# CLASSIFICATION AND TEMPORAL ANALYSIS OF DISTRICT HEATING LEAKAGES IN THERMAL IMAGES

*Amanda Berg[1,2] and Jörgen Ahlberg[1,2]*

[1] Termisk Systemteknik AB, Diskettgatan 11B, SE-583 35 Linköping, Sweden

[2] Dept. of Electrical Engineering, Linköping University, SE-581 83 Linköping, Sweden

## ABSTRACT

District heating pipes are known to degenerate with time and in some cities the pipes have been used for several decades. Due to bad insulation or cracks, energy or media leakages might appear. This paper presents a complete system for large-scale monitoring of district heating networks, including methods for detection, classification and temporal characterization of (potential) leakages. The system analyses thermal infrared images acquired by an aircraft-mounted camera, detecting the areas for which the pixel intensity is higher than normal. Unfortunately, the system also finds many false detections, i.e., warm areas that are not caused by media or energy leakages. Thus, in order to reduce the number of false detections we describe a machine learning method to classify the detections. The results, based on data from three district heating networks show that we can remove more than half of the false detections. Moreover, we also propose a method to characterize leakages over time, that is, repeating the image acquisition one or a few years later and indicate areas that suffer from an increased energy loss.

## INTRODUCTION

Distribution of heat to homes and industries through district heating networks is today one of the most common heating sources in Swedish and Nordic cities. However, the pipes degenerate with time [3] and due to bad insulation or cracks, energy or media leakages might appear. Bad insulation can, for example, be caused by cracks in the outer protective shell, allowing water to enter the insulation layer thus significantly reducing the insulation effect. In addition to being expensive for the network owner, the loss of media or energy also has negative impact on the environment [4]. Therefore, it is of great interest to the owner to have efficient and reliable methods for leakage detection, especially when considering the fact that the pipes generally are placed underground, it is very expensive to dig in the wrong place. Moreover, major leakages of 50 $m^3$ media or more per day may also cause the ground to collapse due to erosion, whereby large amounts of media at boiling temperature are exposed.

This paper presents methods for detection, classification and temporal characterization of such leakages. We have used a commercially available system for large-scale airborne thermal image acquisition, acquiring data from several Nordic cities.

For *detection* of potential leakages, we use the method previously published by Friman et al. [1]. The method is used to analyse the acquired imagery, finding and indicating the areas for which the pixel intensity (temperature) is higher than normal. Apart from the sought-for media and energy leakages, there are several types of objects and phenomena that give rise to such detections. Examples are areas that, for some reason, are warmer than their surroundings, for example, chimneys, cars and heat leakages from buildings. In a large city, there might be several thousands of false detections.

Thus, we want to reduce these false detections as much as possible while maintaining the number of true detections at a fixed level (we use 99%). In order to achieve this goal, we follow a two-step *classification* procedure, as proposed in [2]:

1. Extract building locations from publically available geographic information, and remove all detections located on buildings.
2. Extract image features and use a machine learning method to classify detections as true (media/energy) or false detections.

Next, we propose a novel method for temporal characterization and visualization of the energy loss of the network. Long-term degradation of a pipe might not be detected as a single leakage, but by analysing larger areas and compare the radiated energy from two flights separated by one or a few years, such effects can be detected. The area covering the district heating network is divided into square cells and the comparison of energy loss is done for each cell individually.

### Related work

Over the years, various methods for monitoring of district heating networks have been developed. For example, methods based on change in impedance or frequency response for a thread installed inside the insulation of the pipe. Another common method is to use liquid level switches. They measure the flow of media in the inlet and outlet and if the flow differs, the operator knows that there is a leakage somewhere along that section of the pipe. The major problem with

the above described methods is that it may not be easy to localize the leakage based on the provided information. They detect the presence but not the exact location.

Methods for large-scale monitoring of district heating systems by aerial thermography (that is, using an aircraft equipped by a thermal camera), have been investigated by Ljungberg et al. in the 80's [3],[6],[7] and Axelsson [8]. The results are somewhat antiquated due to the drastic development of thermal cameras during the last two decades. Also, ground-based thermography has been investigated using hand-held cameras [9]. Compared to aerial thermography, this has several drawbacks, such as restricted access to many areas of interest and less scalability.

The first system with automatic image analysis was presented by Friman et al. [1]. The system uses anomaly detection in order to detect abnormally warm areas along the pipes. However, the problem is the large number of false detections. To reduce the number of false detections, buildings are segmented in order to avoid detections due to, e.g., chimneys when the pipes pass under buildings.

Berg and Ahlberg [2], used the detection from [1], proposed a new building segmentation method and reduced the number of false detections even more through classification.

Regarding temporal characterization of remote sensing data, the equivalent within the field is change detection. Change detection is a common usage of remote sensing data which has been extensively studied. Applications include various kinds of environmental monitoring (e.g., land use and land cover (LULC) change, deforestation and crop monitoring; see [10] for a review of such applications), urban change [11], and military target detection [12],[13]. The employed methods usually assume multispectral, sometimes even hyperspectral data, or SAR data [11]. Methods vary greatly, depending of the type of change to be detected, and they can be pixel-based [12]-[14] or object-based [15].

### Contribution

The contribution of this paper is a system incorporating recent and novel advances in thermal monitoring of district heating networks:

1. The detection method invented by Friman et al. [1] and modified by Berg and Ahlberg [2].
2. The machine learning/classification method proposed in [2].
3. A previously unpublished method for temporal characterization and visualization of district heating network energy loss.

## STATE OF THE ART

The state of the art of detecting district heating leakages by airborne thermography is represented by the method by Friman et al. [1] mentioned above, and we use that as a foundation of our work.

Regarding image-based classification of district heating leakages, the only attempt we are aware of is our previous work [2]. The state of the art of object classification in thermal imagery in is presumably represented by pedestrian detection in the automotive industry and target recognition in the defence industry, neither very eager to publish their latest methods. Object classification in general has, however, made significant progress the last decade, with a plethora of new image feature descriptors [18] as well as classification methods such as boosting methods [20] and random forests [21].

Regarding change detection, the state of the art is very application-dependent. While, for example, Theiler [14] compares different pixel-based methods and Blaschke [15] describes the state of the art of object based methods, neither solves our problem. Pixel-based methods find the pixels that have changed since the previous data acquisition in order to point out small targets or, e.g., a land cover attribute change. Object-based methods aim at pointing out that objects have been added to (or removed from) the scene. In our case, we are interested in pointing out an attribute change of one or more objects within a certain area, and the state of the art methods are, therefore, not directly applicable.

## METHODS

The methods used in this work can be divided into four different categories. *Data acquisition*, *detection* of leakages in acquired data, reduction of false detections by *classification* and finally *temporal analysis* of energy loss. Each category is further described below.

### Data acquisition

The thermal infrared images are captured from an aircraft in the night or at dawn to avoid the effect of sun heated objects and cars etc. blocking the view of the street. At this time of day, the ground and buildings will have adopted a homogeneous background temperature and people are not as actively using different kind of vehicles as they are during the day. Furthermore, there should preferably not be any foliage or snow blocking the thermal radiation from reaching the sensor. This provides two windows for data acquisition, one during spring and one during autumn.

During the acquisition, the position, velocity and angles of the aircraft are stored in order to facilitate georeferencing. Weather stations are also placed in the area for camera calibration.

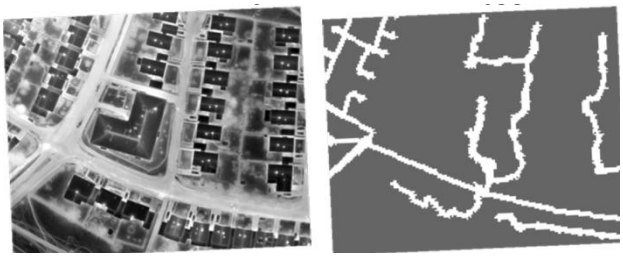

Fig. 1 Thermal image (left) together with its corresponding heat pipe mask (right). White areas within the mask represent the corridor around the network in which the detector will search for pixels with unnaturally high intensity values.

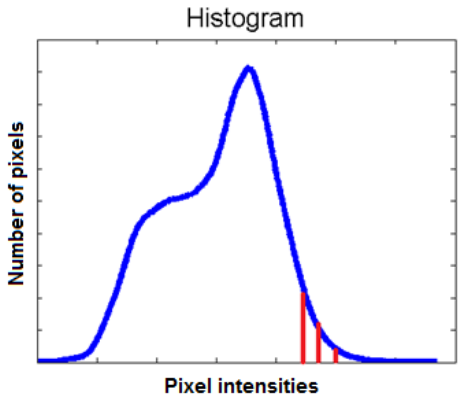


Fig. 2 Histogram of all pixel intensities within the heat pipe mask from one flight. The red, vertical lines illustrate how the histogram is thresholded in order to create several layers of detections.

The data used in this paper was acquired by a cooled mid-wave FLIR SC7000 Titanium with a resolution of 640x512 pixels and a field of view of 11°. At an altitude of 800 m, this yields a pixel footprint of 25x25cm.

The number of flights required in order to cover the whole area depends on the size of the area. For a medium-sized Swedish city, about three flights are needed.

**Detection**

Geographical information on where the heat pipes are buried is provided by the network owner. An assumption that can be made about the leakages is that they will only appear close to a pipe. Therefore, a binary mask is created to be used for detection of unnaturally warm areas. A binary mask corresponding to an image is an image with the same size as the original image and containing binary values (0 and 1) only. The 1's represent the interesting area within the original image, and the 0's represent the uninteresting parts of the image. In this case, the mask represents a corridor around the pipe network with width 2.5 meters. Within this corridor, the detector searches for unnaturally warm areas. In Fig. 1, a captured thermal image together with its heat pipe mask can be seen. Then, statistics of the radiated energy is calculated in the form of a histogram of all pixel intensity values within each heat pipe mask and flight, an example is given in Fig. 2. Since conditions may differ from one flight to another, the captured images from each flight are treated separately.

Fig. 3 Example of visualization of two layers of detections (red and yellow lines). The blue line is the district heating network.

The objective is to find areas within the mask that contain pixels with unnaturally high intensity values. This is achieved by finding the thresholds that generate the upper percentiles (0.95, 0.97, 0.99, 0.995, 0.999, and 0.9995) of the histogram. That is, we find the areas that correspond to the 5%, 3% and so on "warmest" pixels within the heat pipe mask and flight. Each percentile gives rise to a layer of detections, so that in total there are six different layers. It is worth noting that a detection is an extended object that has a shape and an intensity distribution, it is not only a point object.

When the detections are visualized to the operator, he or she can choose different colours for different layers. In Fig. 3, an example of the visualization is provided. The blue line is the district heating network and the red and yellow lines correspond to two different layers of detections where the yellow in this case is the more permissive one. Additionally, the detections are ranked based on the amount of radiated energy providing the operator with a ranking list of all detections.

**Classification**

Classification is the task of deciding to which of a set of predefined categories a new observation belongs. In this case, two categories, or classes, were used: True detections (both media and energy leakages) and false detections. The idea is to use classification to reduce the number of false detections presented to the operator.

*Classification using building masks*

The method for leakage detection described above finds a lot of areas that are unnaturally warm but not true leakages. There are plenty of such objects and phenomena present in a city. For example, ventilation

Fig. 4 An example of a building mask (right) created from a raster map (left). The white areas in the mask correspond to buildings in the map.

outlets, warm car engines, ground heating (to melt snow and ice) and energy leakages from buildings. Sometimes, the heat pipes pass beneath buildings, causing warm chimneys, ventilation outlets and atriums to appear as detections.

In order to minimize the number of false detections caused by objects on top of buildings, a building mask is created from raster map images made available by OpenStreetMap[1]. In these maps, buildings have a certain colour, a fact that can be used for thresholding based on colour and creating a binary mask. An example is given in Fig. 4.

The OpenStreetMap images are stored in GeoTIFF format which means that there is world coordinate information associated to each pixel. This fact facilitates image registration, i.e., the task of connecting pixels in the building masks to pixels in the thermal images.

*Classification using a classifier*

Classification using building masks removes some of the false detections, but far from all. Detections due to other things than buildings are still there. In order to reduce the amount even further, we have evaluated whether or not a classifier, specifically a supervised classifier, can be used for that purpose. We give an overview of the method here, and the details are published in [2].

A supervised classifier is trained by providing it labelled examples to learn from. The training can either be done offline (done once) or online (done continuously). In this work, only offline methods were considered. Part of the success of a supervised classifier is its ability to generalize, that is, draw general conclusions based on only a few observations.

The objects, detections in this case, which are to be categorized into different classes have to be described to the classifier somehow. Therefore, a feature vector is created for each detection. A *feature* is defined as "the specification of an attribute and its value" [16]. For example, a feature for a sample of the object human could be its height, shoe size or hair colour. A *feature vector* is simply a vector containing multiple features.

Table 1 Image features used for classification.

| **Feature** | **Description** |
|---|---|
| Median intensity | Median intensity within the detection. |
| Standard deviation | Standard deviation of the intensity within the detection. |
| Coverage | Ratio of the detection area inside the heat pipe mask. |
| Elongatedness | $\frac{area}{4d^2}$, where $d$ is the number of erosions [19] needed to make the detection disappear. |
| Concentricity | Measurement of how central the maximum intensity is within the detection. |
| Connected components | Number of other detections that lie within a certain radius from the detection. |
| Border average | Mean intensity within an area around the detection. |
| Distance to building | Distance from maximum intensity value to the wall of the closest building. |

If $N$ is the total number of features and the features are scalars, then the features form an $N$-dimensional feature space in which all feature vectors lie. The objective of the classifier is to, by observing labelled examples, find a decision boundary in this space. The decision boundary should then be used to correctly classify previously unseen observations. If the decision boundary is too complex, the classifier will not be able to generalize properly.

The "goodness" of 19 different scalar features were evaluated using the Mahalanobis distance [17] and a final set of 8 image features, provided in Table 1, was used for classifier evaluation. Several types of classifiers were evaluated as well, however, classifier evaluation is not covered here; only the best performing classifier used in the proposed system will be further described. See [2] for details.

The classifier used in the proposed system is the Random forest [21] classifier. Basically, it consists of a forest of decision trees. A decision tree is a hierarchal, decision structure composed of nodes and leaves. At each node, a binary test is made which leads either to another node or to a leaf, see example in Fig. 5. When reaching a leaf, the object is assigned the class label of that leaf. Which test to use in each node is decided when the tree is trained.

[1] http://www.openstreetmap.org

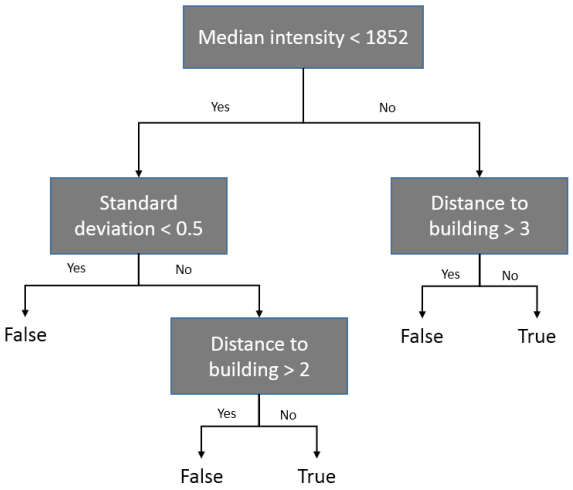


Fig. 5 Example of a decision tree. At each node, a binary test is made (the tests in this tree are examples) and at each leaf, the object is classified as true or false.

A Random forest classifier consists of multiple, random decision trees. Each tree votes on to which class the object belongs. The class assigned to the object is the class with the most votes. We use 120 trees with an average depth of 10 and splitting at nodes based on one randomly selected feature.

**Temporal analysis**

If the acquisition of thermal imagery covering a city is repeated one or a few years later, it is possible to compare the status of the network at the first acquisition with the status at the second acquisition. An automatic comparison method and a visualization technique have been developed for this purpose.

First, a grid consisting of cells, size 50×50 m, is created for the covered area, see Fig. 6. The grid has $M$ rows and $N$ columns (depending on the size of the area). For each acquisition $a$ and cell $(m,n)$, a total radiated power, $P^a_{m,n}$, is calculated ($a = 1,2$; $m = 1,2,\dots,M$ and $n = 1,2,\dots,N$).

Since we know the temperature at the ground, we can compute the radiated power for each detection. For this purpose, Stefan-Boltzmann's law

$$\frac{dQ}{dt} = \varepsilon\sigma(T^4 - T_0^4)A \qquad (1)$$

is used. $\sigma = 5.67 \cdot 10^{-8}\mathrm{W\,m^{-2}K^{-4}}$ is the Stefan-Boltzmann constant, $A$ is the area, and $\varepsilon$ is the emissivity of the object which in this case mainly consists of ground in different forms. Soil, grass and asphalt typically have an emissivity around 0.92. $T$ is the mean temperature of all pixels within the detection. $T_0$ represents the background temperature and is estimated as the mean temperature of all pixels within the current heat pipe mask and flight.

The total radiated power (TRP) of cell $(m,n)$ and acquisition $a$, $P^a_{m,n}$, is calculated as in Eq. (2). For each

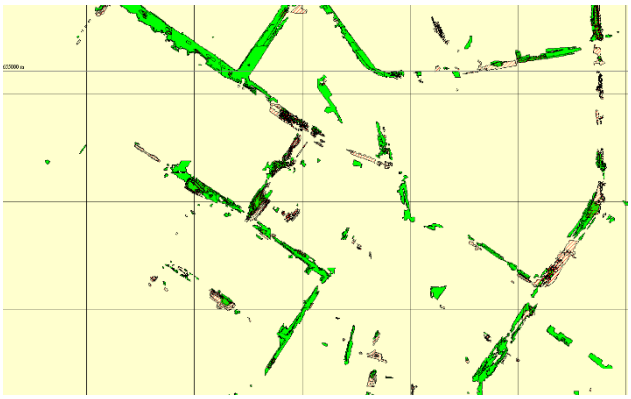

Fig. 6 A part of the grid covering the area. The detections from the old acquisition campaign are marked with red, dashed lines and the new detections are green and filled.

layer, the sum of the radiated power of all detections that have their centroids in cell $(m,n)$ is computed. The TRP of a cell is then defined as the mean of this sum over layers, i.e.,

$$P^a_{m,n} = \frac{\sum_l \sum_k \varepsilon\sigma(T_k^4 - T_0^4)A_k}{L}. \qquad (2)$$

$k \in S_{m,n,l,a}$ where $S_{m,n,l,a}$ is the set of detections in acquisition $a$ and layer $l$ that have their centroids inside the boundaries of cell $(m,n)$. Finally, $l = 1,2,\dots,L$ is the layer and, as mentioned above, the total number of layers in this case is $L = 6$.

The difference of radiated power, $\Delta_{m,n}$, for each cell represents the change in radiated power from the previous acquisition until the current one. It is calculated as the difference of TRP's for the two acquisitions $\Delta_{m,n} = P^1_{m,n} - P^2_{m,n}$.

When the comparison results are visualized to the operator, each cell is colored according to its calculated TRP difference, $\Delta_{m,n}$, and the grid is overlaid on top of a mosaic of the thermal images, see Fig. 7. Red indicates an increase in radiated power and green indicates a decrease. Transparency is used to visualize the degree of change. The cells with the largest increase/decrease are assigned zero transparency and the cells with lowest full transparency. The transparency scale is then linearly distributed for all cells with TRP differences in between. If there is no change at all, i.e. $\Delta_{m,n} = 0$, the cell will be colorless.

**Evaluation methodology**

As mentioned in the description of supervised classifiers above, the methods need labelled examples, or ground truth samples, for training. That is, information on which detections that have proven to be true and false respectively. The system, as described in [1], has been used in over 20 different Nordic district heating networks in recent years, and three of these were chosen for sample collection. These networks were chosen based on their ability to provide ground

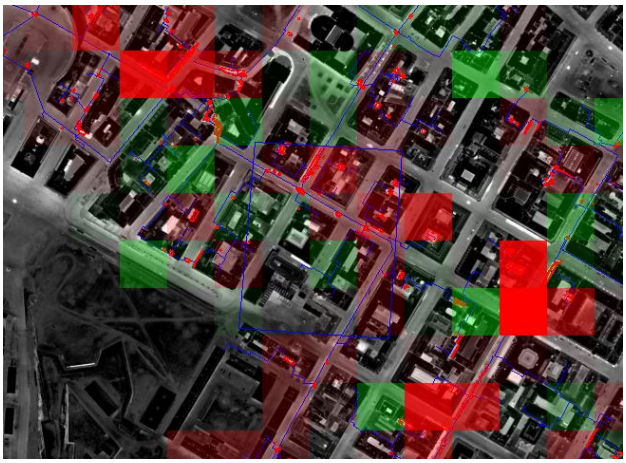

Fig. 7 Visualization of changes in radiated power. A red square indicates that the area suffers from an increased energy loss while a green square means that the energy loss within the area has decreased.

Table 2 Number of ground truth samples for each layer and class.

| **Layer no.** | **1** | **2** | **3** | **4** | **5** | **6** |
|---|---|---|---|---|---|---|
| Threshold | 0.05% | 0.1% | 0.5% | 1% | 3% | 5% |
| True | 34 | 39 | 71 | 89 | 99 | 80 |
| False | 71 | 75 | 148 | 237 | 294 | 348 |

truth samples. The distribution of samples can be seen in Table 2. In order to evaluate the generalization ability of the classifier, a method called 10-fold cross validation [20] was used. The samples are split into 10 different folds. Then, the classifier is trained 10 times, each time using 9 different folds and validated using the 10th. In order to achieve reliable evaluation results, the validation data has to be previously unseen by the classifier.

The confusion matrix [16], Fig. 8, is a common way to describe the different kinds of errors that appear when performing classification. We want to minimize the false positive rate while maintaining a true positive rate of 99%. The limit of 99% comes from the fact that the cost of misclassifying a true detection is much higher than misclassifying a false one.

## RESULTS

### Classification

19% of the false alarms could be removed solely with the use of building masks. In Fig. 9, an example of a false detection that has been removed with the help of the building based segmentation is provided. The false positive rate of the Random forest classifier combined with the building segmentation based classification is 42% when samples from all layers were combined into one dataset. That is, 58% of the false detections can be removed while maintaining a true positive rate of

| | | **Predicted** | |
|---|---|---|---|
| | | **True detection** | **False detection** |
| **Actual** | **True detection** | True positives (true detections correctly classified) | False negatives (True detections that were incorrectly classified as false) |
| | **False detection** | False positives (false detections that were incorrectly classified as true) | True negatives (False detections correctly classified) |

Fig. 8 A confusion matrix, a common way to visualize the performance of a supervised classification method.

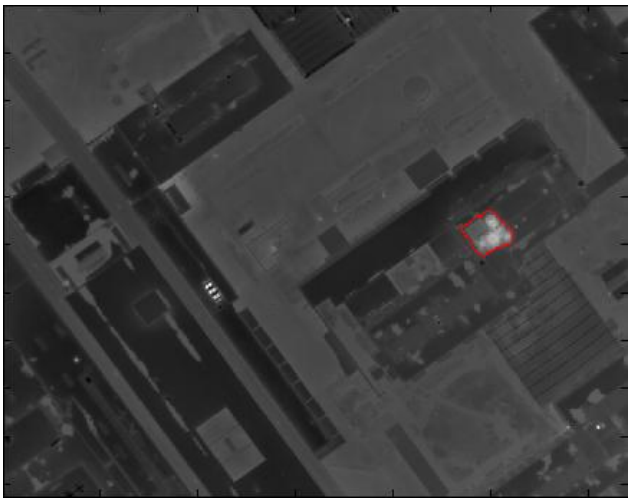

Fig. 9 Example of a false detection (marked with red boundaries) that has been removed with the help of a building mask.

| | | **Predicted** | |
|---|---|---|---|
| | | **True detection** | **False detection** |
| **Actual** | **True detection** | True positives: 408 | False negatives: 4 |
| | **False detection** | False positives: 493 | True negatives: 680 |

Fig. 10 The confusion matrix containing the final results.

99%. The confusion matrix from Fig. 8 has in Fig. 10 been filled out as an illustration of the final results. In a typical medium-sized Swedish city, there are typically around 3000 detections in the 0.995 percentile layer. Among these 3000, only about 10 are true media leakages. Thus, being able to reduce the number of false detections with 58% greatly reduces the workload for the operator.

### Temporal analysis

The temporal analysis had, at the time of writing, been used in the described form in one city. Unfortunately,

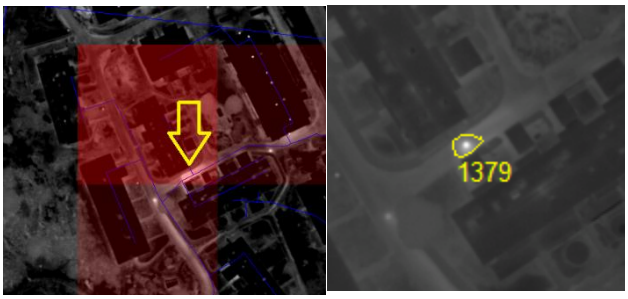


Fig. 11 A cell (left) that has been marked as suffering from an increased energy loss that proved to contain a true media leakage (right). The yellow arrow indicates the position of the leakage within the cell and the yellow number is the ID of the detection.

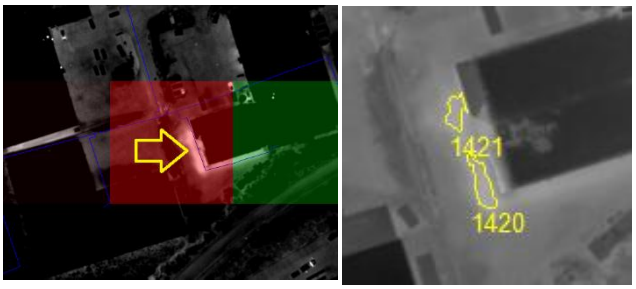


Fig. 12 Another example of a cell (left) that has been marked as suffering from an increased energy loss that proved to contain a true media leakage (right). The yellow arrow indicates the position of the leakage within the cell and the yellow number is the ID of the detection.

that city was not one of the three for which ground truth samples for classifier evaluation had been collected. However, some confirmed leakages have been provided by the network owner allowing us to draw some conclusions about the result. In Fig. 11 and Fig. 12, examples are shown of how the visualization of the comparison clearly indicates cells containing confirmed media leakages as suffering from an increased energy loss.

One particularly interesting example where the comparison acts as a complement to detection ranking based on radiated power is presented in Fig. 13. Here, a major media leakage of 70 $m^3$/day gave rise to some headaches for the network owner who had searched unsuccessfully for the leakage for several years. The district heating pipe laid on top of a bed of gravel and beneath the said pipe was a larger stormwater pipe along which all media from the district heating pipe ran. Due to this choice of path by the media, the temperature difference that could be measured at ground level was only 3°C, placing the detection far down the detection ranking list. Nevertheless, the change in radiated power since the last acquisition was noticeable and in the visualization the cell containing the detection was marked in red.

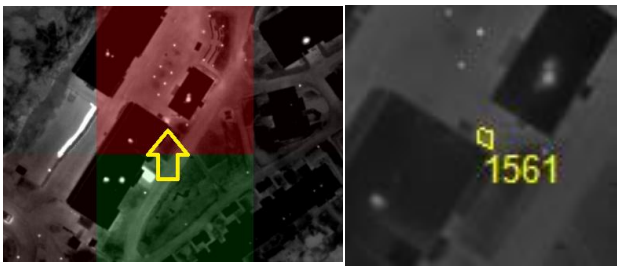


Fig. 13 Example of major leakage made visible by the comparison acting as a complement to detection ranking based on radiated power.

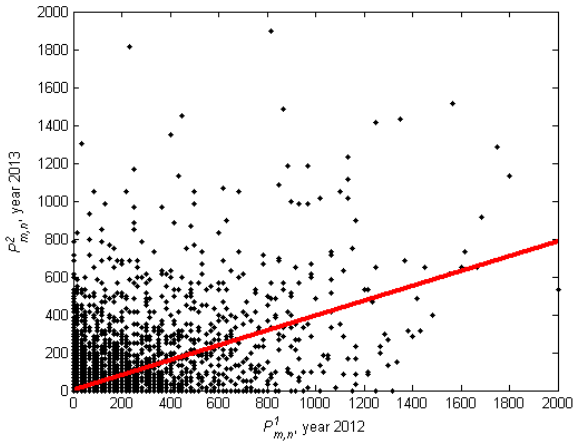


Fig. 14 Scatter plot of the TRP's for each cell (black dots). The red line is a line that has been fitted to the points. It indicates whether or not the overall energy loss of the network has decreased or increased.

In order to make an assessment of the overall status of the network at the current versus the previous acquisition, a scatter plot as the one in Fig. 14, can be used. Each scatter point corresponds to one cell and each cell has two TRP's, $P^1_{m,n}$ and $P^2_{m,n}$, here the x- and y-axis respectively. A red line has been fitted to the points through a least squares fit. If the energy loss has not changed, the angle $\alpha$ between the line and the x-axis is 45°. A smaller $\alpha$, as in Fig. 14, indicates that in general, there has been a decrease in energy loss. In fact, this conclusion coincided with the network owner's feeling about the network's status.

## DISCUSSION

### Classification

With the use of the building segmentation based classification, 19% of the false detections could be removed. There could, however, be a bias present among the ground truth samples since false detections on top of buildings are easier to find than other kind of false detections and thus a larger percentage of such false detections might have been labelled. If so, the false detection reduction rate, in reality, is lower than 19%. However, this has not been further investigated. Also, the number of detections on top of buildings

varies between different cities depending on how much pipes that actually pass beneath buildings.

Furthermore, in the maps generated by OpenStreetMap it has been observed that there sometimes are missing buildings. The opposite error, buildings present in the map but not in reality, has not been observed. Missing buildings is a kind of error that in this case is quite forgiving, since it only leads to false positives, i.e., false detections incorrectly classified as true ones. The opposite, "false" buildings, could result in true detections being classified as false ones (false negatives), which is an unwanted scenario since the cost for false negatives is much higher than for false positives.

### Temporal analysis

Regarding temporal analysis, it should be emphasized that the presented method for temporal analysis is an approximation. It provides a measurement of the radiated power at ground level, but how the heat transfers from the pipe through the soil remains unexplored. It is, however, clear that the properties of the pipe, material, depth, insulation etc., and the soil composition affects how much radiated power that reaches the ground surface.

The presented visualization technique with red and green squares gives the operator a quick overview of the status of the network. He or she can soon pinpoint the most critical areas.

## OUTLOOK

We will continue the work on improving false detection reduction, temporal analysis and quantization of energy loss with the goal of providing the operator with an even more accurate tool for large-scale monitoring of district heating networks.

## CONCLUSIONS

In this paper, a complete system for large-scale monitoring of district heating networks has been presented. Methods for media/energy leakage detection in thermal images and reduction of false detections through classification have been described and a method for temporal analysis of energy losses has been proposed.

The system allows the operator to get a quick overview of the status of the complete network and can be used as a complementing tool for maintenance planning.

The proposed temporal analysis improves usability of the system and the visualization allows the operator to get a quick overview of what areas that should be studied more carefully.

## ACKNOWLEDGEMENT

This work has been partly funded by the Swedish Research Council (Vetenskapsrådet) through the project "Learning systems for remote thermography", grant no. 621-2013-5703.

## REFERENCES

[1] O. Friman, P. Follo, J. Ahlberg, and S. Sjökvist, "Methods for large scale monitoring of district heating systems using airborne thermography," in IEEE Trans. Geoscience and Remote Sensing, Vol. 52 , No. 8, pp. 5175-5182, 2014.

[2] A. Berg and J. Ahlberg, "Classification of leakage detections acquired by airborne thermography of district heating networks," in Proc. 8th International Workshop on Pattern Recognition in Remote Sensing (PRRS 2014), Stockholm, August 2014.

[3] M. Olsson, Long-term thermal performance of polyurethane-insulated district heating pipes, Ph.D. thesis, Chalmers Univ. of Techn. 2001.

[4] M. Fröling, Environmental and thermal performance of district heating pipes, Ph.D. thesis, Chalmers Univ. of Techn., 2002.

[5] S. A. Ljungberg, "Aerial thermography – a tool for detecting heat losses and defective insulation in building attics and district heating networks," in Proc. SPIE Thermosense IX 1987, pp. 257–265.

[6] S. A. Ljungberg, "Thermography for district heating network applications: operational advantages and limitations," in Proc. SPIE Thermosense X 1988, pp. 70–77.

[7] S. A. Ljungberg and M. Rosengren, "Aerial and mobile thermography to assess damages and energy losses from buildings and district heating networks – operational advantages and limitations," in Proc. 16th Congress Int. Soc. Photogrammetry and Remote Sensing 1988, pp. 348–359.

[8] S. R. J. Axelsson, "Thermal modelling for the estimation of energy losses from municipal heating networks using infrared thermography," in IEEE Trans. Geoscience and Remote Sensing, Vol. 26, No. 5, pp. 686-692, 1988.

[9] B. Bøhm and M. Borgström, "A comparison of different methods for in-situ determination of heat losses from district heating pipes," Dept. of Energy Engineering, Technical Univ. of Denmark, 1996.

[10] D. Lu, P. Mausel, E. Brondízio, and E. Moran, "Change detection techniques," Int. Journal of Remote Sensing, Vol. 25, No. 12, pp. 2365-2407, 2003.

[11] X. Li et al., "New approaches to urban area change detection using multitemporal RADARSAT-2 polarimetric synthetic aperture radar (SAR) data," Canadian Journal of Remote Sensing, Vol. 38, No. 3, pp. 253-266, 2012.

[12] J. Meola et al., "Detecting Changes in Hyperspectral Imagery Using a Model-Based Approach," IEEE Trans. Geoscience and Remote Sensing, Vol. 49, No. 7, pp. 2647-2661, 2011.

[13] R. J. Dekker et al., “LWIR Hyperspectral Change Detection for Target Acquisition and Situation Awareness in Urban Areas,” Proc. SPIE, Vol. 8743, paper 874306-1, 2013.
[14] J. Theiler, “Quantitative comparison of quadratic covariance-based anomalous change detectors,” Applied Optics, Vol. 47, No. 28, pp. F12-F26, 2008.
[15] T. Blaschke, “Object based image analysis for remote sensing,” ISPRS Journal of Photogrammetry and Remote Sensing, Vol. 65, pp. 2-16, 2010.
[16] R. Kohavi and F. Provost, “Glossary of terms,” Machine Learning, Vol. 30, No. 2, 1998.
[17] A. Jain and D. Zongker, “Feature selection: Evaluation application and small sample performance,” IEEE Trans. Pattern Analysis and Machine Intelligence, Vol. 19, No. 2, pp. 153-158, 1997.
[18] R. Szeliski, Computer Vision, Algorithms and Applications, Springer, 2011.
[19] R. C. Gonzalez and R. E. Woods, Digital image processing, Pearson Education, 3rd int. ed., 2008.
[20] T. Hastie, R. Tibshirani, and J. Friedman, The elements of statistical learning, 2nd ed., Springer, 2008.
[21] L. Breiman, “Random forests,” Machine Learning, Vol. 45, No. 1, pp. 5-32, 2001.