

 Open access • Journal Article • DOI:10.1109/TGRS.2013.2287238

Methods for Large-Scale Monitoring of District Heating Systems Using Airborne Thermography — [Source link](#)

Ola Friman, Peter Follo, Jörgen Ahlberg, Stefan Sjøkvist

Institutions: Swedish Defence Research Agency

Published on: 01 Aug 2014 - IEEE Transactions on Geoscience and Remote Sensing (IEEE)

Related papers:

- [Thermal modeling for the estimation of energy losses from municipal heating networks using infrared thermography](#)
- [Enhanced analysis of thermographic images for monitoring of district heat pipe networks](#)
- [Environmental and Thermal Performance of District Heating Pipes](#)
- [Long-Term Thermal Performance of Polyurethane-Insulated District Heating Pipes](#)
- [The Elements of Statistical Learning](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/methods-for-large-scale-monitoring-of-district-heating-53me86zbcv>

Methods for Large-Scale Monitoring of District Heating Systems Using Airborne Thermography

Ola Friman, Peter Follo, Jörgen Ahlberg and Stefan Sjökvist

Linköping University Post Print



N.B.: When citing this work, cite the original article.

©2009 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

Ola Friman, Peter Follo, Jörgen Ahlberg and Stefan Sjökvist, Methods for Large-Scale Monitoring of District Heating Systems Using Airborne Thermography, 2014, IEEE Transactions on Geoscience and Remote Sensing, (52), 8, 5175-5182.

<http://dx.doi.org/10.1109/TGRS.2013.2287238>

Postprint available at: Linköping University Electronic Press

<http://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-105285>

Methods for Large-Scale Monitoring of District Heating Systems using Airborne Thermography

Ola Friman, Peter Follo, Jörgen Ahlberg, Stefan Sjökvist

Abstract—District heating is a common way of providing heat to buildings in urban areas. The heat is carried by hot water or steam and distributed in a network of pipes from a central power plant. It is of great interest to minimize energy losses due to bad pipe insulation or leakages in such district heating networks. As the pipes generally are placed underground, it may be difficult to establish the presence and location of losses and leakages. Towards this end, this work presents methods for large-scale monitoring and detection of leakages by means of remote sensing using thermal cameras, so-called airborne thermography. The methods rely on the fact that underground losses in district heating systems lead to increased surface temperatures. The main contribution of this work is methods for automatic analysis of aerial thermal images to localize leaking district heating pipes. Results and experiences from large-scale leakage detection in several cities in Sweden and Norway are presented

I. INTRODUCTION

A district heating system uses a centralized power plant from which heat is distributed to a number of buildings. This solution is commonly found in urban residential and commercial areas on the northern hemisphere as it offers advantages in terms of energy efficiency, pollution control and administrative effort compared to more localized heat generation solutions [1]. The heat is carried from the central power plant by hot water or steam with temperatures in the range 90-150°C (200-300°F). The distribution pipes degenerate with time [2], e.g., due to corrosion, degradation of the insulation material or motions in the ground. As district heating systems in some cities have been in use for many decades, it is of economic and environmental interest to monitor these networks to detect energy losses due to poor insulation or broken pipes where toxic water or steam leaks [3]. Major leakages of 50 m³ to 150 m³ of water or more per day may also cause the ground to collapse due to erosion, whereby large amounts of water at boiling temperature are exposed.

Coarse monitoring to find losses in district heating systems is possible by measuring how much water leaves the centralized power plant and how much arrives at each building, but exact localization of leakages is difficult because the pipes are generally placed underground. Leakages can therefore go unnoticed for long times, or the existence of a leakage may be suspected but its exact location has not been identified. Searching for leakages manually requires digging up the pipes for inspection, a costly and tedious procedure, especially if it has to be repeated several times in order to locate the leakage.

This work presents methods for large-scale monitoring of district heating systems by means of remote sensing from an aircraft using a thermal camera [4], [5], [6]. Such airborne or aerial thermography has found many applications in which

the ground temperature is of interest, e.g., search and rescue operations, fire fighting, mining and geophysical applications [7], [8], studies of urban heat islands [9] and inspection of building insulation and rooftop conditions [10]. In the context of district heating, underground leakages of hot water or steam cause increased temperatures at the surface level that can be measured with the thermal camera, i.e., warmer regions manifest themselves as bright regions in the thermal images [11]. Increases of the surface temperature with more than 20 °C over major leakages have been observed in the thermal images in the present work, but the increase is dependent on the size of the leakage. It has been shown that it is possible to register the surface temperature differences around a leakage using a hand held thermal camera at the ground level [12], [13], but access to non-public areas such as residential gardens and industrial sites may be restricted. In contrast, imaging from an airplane gives both the necessary overview and the possibility to perform large-scale monitoring.

Large-scale monitoring means that an entire city is traversed by the airplane in a predefined pattern while acquiring the thermal images. This results in tens of thousands of aerial images, see Fig. 1. The main contribution of this work is methods for automatic analysis of such images to localize leaking district heating pipes. Results and experiences from large-scale leakage detection in several cities in Sweden and Norway are presented.

II. DATA

This section describes the data acquisition process and pre-processing of the thermal images.

A. Image acquisition

Thermal images covering the district heating system of interest are acquired from an airplane that follows a predefined trajectory, see Fig. 1e. The images in this work were acquired from an altitude of about 800 meters. The camera is triggered at GPS control points so as to ensure full coverage of the area. To minimize the number of irrelevant temperature contrasts, the image acquisition is preferably performed at night or dawn. At this time, the effect of sun heating is minimal and most ground objects have adopted a homogeneous background temperature. For example, cars are generally parked and cool, compare Fig. 1a and Fig. 1d. Some heat sources not related to the district heating system are unavoidable, such as chimneys and heat leakage from houses, see Fig. 1c. At dawn, even before the sun has risen above the horizon, building walls will be heated due to atmosphere scattering. These sources

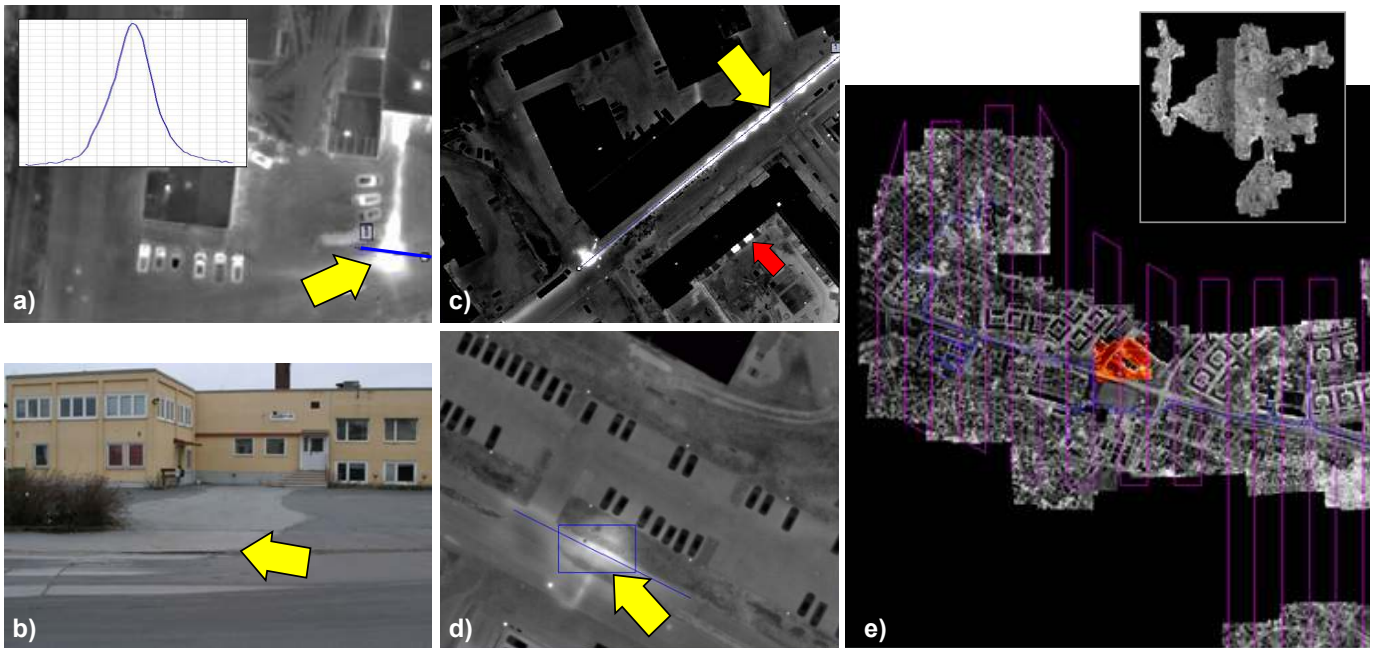


Fig. 1. **a)** An example leakage in an aerial thermal image is marked by the yellow arrow. The temperature profile across the leakage is shown in the inset plot. This thermal image was acquired during daytime and one can note that the parked cars in the image are warm. **b)** A photograph of the leakage area in **a)**. The surface is dry due to the underground leakage. **c)** Another example leakage. The bright regions indicated by the red arrow are heat leakages from a building. **d)** Another example leakage. This image was acquired during night and one can note that the parked cars are cool, as opposed to the cars in **a)**. **e)** The thermal images are acquired along a predefined trajectory (purple line) over a city. After georectification of the images, image mosaics of thousands or tens of thousands of images are generated for simple navigation and exploration of the data. The area shown in red-yellow colorscale shows the size of one image.

of heat may cause false alarm detections. Furthermore, the ground should be free of any insulating layer of snow and there should be no leaves or foliage covering the line of sight to the ground surface. Over the northern hemisphere, where district heating is mostly used, this means that there are two time windows during the spring and the fall where there is no snow and no foliage. Ideally, the ground should also be dry and the wind speed low.

If the area of interest is large, it may not be possible to cover it in one night. Several acquisitions are then made over several nights. As the environmental conditions will be different for these acquisitions, e.g., different background temperatures, each acquisition must be treated separately in the ensuing automatic analysis. In the inset mosaic image in Fig. 1e, the striping effect stems from different acquisitions with slightly different image contrasts.

At the beginning of each acquisition, the airplane passes over calibration panels placed on the ground. The emissivity of the panels is known and their surface temperature is measured with a contact method. Weather parameters such as temperature, humidity and wind speed are also measured. Using this information, atmospheric effects can be estimated and the image intensities can be translated to a temperature scale¹. While the methods for leakage detection presented below are invariant to intensity scaling, such calibration enhances the interpretability of the image contrasts and it facilitates change detection and quantification between several overflights of the

same area.

The thermal camera used in this work is a cooled FLIR SC7000 Titanium which operates in the mid-wave IR domain ($3.0 \mu\text{m} - 5.0 \mu\text{m}$). The sensor size is 640×512 pixels and the FOV 11° . At an acquisition altitude of 800 m, this yields a pixel footprint size of 24 cm.

B. Georeferencing

The thermal images are georectified using standard methods based on GPS and IMU measurements acquired simultaneously with the images. An image mosaic is built from the georectified images. This mosaic is stored in a scale-space fashion to facilitate interactive exploration by allowing quick navigation and zooming in the images. Detection of leakages is always performed in the original images, however, as they have the highest resolution and quality.

C. Auxiliary GIS data of the district heating network

Cities with district heating systems usually maintain a blueprint of the pipes in a Geographic Information System (GIS) format. This information greatly facilitates the detection of leakages as the search area can be narrowed down to an area above the pipes. A large portion of the false alarm detections is thereby rejected. Using the georeferencing information of the images, the pipes can be projected into the original images and the search area defined by rasterizing an image mask around the pipes that extends 5 pixels on each side, see Fig. 2.

¹The software *Altair* from *FLIR Systems, Inc* was used for this purpose in this work.

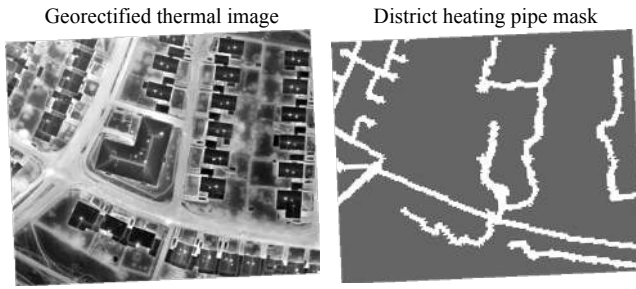


Fig. 2. An example of a georectified thermal image and a pixel mask of the district heating pipes obtained by rasterizing a GIS layer of the district heating system.

III. METHODS

A. Modeling leakages

The detection of district heating leakages relies on a surface temperature contrast induced by the underground leakages. It is possible to model the temperature field and heat flux around a buried heating pipe, and thereby predict the temperature profile at the surface level above a leakage [11]. Such modelling and empirical observations show that an underground heat loss manifests itself as an area with a smooth bell-shaped temperature variation profile at the surface level, see Fig. 1a. This information is potentially useful as a distinguishing feature of district heating leakages. However, the exact temperature profile shape depends on a number of parameters and assumptions, such as the thermal conductivity and humidity of the soil, a homogeneous ground volume around the pipe, and surface conditions such as air temperature, wind and surface cover. In addition, atmospheric effects and the surface emissivity affect the radiance measured by the thermal camera. Finally, a leakage may be partially occluded by a car or any other object, introducing discontinuities in the measured temperature profiles from the airborne platform.

For the above reasons, it is very difficult to formulate a useful compact model of the spatial shape and temperature profile induced by a real leakage, and to use it to discriminate against other warm areas seen in the images. The problem of detecting leakages is therefore treated as an anomaly detection problem instead. In this approach, a model of normal temperature variations is created from the acquired images and outliers with respect to this model are detected as potential leakages that should be investigated further. The pixel intensity probability distribution $p_T(x)$, i.e., the probability that the pixel value T falls in an infinitesimal interval around the value x , for the ground surface above the district heating pipes is used as a model of normal variations, see Fig. 3. For the purpose of detection, it does not matter if $p_T(x)$ is expressed in the raw radiance unit delivered by the camera or if the intensity scale of the images has been calibrated to apparent temperature.

To estimate $p_T(x)$, the aerial thermal images and a pixel mask of the pipe locations are required. The latter is obtained by projecting the GIS layer of the district heating system into the acquired images as described in Section II-C. Real leakages are rare events so that $p_T(x)$ can be non-parametrically estimated, without significant bias, by the histogram of all




Fig. 3. Distributions of all pixel intensities (dash-dotted blue line) across thousands of images in one acquisition, and the distribution $p_T(x)$ of the pixels in a mask above the district heating pipes (solid red line). Thresholds to detect candidate leakages are found as percentiles of $p_T(x)$, i.e., as abnormal temperatures, which is illustrated by the filled areas. Histograms for acquisitions on two consecutive nights are shown and the thresholds are adapted to the acquisition conditions.

pixels under the pipe mask across all images in one acquisition. As it is generally not possible to hold all thousands of images in the computer memory simultaneously, the histogram is estimated iteratively in a two-pass process: In a first pass over the images, the pixel intensity mean and standard deviation are determined iteratively using one image at the time. Based on this information, suitable histogram bin centers are derived. In this work, 1000 bins evenly distributed in a range ± 5 standard deviations from the mean are defined. The second pass over the images then accumulates the total histogram image by image.

The estimation of $p_T(x)$ should be made using images acquired under similar conditions. A city is typically imaged in a number of acquisitions over several days. Each acquisition requires its own $p_T(x)$ as the atmospheric conditions and air temperature generally will differ, leading to different distributions of the normal surface temperature. If the conditions change within an acquisition, it is possible to estimate different $p_T(x)$ for images within the acquisition, but this is usually not necessary. As a final note, $p_T(x)$ will differ somewhat from the distribution of *all* pixel intensities, as $p_T(x)$ is restricted to the area above the district heating pipes, and some dissipation from the pipes is unavoidable and normal, and because the pipes are usually drawn under streets which tend to be somewhat warmer than vegetation and building roofs. This is illustrated in Fig. 3.

B. Detecting leakages

District heating leakages are assumed to cause surface temperatures that are larger than normal surface temperatures, i.e., leakages are found in the upper tail of $p_T(x)$. To detect potential leakages, a threshold defined by a certain percentile of $p_T(x)$ is therefore used. For example, one may mark the warmest 1% of all pixels over the pipes as outliers, representing candidate leakage regions. The threshold obtained in this fashion is objective and adapted to the imaging conditions, which is exemplified in Fig. 3. This is in contrast to human screening of individual images, which is superior in terms of recognition of local context and scene analysis, but which lacks the objective information across all tens of thousands of images that is encoded in $p_T(x)$.

The pixels with extreme temperatures identified with the percentile threshold are typically clustered in regions. Further region-based analysis is therefore applied to filter out warm regions that are unlikely to stem from leakages. One may incorporate discriminating features describing the temperature profile over each detected region using for example shape, blobness and intensity gradient features. However, as discussed in Section III-A, due to the large variation of leakage shapes and locations, e.g., a leakage under a parked car will appear to have both an uncharacteristic shape as well as apparent steep edge in the temperature profile, it is difficult to find features that uniquely characterize real leakages. To avoid that real leakages are filtered out, the only additional feature considered in this work is the area of the detected candidate leakage regions. A ranking of the candidate regions is produced as the product of the region intensity and region area, which gives a notion of the energy loss in the region.

The ranked candidate leakage regions are manually screened by a human operator who is familiar with the district heating system at hand. Most of the candidate detections will be false alarms but these are generally easily rejected by a human using the scene context, e.g., heat leakage from houses, wells, garages etc. The expert operator also has background knowledge of where leakage detections are normal and expected (culverts, constructions sites etc.) and where they can be suspected due to the age of the pipes or losses indicated by coarse monitoring systems. It should be emphasized that the automatic leakage detection is a support tool for an efficient detection process compared to visually screening tens of thousands of images without any hints of the locations of potential leakages.

C. Automatic building segmentation

Many false alarm detections arise around buildings, for example from warm chimneys, vents and entrances. Buildings are therefore important context carriers for a human operator when rejecting false alarms, and an important step to further improve the automated analysis is to segment buildings in the images. Automatic building segmentation in remote sensing images has been investigated previously in the research literature [14], frequently with the aid of an additional 3D laser scan that provides an accurate digital surface model [15]. An additional 3D scanning sensor facilitates the segmentation greatly but also incurs higher costs and complexity of the data acquisition.

Here, an approach to segmenting buildings based only on the airborne thermal images is presented. The building rooftops cool off faster than the ground at night and therefore appear dark in the images, see Figs. 1, 2, 4 and 7. The overall approach to the building segmentation is to first divide an image into homogeneous regions and then to classify each region as representing a building or not. The algorithm consists of the following three basic steps:

- 1) Calculate an edge image.
- 2) Apply a Watershed transform to produce image regions.
- 3) Extract features and classify the regions.




Fig. 5. Histograms of the image gradient magnitude for all pixels in two acquisitions. Adaptive thresholds are found at the 75% percentiles as shown by the shaded areas.

The steps are visualized in Fig. 4. In step 1, there are many different methods for calculating an edge image, e.g., gradient-based methods, the Canny edge detector or phase congruency methods [16]. Any of these methods would be applicable. In this work, the magnitude of the image gradient calculated using a Derivative of Gaussian filter is used to find edges. The main problem in step 1 is to automatically decide when the gradient in a pixel is large enough to stand out from the noise, i.e., when a pixel should be classified as an edge. Using same idea as for the image intensity above, this is solved by calculating the histogram of the gradient magnitude across all thousands of images in an acquisition, see Fig. 5. The temperature difference between buildings and the ground is larger than most other gradients in the images, so that a threshold can be derived automatically from this histogram, e.g., as the 75% percentile. This approach also adapts to scaling and the unit of the pixel intensity in the images.

In step 2, a Watershed transform is applied to the edge image to decompose it into a large number of regions [17]. In the last step 3, the goal is to classify each region as a building or not. To this end, the following image intensity and shape features are extracted from each region:

- Median pixel intensity in the region
- Standard deviation of the pixel intensities in the region
- Area of the region
- Region compactness
- Border regularity

Buildings in the thermal images are characterized as larger homogeneous areas (i.e. low intensity variation) with low intensity and a high border regularity. The region compactness is the area of the region divided by its squared perimeter length, also known as the isoperimetric quotient. This shape feature has its maximum for a circular region, but also quadratic building regions respond strongly compared to vegetation regions, which typically are more irregular and less compact. The border regularity feature is based on the curvature of the region perimeter. Buildings generally consist of straight walls with a few corners, so that the edge curvature is low except at a few places. The feature is calculated by considering edge normal direction (i.e., the image gradient direction) $\hat{\mathbf{g}}_i, i = 1, \dots, N$ at N regularly spaced points along the region contour. Let the angle between two vectors be denoted by

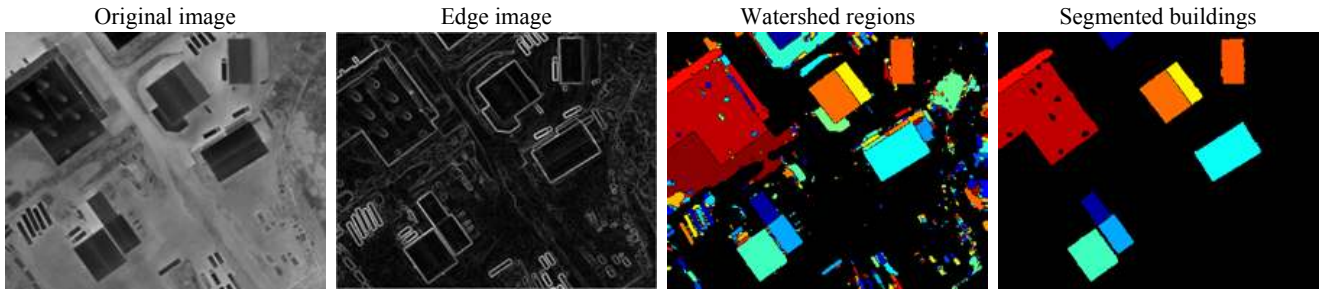


Fig. 4. Overview of the steps from the original thermal image to a building mask. The first step is to calculate an edge image of the original image. A Watershed transform is then applied to the edge image to divide the image into homogeneous image regions. Features are extracted from each of these regions and they are classified into the categories building or non-building.

$\angle(\hat{\mathbf{g}}_i, \hat{\mathbf{g}}_j)$ and define the indicator function

$$I(\hat{\mathbf{g}}_i, \hat{\mathbf{g}}_j) = \begin{cases} 1 & \text{if } \angle(\hat{\mathbf{g}}_i, \hat{\mathbf{g}}_j) < \alpha, \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

The parameter α defines how much two neighboring edge gradients along the region border are allowed to differ and still be considered as regular. In this work we use $\alpha = 10^\circ$. Finally, the border regularity feature is calculated as

$$\frac{1}{N} \sum_{i=1}^N I(\hat{\mathbf{g}}_i, \hat{\mathbf{g}}_{i+1}). \quad (2)$$

Put in words, the border regularity feature is a number between 0 and 1 that equals the fraction of the region border that is considered regular.

Finally, a classifier must be trained to classify the regions based on the extracted features. Training data was obtained by manually delineating about 500 buildings in a batch of training images. Regions in the Watershed transform of these images that coincide with the manually drawn areas were taken as examples of building regions. Training examples for non-building regions were similarly taken as Watershed regions outside the manually drawn areas. There were many more examples of non-building regions than of buildings, and this imbalance was accounted for in the training of the classifier. Distributions of some of the features for the building and non-building training examples are shown in Fig. 6. There are several possible non-linear classifier types that can be trained using this data, e.g., Support Vector Machines, neural networks and boosting-based methods [18]. There is no particular reason to believe that either of these should perform significantly better or worse. After some experimentation with Random Forest and Adaboost classifiers [19], with similar results, the Adaboost classifier was used due to its slightly simpler construction.

IV. RESULTS

Airborne thermography acquisitions over 15 cities in Sweden and Norway were performed to optimize the acquisition methodology and develop the algorithms presented in this work. The image analysis methods presented above were implemented in standalone software for image analysis, presentation and interaction. The imaged cities were of different sizes with populations between 5,000 to 500,000. The district

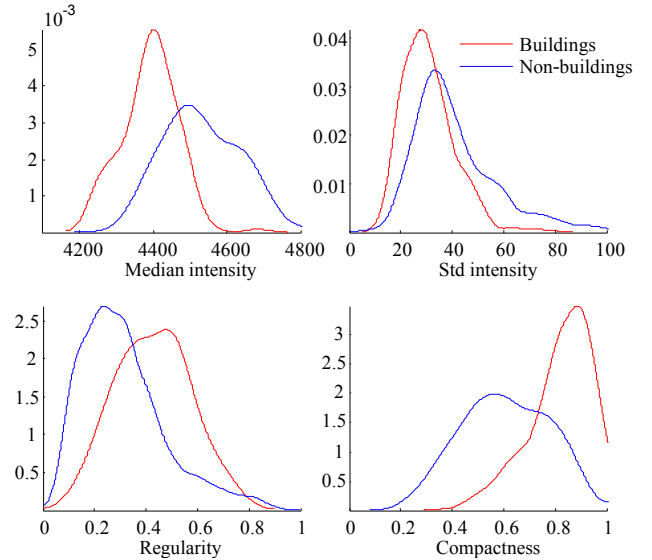


Fig. 6. Distributions of feature values for buildings and other regions respectively. Buildings generally have a low homogeneous temperature and high regularity compared to other objects in the images.

heating systems in these cities comprise tens of kilometers of pipes each. Up to 50,000 thermal images were acquired to cover the district heating system in the larger cities. Comprehensive ground truth results to validate candidate leakage detections in the images are not available for all cities, as this requires extensive digging and feedback from network operators and city planning departments which are not part of the current research project. Partial ground truth results are available though, as presented below.

A. Leakage detection

Figures 7 and 8 illustrate typical detection results. The experience has shown that three detection thresholds corresponding to the 0.005%, 0.05% and 0.5% of the warmest pixel percentiles are suitable for the interactive exploration of the images. Table I presents results for three cities from which detailed feedback and ground truth analysis have been obtained. There are about 1 to 2 candidate detections per kilometer pipe at the 0.05% significance level. Most of these are either false alarms or detections caused by bad pipe insulation that may need repair but which do not represent

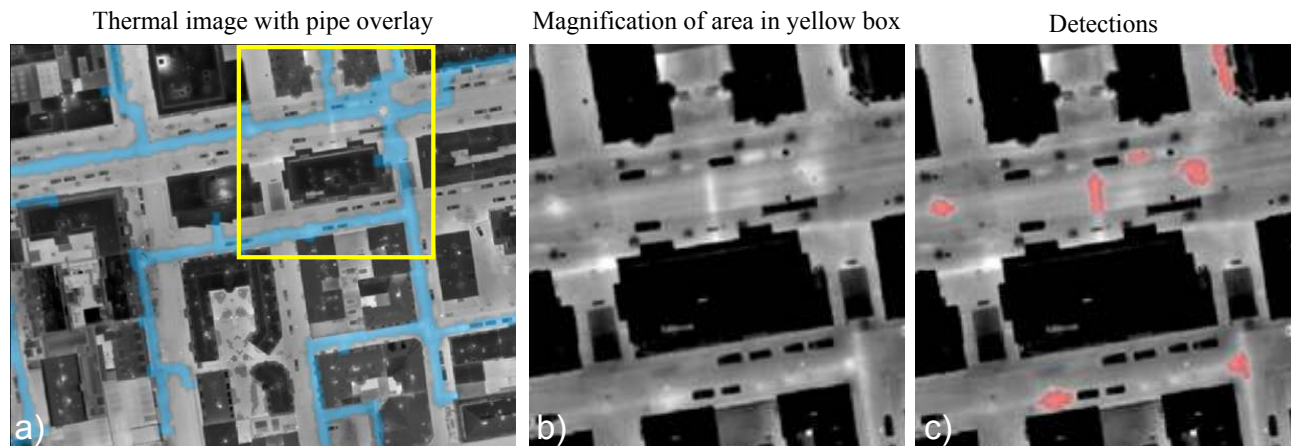


Fig. 7. a) A thermal image (640×512 pixels) of a typical city block. The blue overlay indicates a mask above the district heating pipes obtained from a Geographic Information System. At night, buildings and cars are generally colder than streets. b) Magnification of the region marked with yellow in a). c) Detected areas. The gray scale mapping in b) and c) is different than in a) to better highlight warm regions.



Fig. 8. Another detection example over a larger area visualized as an image mosaic. Candidate detections are shown in red and district heating pipes as blue lines. One major leakage is found in the image with a brighter color scale. A few false alarms are also present, typically around buildings.

TABLE I
DETECTION RESULTS IN THREE CITIES AT THE 0.05% LEVEL.

	City A	City B	City C
Network length (km)	400	150	150
Nbr. detections	549	193	283
Nbr. detections / km	1.4	1.3	1.9
Nbr. confirmed leakages	20	7	7

immediate danger. In the cities in Table I, between 7 and 20 previously unknown leakages requiring immediate action, i.e., large amounts of water were leaking from the pipes, were confirmed by digging at locations indicated as candidate leakages.

B. Building segmentation

The AdaBoost classifier that classifies image regions as buildings or non-buildings was trained on 250 of the 500 manually delineated buildings and evaluated using the remaining 250 examples. A classification accuracy of 97% on the training data and 86% on the test data was obtained. Some example segmentations are shown in Fig. 9. A few buildings or building parts are not detected, either because they were not identified as one object by the Watershed transform or because they were falsely classified by the AdaBoost algorithm. A ranking of the importance of the extracted features can also be obtained by studying the trained AdaBoost classifier. Overall, the median intensity of a region, its compactness and its border regularity are the most characteristic features. However, all extracted features add to the classification accuracy.

V. DISCUSSION

The methods presented above have been applied to find district heating leakages in a number of cities, and they have helped locating numerous of real leakages. The cost of the data collection and analysis is currently of the order of \$500 per kilometer pipe, which well offsets the costs generated by leakages. The stand-alone software makes it possible for city representatives and operators of the district heating systems to efficiently browse the large amount of image data and to evaluate candidate leakages detected by the automatic algorithm. This approach saves tremendous amounts of time compared to screening each individual image manually.

The automatic detection of candidate leakage regions is purposely kept quite simple, relying mainly on intensity (temperature) and area characteristics. The reason for this is to keep the automatic analysis as transparent as possible to the user. This means that the user understands why regions are highlighted as candidate detections and he or she can quickly classify them as suspected leakages or false alarms based on contextual scene information or on additional background knowledge about the district heating network. Thus is the objective analysis of the computer, which ensures that no warm regions are missed, combined with the superior cognitive skills and background knowledge of the human expert. The automatic image analysis system could try to further reject false alarms through more complex scene analysis or by extracting more intricate features from the candidate regions




Fig. 9. Example detections of buildings in four images. The yellow arrows show missed building parts.

and try to train a classifier to recognize them. However, such solutions tend to become a black-box to the human operator, who may not understand why some warm regions are detected as possible leakage candidates while others are not. Such obscurity affect the acceptance and trust in the system. There is also an asymmetric cost of a false positive and a false negative detection that must be considered: a false leakage detection generally only incurs a small cost in terms of the time it takes for the human operator to reject it, whereas a missed real leakage will incur a large economic costs in terms of energy losses and potentially also damages to the environment and persons.

One relatively transparent way of trying to reduce the number of false alarms is to remove candidate leakage detections that border to buildings. A rough estimation is that 20% of the false alarms are found next to buildings. A method for automatically segmenting buildings have been presented to this end. In general, building segmentation in remote sensing images is a challenging task as it requires understanding of the scene, e.g., the relation between buildings, roads, cars, vegetation, for a correct result. Even a human cannot always delineate building structures in the thermal images used in this work. It can be noted that a perfect segmentation is not required for the current application: if a building is missed it may possibly generate an extra false alarm detection due to some heat leakage from that building. Single additional false alarms add only little extra time in the screening of the images, as discussed above. In contrast, it is more important not to misclassify a non-building as a building, because this may remove a real leakage from the candidate detection map. It should also be stressed that one cannot exclude the possibility of a real leakage next to a building, sometimes the district heating pipes even run under buildings. It is therefore advisable that a human expert looks also at candidate detections close to buildings to decide whether they should be further evaluated or not. The building segmentation is for this reason currently intended as a priority indicator in the visual screening process that can be used in conjunction with the ranking of the candidate regions based on image intensity and area. It should be emphasized that the building segmentation can be improved by acquiring an additional 3D laser measurement that maps the surface height for each pixel. Some cities also maintain a GIS map of all buildings that can be projected as masks into the images in the same way as the district heating pipes. Compared to this solution, however, a reliable segmentation based on the images themselves has the advantage that there are no registration errors.

The main hurdle towards a fully automatic leakage detection system is improved analysis and understanding of context and patterns in the scene. In computer vision and image analysis research in general, this has proven a very difficult problem. It should be stressed that the camera hardware, e.g., in terms of resolution, temperature sensitivity or signal-to-noise, is currently not a limiting factor. More near-time goals are to include additional GIS sources, such as layers of roads and buildings, as they are becoming commonplace in city planning. Another relevant problem is to estimate the energy loss at a leakage based on the image information.

VI. CONCLUSIONS

Computerized image analysis methods for the detection of leakages in district heating networks by means of airborne thermography have been presented. Leakages are detected as temperature anomalies, i.e., as regions with unusually high temperature compared to other regions in the images. A method to adaptively set the threshold temperature for detecting leakages has been presented. This is important as the image contrasts depend on the environmental conditions during the acquisition. Several large-scale real life studies with image acquisitions covering entire cities have been carried out, and the presented methods have proven valuable tools when screening the large amounts of acquired images, both by pinpointing real leakages and by saving large amounts of tedious manual work. A method for segmenting buildings in the thermal images has also been suggested as a way of reducing false alarms.

ACKNOWLEDGMENT

The authors would like to thank the Swedish Governmental Agency for Innovation Systems (VINNOVA) for project funding through the FOCUS Institute Excellence Center. The work has been carried out in cooperation with several Scandinavian district heating companies, the Swedish District Heating Association (Svensk Fjärrvärme) and Sky Movies AB.

REFERENCES

- [1] A. Poredos and A. Kitanovski, "District heating and cooling for efficient energy supply," in *Proceedings of the 2011 International Conference on Electrical and Control Engineering*, 2011, pp. 5238–5241.
- [2] M. Olsson, "Long-term thermal performance of polyurethane-insulated district heating pipes," Ph.D. dissertation, Chalmers University of Technology, 2001.
- [3] M. Fröling, "Environmental and thermal performance of district heating pipes," Ph.D. dissertation, Chalmers University of Technology, 2002.
- [4] S.-A. Ljungberg, "Aerial thermography - a tool for detecting heat losses and defective insulation in building attics and district heating networks," in *Proceedings of SPIE Thermosense IX: Thermal Infrared Sensing for Diagnostics and Control*, 1987, pp. 257–265.
- [5] —, "Thermography for district heating network applications: operational advantages and limitations," in *Proceedings of SPIE Thermosense X: Thermal Infrared Sensing for Diagnostics and Control*, 1988, pp. 70–77.
- [6] 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 *Proceedings of the 16th Congress of the International Society for Photogrammetry and Remote Sensing*, 1988, pp. 348–359.
- [7] D. Turton, "Aerial thermography - applications in mining," AAMHatch Pty Limited, Tech. Rep., 2008.
- [8] K. McManus, "Airborne thermography and ground geophysical investigation for detecting shallow ground disturbance under vegetation," Ph.D. dissertation, Durham University, 2004.
- [9] R. Gluch, "Using high spatial resolution airborne thermal-IR data to map the urban thermal environment," in *Proceedings of the 2002 IEEE International Geoscience and Remote Sensing Symposium*, 2002, pp. 1923–1925.
- [10] D. Allinson, "Evaluation of aerial thermography to discriminate loft insulation in residential housing," Ph.D. dissertation, University of Nottingham, 2007.
- [11] S. Axelsson, "Thermal modeling for the estimation of energy losses from municipal heating networks using infrared thermography," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 26, no. 5, pp. 686–692, 1988.

- [12] H. Zinko, J. Björklev, H. Bjurström, M. Borgström, B. Bohm, L. Koskelainen, and G. Phetteplace, "Quantitative heat loss determination by means of infrared thermography - the TX model," International Energy Agency, Tech. Rep., 1996.
- [13] M. Borgström and B. Bohm, "A comparison of different methods for in-situ determination of heat losses from district heating pipes," Technical University of Denmark, Tech. Rep., 1996.
- [14] C. Benedek, X. Descombes, and J. Zerubia, "Building detection in a single remotely sensed image with a point process of rectangles," in *Proceedings of the 20th International Conference on Pattern Recognition*, 2010, pp. 1417–1420.
- [15] B. Matei, H. Sawhney, S. Samarasekera, J. Kim, and R. Kumar, "Building segmentation for densely built urban regions using aerial LIDAR data," in *Proceedings of the 2008 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2008, pp. 1–8.
- [16] D. Ziou and S. Tabbone, "Edge detection techniques: An overview," *International Journal of Pattern Recognition and Image Analysis*, vol. 8, no. 4, pp. 537–559, 1998.
- [17] J. Roerdink and A. Meijster, "The watershed transform: definitions, algorithms, and parallelization strategies," *Fundamenta Informaticae*, vol. 41, pp. 187–228, 2000.
- [18] S. Haykin, *Neural Networks: A Comprehensive Foundation*. Prentice Hall, 2009.
- [19] J. Friedman, T. Hastie, and R. Tibshirani, "Additive logistic regression: a statistical view of boosting," *The Annals of Statistics*, vol. 38, no. 2, pp. 337–374, 2000.