

Automatic analysis of UAS-based thermal images to detect leakages in district heating systems

Elena Vollmer , Rebekka Volk  and Frank Schultmann 

Institute for Industrial Production, Karlsruhe Institute for Technology, Karlsruhe, Germany

ABSTRACT

The mostly subterranean nature of district heating system pipelines makes pinpointing any occurring leakages a challenge. Airborne thermography offers a means for widespread monitoring, allowing thermal anomalies to be identified within multitudes of infrared images. This paper details a program developed to automate the entire image analysis process using mainly open-source software. Thermal images are acquired via unmanned aircraft system (UAS), pre-processed, and georeferenced individually or combined to orthomosaics. The search space is minimized to areas around the pipelines. Regions of interest are determined by image segmentation via tailored triangle histogram thresholding. The majority of resulting false alarms are removed by comparison with characteristic traits and results classified by their severity. The algorithm is applied to images newly acquired in Germany as part of a case study. The implemented methodology allows for a reduction of between 92 and 99% of thermal anomalies to a manageable amount of potential leakages for network operators to view. The use of orthomosaicking software in this context, though helpful in coalescing data, is found to lack robustness, precision and therefore reliability. Despite some limitations, the developed program is able to confidently detect and categorize leakages of varying severity and can be used directly by network operators. Future research will focus on further data pre-processing to eliminate thermal drift and remove the remaining false alarms, which mostly pertain to common urban features.

KEYWORDS

District heating systems; image processing; leakage detection; remote thermography; software development; unmanned aircraft systems

1. Introduction

District heating systems (DHSs) are networks of underground pipelines, which distribute thermal energy in the form of hot water or steam to end users (Krist et al. 2017, Werner 2017). They are capable of incorporating new energy sources and therefore play an important role in the integration of renewables in the heat sector (Krist et al. 2017, Werner 2017). Compared to local heat generators, DHSs are eco-friendlier, highly energy and resource efficient, provide means for easy pollution control, and are accompanied by a reduced administrative effort (Krist et al. 2017, Friman et al. 2014).

However, after decades of use, such subterranean pipelines fatigue. Leakages develop owing to corrosion, the deterioration of insulation material, or motion in the ground (Olsson 2001). If left unattended, these leakages can turn into public safety hazards, eroding the surrounding soil and causing the ground to collapse (Friman et al. 2014). The results include damages of considerable cost to infrastructure, the environment, involved personnel, and system efficiency (El-Zahab and Zayed 2019). Any loss of energy or fluid with such a potentially catastrophic impact needs to be mitigated.

Effective monitoring and maintenance are key to ensuring these systems stay efficient and safe. The subterranean nature of the networks, however, complicates the matter. In- and outlet flow sensors can generally identify system changes, but pinpointing exact leakage locations becomes a difficult and cost-intensive undertaking when the pipes need to be dug free for inspection (El-Zahab and Zayed 2019, Friman et al. 2014). Therefore, various alternatives for large-scale network monitoring have emerged, including the use of remote sensing and infrared thermography. This method is based on the premise that a pipe leakage causes an increase in temperature at the surface above, which may then be detected as an anomaly within thermal images (Axelsson 1988). However, manually screening the vast number of images resulting from such an aerial thermography process would be extremely laborious and time-consuming for network operators (Friman et al. 2014). As a result, some researchers have already presented methods for automated analyses of thermal images and leakage detection.

1.1. Related work

After Axelsson (1988) and Ljungberg and Rosengren (1988) suggest using airborne thermography for large-scale monitoring of district heating systems, Friman et al. (2014) are the first to describe a method to automatically analyse images and detect pipeline leakages. They acquire thermal infrared (TIR) images via manned flight and georectify them using recorded GPS and inertial measurement unit (IMU) measurements. The search area is minimized by masking the images with a network blueprint and additional buffer in geographic information system (GIS) format. The anomaly detection problem is solved by identifying those pixels which exceed a certain threshold of the pixel intensity probability distribution. False alarms caused by buildings are removed through segmentation via watershed transform. Berg, Ahlberg, and Felsberg (2016) expand on this methodology by extracting and evaluating distinguishing image features and using machine learning to classify these previously found detections as true or false. Building segmentation is improved upon by implementing a building mask created using OpenStreetMap (OSM).

Xu et al. (2016) implement saliency computation models based on colour, orientation, and intensity features on TIRs acquired during manned flights. Georeferencing the images allows the search space to be minimized to areas close to the DHS. This is achieved with a network blueprint in GIS format manually buffered using the commercial software ArcGIS. Saliency maps of the masked TIR images are created with a simplified Itti model and leakages detected by threshold segmentation. Zhong et al. (2019) expand on this method by fusing a local and global saliency map to include a comparison to local pixel neighbourhoods as well as overall feature rarity. The authors use maximum entropy segmentation to identify leakages in the final saliency map. Thermal images are acquired during both manned and unmanned aircraft

system (UAS) flights. When pipeline GIS data is unavailable, the authors assume the DHS to be located underneath roads and segment streets from RGB images as an alternative.

Sledz, Unger, and Heipke (2020) manually create orthomosaics¹ from UAS-based images using the software Agisoft Metashape and mask these with buffered DHS GIS data. Anomalies are detected as elliptical thermal hot spots by means of an Laplacian of Gaussian (LoG) blob detector. Blobs found by the detector are divided into regions and the temperature difference from warmest region to surrounding area is thresholded to identify which depict a leakage.

Hossain, Villebro, and Forchhammer (2020) demonstrate a machine learning approach on images acquired via UAS to find potential leakages. A region extraction algorithm is used for pre-processing, which detects warm regions and large gradients in an image, then combines both to an output image. A human operator identifies and labels the thus identified regions as leakages or not using ground truth data and expert knowledge. Some images are subsequently used to train a deep learning convolutional neural network and several different conventional machine learning classifiers.

1.2. Aim and scope

This paper showcases novel implementations of various methodology to the end of identifying leakages in district heating networks. The problem at hand is solved while simultaneously focusing on automation and generalization of the thermal image analysis procedure. The work thereby features a number of contributions to the fields of remote sensing, automation and image analysis.

While the above mentioned papers present different approaches, few elaborate on the extent of achieved automation or any actual software created for a widespread application. Oftentimes, third-party software is mentioned to have manually been used for steps like georeferencing or adding buffers. Some steps are not detailed at all and the means of creating and cost of running any mentioned or built software is unknown. Neither source code nor any thermal image data sets are provided, thus prohibiting tests, comparisons, or improvements. All these factors substantially impede the use or recreation of these methods and make the adoption of the proposed methodology near impossible.

The first contribution of this work therefore lies in the detailed description of a fully automated, implementable and replicable program for DHS leakage detection. In contrast to related work, the source code will be provided with the publication of this paper and is available on Zenodo (Vollmer, Volk, and Vogl 2023). The focus of the program design lay on using mostly open-source software to allow as large a user-ship as possible. This new form of implementation requires novel adaptations as well as the introduction of new methods.

Furthermore, this paper showcases various new methodological implementations. Following the example of only the most recent papers in this field, the thermal images are acquired via UAS. This is preferable to manned aircraft because flight planning is much more flexible, cheaper, and the process yields images with a higher ground sampling distance (GSD) (Sledz, Unger, and Heipke 2020). However, the close proximity to the ground complicates the georeferencing process because buildings have a more prominent distortive effect. The authors of the previously discussed work do not detail the

means by which they realized georeferencing or did not automate the process at all, and they do not address the problem of said distortions.

Another key contribution of this work therefore lies in the incorporation of three different approaches of handling the georeferencing problem automatically, varying in complexity and precision. The first is an individual image georeferencing method by means of affine transformation matrix using global navigation satellite system (GNSS) and IMU data and its adaption to remove general distortions by including building and topography data. Additionally, a novel indicator for false attitude recordings is identified and according automatic sorting out developed. This allows for new insights to be gathered concerning best practice rules for image acquisition via UAS. The other two implemented georeferencing methods demonstrate and compare the utilization of open- and closed-source photogrammetry software for automated orthomosaic generation. Additionally, because photogrammetry software itself can only handle input images of high quality, another contribution of this work is the development of an automatic thermal image quality indicator using the Canny edge detector. Finally, the novel use of these three different georeferencing approaches in parallel allows for a first direct comparison between these methods in their automatically implemented form.

Further contributions of this work revolve around the leakage identification itself. Previously proposed detection methods in the afore-described papers either failed to yield satisfactory results on the given real-world data² or proved to be inapplicable owing to a temperature drift in the given case study images. Therefore, the focus of this work lay in identifying a generally applicable method independent of data drifts and without requiring extensive preliminary effort. This new method for image segmentation implements the triangle histogram algorithm for thresholding as well as essential adaptations for its use in this new application. False positives are sorted out based on a novel combination of three characteristic leakage traits, as separately described in aforementioned work. Additionally, a simple means of result classification by severity via relative temperature difference is introduced. Lastly, the results are automatically prepared for user-friendly reviewing.

The automation of these aforementioned steps in their entirety as a focus of this work limits the required user interaction to the beginning (for data input) and the end (for result review). The program and method functionality are demonstrated and discussed using newly acquired images for a case study in Germany, which were used as a basis for design and testing. Contrary to predecessors and in the spirit of open data exchange, a dataset of thermal images will be made publicly available and can be viewed on Zenodo (Vollmer, Volk, and Vogl 2023).

1.3. Outline

The rest of this paper is organized as follows. Section 2 discusses the required data for the program to function. This includes details on the thermal image acquisition process, required prevalent conditions, and necessary auxiliary data. Section 3 describes the five key methodological steps implemented for leakage detection within the framework of the developed program. Section 4 details its application on images of a new case study, while Section 5 discusses the results and limitations. Section 6 concludes the paper with a summary and outlook.

2. Thermal image acquisition and auxiliary program data

The required data for program usage are thermal images acquired via UAS, GIS data of the DHS, and surface information, which include a level of detail 1 (LOD1) model of buildings in the area and a digital terrain model (DTM).

2.1. Thermal images and their acquisition

Image acquisition via UAS is preferable as it enables a flexibility, cost and ground sampling distance unobtainable by manned aircraft (Sledz, Unger, and Heipke 2020). Depending on the size of a DHS, it can take several days and the subdivision of the network into multiple regions to cover the entire area (Friman et al. 2014). Certain conditions must prevail during all flights to ensure sufficient image quality. These can be summed up as follows:

- Image acquisition should take place at night, preferably a few hours before dawn so as to ensure most irrelevant heat sources caused by sun and people have been eliminated (Friman et al. 2014, Heipke and T'odter 2020).
- Outdoor air temperatures should not exceed 10°C. In the northern hemisphere, the acceptable time frame is given as October to April (Heipke and T'odter 2020).
- The ground should be free of foliage or snow to give the thermal camera an unhindered view of the ground above the DHS. Not only may leakages otherwise be overlooked, but also snow reflectance can falsify the thermal imagers' temperature recording (Friman et al. 2014, Heipke and T'odter 2020).
- It should not be snowing, raining, or foggy during image acquisition and wind speeds should be comparatively low.³ Water particles absorb some of the IR radiation before it can reach the thermal camera, resulting in grainy images lacking contrast and detail. Long-lasting rain and high winds cause the ground surface to cool down and puddles to form, which falsify the recorded temperature data (Friman et al. 2014, Heipke and T'odter 2020).

Several characteristics of the DHS in question need to be considered as well. The pipes cannot be placed too deep below the surface for thermography to work reliably. Heipke and T'odter (2020) set the limit at 1 m and stress that the depth, just like the temperature of the transported medium, should stay constant throughout the system. Additionally, the method may not function reliably on unconventional types of pipe casings, meaning these must be tested in advance.⁴ Lastly, local laws pertaining to UAS usage on private or company property might call for special permits or prevent data acquisition in specific areas (Heipke and T'odter 2020).

The recorded images must contain extractable thermal data and specific metadata such as a GNSS coordinate, flight height, camera parameters, and attitude information in the form of yaw, pitch, and roll angles of flight and gimbal. A fitting file format is therefore Radiometric Joint Photographic Experts Group (RJPEG) as these images contain thermal matrices with a temperature value for every recorded pixel.

2.2. Auxiliary data

Friman et al. (2014) already recognized the advantage of reducing the search space within the images to areas close to the DHS. It allows distant objects, which otherwise would cause false alarms,⁵ to be removed from consideration before leakage detection has even begun. This same step is therefore also implemented within the developed program and requires a blueprint of the pipeline network in GIS format, specifically a shapefile. These are obtainable from the local authorities or network operators (Friman et al. 2014).

This step can only be performed if the images themselves are georeferenced in advance. Georeferencing refers to the conversion of the image coordinate system to a chosen geographic coordinate system, thus correctly placing an image on a map (Jayapalan et al. 2021). Images acquired during UAS flights suffer from high amounts of relief displacement (Whitehead and Hugenholtz 2014), which is the distortion and positional error induced by objects of different heights in the camera's field of view (Government of Canada 2015). To mitigate this somewhat, building and topography data are taken into account. A LOD1 model of surrounding buildings should be provided in shapefile format, topographic information in the form of a DTM of points and corresponding ground elevation values in text file format.

3. Methodology

Using a similar methodology to Friman et al. (2014)'s, the leakage detection process is divided into five main steps within the program.

- (1) preliminary image processing
- (2) image georeferencing and/or orthomosaic generation
- (3) search space minimization to areas above and around the DHS
- (4) thermal anomaly detection via image segmentation
- (5) leakage identification by reduction of false alarms

The procedure is visualized as a flowchart in Figure 1. The acquired thermal images are analysed in sets, which are groups of images collected in one flight. Pre-processing occurs both in preparation for the subsequent step as well as for later leakage identification. The images are placed on a map or alternatively combined to generate orthomosaics using third-party software. By overlaying these with a DHS pipeline blueprint the search space can be minimized to areas above and around the network, thus excluding those of no importance and ensuring considerably fewer false alarms ensue. Areas of interest within the resulting masked images are defined as those pixels that exceed a specific threshold. These thermal anomalies are evaluated by size, shape, and temperature to identify those which are true leakages and classify the results into three categories of severity. The leakage positions are additionally evaluated.

The program is designed to run on a ground workstation with at minimum 8 GB of RAM, an Intel i5 processor, and either a Windows 10 Pro or Linux Ubuntu 22.04 operating system. For the purpose of automation, the open-source programming language Python is utilized in the versions 3.6.7 and 3.9.1. Python has several advantages for its use in this application, such as its widespread availability, large support network, and variety of third-party modules in addition to extensive support libraries. Other programs are also easily accessible via Python, such as the

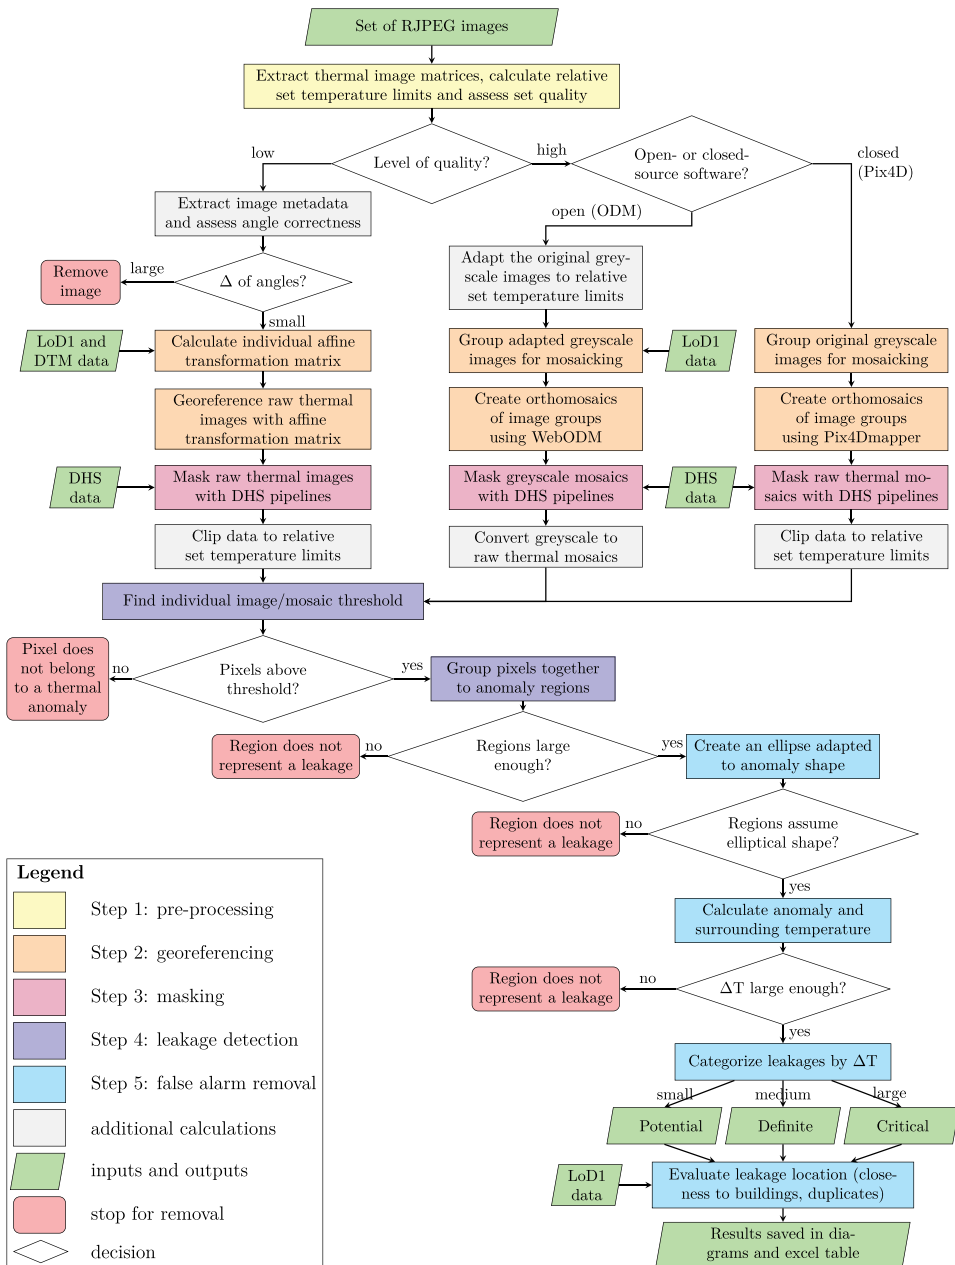


Figure 1. Leakage detection procedure implemented in the developed Python program.

aerial image mapping tools OpenDroneMap (ODM) (OpenDroneMap 2020) and Pix4D (Pix4D 2020) that are employed in this case (Section 3.2). The code can be viewed in full on Zenodo (Vollmer, Volk, and Vogl 2023).

It is run via command line and utilizes user prompts to attain all relevant inputs, analysing as many sets of images as desired. The final results are saved into excel files and different diagrams depict the stages of the process. Users therefore need only interact

with the program at the beginning of a run-through by providing all required data. The results can then be viewed at the end to determine which identified leakages merit closer observation or even direct intervention.

3.1. Image pre-processing

Preliminary image processing encompasses the calculation of the raw thermal data, the definition of a set's temperature limits in preparation for subsequent steps, as well as a quality assessment of the entire set. The temperature data is imperative for later leakage identification and can only be obtained from the original images.

For later data conversions, a credible temperature range for each set of images is determined. The maximum and minimum of all extracted values cannot represent this range because thermal cameras can incorrectly record a temperature given adverse surrounding conditions or specific material types.⁶ Relative minimum and maximum temperature limits of a set of images are instead defined using the arithmetic mean μ and standard deviation σ of all temperature values:

$$T_{\text{set}_{\text{min}_{\text{rel}}}} = \mu_{\text{set}} - 3\sigma_{\text{set}} \quad (1)$$

$$T_{\text{set}_{\text{max}_{\text{rel}}}} = \mu_{\text{set}} + 9\sigma_{\text{set}} \quad (2)$$

In a normal distribution, the area between $\mu - 3\sigma$ and $\mu + 3\sigma$ encompasses 99.7% of all data (Merziger et al. 2013). Although a histogram of typical set values does not directly equate such a distribution, they are akin in shape. Therefore, these values are chosen as initial limits. The upper limit is extended and defined by triple the distance to the set's mean temperature value because the usage of this data for leakage detection purposes inherently requires a focus shift to higher temperature ranges. Leakages may display a more than 20°C higher temperature in comparison to their surroundings (Friman et al. 2014) and thus extend well beyond such an average set temperature.

Due to the vast amount of data in a set, the values of all temperature matrices cannot be handled simultaneously. Instead, the mean and standard deviation are computed for each thermal image individually and then combined. The combined arithmetic mean of all temperature matrices is simply the average of all single image means. To calculate the combined standard deviation, Burton (2013)'s code using analysis of variance is applied. Owing to an extensive image overlap, not all image pixels are considered equally. However, the mathematical error is minor and the assumption of using all temperature data still valid for determining a general range of relevant values.

Lastly, the general image quality of a set is evaluated because it defines the way in which georeferencing can be implemented. An analysis of both photogrammetry software shows that low quality images do not yield usable results (Figure 2). Low quality hereby refers to having few distinguishable features, in other words being particularly grainy or depicting large amounts of vegetation. Resulting mosaics are blurry and contain large gaps of data missing (Figure 2(d)).

The quality of a set is determined using the Canny edge detector, an algorithm that can optimally balance precise edge localization and noise influence (Distante and Distante 2020, University Consortium for Geographic Information Science 2020). It returns a binary image

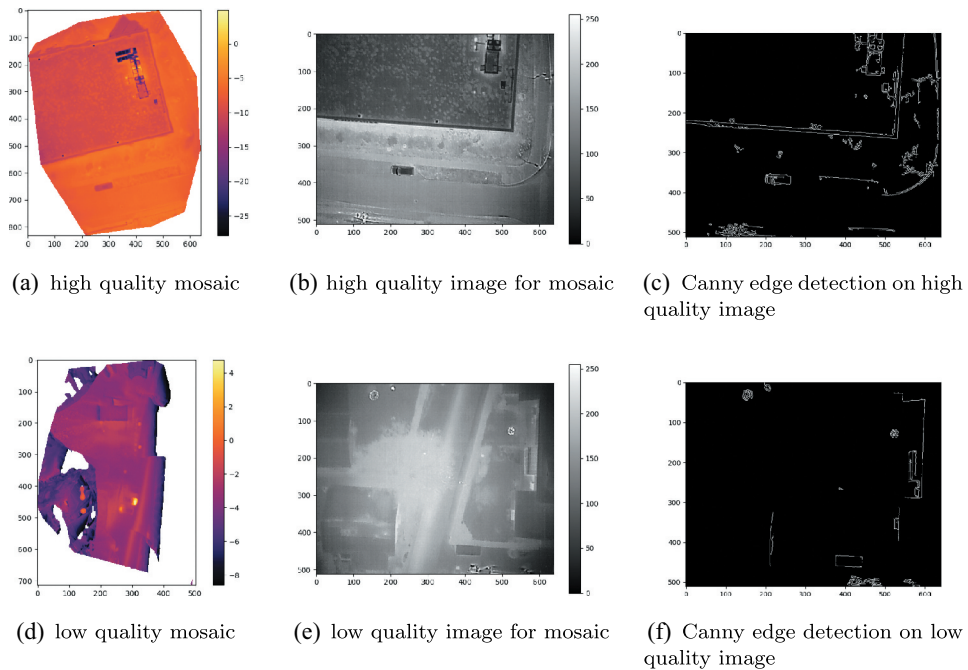


Figure 2. Comparison of mosaics generated from high and low quality images. Quality is defined by the existence of distinguishable features or lack thereof and visualised using the Canny edge detector algorithm applied to the images.

displaying the dominant edges as white pixels, the pixel count of which ultimately defines set quality. The algorithm is applied to every image (Figure 2(c,f)) to determine a set's mean edge pixel count. Each set is then classified as being of either high or low quality via thresholding. The optimal threshold for this case study was found to be 4000 pixels.

3.2. Image georeferencing and orthomosaic creation

In order to minimize the search space to areas of relevance, the images are georeferenced first by positioning them on a map and thus linking internal coordinate system to geographic counterpart (Leidner 2016). The step is crucial to overall accuracy, as only precisely positioned images allow for exact masking results. The difficulty, as mentioned in Section 2.2, lies in mitigating the relief displacement and geometric distortions inherent in remote sensing applications (Government of Canada 2015). Three methods are implemented within the program: (I) georeferencing of individual images, (II) using aerial image mapping software WebODM (OpenDroneMap 2020) and (III) Pix4Dmapper (Pix4D 2020). Both applications create mosaics of several images at a time while simultaneously removing the inherent distortions in a process called geo- or orthorectification (ArcGIS Pro 2021a, University Consortium for Geographic Information Science 2020). Since the mapping software requires high image quality for successful processing, the first method comes into play when the images have a low quality.

The low quality images are georeferenced individually by calculating an affine transformation matrix for each (Garrard 2016). Applying this matrix will shift, scale, and rotate the image

raster data to its correct position within a chosen coordinate reference system (CRS) (ArcGIS Pro 2021b). It can be calculated with a set of geographical coordinates, the GSD, and the drone heading angle (Garrard 2016). For this, the following parameters are extracted from the image metadata: focal length f , sensor width s_w , image width i_w and height i_h , flight height h , longitude lon and latitude lat of the image centre point, and yaw or so-called heading angles κ of both unmanned aircraft (UA) and camera. The intrinsic camera parameters shown in Figure 3 (left side) are used to calculate the size of the image's footprint g_w and g_h on the earth's surface using the following equations (Pix4D 2019):

$$GSD = \frac{s_w \times h \times 100}{f \times i_w} \quad \left[\frac{mm}{pxl} \right] \quad (3)$$

$$g_w = \frac{GSD \times i_w}{100} \quad [m] \quad \text{and} \quad g_h = \frac{GSD \times i_h}{100} \quad [m] \quad (4)$$

The coordinates of each image corner are calculated by combining an image's footprint with its centre point coordinate's easting and northing values⁷ and heading angle κ ⁸ as depicted in Figure 3 (right side). The following trigonometric equations (Merziger et al. 2013) exemplify the procedure for the upper left corner:

$$d = \sqrt{\left(\frac{g_w}{2}\right)^2 + \left(\frac{g_h}{2}\right)^2} \quad \text{and} \quad a = \arctan\left(\frac{g_h}{g_w}\right) \quad (5)$$

$$x'_{UL} = Easting_{UL} = Easting_{center} - d \times \cos(a - \kappa) \quad (6)$$

$$y'_{UL} = Northing_{UL} = Northing_{center} + d \times \sin(a - \kappa) \quad (7)$$

By calculating the coordinates of the other three image corners in analogous fashion, Python's Geospatial Data Abstraction Library (GDAL) (GDAL/OGR contributors 2022) is capable of constructing the affine transformation matrix.

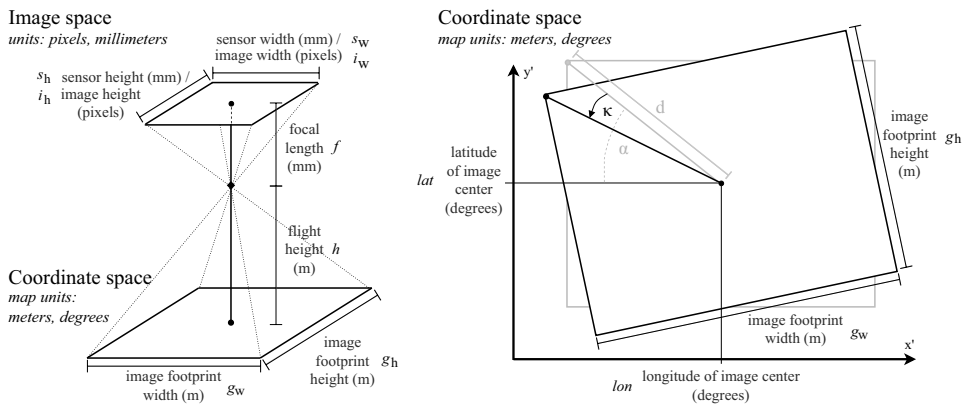


Figure 3. Visualisation of the image parameters required for georeferencing. Left side: parameters relevant for calculating the GSD and image footprint. Right side: combination with the UA heading angle to calculate ground coordinates of the image corners (Pix4D 2019).

Table 1. Overview and comparison of photogrammetry software WebODM and Pix4Dmapper (information from OpenDroneMap [2020], Pix4D [2020], Pix4D [2021], and , Groos et al. [2019]).

	WebODM	Pix4Dmapper
Relevant inputs	JPEG images including all original metadata	JPEG, RJPEG images including all original metadata
Useful outputs	orthomosaic	orthomosaic, thermal index map
Installation effort	high - prerequisites: Docker, Git, if Windows: 10 Pro OS	medium - prerequisites: Pix4Dengine, Python version 3.6.7
Python access	directly	indirectly - via Pix4Dengine
Cost	none	high - licenses for Pix4Dmapper (1500 for research purposes) and -engine

Some adjustments are necessary for the application of this method on images recorded via UAS. Relief displacement caused by the small flight height is prominent in images recorded above buildings and mitigated by subtracting the building height from the flight height when the centre coordinate lies above a building. A LOD1 model of the buildings is used for this purpose. Additionally, ground elevation information from a given DTM is incorporated, as terrain following was not part of the flight program.

A fail-safe is included to account for incorrect flight attitude, as the recorded angles are not always accurate (Prof. Grant Petty, pers. comm., April 14, 2021). As both the UA and camera save a yaw angle per image, inaccurate measurements can be identified by a large delta between the two.⁹ In these cases, the images are removed from consideration because the calculated corner points lead to incorrectly placed images.

Alternatively, orthomosaics are generated from high quality thermal images by either or both WebODM (OpenDroneMap 2020) and Pix4Dmapper (Pix4D 2020). These are chosen owing to their capacity for automation and advantages (Table 1).

The images are divided into groups (by size and proximity to surrounding buildings) prior to processing. For WebODM, an error handling mechanism is included to increase the general groups of between 15 and 30 images to a greater size if no usable mosaic is generated from the original group. The groups for Pix4Dmapper are generally larger at around 30 images as those result in the highest quality mosaics. It should be noted that image groups cannot be extended endlessly as the resulting mosaics will display a too low resolution and quality for effective leakage detection (Friman et al. 2014), which is why the fail-safes are limited in their extension capabilities.

WebODM requires greyscale equivalents of the raw thermal images as inputs, which are calculated using the relative set temperature limit defined during pre-processing (Section 3.1). This is necessary because, unlike Pix4Dmapper which can generate thermal index maps, WebODM is incapable of creating orthomosaics from the thermal RJPEG themselves.

3.3. Minimizing the search space

Georeferenced images and orthomosaics are overlaid with the network blueprint in GIS format to remove pixels that are too far away from the DHS. Overlaying image raster data with network pipeline vector data is achieved with Python's 'Rasterio' (Gillies et al. 2013), a library also used to extend the DHS shapefile by a 3.5 m buffer on all sides. Figure 4 exemplifies the described process using the georeferencing results of the previous step. The WebODM mosaic data are finally converted back to thermal values using the same relative set temperature limits and all other image and mosaic data are clipped to the same range.

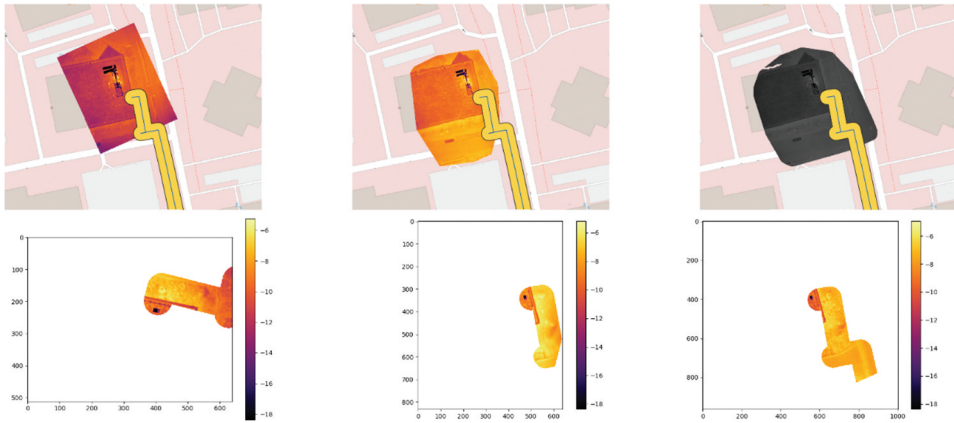


Figure 4. Comparison of masking and post-processing procedures of individual image georeferencing (first), Pix4Dmapper (second) and WebODM.

3.4. Detection of thermal anomalies

The masked images or mosaics are segmented to find pixels of interest and thus thermal anomalies. Similarly to Friman et al. (2014), pixels of interest are hereby assumed to reside in the upper tail of a histogram of all values. Instead of choosing the threshold as an upper percentile at random, the triangle threshold histogram algorithm is implemented to identify a fitting limit. Owing to a temperature calibration drift within the case study data, the method cannot be used to identify an overall set threshold and is instead repeated for each individual image. Figure 5 exemplifies this non-linear drift in a dataset of case study images, showing an increase of temperatures without cause.

The algorithm uses the hypotenuse that links the histogram peak P_1 to the right-side edge P_2 of the set temperature range (Figure 6). The longest orthogonal connection d_{max} between hypotenuse and histogram bin is determined iteratively and identifies the threshold as the corresponding bin's temperature value (Zack, Rogers, and Latt 1977).

Here, an adjustment is made to account for histograms consisting of multiple distinctive peaks when large uniform areas of different temperatures are depicted. This can happen when a building and street take up near equal areas within an image. The peak is then not simply defined as the bin of maximum height, but instead as the warmest amongst all distinctive local maxima. Only if no other peaks of considerable size are detected is the bin with the maximum percentage selected.

Experimental analyses showed that the threshold may at times be too conservative, choosing too many pixels as being of interest. This can happen in very varied images, which include a broad range of temperatures occurring multiple times. For this reason, the pixel percentage of the chosen threshold bin is taken into account. If the percentage exceeds 1%, the threshold is moved further to the right to a bin that does not exceed said experimentally determined limit. The Appendix A exemplifies these described amendments to the traditional triangle histogram thresholding method, thus demonstrating the reasons for how it is tailored to match the specific requirements of thermal leakage detection in the urban context.




Figure 5. Temperature drift in a case study set, with acquisition from top right to bottom left. Warm areas do not stem from underground heat (SWM).

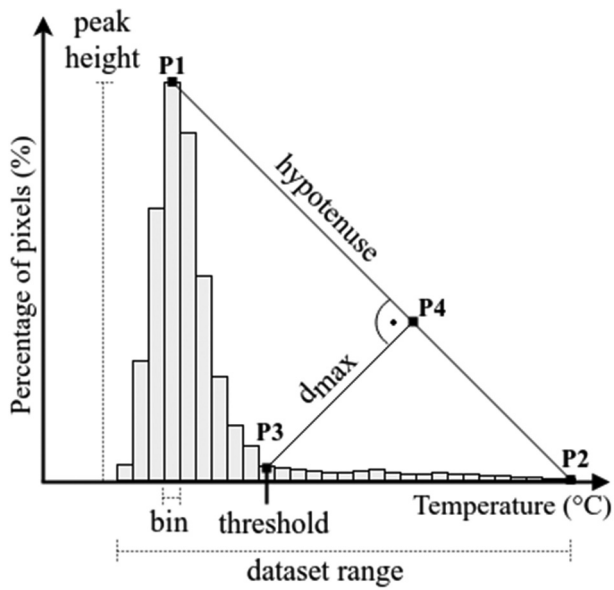


Figure 6. Visualisation of the triangle histogram thresholding method and involved parameters (Zack, Rogers, and Latt 1977).

The described method results in an array for each image consisting of pixels labelled according to their relation to the found threshold. The pixels of interest are grouped to regions of interest. This is accomplished by means of Python's 'scikit-image' (Van der Walt et al. 2014), a library that calculates various parameters for pixels with identical labels. First, the labelled arrays are assessed to determine whether more than one region of interest exists by thresholding the so-called extent (ratio of labelled pixels to all within their bounding box). If the pixels are all in the same area, they are said to comprise a single region, while sparse and widely spread out pixels are recognized as multiple clusters. The clusters are identified using the flood fill algorithm. Since the algorithm only connects direct neighbours, small regions are additionally joined with larger ones if they are in close proximity to one another.

3.5. False alarm removal

Even after search space minimization and segmentation, many false alarms can be found amongst the detected thermal anomalies. Therefore, only those anomalies that satisfy specific size, form, and temperature constraints are classified as leakages.

- (1) Size: A visual analysis of the thermal images and correspondence with local municipal utilities companies show that very small thermal anomalies¹⁰ cannot represent actual pipeline leakages and are therefore sorted out.
- (2) Form: The form constraint is defined by elliptical or circular shape.¹¹ The thermal anomaly should, to a certain degree, exhibit such a form to be considered a leakage. The anomaly's centre pixel coordinate, major and minor axis lengths, and rotational angle are calculated to define the ellipse's size, shape and orientation (Merziger et al. 2013). The ratio of anomaly pixels to the amount constituting the ellipse is thresholded with an experimentally found limit to determine which thermal hotspots can be interpreted as leakages.
- (3) Temperature: True leakages are identified in a comparison of anomaly temperature to its surroundings.¹² A ring is constructed at a distance around an anomaly to calculate the average temperature of the surrounding area. The anomaly temperature is defined by the mean of all pixels in small hotspots, the warmest 50 in medium-sized, and the top 100 pixels in large anomalies. If the temperature difference exceeds 5°C¹³ the anomaly is declared a leakage.

The identified leakages are categorized as either potential, definite, or critical depending on the severity of said temperature difference. Potential leakages have a temperature delta between 5°C and 10°C, definite ones display a difference of 10°C to 15°C, while critical leakages are defined as exceeding the 15°C mark.

To remove redundancy and give insight into further potential false alarms, a location analysis is preformed. The leakage positions are compared to surrounding buildings as well as each other. This helps determine whether a found region of interest is situated on a building roof (e.g. caused by faulty insulation or a warm chimney) or if any duplicates are present (e.g. the same thermal anomaly is displayed within multiple subsequent images). Additionally, anomalies in close proximity to a building may be caused by vents or

entrances (Friman et al. 2014). To identify leakages on top of or close to buildings, the coordinate of each leakage centre point is compared to the buildings LOD1 model.

In theory, the procedure works well. However, practical application has shown the calculated geotransform (Section 3.2) to lack precision, thus prohibiting the inclusion of this positional information in the leakage detection process itself. Appendix B exemplifies how in particular the calculated mosaic placement can deviate from its true positioning. Furthermore, true leakages may appear near buildings as pipeline networks are often situated in their close proximity (Friman et al. 2014). Therefore, this information is provided as supplementary material to the list of leakages to aid in facilitating the result evaluation.

4 . Results

4.1. Case study description

The case study examines parts of the DHS of suburban areas close to Munich in Germany. It encompasses several regions including the towns Taufkirchen, Ottobrun, and Neubiberg, which in total contain approximately 43 km of DHS pipelines. The water transported by these pipelines ranges from 80°C to 130°C. The installation depth is unknown. The level of urbanization varies throughout the given region, being mostly suburban but including many parks and forests. The area was split into 49 zones, with every flight requiring approximately 25 min. Image acquisition took place at night-time between 7 pm and 6 am in December of 2019. The ambient temperature ranged from 2° to -5°C with maximum wind speeds at 20 km/h. The process produced approximately 55,000 images.

While many of the resulting images have a high quality, there were some disadvantages to the chosen time frame, which lead to low quality, temperature drift or the depiction of features that falsify leakage detection. In some images, snow coverage is distinctly visible while ongoing snowfall during the acquisition of others produced extremely grainy pictures. The comparatively early starting time allowed for a large amount of human activity, which resulted in the presence of people as well as vehicles exuding heat. These scenarios are exemplified in Figure 7. Such problematic occurrences naturally reduce the efficiency of any leakage detection algorithm owing to either an increase of false alarms or loss of valuable information.




Figure 7. Select case study TIR images examples of high and low quality showcasing various aspects of image acquisition.

DJI's Matrice 600 Pro UA was utilized for nearly fully automated image acquisition. The routes for image acquisition were pre-programmed into the control software and only required a human operator for takeoff. The UA's flight height was consistent at 60 m above the starting point. An image overlap of approximately 88% was achieved. The camera technology was provided in form of the Zenmuse XT2, a combination of FLIR's Duo Pro R thermal and visual camera and DJI's gimbal. The thermal camera contains an uncooled VOx microbolometer and therefore works in the long-wavelength infrared (LWIR) range between 7.5 and 13.5 μm . It has a focal length of 13 mm, a sensor width of 10.88 mm, and a spatial resolution of 640 \times 512 pixels. (FLIR Systems, Inc. 2017, SZ DJI Technology Co. Ltd. 2018a, 2018b)

The thermal camera can output images in JPEG, TIFF or FLIR's own RJPEG file format (FLIR Systems, Inc. 2017). Because RJPEG contains extractable thermal data saved to each image pixel, they were chosen as the preferred output. Every image stored all relevant metadata including the camera parameters, GNSS coordinate, relative height, and positional angles. It should be noted that the recorded absolute altitude values were incorrect, as they placed the images below the ground. For this reason, the relative values were utilized in combination with a DTM. Ordinary RGB images (in JPEG format) were recorded as well, although they are not utilized. Both the DTM and LoD1 model were provided by the city of Munich.

4.2. Case study results

Four sets out of the 49 (Figure 8) were evaluated to demonstrate the program's functionality, allude to its effectiveness, and discuss potential for improvement. They were chosen for the presence of confirmed ground truth data and varying image quality (Table 2).

Applying the developed leakage detection algorithm to the four select sets lead to the results given in Table 3. Between 92 and 99% of false alarms were sorted out using the various methods described in Section 3.5. The remaining false positives were mostly found to stem from common urban features, such as warm vehicles, buildings or manholes.¹⁴

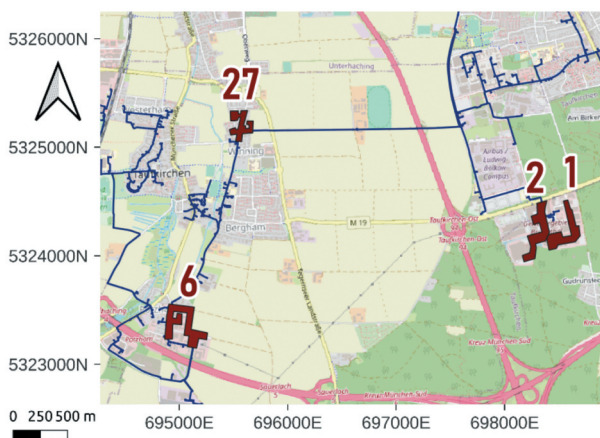


Figure 8. Locations of four case study sets.

Table 2. General information on the four select case study sets.

set #	1	2	6	27
date	02.12.	03.12.	02.12.	11.12.
time	11.15–11.50 pm	12.05–12.40 am	10.00–10.45 pm	10.30–11.00 pm
image count	875	868	982	640
Canny pixel count	8250	10706	8444	2163
quality	high	high	high	low

Individual image evaluation leads to the largest amount of detected leakages, mainly due to the high image overlap causing regions of interest to be identified multiple times. The leakage location evaluation helps determine these cases and remove duplicates. A comparative analysis of leakage groupings shows the chosen methodology to be effective in grouping together true positives, meaning identical thermal anomalies.¹⁵ The false positive rate, in other words the amount of leakages defined as being a replica although they are not, equalled approximately 8% in both mosaics and individual images. The location information also gives insight into false alarms that are caused by buildings. An evaluation of the computed leakage locations showed that thermal anomalies close or above buildings were correctly found in 58% of cases within mosaics and 68% within single images. Of the anomalies determined to be in proximity to buildings, a small percentage¹⁶ was found to stem from objects close by (such as vehicles) instead of the building itself.

This requires different computational effort depending on the method involved. Single image georeferencing takes the least amount of time and depends solely on the amount of images in a set. Orthomosaic generation is more time-intensive owing to the need to access the third party software. WebODM outstrips both alternatives by far because of the included fail-safe mechanism, which reruns faulty image groups.

5. Discussion

The results of the developed leakage detection algorithm were compared to a manual assessment of the select case study sets, an overview of which is given in Table 4. This manual evaluation was performed on a mosaic created from all images within a set, as little actual ground truth data was available. The algorithm returns only a small portion of the manually identified regions of interest, the reasons for which are provided in last row of Table 4.

In approximately 70% of the cases (denoted by ' ΔT ') manually defined areas of interest were detected by the algorithm, yet sorted out owing to a low temperature difference to their surroundings. The vast majority of manually identified thermal hotspots therefore did not differ enough from their setting to be classified as leakages. Almost 15% of the visually identified leakages (denoted by 'mask') were eliminated from consideration due to their increased distance from the DHS pipelines. The algorithm is therefore found to greatly support human analysis by removing those manually identified regions which do not fulfil key leakage criteria.

The other missed leakages shed light upon two shortcomings of the leakage detection algorithm in its current implementation. The first is caused by the automation of the georeferencing and mosaicking step (denoted by 'no mos' or 'mos wrong'), the second by

Table 3. Overview of the results of the leakage detection algorithm applied to four select sets

set #	1				2				6				27								
	Pix4D mapper	Web ODM	single images	Pix4Dmapper	Web ODM	single images	Pix4Dmapper	Web ODM	single images	Pix4Dmapper	Web ODM	single images	single images	single images							
leakages before removal	170	1488	4439	150	2083	5014	137	2275	5922	17	2083	5014	137	2275	5922	17	2083	5014	137	2275	5922
leakages of fitting size	73	171	727	75	218	541	61	226	942	7	218	541	61	226	942	7	218	541	61	226	942
leakages of fitting shape	31	55	376	30	62	200	33	101	495	31	62	200	33	101	495	31	62	200	33	101	495
leakages of high ΔT	7	14	117	12	29	90	10	30	141	7	29	90	10	30	141	7	29	90	10	30	141
unique leakages	6	9	17	7	12	13	8	20	41	6	12	13	8	20	41	6	12	13	8	20	41
	pot: 6	pot: 6	pot: 13,	pot: 4,	pot: 7,	pot: 7,	pot: 5,	pot: 17,	pot: 37,	pot: 6	pot: 7,	pot: 7,	pot: 5,	pot: 17,	pot: 37,	pot: 6	pot: 7,	pot: 7,	pot: 5,	pot: 17,	pot: 37,
	def: 2,	def: 2,	def: 3,	def: 2,	def: 4,	def: 2,	def: 3	def: 2,	def: 3,	def: 2,	def: 2,	def: 2,	def: 3	def: 2,	def: 3,	def: 2,	def: 2,	def: 2,	def: 3	def: 2,	def: 3,
	crit:	crit:	crit: 1	crit: 1	crit: 1	crit: 4	crit: 1	crit: 1	crit: 1	crit:	crit: 1	crit: 4	crit: 1	crit: 1	crit: 1	crit:	crit: 1	crit: 4	crit: 1	crit: 1	crit: 1
false positives	2x veh,	1x veh,	2x veh,	1x veh,	3x mh,	3x mh,	2x veh,	8x veh,	9x veh,	2x veh,	3x mh,	3x mh,	2x veh,	8x veh,	9x veh,	2x veh,	3x mh,	3x mh,	2x veh,	8x veh,	9x veh,
	3x mh,	3x mh,	1x mh,	1x mh,	5x bldg,	6x bldg,	5x bldg,	1x mh,	2x mh,	1x mh,	6x bldg,	6x bldg,	1x mh,	1x mh,	2x mh,	2x mh,	1x mh,	1x mh,	2x mh,	1x mh,	2x mh,
	1x mos	1x bldg,	10x bldg,	1x bldg,	1x metal	1x metal	1x bldg,	8x bldg,	23x bldg,	1x metal	1x metal	1x metal	8x bldg,	8x bldg,	23x bldg,	2x metal	23x bldg,	23x bldg,	2x metal	23x bldg,	2x metal
actual leakages	0	pot: 1	pot: 1	pot: 1,	pot: 2,	pot: 2,	0	pot: 2	pot: 5	pot: 1	pot: 2,	pot: 2,	0	pot: 2	pot: 5	pot: 1	pot: 2,	pot: 2,	0	pot: 2	pot: 5
run time [h]	1.57	4.49	1.25	1.75	4.34	1.26	1.99	3.96	1.48	0.93	1.26	1.99	3.96	1.48	0.93	0.93	1.26	1.26	1.99	3.96	1.48
\sum run time [h]	6.08	6.08	6.02	6.02	6.02	6.02	6.02	6.02	6.02	6.02	6.02	6.02	6.02	6.02	6.02	6.02	6.02	6.02	6.02	6.02	6.02

Abbreviations: pot = potential, def = definite, crit = critical, bldg = building, mh = manhole, mos = mosaic, veh = vehicle.

Table 4. Comparison of the leakage detection algorithm results with the manual evaluation of the select case study sets.

set #	1		2		6		27
method	Pix4D mapper	Web ODM	Pix4D mapper	Web ODM	Pix4D mapper	Web ODM	single images
unique leakages	pot: 6 def: 2, crit: 1	pot: 6, def: 2, crit: 1	pot: 4, def: 2, crit: 1	pot: 7, def: 4, crit: 1	pot: 5, def: 3	pot: 17, def: 2, crit: 1	pot: 37, def: 3, crit: 1
true leakages	0	pot: 1	pot: 1, crit: 1	pot: 2, crit: 1	0	pot: 2	pot: 5
manually identified leakages		pot: 4	pot: 3, crit: 1	pot: 3, crit: 1		pot: 5, def: 2	pot: 5
manually identified and true leakages	0	0	pot: 1, crit: 1	pot: 2, crit: 1	0	0	0
reasons manually identified leakages are missed	1x ΔT , 1x mask, 1x T_{th} , 1x no mos	2x ΔT , 1x mask, 1x T_{th}	1x ΔT , 1x mos wrong	1x ΔT	4x ΔT , 2x T_{th} , 1x mos wrong	5x ΔT , 1x mask, 1x T_{th}	4x ΔT , 1x T_{th}

Abbreviations: pot - potential, def - definite, crit - critical, mos - mosaic, th - threshold.
Reasons for oversight: incorrect manual identification: ΔT - delta too small, mask - too far from DHS incorrect removal by algorithm: T_{th} - incorrect threshold, mos - mosaic failure

Table 5. Comparison of the select case study sets illustrating the methodology- dependent amounts of unevaluated image data. Files refer to either mosaics or thermal images. Different file counts pertain to the originally given or planned number of files versus those resulting after georeferencing. Percentages are calculated in terms of images, although overlapping mosaics are not considered.

set #	1			2			6			27
method	Pix 4D	Web ODM	single	Pix 4D	Web ODM	single	Pix 4D	Web ODM	single	single
original file count	28	75	875	28	82	868	32	109	982	640
resulting file count	22	75	694	26	81	651	30	111	796	448
missing images	21%	0%	21%	7%	0.8%	25%	6%	0.9%	18.9%	30%
relevant missing images	21%	0%	0%	7%	0.8%	0.5%	6%	0.9%	0%	0.2%

the thresholding method being applied individually to images and mosaics (denoted by ' T_{th} ') due to the mentioned calibration drift in the case study sets (Appendix A).

Automating the georeferencing step – be it as orthomosaics or individually – means not all images can be processed. Table 5 shows the extent of this fact. As discussed in Section 3.2, some images include incorrect heading angles and must therefore be removed from consideration during individual georeferencing. Owing to the high image overlap of nearly 90%, this only poses a problem when large amounts of consecutive images fail the angle evaluation – an extremely rare situation (“single” columns in Table 5). Figure 9 depicts the source of most incorrect angles: they occur when the UA abruptly changes direction. This observation can be used as a newly identified best practice rule and should be taken into account during future image acquisition and flight planning to ensure the usability of all images.

This is more problematic in the automated mosaicking process. Both WebODM and Pix4Dmapper were at times found to fail to generate mosaics, leading to entire image groups being passed over. This happens in WebODM despite implemented quality and failure control mechanisms when an image group still yields no usable results after considerable size extension. In Pix4Dmapper, even extending a previously failed image group to double its amount of images may not yield any

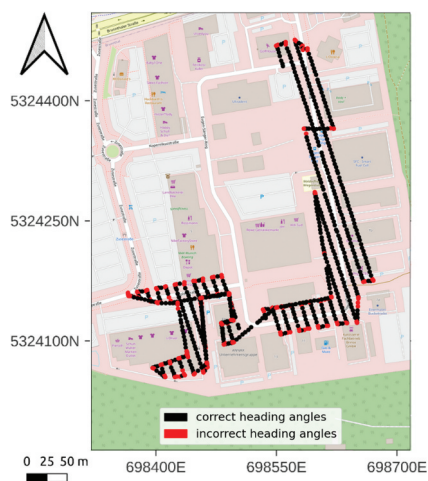
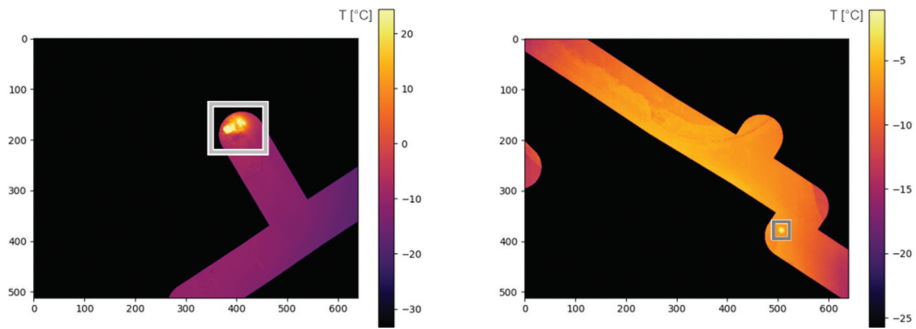


Figure 9. GNSS coordinates of all images in a case study set displayed on an OSM background. Images with incorrect heading angles are red.



(a) Critical leakage (set 2) with a temperature difference of 36.4°C (b) Potential leakage (set 6) with a temperature difference of 5.1°C

Figure 10. Example leakages detected within the select case study image sets.

result and forces the original group to be skipped. This explains why the automated Pix4mapper software misses most images (Table 5) and one of the manually identified leakages in Set 1 (Table 4). While theoretically more adept at removing distortions and image vignetting, evaluation of the case study datasets showed the georeferencing of both WebODM and Pix4Dmapper mosaics to be imprecise (Appendix B).

Despite all specified processing options, the mosaics sometimes failed to achieve high quality. In the selected image sets, the WebODM mosaics included holes or small areas missing in approx. 7% of the cases. Pix4Dmapper displayed the more serious peculiarity of sometimes producing near (approx. 17%) or completely (approx. 12%) unusable mosaics. While a cause for this might be the previously mentioned conflicting image metadata, it highlights the necessity for a more in-depth testing, particularly of Pix4Dmapper, to ensure the usability for applications such as leakage detection. Additionally, the importance of high image quality must be stressed: The optimal acquisition conditions described in Section 2.1 should always be adhered to. The comparative analyses of all implemented georeferencing methodology has shown the single image georeferencing process, while not as adept at removing distortions, to be the most consistent, robust and reliable. Moreover, due to the large number of images and various viewpoints of similar scenery, the scaling distortions do not weigh quite as heavily. It also requires the least processing time (Table 3).

The second shortcoming of the implemented leakage detection algorithm comes from the necessity to individually threshold images and mosaics due to the observed calibration drift (Section 3.4). While this individual segmentation method works well in detecting locally concentrated leakages of various sizes, it cannot identify more widespread variants that encompass large areas. Extending the algorithm to perform both individual as well as the combined thresholding demonstrated by Friman et al. (2014) could solve this problem and allow for the detection of long pipeline sections of interest. Any temperature drift similar to this case study's would, however, cause a vast increase in false positives.

Ultimately, an extensive analysis of the thermal camera is required to ensure such a drift does not occur during future acquisitions. Alternatively, calibration references such as black-bodies with known temperatures can be placed within the flight zone to calculate it (Prof. Grant Petty, pers. comm., April 14, 2021).

Owing to the algorithm design and non-automatic removal of duplicates as well as little actual ground truth data available, an exact comparison to methods from related work cannot properly be performed. However, despite the revealed problems, the implemented program is capable of identifying the most critical and important leakage in the entire case study within Set 2, which ruptured two weeks after image acquisition. It was detected within all variations of the georeferencing process, characterized by a 36.4°C temperature difference to its surroundings, and therefore categorized correctly as critical in all cases (Figure 10(a)). Various smaller, potentially existent leakages were reliably identified as well (Figure 10(b)). Most importantly, the developed algorithm was shown to be capable of automatically evaluating thousands of images, detecting the important leakages and returning concise and easily manageable results completely autonomously.

Finally, some limitations of the program in its current form need to be addressed. For adequate functionality, thermal images must be acquired in a similar fashion as those of the case study. For instance, terrain following should not be implemented because the elevation data used to calculate the flight height would otherwise have an adverse effect and cause imprecision. Changes in flight height or choice of thermal camera affect the correctness of the thermal image quality assessment, buffer size around the DHS network, and false alarm removal by anomaly size and may therefore cause subpar functionality.

Incorporating building data into false alarm identification requires extremely precise georeferencing. Because none of the three presented methods can offer this to the desired extent, the location evaluation step could not be included in the overall process. Further enhancement is necessary to do so.

6. Conclusion

6.1. Summary

The aim of this study was to develop and demonstrate a simple, automated way for network operators to monitor and assess their DHSs. Aside from making use of the cost- and time-efficient UAS-based acquisition process, this meant the development of an program suitable for automated thermal image analysis. The designed software is functional, completely automated, and capable of correctly identifying leakages of varying importance. Considerable value was placed upon the development of a freely implementable program to ensure the program's future usability and allow for effortless modifications of both code and general methodology.

The developed algorithm comprises five fundamental steps modified for automation. The images are pre-processed, three variations of georeferencing are applied and compared, results are masked with the DHS blueprint, and thermal anomalies are identified within the minimized search space. Various false positives are sorted out based on characteristic size, shape, and temperature of thermal anomalies, and the final results are categorized as leakages of different severity. Lastly, information is supplied on leakage locations, their proximity to buildings, and potentially existing duplicates.

The results are presented in an easily comprehensible and transparent fashion using excel data sheets and the plotting of leakage locations on maps to give operators an effortless and quick overview. Additionally, the identification of identical leakages and anomalies in proximity to buildings can help reduce the amounts of remaining false alarms.

Result evaluation exposed two drawbacks of the algorithm in its current implementation, with suggestions for improvement discussed in Section 5. It highlighted the importance of capturing high quality images by adhering to best-practice acquisition rules (Section 2.1). The camera should be tested to ensure no discernible temperature drift occurs. Additionally, flight speeds or patterns should be altered to record correct heading angles and avoid sharp curves. In order to alleviate the misinterpretation of common urban features such as vehicles and people, it would be prudent to acquire images during the night and early morning.

6.2. Outlook

An important extension to a leakage detection program would be the inclusion of temporal analyses as described by Berg, Ahlberg, and Felsberg (2016). Repeating the image acquisition process after one or more years under similar conditions allows for a comparison to the initial state and gives further insight into the general network status and its degradation. Critical changes can be identified and initially overlooked leakages discovered (Berg, Ahlberg, and Felsberg 2016).

Additional insight can also be gained by integrating simultaneously recorded RGB images as suggested by Sledz, Unger, and Heipke (2020), which would allow for a parallel analysis of both thermal and optical images. For instance, it would be possible to remove common false alarms in the given case study data sets caused by streetlamps or manholes through their identification within the corresponding optical images. However, since the optimum time frame is image acquisition at night-time, the inherent RGB information that can be extracted from optical images will be limited.

Future work may also profit from the integration of network characteristics like age, repairs, material information, and pipeline depth. Including these can help improve the quality of leakage forecasts. Data gained from collaborations with local utility companies would also allow for more precise analyses by including any identified true ground truth information, something that studies so far lack.

One of the main gaps in current literature is a true comparison of already developed solutions. The lack of published thermal datasets and source code makes this impossible up to now. We hope that by sharing our data as part of this paper's publication, the open-source approach will become common-place amongst these kinds of case studies and thus allow for genuine analyses and a collaborative exchange of knowledge.

Abbreviations

CRS	coordinate reference system.
DHS	district heating system.
DTM	digital terrain model.

GDAL	Geospatial Data Abstraction Library.
GIS	geographic information system.
GNSS	global navigation satellite system.
GSD	ground sampling distance.
IMU	inertial measurement unit.
LOD1	level of detail 1.
LoG	Laplacian of Gaussian.
LWIR	long-wavelength infrared.
ODM	OpenDroneMap.
OSM	OpenStreetMap.
RJPEG	Radiometric Joint Photographic Experts Group.
SWM	Stadtwerke München.
TIR	thermal infrared.
UA	unmanned aircraft.
UAS	unmanned aircraft system.

Notes

1. Orthomosaics are georectified mosaics, meaning mosaics that are correctly mapped and have had their inherent distortions removed (ArcGIS Pro 2021a).
2. Preliminary testing showed saliency maps and LoG blob detectors to be ineffective or require the setting of multitudes of dataset-dependent parameters.
3. Heipke and Todter (2020) define 2 m/s as the maximum wind speed yet exceeded that limit frequently during their own image acquisition.
4. Heipke and Todter (2020) name the following as unproblematic: fibre cement, plastic, protective, steel, and flexible pipes, independent of overhead or underground placement.
5. For instance faulty building insulation, streetlamps, people, vehicles, and metal surfaces.
6. Metal objects, for instance, are depicted with much colder temperatures than are actually measured on their surface (Gade and Moeslund 2014).
7. Latitude and longitude are translated from the geographical coordinate system (unit: degrees, minutes and seconds) into easting and northing values in a Cartesian CRS (unit: metres).
8. Ordinarily, complex angle transformations would be necessary to convert the recorded navigational attitude yaw, pitch, and roll to the photogrammetric equivalents' omega, phi, and kappa used in georeferencing (Bäumker and Heimes 2002). These can be omitted here because the implemented 3-axis gimbal (SZ DJI Technology Co. Ltd. 2018b) stabilizes the UA, ensuring negligibly small roll and pitch angles (Prof. Grant Petty, pers. comm., April 14, 2021). Owing to nadir UA alignment, the yaw angle is directly assumed to be κ .
9. An 8° difference is experimentally determined as the threshold for the given UA Matrice 600 Pro (SZ DJI Technology Co. Ltd. 2018a).
10. In the case of a UA's flight height, a conservative threshold is chosen at 10 pixels.
11. This form is based upon Sledz, Unger, and Heipke (2020)'s observation that an elliptical shape best depicts leakages. Berg, Ahlberg, and Felsberg (2016) consider form as an important factor, but only take circularity into account as a feature for distinguishing leakages. However, as Friman et al. (2014) state, a leakage is not always perfectly round and may also be partially covered, making an ellipse the more obvious choice.
12. Berg, Ahlberg, and Felsberg (2016) propose a similar method of comparing the intensity of a hotspot and its surroundings. However, using temperature data has the advantage of enabling postliminary leakage categorization with empirical values.
13. Sledz, Unger, and Heipke (2020) define a delta of 5°C already sufficient to indicate a leakage.

14. Munich's municipal utilities company SWM states some manholes may indicate DHS leakages and should therefore not be disregarded completed, although the majority are marked as false alarms in Table 3.
15. In mosaics approximately 85% and in individual images 96% of thermal hotspots.
16. 7% within mosaics and 10% within single images.
17. In the Universal Transverse Mercator (UTM) projection, zone '32N' (denoting European territory), and European Terrestrial Reference System 1989 (ETRS89) map datum.

Acknowledgements

The thermal images for the case study in Germany were acquired in collaboration with the Air Bavarian GmbH and Munich's municipal utilities company Stadtwerke München.

Disclosure statement


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


Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

ORCID

Elena Vollmer  <http://orcid.org/0000-0002-8805-3726>

Rebekka Volk  <http://orcid.org/0000-0001-9930-5354>

Frank Schultmann  <http://orcid.org/0000-0001-6405-9763>

CRedit author statement

Elena Vollmer: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Writing – Original Draft, Writing – Review & Editing, Visualization.

Rebekka Volk: Conceptualization, Data Curation, Writing – Review & Editing, Supervision.

Frank Schultmann: Writing – Review & Editing, Supervision. All authors have agreed to the published version of the manuscript.

References

- ArcGIS Pro. 2021a. "Documentation: Generate an Orthomosaic Using the Orthomosaic Wizard." Accessed February 8, 2022. <https://pro.arcgis.com/en/pro-app/latest/help/data/imagery/generate-an-orthomosaics-using-the-orthomosaic-wizard.htm>.
- ArcGIS Pro. 2021b. "Documentation: Overview of Georeferencing." Accessed May 25, 2021. <https://pro.arcgis.com/en/pro-app/latest/help/data/imagery/overview-of-georeferencing.htm>.
- Axelsson, S. 1988. "Thermal Modeling for the Estimation of Energy Losses from Municipal Heating Networks Using Infrared Thermography." *IEEE Transactions on Geo-Science and Remote Sensing* 26 (5): 686–692. <https://doi.org/10.1109/36.7695>
- Bäumker, M., and F.-J. Heimes. 2002. "New Calibration and Computing Method for Direct Georeferencing of Image and Scanner Data Using the Position and Angular Data of an Hybrid Inertial Navigation System." Proceedings of OEEPE Workshop on Integrated Sensor Orientation.

- https://www.researchgate.net/publication/254001148_New_Calibration_and_Computing_Method_for_Direct_Georeferencing_of_Image_and_Scanner_Data_Using_the_Position_and_Angular_Data_of_an_Hybrid_Inertial_Navigation_System
- Berg, A., J. Ahlberg, and M. Felsberg. 2016. "Enhanced Analysis of Thermographic Images for Monitoring of District Heat Pipe Networks." *Pattern Recognition Letters* 83:215–223. <https://doi.org/10.1016/j.patrec.2016.07.002>.
- Burton, D. A. 2013. "Composite Standard Deviations." Cary, North Carolina, USA. Accessed April 23, 2021. http://www.burtonsys.com/climate/composite%7B%5C_%7Dsd.php%7B%5C#%7Dpython.
- Distante, A., and C. Distant. 2020. *Handbook of Image Processing and Computer Vision: From Image to Pattern*. Vol. 2. Cham, Switzerland: Springer International Publishing. <https://doi.org/10.1007/978-3-030-42374-2>.
- El-Zahab, S., and T. Zayed. 2019. "Leak Detection in Water Distribution Networks: An Introductory Overview." *Smart Water* 4 (1): 1–23. issn: 2198-2619. <https://doi.org/10.1186/s40713-019-0017-x>.
- FLIR Systems, Inc. 2017. *Duo Pro R: User Guide*. Product information. Version 1.0. FLIR® Systems, Inc. <https://www.flir.com/products/duo-pro-r/>.
- Friman, O., P. Follo, J. Ahlberg, and S. Sjökvist. 2014. "Methods for Large-Scale Monitoring of District Heating Systems Using Airborne Thermography." *IEEE Transactions on Geoscience & Remote Sensing* 52 (8): 5175–5182. <https://doi.org/10.1109/TGRS.2013.2287238>.
- Gade, R., and T. B. Moeslund. 2014. "Thermal Cameras and Applications: A Survey." *Machine Vision and Applications* 25 (1): 245–262. <https://doi.org/10.1007/s00138-013-0570-5>.
- Garrard, C. 2016. *Geoprocessing with Python*. Shelter Island, NY: Manning Publications Co.
- GDAL/OGR contributors, 2022. *Open Source Geospatial Foundation "GDAL/OGR Geospatial Data Abstraction Software Library"* Accessed February 20, 2022. <https://gdal.org/10.5281/zenodo.5884351>.
- Gillies, S. 2013. "Rasterio: Geospatial Raster I/O for Python Programmers." Mapbox. Accessed May 30, 2021. <https://github.com/rasterio/rasterio>.
- Government of Canada. 2015. "Geometric Distortion in Imagery." Accessed May 19, 2021. <https://www.nrcan.gc.ca/maps-tools-publications/satellite-imagery-air-photos-remote-sensing-tutorials/satellites-sensors/geometric-distortion-imagery/9401>.
- Groos, A. R., T. J. Bertschinger, C. M. Kummer, S. Erlwein, L. Munz, and A. Philipp. 2019. "The Potential of Low-Cost UAVs and Open-Source Photogrammetry Software for High-Resolution Monitoring of Alpine Glaciers: A Case Study from the Kanderfirn (Swiss Alps)." *Geosciences* 9 (8): 356. 10.3390/geosciences9080356.
- Heipke, C., and J. Tödter. 2020. *Drohnen-gestützte Thermografie als Basis der Asset- und Instandhaltungsstrategie von Fernwärmenetzen [Drone-based thermography as the basis for asset and maintenance strategies of district heating networks]*. Final report 19768 N. IGF, Fernwärme-Forschungsinstitut in Hannover e.V. [District heating research in Hannover e.V.], and Leibniz University Hannover.
- Hossain, K., F. Villebro, and S. Forchhammer. 2020. "UAV Image Analysis for Leakage Detection in District Heating Systems Using Machine Learning." *Pattern Recognition Letters* 140:158–164. <https://doi.org/10.1016/j.patrec.2020.05.024>.
- Jayapalan, L., C. Murray, F. Chen, and Q. Xie. 2021. Oracle® Spatial: Spatial GeoRaster Developer's Guide. 21c. Oracle®. <https://docs.oracle.com/en/database/oracle/oracle-database/21/geors/spatial-georaster-developers-guide.pdf>.
- Krist, H., B. Reinelt, A. Frisch, S. Kreidenweis, A. Dalsass, and A. Hofmann. 2017. *Wärmenetze in Kommunen - in Zehn Schritten Zum WE4rmenetz, Leitfaden [Heating Networks in Municipalities - ten Steps to a Heat Network, a Guide]*. Tech. Rep. 4th ed. Kempten, Germany: Bayerisches Landesamt für Umwelt (LfU) [Bavarian state office for the environment] and Bayerisches Staatsministerium für Wirtschaft, Landesentwicklung und Energie (StMWi) [Bavarian state ministry for economic affairs, state development and energy].
- Leidner, J. L. 2016. "Georeferencing: From Texts to Maps." In *International Encyclopedia of Geography: People, the Earth, Environment and Technology*, edited by D. Richardson, N. Castree, M. F. Goodchild, A. Kobayashi, W. Liu, and R. A. Marston, 1–10. Oxford: John Wiley & Sons, Ltd. <https://doi.org/10.1002/9781118786352.wbieg0160>.

- Ljungberg, S.-A., and M. Rosengren. 1988. "Aerial and Mobile Thermography to Assess Damages and Energy Losses from Buildings and District Heating Networks - Operational Advantages and Limitations." *International Arch Photogrammetry Remote Sensing XXVII- B7 (27-B7)*: 348–359.
- Merziger, G., G. Mühlbach, D. Wille, and T. Wirth. 2013. *Formeln Und Hilfen: Höhere Mathematik. 7th. [Formulas and Equations: Higher Mathematics]*. 7th ed. Barsinghausen, Germany: Binomi.
- Olsson, M. E. 2001. "Long-Term Thermal Performance of Polyurethane-Insulated District Heating Pipes." PhD thesis, Göteborg, Sweden: Chalmers University of Technology.
- OpenDroneMap. 2020. "WebOdm". Accessed May 25, 2021. <https://www.opendronemap.org/webodm/>.
- Pix4D. 2019. "What is Accuracy in an Aerial Mapping Project?" Accessed January 25, 2021. <https://www.pix4d.com/blog/accuracy-aerial-mapping>.
- Pix4D. 2020. Pix4Dmapper. Accessed May 25, 2021. <https://www.pix4d.com/product/pix4dmapper-photogrammetry-software>.
- Pix4D. 2021. "Pix4dmapper: Processing Steps." Accessed May 25, 2021. <https://support.pix4d.com/hc/en-us/articles/115002472186-Processing-steps>.
- Sledz, A., J. Unger, and C. Heipke. 2020. "UAV-Based Thermal Anomaly Detection for Distributed Heating Networks." *ISPRS - International Archives of the Photogrammetry Remote Sensing and Spatial Information Sciences XLIII-B1-2020:499–505*. <https://doi.org/10.5194/isprs-archives-XLIII-B1-2020-499-2020>.
- SZ DJI Technology Co. Ltd. 2018a. *Matrice 600 Pro: User Manual*. Product information. Version 1.0. SZ DJI Technology Co. Ltd. <https://www.dji.com/de/matrice600-pro>.
- SZ DJI Technology Co. Ltd. 2018b. *Zenmuse XT 2: User Manual*. Product information. Version 1.0. SZ DJI Technology Co. Ltd. <https://www.dji.com/zenmuse-xt2>.
- University Consortium for Geographic Information Science. 2020. "GIS&T DC-30- Georeferencing and Georectification." Accessed May 19, 2021. <https://gistbok.ucgis.org/bok-topics/georeferencing-and-georectification>.
- Van der Walt, S., J. L. Schönberger, J. Nunez-Iglesias, F. Boulogne, J. D. Warner, N. Yager, E. Gouillart, and T. Yu. 2014. "Scikit-Image: Image Processing in Python." *PeerJ*. 2:e453. <https://doi.org/10.7717/peerj.453>.
- Vollmer, E., R. Volk, and M. Vogl. 2023. "Automatic Analysis of UAS-Based Thermal Images to Detect Leakages in District Heating Systems: Source Code and an Example Thermal Image Dataset." *Zenodo V1*. <https://doi.org/10.5281/zenodo.7851726>.
- Werner, S. 2017. "District Heating and Cooling in Sweden." *Energy*. 126:419–429. <https://doi.org/10.1016/j.energy.2017.03.052>.
- Whitehead, K., and C. H. Hugenholtz. 2014. "Remote Sensing of the Environment with Small Unmanned Aircraft Systems (UASs), Part 1: A Review of Progress and Challenges." *Journal of Unmanned Vehicle Systems* 2 (03): 69–85. <https://doi.org/10.1139/juvs-2014-0006>.
- Xu, Y., X. Wang, Y. Zhong, and L. Zhang. 2016. "Thermal Anomaly Detection Based on Saliency Computation for District Heating System". In *2016 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, pp. 681–684. <https://doi.org/10.1109/IGARSS.2016.7729171>.
- Zack, G. W., W. E. Rogers, and S. A. Latt. 1977. "Automatic Measurement of Sister Chromatid Exchange Frequency." *The Journal of Histochemistry and Cytochemistry: Official Journal of the Histochemistry Society* 25 (7): 741–753. issn: 0022-1554. <https://doi.org/10.1177/25.7.70454>.
- Zhong, Y., Y. Xu, X. Wang, T. Jia, G. Xia, A. Ma, and L. Zhang. 2019. "Pipeline Leakage Detection for District Heating Systems Using Multisource Data in Mid- and High-Latitude Regions." *ISPRS Journal of Photogrammetry and Remote Sensing* 151:207–222. <https://doi.org/10.1016/j.isprsjprs.2019.02.021>.


Appendix A

The following figures exemplify the amendments made to the basic triangle histogram thresholding method, showing why these are necessary in form of a cursory, visual ablation study. Each figure depicts in (a) a thermal image after mapping and masking, in (b) a distribution of the thermal data and threshold automatically selected by the thresholding method, and in (c) the results of applying that threshold to the image to obtain a binary matrix with white pixels being those of interest.

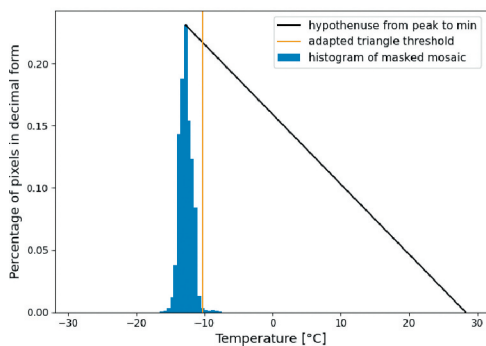
Figure A1 shows a regular thermal mosaic and its threshold, which is used without further changes to create the resulting binary image in (c). No amendments are necessary for the majority of images such as this one.

Figure A2 showcases an example where the standard choice of threshold lies between peaks, making an amendment necessary to choose the correct threshold according to peak distribution. Figures (d) and (e) display the effects of the methodological change on the choice of threshold and subsequent pixels of interest.


Figure A3 demonstrates a too conservatively chosen threshold and the implemented correction to mitigate the problem. Figures (d) and (e) display the effects of the methodological change on the choice of threshold and subsequent pixels of interest.



(a) Standard thermal mosaic (temperature colourbar in °C)

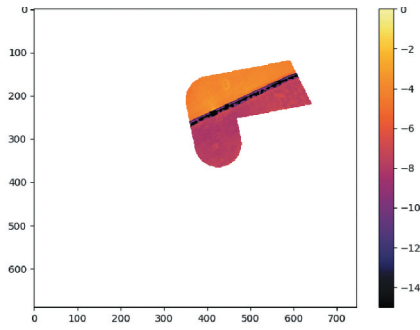


(b) Histogram and triangle threshold

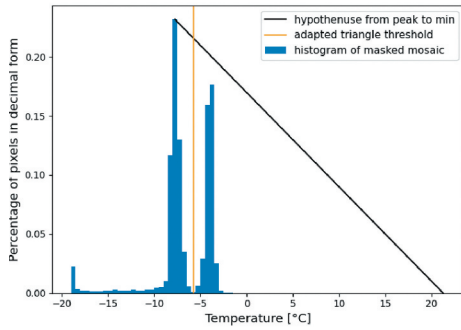


(c) Resulting binary thresholded mosaic


Figure A1. Triangle histogram thresholding demonstrated using a standard thermal mosaic.



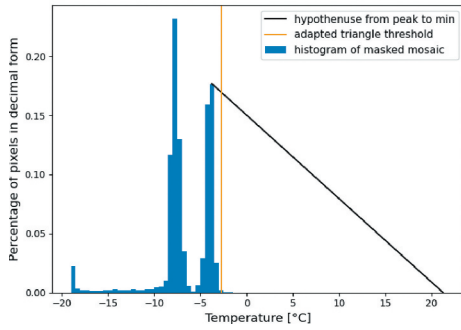
(a) Thermal mosaic with two distinct areas (temperature colourbar in $^{\circ}\text{C}$)




(b) Histogram and triangle threshold of thermal mosaic without peak consideration



(c) Resulting binary thresholded mosaic with too many pixels of interest

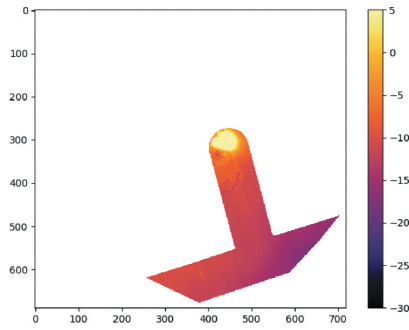


(d) Histogram and triangle threshold of thermal mosaic with peak consideration




(e) Amended resulting binary thresholded mosaic

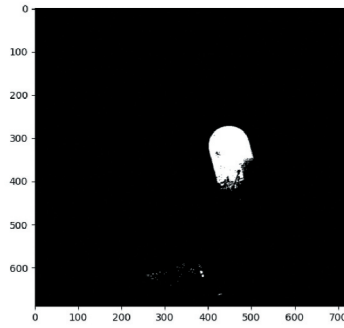
Figure A2. Triangle histogram thresholding demonstrated using a thermal mosaic with two distinct areas causing two peaks. The peak adjustment allows the selection of only relevant pixels within the warmer region.




(a) Varied thermal mosaic with a large leakage (temperature colourbar in $^{\circ}\text{C}$)




(b) Histogram and triangle threshold of thermal mosaic without consideration of bin magnitude



(c) Resulting binary thresholded mosaic with too many pixels of interest



(d) Histogram and triangle threshold of thermal mosaic with consideration of bin magnitude




(e) Amended resulting binary thresholded mosaic


Figure A3. Triangle histogram thresholding demonstrated using a varied thermal mosaic with a large leakage causing a too conservative threshold. The adjusted threshold ensures the region of interest is reduced to the appropriate size.

Appendix B

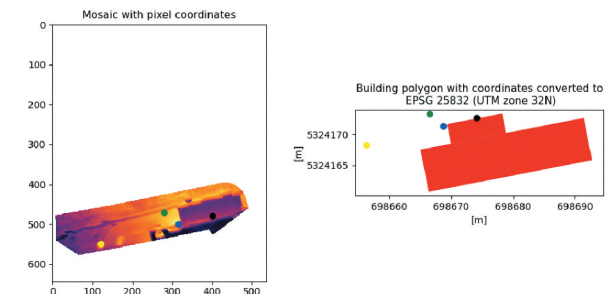
Figure B1 comparatively visualises how the geotransform calculated during (a) individual image georeferencing and georectification via (b) WebODM or (c) Pix4Dmapper can deviate from the true positioning. Particularly in orthomosaics it is found to be too inconsistent to use reliably in classification. Figure B1 demonstrates this with example points, marked in relation to the depicted building,¹⁷ to illustrate the shift in comparison to the building's equivalent geographically accurate polygon.



(a) Georeferencing discrepancies in an exemplary georeferenced single image



(b) Georeferencing discrepancies in an exemplary WebODM orthomosaic



(c) Georeferencing discrepancies in an exemplary Pix4D orthomosaic

Figure B1. Possible discrepancies in georeferencing exemplified for all variations of georeferencing. The left side shows the georeferenced image or mosaic with points plotted in relation to the depicted building. The right side draws those same geographic points in relation to the building's equivalent geographic polygon.

Legend: black - on building, blue - at building border, green - close to building, yellow - far from building