Bjørk, A. Schomacker, and K. H. Kjær. 2016. "Digital Elevation Model and Orthophotographs of Greenland Based on Aerial Photographs From 1978–1987." *Scientific Data* 3:160032. <https://doi.org/10.1038/sdata.2016.32>.
- Kunz, M., M. A. King, J. P. Mills, P. E. Miller, A. J. Fox, D. G. Vaughan, and S. H. Marsh. 2012. "Multi-Decadal Glacier Surface Lowering in the Antarctic Peninsula: Peninsula Glacier Surface Lowering." *Geophysical Research Letters* 39 (19). L19502 <https://doi.org/10.1029/2012GL052823>.
- Lindenberger, P., P.-E. Sarlin, and M. Pollefeys. 2023. "LightGlue: Local Feature Matching at Light Speed." ICCV. <https://arxiv.org/abs/2306.13643v1>.
- Maiwald, F., D. Feurer, and A. Eltner. 2023. "Solving Photogrammetric Cold Cases Using AI-based Image Matching: New Potential for Monitoring the Past with Historical Aerial Images." *ISPRS Journal of Photogrammetry and Remote Sensing* 206:184–200. <https://doi.org/10.1016/j.isprsjprs.2023.11.008>.
- McNabb, R., L. Girod, C. Nuth, and A. Käb. 2020. "An Open-Source Toolset for Automated Processing of Historic Spy Photos: sPyMicMac." In *EGU General Assembly*. (Online, <https://doi.org/10.5194/egusphere-egu2020-11150>).

- Morgenstern, A., P. P. Overduin, F. Günther, S. Stettner, J. Ramage, L. Schirrmeister, M. N. Grigoriev, and G. Grosse. 2021. "Thermo-Erosional Valleys in Siberian Ice-Rich Permafrost." *Permafrost and Periglacial Processes* 32 (1): 59–75. <https://doi.org/10.1002/ppp.v32.1>.
- Moussavi, M., A. Pope, A. R. W. Halberstadt, L. D. Trusel, L. Cioffi, and W. Abdalati. 2020. "Antarctic Supraglacial Lake Detection Using Landsat 8 and Sentinel-2 Imagery: Towards Continental Generation of Lake Volumes." *Remote Sensing* 12:134. <https://doi.org/10.3390/rs12010134>.
- Nicu, I. C., and C. C. Stoleriu. 2019. "Land Use Changes and Dynamics Over the Last Century Around Churches of Moldavia, Bukovina, Northern Romania – Challenges and Future Perspectives." *Habitat International* 88:101979. <https://doi.org/10.1016/j.habitatint.2019.04.006>.
- Perreault, N., E. Lévesque, D. Fortier, D. Gratton, and L. J. Lamarque. 2017. "Remote Sensing Evaluation of High Arctic Wetland Depletion Following Permafrost Disturbance by Thermo-Erosion Gully Processes." *Arctic Science* 3 (2): 237–253. <https://doi.org/10.1139/as-2016-0047>.
- Polar Geospatial Center. 2023. "Aerial Photography." Accessed September 5, 2023. <https://www.pgc.umn.edu/data/aerial/>.
- Pope, A., W. Rees, A. Fox, and A. Fleming. 2014. "Open Access Data in Polar and Cryospheric Remote Sensing." *Remote Sensing* 6:6183–6220. <https://doi.org/10.3390/rs6076183>.
- Rodrigo, C., A. Varas-Gomez, A. Bustamante-Maino, and E. Mena-Hodges. 2021. "High-Concentration Sediment Plumes, Horseshoe Island, Western Antarctic Peninsula." *Antarctic Science* 33 (1): 1–4. <https://doi.org/10.1017/S0954102021000055>.
- Sarlin, P.-E., D. DeTone, T. Malisiewicz, and A. Rabinovich. 2020. "SuperGlue: Learning Feature Matching With Graph Neural Networks." In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4937–4946. Seattle, WA: IEEE. <https://doi.org/10.1109/CVPR42600.2020.00499>.
- Smith, I. F. 1973. "Gravity Survey on Shoesmith Glacier, Horseshoe Island, Graham Land." *British Antarctic Survey Bulletin* 33:77–82.
- Tyszkiewicz, M., P. Fua, and E. Trulls. 2020. "DISK: Learning Local Features with Policy Gradient." In *Advances in Neural Information Processing Systems*, edited by H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Vol. 33, 14254–14265. Curran Associates, Inc. NeurIPS 2020 (Online).
- Usmanov, B., I. C. Nicu, I. Gainullin, and P. Khomyakov. 2018. "Monitoring and Assessing the Destruction of Archaeological Sites From Kuibyshev Reservoir Coastline, Tatarstan Republic, Russian Federation. A Case Study." *Journal of Coastal Conservation* 22 (2): 417–429. <https://doi.org/10.1007/s11852-017-0590-9>.
- Wang, Z., A. Bovik, H. Sheikh, and E. Simoncelli. 2004. "Image Quality Assessment: From Error Visibility to Structural Similarity." *IEEE Transactions on Image Processing* 13 (4): 600–612. <https://doi.org/10.1109/TIP.2003.819861>.
- Wang, J., D. Li, W. Cao, X. Lou, A. Shi, and H. Zhang. 2022. "Remote Sensing Analysis of Erosion in Arctic Coastal Areas of Alaska and Eastern Siberia." *Remote Sensing* 14:589. <https://doi.org/10.3390/rs14030589>.
- Yu, Z., Z. Cao, C. Yu, G. Qiao, and R. Li. 2022. "Ice Flow Velocity Mapping in Greenland Using Historical Images From 1960s to 1980s: Scheme Design." *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* XLIII-B3-2022:799–804. <https://doi.org/10.5194/isprs-archives-XLIII-B3-2022-799-2022>.
- Yuan, X., X. Yuan, J. Chen, and X. Wang. 2022. "Large Aerial Image Tie Point Matching in Real and Difficult Survey Areas Via Deep Learning Method." *Remote Sensing* 14:3907. <https://doi.org/10.3390/rs14163907>.
- Zhang, L., E. Rupnik, and M. Pierrot-Deseilligny. 2021. "Feature Matching for Multi-Epoch Historical Aerial Images." *ISPRS Journal of Photogrammetry and Remote Sensing* 182:176–189. <https://doi.org/10.1016/j.isprsjprs.2021.10.008>.
- Zhao, C., M. King, C. Watson, V. Barletta, A. Bordoni, M. Dell, and P. Whitehouse. 2017. "Rapid Ice Unloading in the Fleming Glacier Region, Southern Antarctic Peninsula, and Its Effect on Bedrock Uplift Rates." *Earth and Planetary Science Letters* 473:164–176. <https://doi.org/10.1016/j.epsl.2017.06.002>.

Appendix

Table A1. Overview of all flight paths located inside the research area.

Flight path	Nr. of images	Nr. of georef. images	information content	length error vector (m)
1644	67	66	0.1	26,939
1741	18	0	0.07	n.a.
1742	36	0	0.04	n.a.
1743	28	28	0.02	51,102
1744	10	10	0.07	28,265
1745	16	15	0.03	19,384
1746	19	18	0.12	18,507
1747	15	14	0.08	10,705

(Continued)

Table A1. Continued.

Flight path	Nr. of images	Nr. of georef. images	information content	length error vector (m)
1800	76	76	0.12	20,709
1801	85	85	0.18	11,751
1803	14	0	0.52	n.a.
1804	20	2	0.18	92,024
1805	41	41	0.05	24,839
1807	41	41	0.07	27,766
1809	36	0	0.02	n.a.
1810	61	0	0.03	n.a.
1811	21	0	0.03	n.a.
1812	80	80	0.08	22,013
1813	123	123	0.15	18,837
1814	45	45	0.04	23,785
1815	76	76	0.16	26,304
1816	182	182	0.06	20,490
1817	35	35	0.11	14,912
1818	70	70	0.09	15,613
1821	188	188	0.07	7,203
1822	114	114	0.13	18,893
1823	103	103	0.19	12,645
1824	109	109	0.3	28,286
1825	115	114	0.09	59,068
1826	87	0	0.05	n.a.
1827	62	10	0.24	57,809
1829	45	2	0.06	40,189
1830	54	0	0.05	n.a.
1831	78	78	0.13	79,129
1832	81	81	0.08	57,454
1833	156	156	0.17	17,072
1834	50	50	0.46	32,106
1835	93	93	0.34	17,655
1843	89	89	0.12	13,001
1844	97	97	0.19	56,554
1845	86	86	0.2	91,998
1846	108	107	0.2	9,441
1847	95	95	0.19	13,301
1848	88	88	0.24	5,082
1849	101	101	0.16	160,741
1962	53	53	0.19	66,112
1963	11	0	0.08	n.a.
1965	20	20	0.21	44,150
1966	41	41	0.09	11,386
1967	46	1	0.08	13,889
1968	36	36	0.1	20,710
1969	34	0	0.07	n.a.
1993	26	1	0.14	20,039
2073	89	89	0.02	36,656
2074	73	73	0.02	41,598
2075	91	0	0.0	n.a.
2121	76	75	0.09	36,459
2123	54	54	0.07	5,665
2124	66	66	0.04	6,243
2133	15	1	0.38	55,343
2134	13	13	0.37	6,242
2135	25	0	0.35	n.a.
2141	53	53	0.12	15,077
2142	75	75	0.24	13,084
2160	28	28	0.5	16,189
2161	38	38	0.53	35,423
2163	36	0	0.94	n.a.
2164	48	47	0.52	18,065
2165	26	7	0.76	27,797
2166	67	67	0.6	20,396
2167	37	37	0.23	3,338