Automated location of active fire perimeters in aerial infrared imaging using unsupervised edge detectors

M M Valero, O Rios, E Pastor* and E Planas

Centre for Technological Risk Studies, Universitat Politècnica de Catalunya – BarcelonaTech, Eduard Maristany 10-14, E-08019 Barcelona, Spain.

*E-mail (corresponding author): elsa.pastor@upc.edu

Abstract

A variety of remote sensing techniques have been applied to forest fires. However, there is at present no system capable of monitoring an active fire precisely in a totally automated manner. Spaceborne sensors show too coarse spatio-temporal resolutions and all previous studies which extracted fire properties from infrared aerial imagery incorporated manual tasks within the image processing workflow. As a contribution to this topic, this article presents an algorithm to automatically locate the fuel burning interface of an active wildfire in georeferenced aerial thermal infrared (TIR) imagery. An unsupervised edge detector, built upon the Canny method, was accompanied by the necessary modules for the extraction of line coordinates and the obtention of the total burned perimeter. The system was validated in different scenarios ranging from laboratory tests to large scale experimental burns performed under extreme weather conditions. Output accuracy was computed through three common similarity indices and proved acceptable. Computing times were below 1 second per image on average. The produced

information was used to measure the temporal evolution of the fire perimeter location and automatically generate rate of spread (ROS) fields. Information products were easily exported to standard Geographic Information Systems (GIS) such as GoogleEarth and QGIS. Therefore, this work contributes towards the development of an affordable and totally automated system for tactical wildfire surveillance.

Keywords

Thermal infrared (TIR), remote sensing, image analysis, wildfire evolution, forest fire, monitoring, segmentation, perimeter tracking, rate of spread (ROS), emergency response

Introduction

Wildland fires constitute a widespread problem with important social, economic and ecological consequences which has shown growing rates of occurrence in past decades. Different authors have reported a clear increase in the annual number of fires and area burned in the Mediterranean region (Pausas 2004; Shakesby 2011) and a further rise in fire seasons length and severity is to be expected worldwide in the near future (Flannigan *et al.* 2009, 2013). Furthermore, fire regimes are shifting and their dynamics are not yet fully understood (Finney *et al.* 2013; Duane *et al.* 2015). At present, fire spread cannot be successfully forecasted given a set of initial and boundary conditions. A number of propagation models have been published in the past (Pastor *et al.* 2003; Sullivan 2009a; b; c) but the simulators built upon them have shown low accuracy when estimating fire rates of spread (Cruz and Alexander 2013). In this context, there is a need for new means of obtaining accurate detailed information on the evolution of wildland fires, with a triple goal: firstly, to enhance the understanding of fire dynamics; secondly, to assist firefighting tasks if such data can be acquired and provided in real time during a fire event; and thirdly, to tune data-driven fire propagation simulators on line and improve the accuracy of their forecast, also in real time. Recent studies developed data

2

assimilation techniques with this aim (Mandel *et al.* 2008; Rios *et al.* 2014, 2016, Rochoux *et al.* 2014, 2015). One of the key pieces of information with highest value for fire management teams is the location of the fire perimeter and its rate of spread. Awareness of fire position and rate of spread would reduce the risk of firefighter entrapment and boost the effectiveness of firefighting activities.

Remote sensing has shown a great applicability to the study of forest fires. A variety of sensors and platforms have been tested and several systems are currently operational. Most systems focus on pre- and post-fire applications such as estimating surface and crown fuel loading, early fire detection, remote measurement of burned areas and estimation of carbon dioxide and trace gas emissions (Lentile *et al.* 2006). Although others collect active-fire data, automatic surveillance of fire evolution has not yet been achieved. Spaceborne sensors exhibit too coarse spatial (500m at best) and temporal (2 - 4 overpasses a day) resolutions to be suitable for tactical wildfire monitoring, and they are affected by cloud cover (Csiszar *et al.* 2006; Boschetti *et al.* 2010; Veraverbeke *et al.* 2014). A number of airborne sensors with much better properties are currently employed for fire evolution surveillance (Riggan *et al.* 2003; Ambrosia and Wegener 2009; Ambrosia, Wegener, *et al.* 2011; Ambrosia, Sullivan, *et al.* 2011), but the extraction of fire properties is only partially automated. At present, images are in general analysed merely qualitatively and fire front locations, when computed, are delineated manually (Pérez *et al.* 2011; Manzano-Agugliaro *et al.* 2014; Stow *et al.* 2014; Dickinson *et al.* 2016; Zajkowski *et al.* 2016).

Image processing algorithms so far applied to wildfires are rather simple and those with greater capabilities work exclusively under very specific conditions. A large amount of research has been conducted into the segmentation of visual imagery (i.e. classification of image pixels as fire or non-fire pixels). However, fire detection in the visible spectrum is problematic due to the presence of smoke and the non-uniformity of the background. Çetin *et al.* (2013) provided a detailed review of the state of the art. Although very important advancements related to forest fires have been published by Çelik and Demirel (2009), Ko *et al.* (2009), Rudz *et al.* (2009a,

2013), Rossi *et al.* (2010, 2011, 2013), Borges and Izquierdo (2010) and Toulouse *et al.* (2015), these algorithms are far from reaching the needed robustness and maturity. On the other hand, Thermal Infrared (TIR) cameras offer important advantages when compared to visible imaging. The TIR range is generally smoke free (for fire scenarios with scattering limited to particles under 1 µm) and fire is usually easy to distinguish from the cold background. These cameras, which are becoming more and more affordable, have proved to be useful for active fire mapping even when handheld and viewing far off-nadir (Plucinski and Pastor 2013; Paugam *et al.* 2013). TIR images have also seen important advancements concerning fire segmentation, although no completely autonomous system has been developed yet for automated forest fire monitoring. Existing processing algorithms for TIR images rely on either special working conditions or additional support systems. Pastor *et al.* (2006b), for instance, computed the location of linear flame fronts and their rate of spread using a laboratory set-up, whereas Martínez-de Dios *et al.* (2011) needed close views of the fire from different angles simultaneously and combined TIR and visible cameras. Moreover, most existing segmentation algorithms apply intensity thresholding methods which can neither distinguish between flames and fire base nor detect which part of the fire contour is actually the most active. Pérez *et al.* (2011) defined several intensity thresholds so as to discriminate between different characteristic fire zones, but such classification did not allow for automatic location of the active fire front. Paugam *et al.* (2013) used time-of-arrival (t_a) isocontours to locate the fire front in middle infrared images, defining the time of arrival as the time at which each pixel reached a brightness temperature of 600 K. However, their approach relied on accurate brightness temperature measurements, which requires a high-end camera. Furthermore, this method was based on a time-series analysis and it would not allow the detection of the fire front in an isolated image. Finally, Ononye *et al.* (2007) made use of multispectral images, thus taking advantage of the difference in emission profiles between burned and unburned vegetation. Unfortunately, this type of sensors are complex, heavy and expensive and cannot be widely implemented operationally at present.

In this context, this article presents a fire tracking algorithm designed for georeferenced TIR images acquired using an airborne thermal camera. Images are segmented individually to detect the position of the fuel burning interface in each video frame. Subsequently, time-series information is used to obtain the fire perimeter evolution. This allows the extraction of secondary information products such as ROS fields. Outputs can be imported into common GIS platforms. Several standard edge detection methods were analysed and the one which showed the best performance was selected and adapted. The core segmentation algorithm was complemented with the needed modules for parameters estimation and extraction of derived information products. The complete system was tested on a set of experimental scenarios at different scales including large-scale burns. Needed input images may be captured by low-cost TIR cameras as long as the imagery can be georeferenced. Corrections that account for bodies' emissivity or atmospheric transmissivity are not required.

Background

Edge detection methods

Edge detection methods exhibit important advantages over simpler segmentation algorithms based on intensity thresholding. Active fire areas appear in IR images as most intense regions. However, the divergence in intensity is usually slight between fuel burning areas and other hot zones such as flames, which makes intensity thresholding methods not robust enough for this purpose. In fact, the region that is to be detected corresponds to the interface between the unburned (cold) fuel and the most intensively burning (hottest) zone. For this reason, it seems a better approach to compute and threshold intensity gradients, which is the basis of a series of well-known edge detection algorithms (Gonzalez and Woods 2008).

We assessed the suitability of such techniques for this particular purpose by comparing different common methods. Both simple gradient masks and more complex edge detection algorithms were applied to various sets of IR images. The studied gradient masks were Sobel's and Prewitt's (Prewitt 1970; Sobel 1970; Gonzalez and Woods 2008), whereas the Laplacian of a

Gaussian (LoG) edge detector (Marr and Hildreth 1980) and the Canny edge detector (Canny 1986) were chosen among more sophisticated algorithms.

Gradient masks are a useful tool to compute the gradient magnitude of a two-dimensional field variable with high computational efficiency. As such, they can be applied to intensity images to compute intensity gradients. The most commonly employed masks are Sobel's and Prewitt's**¡Error! No se encuentra el origen de la referencia.**. In both cases, the convolution of the original image with masks of size 3 x 3 pixels approximates the intensity partial derivatives in the vertical and horizontal direction. Once these partial derivatives are known, the magnitude and direction of the gradient can be computed for every pixel in the image. Finally, a threshold can be applied in order to consider only the pixels with highest gradient magnitudes.

The Laplacian of a Gaussian (LoG) detector is based on the expression shown in equation 1, which corresponds to the application of the Laplacian operator to a 2-D Gaussian function of standard deviation σ (*x, y* represent pixel horizontal and vertical coordinates, respectively). The Gaussian function smooths the image reducing noise while the second derivatives of the Laplacian operator detect the most abrupt changes in intensity. Pixels with the highest values in intensity gradient show peaks in the first derivative and, consequently, zero-crossings in the second derivative.

$$\nabla^2 G(x,y) = \left[\frac{x^2+y^2-2\sigma^2}{\sigma^4}\right] e^{-\frac{x^2+y^2}{2\sigma^2}} \qquad (1)$$

The Canny edge detector, whose detailed mathematical description can be found in (Canny 1986; Gonzalez and Woods 2008), consists of four basic steps. Firstly, the input image is smoothed with a Gaussian filter; next, the gradient magnitude is computed; afterwards, an algorithm of non-maxima suppression is applied; lastly, edges are detected and linked by means of double thresholding and connectivity analysis. Its basic interpretation is the following: after noise is reduced through the Gaussian filter, main edges are detected and non-edges discarded. Considered edge points are then classified into two categories: strong and weak edges. Two

threshold values, usually named hysteresis thresholds, must be defined for this purpose. Pixels with gradient values above the highest threshold are considered strong edges. Pixels with gradient values above the lowest threshold but below the highest one and connected to strong edges are considered weak edges. The final edge ridges are composed by strong edge points plus the weak edge points.

Evaluation criteria

A good number of similarity metrics are available in the literature for segmentation quality assessment. Three common indices were used here, namely the Figure of Merit (FOM), first published by Pratt (1978) and defined as in equation 2, the mean distance as originally defined by Peli and Malah (1982) (eq. 3) and the Baddeley distance (Baddeley 1992) (eq. 4).

$$FOM\left(I, I_{gt}\right) = \frac{1}{\max\{\text{card}(I), \text{card}(I_{gt})\}} \sum_{k \in I} \frac{1}{1 + \alpha \cdot d(k, I_{gt})^2} \tag{2}$$

$$MD\left(I, I_{gt}\right) = \frac{1}{\text{card}(I)} \sum_{k \in I} d(k, I_{gt}) \tag{3}$$

$$BAD\left(I, I_{gt}\right) = \left[\frac{1}{\text{card}(I) + \text{card}(I_{gt})} \sum_{k \in I \cup I_{gt}} \left|d\left(k, I_{gt}\right) - d(k, I)\right|^P\right]^{\frac{1}{P}} \tag{4}$$

In all of them, *I* is the segmentation result, I_{gt} its corresponding ground truth, *card()* represents the cardinality of a curve (i.e., the number of points it consists of) and $d(k, I_{gt})$, the minimum distance between the point *k* and the curve I_{gt}. α is a constant with an assigned value of 1/9 by Pratt and *P* another constant that may take any value ≥ 1. Not having found in the literature any justified recommendation for the value of P, here *P* = 1 was taken in order to keep a formulation similar to that of the mean distance.

These three indices were selected among all existing options for the following reasons. Pratt's Figure of Merit was chosen by Chabrier *et al.* (2008) over numerous evaluation criteria for its best performance. Its use was recommended for contour similarity assessment (Hemery *et al.* 2010) and it was implemented for comparison of two segmentation algorithms applied to wildfire front localisation in visual images (Rudz *et al.* 2009b). In addition to the Figure of Merit, which is dimensionless, dimensional distances were computed in meters to provide additional understanding of the algorithm performance. Both the mean distance (eq. 3) and the Baddeley distance (eq. 4) were as well recommended by Hemery *et al.* (2010), and the latter was also chosen by other authors for its high-quality performance (Fernández-García *et al.* 2004; Medina-Carnicer *et al.* 2011a).

Although all three indices quantify the accuracy of the tested methods, they present some differences that should be noticed. The Figure of Merit is dimensionless and takes values in the restricted range [0, 1], 1 being the best score and 0 the worst. The mean distance provides an average value of the offset existing in each instance between the automatically segmented edge and the ground truth edge. The Baddeley distance, despite being dimensional, does not provide easily interpretable information but includes an important aspect: both the FOM and the mean distance account only for the points which were wrongly considered as edges by the segmentation algorithm but are not in the ground truth, i.e., the false positives. Neither penalises the algorithm for those points that should be considered edges but were not, i.e., false negatives. The Baddeley distance takes this into consideration and accounts for false positives and false negatives alike.

Automatic algorithm for fire perimeter tracking

The main aim of the developed system is to monitor the evolution of an active wildfire using aerial TIR images. This goal is achieved following three steps. Firstly, the algorithm makes use of image processing tools to detect which pixels in each frame belong to the fuel burning interface. Secondly, the coordinates of these pixels are ordered to build a connected line, and the

total burned perimeter is updated. Finally, the perimeter evolution is imported into a GIS and used to compute rates of spread. The following subsections detail each of these steps, whereas figure 1 shows a block diagram of the whole process.

Unsupervised fire front detection

By *fire front* or *fuel burning interface* we refer in this article to the part of the fire perimeter which shows active combustion at a given moment. To detect it, the suitability of the four presented edge detection algorithms was assessed using the metrics described above. Following the comparative study, whose outcome is discussed in the results section, the Canny method was selected and subsequently totally automated. During the comparative study, all detectors were tested at their best possible performance by the manual adjustment and optimisation of all required parameters. However, such manual tuning is to be avoided in a completely unsupervised framework like the one proposed here.

Therefore, Canny's parameter selection was automated following the algorithm proposed by Medina-Carnicer *et al.* (2011b). This algorithm allows the unsupervised estimation of the optimum hysteresis thresholds for the Canny detector given a set of candidates by using only the information available in the image being analysed. I.e., each frame is processed independently and the time series, as well as its frequency, has no impact on the detected fire front. The correct determination of the hysteresis thresholds is crucial since pixels with gradient values comprised between both thresholds will only be classified as edges if they meet certain connectivity criteria. In contrast to previous work on this issue where only gradient information had been used to estimate optimum threshold values, this approach takes into consideration the number of times each pixel is added to the edge during the hysteresis process. This is accomplished by searching for the *instability zone*, defined as the interval of gradient levels between both hysteresis thresholds. For instance, a pixel with a low gradient value is unlikely to become an edge point. However, if many pixels with the same gradient value are added during the hysteresis process, this gradient value should probably be a member of such instability zone.

9

Similarly, pixels with a high gradient value will most probably be part of edges. However, if few pixels with the same gradient value are added during the hysteresis process, such gradient level should probably not fall within the instability zone. Therefore, the determination of the instability zone depends not only on the gradient values but also on the way these gradient values are treated during the hysteresis process. In consequence, how gradient values are distributed in the image is relevant. The proposed algorithm introduces connectivity in the analysis and uses it for optimum parameter estimation.

Figure 2 shows the block diagram of this sub-module. In summary, given an input image I and the corresponding gradient image $G(I)$, a set of different edge maps $G_C(I)$ can be obtained using a set of possible hysteresis thresholds $C(I)$. If a set with a large enough number of candidates is used, then the set $\{\Delta G_{low,high}(I), (low, high) \in C\}$ provides information about how many times each pixel is added during the hysteresis process. Such information is consolidated in the image SM_{H_C}. The grey level of each pixel in the image $Prob(SM_{H_C})$ represents the probability of each pixel being added as an edge point during the hysteresis process. If $Prob(SM_{H_C})$ is thresholded with $x \in (0,1)$, only the pixels with a probability equal to or greater than x of being added by the hysteresis process will appear in the image $F_I(x) = G(I) \cdot Prob_x(SM_{H_C})$, and their grey level will be the corresponding gradient value contained in image $G(I)$. Pixels with a probability lower than x of being added by the hysteresis process will have the grey level zero in the image $F_I(x)$. Finally, the distribution $P(F_I(x))$ (eq. 5) provides information about the probability that a pixel has a gradient value x if its probability of being added during the hysteresis process is equal to or greater than x. Whereas $P(F_I(x))$ is expected to take the value zero in the majority of the points, gradient levels x that meet the condition $P(F_I(x)) \neq 0$ can be considered instable. The instability zone is defined by the instable gradient levels.

$$P(F_I(x)) = \begin{cases} \dfrac{|F_I(x)|}{|Prob_x(SM_{H_C})|}, & |Prob_x(SM_{H_C})| > 0 \\ 0, & |Prob_x(SM_{H_C})| = 0 \end{cases} \qquad (5)$$

From a practical point of view, the proposed method only needs to be fed a set of threshold candidates and the standard deviation of the Gaussian filter used for noise removal. Input parameters can generally be kept constant as we did here for all study cases (see Table 1). Furthermore, there is no need to provide a subset of ground-truth fire fronts in order to calibrate the Canny detector. The system itself performs this calibration using only the information contained in the thermal images.

Note that, according to Medina-Carnicer *et al.* (2011b), the set of candidates for the hysteresis thresholds could take the form $c = \{(low, high)|low, high \in (0, 1)\}$. However, this range was narrowed here attending to the expectable values of the optimum thresholds. *low* will always be lower than *high*. Furthermore, based on previous experience with this type of images, *low* was expected to take values between 0.2 and 0.7 and *high*, between 0.5 and 0.9. While simple, these assumptions allow an important reduction in computing times since each candidate pair entails the computation of two edge maps for each input frame. The results shown below proved the validity of these assumptions. Note however that this candidates set is to some degree arbitrary and can be modified attending to computational constraints or accuracy requirements.

Extracting the fire front coordinates and building the fire perimeter

Once the pixels comprising the fire front have been detected in each frame, their coordinates must be read and ordered to form a line suitable for further processing. Previously, the thickness of the detected fire front is to be reduced to the minimum possible. This was performed here through a medial axis transformation, which allows reducing a plane region to a line. For every point in the region, this thinning algorithm finds the closest point in the region border. If a point is found to be equidistant to two or more points in the border, it is considered to belong to the medial axis of the region. The output can be understood as the *skeleton* of the region (Gonzalez and Woods 2008).

Since input images were already ortho-normalised and geo-referenced, the task of computing the geographic coordinates of the pixels in the fire front was rather straight-forward. However, ordering these points required further attention. A set of randomly distributed points must be ordered in either clockwise or counter clockwise direction considering their connectivity. To achieve this, the line endpoints had to be previously located. Although edge detection algorithms try to maximise the connectivity of resulting edges, these are usually composed of more than one connected line. Thus, intermediate endpoints may easily be confused with the actual endpoints of the fuel burning interface. The task of identifying the latter in a general case and computing them in an automatic manner may be difficult to implement. The approach followed here relies upon the assumption that the actual endpoints will generally lie the farthest from the centroid of the complete collection of points. Note that this centroid does not refer to the whole fire perimeter but to the fire front detected in each frame. Once one of these endpoints was located, the curve could be read from a binary image in the correct direction.

The next processing step consists in updating the burned perimeter with the new information provided by the edge detector in each time step as exemplified in figure 3. The points comprising the detected burning interface at t_k are ordered clockwise. Subsequently, its endpoints are connected to the closest points in the burned perimeter computed at t_{k-1}. Depending on the time elapsed between processed frames, the newly detected active front might be separated from the latest built perimeter. Afterwards, the points in the burned perimeter at t_{k-1} comprised between those endpoints are read clockwise and added to the updated perimeter at t_k. Finally, the updated fire mask at t_k is applied a union operation ($\cup$) with the fire scar at t_{k-1} to ensure the fire perimeter at t_k will always contain the fire perimeter at t_{k-1}.

Derived information products

As long as input footages are georeferenced, fire perimeter locations can be expressed in geographic coordinates and imported into a GIS. We developed the necessary modules to export results in the formats used by GoogleEarth and QGIS, two widely used and free tools. The

exportation to GoogleEarth is performed by writing a *.kml* file, whereas the QGIS exportation module creates an ASCII points file which is later on loaded by the GRASS function *v.in.ascii.points*. This GIS integration allows coupling the information produced by the presented methodology with topographic maps, digital elevation models and crew teams' position, just to name a few.

Finally, geo-located fire fronts and burned perimeters are used to derive second-level information products such as ROS fields. The fire rate of spread is computed at each point of the perimeter following the methodology used in previous studies such as (Ononye *et al.* 2007; Pastor *et al.* 2010; Planas *et al.* 2011; Paugam *et al.* 2013), i.e. measuring the distance that point would travel in a direction perpendicular to the perimeter until intersecting the perimeter location corresponding to the next time instant, and dividing this distance by the time needed to travel it. If a digital elevation model of the area is available, the actual surface ROS is computed in addition to the horizontal ROS. Whereas the latter is important to assess fire growth velocity, it is the former which possesses physical meaning. Furthermore, the surface ROS magnitude may be used to estimate fireline intensity (Johnston *et al.* 2017) and it is essential to correctly determine fire ecological impact.

Experimental data

The algorithm was applied to around 6600 frames from five different video sequences captured under dissimilar circumstances. Table 2 summarises their most relevant properties whereas sample frames are displayed in figure 4.

Footage 1 was acquired during a laboratory burning over a combustion table of about 1.5 meters width and 3 meters length, using straw as fuel. Footage 2 was recorded during field burning experiments conducted in a 17 x 8 meters plot covered with thinned Aleppo pine saplings stands (Domènech 2011). Footages 3 to 5 were captured during various field experiments conducted in the Ngarkat Conservation Park, South Australia (Cruz *et al.* 2010; Plucinski and Pastor 2013).

These experiments consisted of a series of controlled burns in horizontal mallee-heath shrub plots with areas ranging from 4 to 25 ha. Despite being part of a scientific experimental burning programme, fires exhibited real behaviour patterns as were performed in large plots under extremely severe weather conditions. Furthermore, they were recorded similarly to a real emergency situation.

All images were acquired with a thermal IR camera whose properties are listed in Table 3. Always operated manually, it was fixed at elevated oblique positions in experiments 1 and 2 and mounted on a hovering helicopter for footages 3 to 6.

The spatial resolution of footages depended on the distance between the fire and the camera and the angle of view in each case. Pixels in footages 1 and 2 had a side length of approximately 0.008m and 0.04m, respectively. Images corresponding to large burning plots exhibited pixels of between 1 and 3 meters per side. All TIR footages were orthorectified through a direct linear transformation algorithm similar to that proposed by Pastor *et al.* (2006b). Furthermore, large-scale imagery was georeferenced using ground control points. Subsequently, fire perimeters were delineated manually every 10 seconds. These manual isochrones were taken as ground truth for accuracy assessment. In addition to constituting the most reliable reference data available, they represent the output of the current methodology that we aim to automate. Therefore, it can be assumed that the closer an algorithm's output is to this reference data, the higher its performance should be rated.

Results and discussion

Comparative study of edge detection methods

The four edge detection algorithms detailed in the background section were applied to the different geo-corrected footages described in the experimental data section. Figure 5 shows the results for some representative frames for qualitative assessment. These results represent the best possible output obtained with each method, i.e. edge detection parameters were optimised

for each frame independently. In addition to the edge detection algorithms, the result of applying global intensity thresholding is also shown. Since the latter is the approach most commonly used to date, comparison to its output allows the assessment of the improvements brought by our proposal.

In order to quantitatively assess the performance of all methods, the three similarity metrics discussed in the background section were computed for all frames with available ground truth. Results are shown in figure 6.

Figures 5 and 6 show the best possible output given by every method in each sequence. Fire fronts were obtained after manually adjusting the parameters independently for each method and each footage, so that the output was always the closest to the desired one. Sobel's and Prewitt's masks behaved similarly, which could be anticipated considering their similar formulation. Their application proved to be unsatisfactory, as did the Laplacian of a Gaussian (LoG) detector. Also poor results were obtained by direct intensity thresholding. None of these methods was able to detect the whole burning interface (omission errors) and they included in the output points that were obviously not part of the targeted boundary (commission errors). More importantly, commission and omission errors showed opposite tendencies when modifying the thresholding value. The threshold can be adjusted to optimize segmentation, and such tuning was performed here. However, increasing its value raised omission errors, whereas reducing the threshold resulted in higher commission errors. This behaviour implies the impossibility of obtaining satisfactory outputs with these methods by adjusting the threshold alone.

On the contrary, the Canny method gave satisfactory results. Not only did it successfully recognise the points comprising the fuel burning interface but it also omitted the flames in most cases, adjusting the output line to the fire base. This fact shows an important improvement when compared to the global intensity thresholding technique, which cannot distinguish between burning fuel (i.e. fire base) and flames. This problem was exacerbated by the lengthening and distortion of flames in the geocorrected frames (see for instance frames of footage 5 in figure 5). Flame distortion appears during the homography transformation used for ortho-rectification,

which projects pixels in the camera plane to the ground. Since flames are contained in neither of these planes, they become distorted. Furthermore, they are not part of the fuel burning interface and must be filtered out. The only method able to do this to some extent was Canny's. Another relevant aspect to note is the behaviour with respect to internal hot zones which, despite still being actively burning (and therefore exhibiting flames), are not part of the fire front and should not be present in the output line. No method but Canny's could omit such internal burning zones (see figure 5). Lastly, Canny output lines are highly connected (few gaps appear), which constitutes an important advantage for further processing.

Accuracy of the automatic algorithm

For the aforementioned reasons, the Canny method was selected and the estimation of its optimum parameters was automated as explained above. Qualitative assessment of the totally unsupervised edge detector is shown in figure 7. Moreover, the same similarity indices were applied in the same footages (figure 8). Although slightly lower than when tuning the parameters manually, the quality was still found to be acceptable in this case. Values for the Figure of Merit stayed above 0.6 in most cases and mean distances fell within centimetres in the small-scale burns and took values between 1 and 18 metres in the large scale scenarios. The fact that the values for the Baddeley distance were not much higher than the average distance proves than the algorithm did not present an important amount of false negatives.

The most significant dissimilarities correspond to frames in which the Canny method was affected by flame projections and could not filter them out. Figure 9 shows examples of successful and unsuccessful behaviour when considerable flames are present. Higher errors occur when flames are particularly concentrated. This incorrect behaviour was not easily foreseeable and had an important impact in the quantitative metrics computed.

Derived information products

The resulting rate of spread fields are displayed in figure 10, where a subset of the detected fire isochrones has been superimposed. ROS values computed for footage 1 were close to 0.3 m/min. These results are in accordance with measurements obtained by Pastor *et al.* (2006a) in the same laboratory setting with similar fuel and ambient conditions. In footage 2, although ROS were not explicitly measured, the fire travels around 17m in approximately 4.5 minutes, which gives an average ROS of almost 4 m/min, also in accordance with the results obtained here. Finally, ROS values obtained for footages 4 and 5 can be compared to measurements performed by Planas *et al.* (2011). In the experimental burn corresponding to footage 4 (Plot A), ROS was estimated to take values between 20 m/min and 143 m/min. During the burn corresponding to footage 5 (Plot AS2E), ROS measurements fell between 20 m/min and 166 m/min. There were no experimental data available for footage 3.

Analysis of computing times

We measured the computing time needed by the presented system in all the studied scenarios. An Intel i7 Quad Core desktop computer was used and no code parallelisation was implemented. In average, the algorithm needed 0.977s to obtain the fire perimeter location corresponding to each TIR image.

Computing speed may be further increased owing to the stability detected in the values of estimated optimum thresholds. Figure 11 shows these optimum values for the five considered scenarios. Their minimal dispersion suggests actual values might be approximated by means or medians. This hypothesis was tested by computing, for each footage, all fire fronts with constant thresholds obtained as an average of the optimum values for the first 50 frames. The quality of the fronts was similar to the quality of fire fronts obtained with optimum thresholds computed every frame (table 4).

When constant average thresholds were used instead of continuously computing the optimum thresholds for every single frame, accuracy did not drop significantly. Meanwhile, the computation time per frame could be reduced to 0.017s in average following this approach, i.e. more than 50 times. Therefore, this algorithm may be deployable even in scenarios where computational resources are extremely limited (e.g. in the field). In such cases, optimum thresholds may be approximated by average values computed during an initial calibration step.

Conclusions and future work

This article proposes the use of edge detection algorithms applied to aerial thermal infrared images in order to automatically locate the active fuel burning interface of an active wildland fire. Four existing edge detection techniques were assessed and their performance compared to each other's and to the current approach in the literature, i.e. direct intensity thresholding. Three widely accepted similarity indices were computed to measure the quality of each algorithm's output taking as reference the manually annotated ground truth. In all cases, the Canny method outperformed the others when the parameters needed for all of them were manually optimised. Moreover, it showed important additional advantages such as omitting undesired flame projections in most cases, adjusting the output flamefront to the fire base and providing highly connected lines with very few gaps. Furthermore, the Canny method could be completely automated and its parameters estimated without supervision using only the information contained in each input TIR frame. Neither ground truth isochrones nor manual calibration were needed. Performance of the Canny method with automatically defined parameters was reduced slightly but remained acceptable. The totally automated location of the fuel burning interface at each moment was imported into standard geographic information systems such as GoogleEarth and QGIS, thus allowing integration with other spatial awareness tools and decision support systems. Generated data also permitted measuring the evolution of the fire perimeter with time and extracting maps of the fire rate of spread.

The system we present was successfully tested over a set of thermal footages captured under different conditions ranging from laboratory tests to large-scale experimental burns. The latter significantly resembled an actual emergency response scenario since fires exhibited real behaviour patterns and were recorded by the crew aboard a hovering helicopter. Although in these situations there is a significant production of smoke, this fact did not interfere with the algorithms exposed here since TIR imagery is usually smoke free. Furthermore, the nature of the algorithms makes them independent of the availability of temperature values so that cameras without thermogrammetry capacity may be used and corrections that account for bodies' emissivity or atmospheric transmissivity are not required. The fire front detection algorithm is also time-series-independent since the fire front is detected separately in each frame. Therefore, the recording frequency has no effect on the detected fire fronts. Subsequent outputs such as the perimeter evolution and its rate of spread do depend on the quality of the time series, but the core algorithm does not. This fact allows the easy handling of faulty video sections.

Our results suggest an important applicationof aerial thermal infrared imaging combined with edge detection methods to automatic real-time monitoring of active forest fires. The steady decrease in price, size, weight and complexity that TIR cameras have recently shown will support and encourage developments such as the one examined here as well as their later implementation. Nonetheless, in order to ensure a high level of robustness and interoperability, this algorithm should be applied to a larger amount of images captured by different cameras, in more complex scenarios and corresponding to various fire behaviour regimes. Due to the inherent difficulty to obtain this type of imagery, available datasets are limited. After the acquisition of images, there are already algorithms which use the aircraft navigation systems (GPS, IMU) to automatically geo-reference the images taken onboard. However, the majority of these algorithms rely on the visible spectrum. Accurate and fast geocorrection methodologies might as well be developed ad-hoc for thermal imagery so that system's weight, cost and complexity can be reduced to the minimum. Finally, Geographic Information Systems (GIS) are currently powerful enough to integrate the output data in a decision support system suitable for

operational use, where real-time accurate information is crucial. In addition, data sensed remotely in real time can be used to feed data driven simulators which are able to provide fast, reliable forecasts of the subsequent fire spread. End user collaboration is essential in the development of intuitive and informative user interfaces that allow for the rapid delivery of meaningful information. Ultimately, emergency responders could benefit from state-of-the-art technology which would reduce risks for firefighters and boost the impact, efficacy and efficiency of suppression activities.

Acknowledgements

References

Ambrosia VG, Wegener S, Zajkowski T, Sullivan DV, Buechel S, Enomoto F, Lobitz B, Johan S, Brass J, Hinkley E (2011) The Ikhana unmanned airborne system (UAS) western states fire imaging missions: from concept to reality (2006 – 2010). *Geocarto International* **26**, 85–101. doi:10.1080/10106049.2010.539302.

Ambrosia VG, Sullivan DV, Buechel SW (2011) Integrating sensor data and geospatial tools to enhance real-time disaster management capabilities: Wildfire observations. *Special Paper of the Geological Society of America* **482**, 1–12. doi:10.1130/9780813724829.

Ambrosia VG, Wegener SS (2009) Unmanned airborne platforms for disaster remote sensing support. In 'Geoscience and Remote Sensing'. (Ed P-G Ho)(InTech) doi:10.5772/8302.

Baddeley A (1992) An error metric for binary images. *Robust Computer Vision* 59–78. doi:10.1.1.52.3879.

Borges PVK, Izquierdo E (2010) A Probabilistic Approach for Vision-Based Fire Detection in Videos. *IEEE Transactions on Circuits and Systems for Video Technology* **20**, 721–731. doi:10.1109/TCSVT.2010.2045813.

Boschetti L, Roy DP, Justice CO, Giglio L (2010) Global assessment of the temporal reporting accuracy and precision of the MODIS burned area product. *International Journal of Wildland Fire* **19**, 705. doi:10.1071/WF09138.

Canny J (1986) A computational approach to edge detection. *IEEE transactions on pattern analysis and machine intelligence* **8**, 679–698. doi:10.1109/TPAMI.1986.4767851.

Çelik T, Demirel H (2009) Fire detection in video sequences using a generic color model. *Fire Safety Journal* **44**, 147–158. doi:10.1016/j.firesaf.2008.05.005.

Çetin AE, Dimitropoulos K, Gouverneur B, Grammalidis N, Günay O, Habiboğlu YH, Töreyin BU, Verstockt S (2013) Video fire detection - Review. *Digital Signal Processing: A Review Journal* **23**, 1827–1843. doi:10.1016/j.dsp.2013.07.003.

Chabrier S, Laurent H, Rosenberger C, Emile B (2008) Comparative Study of Contour Detection Evaluation Criteria Based on Dissimilarity Measures. *EURASIP Journal on Image and Video Processing* **2008**, 1–13. doi:10.1155/2008/693053.

Cruz MG, Matthews S, Gould J, Ellis P (2010) Fire dynamics in mallee-heath: fuel, weather and fire behaviour prediction in south Australian semi-arid shrublands. *Bushfire Cooperative Research Centre, Technical Report A1001*. http://www.bushfirecrc.com/sites/default/files/firedynamicsinmalleeheathreport.pdf.

Cruz MG, Alexander ME (2013) Uncertainty associated with model predictions of surface and crown fire rates of spread. *Environmental Modelling and Software* **47**, 16–28. doi:10.1016/j.envsoft.2013.04.004.

Csiszar IA, Morisette JT, Giglio L (2006) Validation of active fire detection from moderate-resolution satellite sensors: the MODIS example in northern eurasia. *IEEE Transactions on Geoscience and Remote Sensing* **44**, 1757–1764. doi:10.1109/TGRS.2006.875941.

Dickinson MB, Hudak AT, Zajkowski T, Loudermilk EL, Schroeder W, Ellison L, Kremens RL, Holley W, Martinez O, Paxton A, Bright BC, O'Brien JJ, Hornsby B, Ichoku C, Faulring J, Gerace A, Peterson D, Mauceri J (2016) Measuring radiant emissions from entire prescribed fires with ground, airborne and satellite sensors - RxCADRE 2012. *International Journal of Wildland Fire* **25**, 48–61. doi:10.1071/WF15090.

Domènech R (2011) Efectivitat dels tractaments d'aclarida en la reducció del risc de propagació d'incendis en regenerats de pi blanc. Universitat Politècnica de Catalunya - BarcelonaTech.

Duane A, Piqué M, Castellnou M, Brotons L (2015) Predictive modelling of fire occurrences

from different fire spread patterns in Mediterranean landscapes. *International Journal of Wildland Fire* **24**, 407–418. doi:10.1071/wf14040.

Fernández-García NL, Medina-Carnicer R, Carmona-Poyato A, Madrid-Cuevas FJ, Prieto-Villegas M (2004) Characterization of empirical discrepancy evaluation measures. *Pattern Recognition Letters* **25**, 35–47. doi:10.1016/j.patrec.2003.08.011.

Finney MA, Cohen JD, McAllister SS, Jolly WM (2013) On the need for a theory of wildland fire spread. *International Journal of Wildland Fire* **22**, 25–36. doi:10.1071/WF11117.

Flannigan M, Krawchuk M, De Groot W, Wotton B, Gowman L (2009) Implications of changing climate for global wildland fire. *International Journal of Wildland Fire* **18**, 483–507. doi:10.1071/wf08187.

Flannigan M, Cantin AS, De Groot WJ, Wotton M, Newbery A, Gowman LM (2013) Global wildland fire season severity in the 21st century. *Forest Ecology and Management* **294**, 54–61. doi:10.1016/j.foreco.2012.10.022.

Gonzalez RC, Woods RE (2008) 'Digital Image Processing, 3rd ed.' (Pearson Prentice Hall)

Hemery B, Laurent H, Rosenberger C (2010) Comparative study of localization metrics for the evaluation of image interpretation systems. *Journal of Electronic Imaging* **19**, 23017. doi:10.1117/1.3446803.

Johnston JM, Wooster MJ, Paugam R, Wang X, Lynham TJ, Johnston LM (2017) Direct estimation of Byram´s fire intensity from infrared remote sensing imagery. *International Journal of Wildland Fire* **26**, 668–684. doi:http://dx.doi.org/10.1071/WF16178.

Ko BC, Cheong K-H, Nam J-Y (2009) Fire detection based on vision sensor and support vector machines. *Fire Safety Journal* **44**, 322–329. doi:10.1016/j.firesaf.2008.07.006.

Lentile LB, Holden ZA, Smith AMS, Falkowski MJ, Hudak AT, Morgan P, Lewis SA, Gessler PE, Benson NC (2006) Remote sensing techniques to assess active fire characteristics and post-fire effects. *International Journal of Wildland Fire* **15**, 319. doi:10.1071/WF05097.

Mandel J, Bennethum LS, Beezley JD, Coen JL, Douglas CC, Kim M, Vodacek A (2008) A wildland fire model with data assimilation. *Mathematics and Computers in Simulation* **79**, 584–606. doi:10.1016/j.matcom.2008.03.015.

Manzano-Agugliaro F, Pérez-Aranda J, De La Cruz JL (2014) Methodology to obtain isochrones from large wildfires. *International Journal of Wildland Fire* **23**, 338–349. doi:10.1071/WF13166.

Marr D, Hildreth E (1980) Theory of Edge Detection. *Proceedings of the Royal Society B: Biological Sciences* **207**, 187–217. doi:10.1098/rspb.1980.0020.

Martínez-de Dios JR, Merino L, Caballero F, Ollero A (2011) Automatic forest-fire measuring using ground stations and Unmanned Aerial Systems. *Sensors (Basel, Switzerland)* **11**, 6328–53. doi:10.3390/s110606328.

Medina-Carnicer R, Muñoz-Salinas R, Carmona-Poyato A, Madrid-Cuevas FJ (2011) A novel histogram transformation to improve the performance of thresholding methods in edge detection. *Pattern Recognition Letters* **32**, 676–693. doi:10.1016/j.patrec.2010.12.012.

Medina-Carnicer R, Muñoz-Salinas R, Yeguas-Bolivar E, Diaz-Mas L (2011) A novel method to look for the hysteresis thresholds for the Canny edge detector. *Pattern Recognition* **44**, 1201–1211. doi:10.1016/j.patcog.2010.12.008.

Ononye A, Vodacek A, Saber E (2007) Automated extraction of fire line parameters from multispectral infrared images. *Remote Sensing of Environment* **108**, 179–188. http://www.sciencedirect.com/science/article/pii/S0034425706004639.

Pastor E, Zarate L, Planas E, Arnaldos J (2003) Mathematical models and calculation systems for the study of wildland fire behaviour. *Progress in Energy and Combustion Science* **29**, 139–153. doi:10.1016/s0360-1285(03)00017-0.

Pastor E, Àgueda A, Andrade-Cetto J, Muñoz M, Pérez Y, Planas E (2006a) Computing the rate of spread of linear flame fronts by thermal image processing. *Fire Safety Journal* **41**, 569–579. doi:10.1016/j.firesaf.2006.05.009.

Pastor E, Àgueda A, Andrade-Cetto J, Muñoz M, Pérez Y, Planas E (2006b) Computing the rate of spread of linear flame fronts by thermal image processing. *Fire Safety Journal* **41**, 569–579. doi:10.1016/j.firesaf.2006.05.009.

Pastor E, Pérez Y, Cubells M, Planas E, Plucinski M, Gould J (2010) Quantifiable assessment of aerial suppression tactics in wildland fires using airborne infrared imagery. In 'In 'VI International Conference on Forest Fire Research'',

Paugam R, Wooster MJ, Roberts G (2013) Use of Handheld Thermal Imager Data for Airborne Mapping of Fire Radiative Power and Energy and Flame Front Rate of Spread. *IEEE Transactions on Geoscience and Remote Sensing* **51**, 3385–3399. doi:10.1109/TGRS.2012.2220368.

Pausas JG (2004) Changes in fire and climate in the eastern Iberian Peninsula (Mediterranean Basin). *Climatic Change* **63**, 337–350. doi:10.1023/b:clim.0000018508.94901.9c.

Peli T, Malah D (1982) A study of edge detection algorithms. *Computer Graphics and Image Processing* **20**, 1–21.

Pérez Y, Pastor E, Planas E, Plucinski M, Gould J (2011) Computing forest fires aerial suppression effectiveness by IR monitoring. *Fire Safety Journal* **46**, 2–8. doi:10.1016/j.firesaf.2010.06.004.

Planas E, Pastor E, Cubells M, Cruz M, Greenfell I (2011) Fire behaviour variability in mallee-heath shrubland fires. *5th International Wildland Fire …*. http://scholar.google.com/scholar?hl=en&btnG=Search&q=intitle:Fire+behavior+variability+in+mallee-heath+shrubland+fires#7.

Plucinski M, Pastor E (2013) Criteria and methodology for evaluating aerial wildfire suppression. *International Journal of Wildland Fire* **22**, 1144–1154. doi:10.1071/WF13040.

Pratt WK (1978) 'Digital image processing.'

Prewitt J (1970) Object enhancement and extraction. *In 'Picture processing and Psychopictorics'* **10**, 15–19.

Riggan PJ, Tissell RG, Hoffman JW (2003) Application of the firemappertm thermal-imaging radiometer for wildfire suppression. *IEEE Aerospace Conference Proceedings* **4**, 1863–1872. doi:10.1109/AERO.2003.1235116.

Rios O, Jahn W, Rein G (2014) Forecasting wind-driven wildfires using an inverse modelling approach. *Natural Hazards and Earth System Sciences* **14**, 1491–1503. doi:10.5194/nhess-14-1491-2014.

Rios O, Pastor E, Valero MM, Planas E (2016) Short-term fire front spread prediction using inverse modelling and airborne infrared images. *International Journal of Wildland Fire* **25**, 1033–1047. doi:10.1071/WF16031.

Rochoux MC, Ricci S, Lucor D, Cuenot B, Trouvé A (2014) Towards predictive data-driven simulations of wildfire spread - Part I: Reduced-cost ensemble Kalman filter based on a polynomial chaos surrogate model for parameter estimation. *Natural Hazards and Earth System Sciences* **14**, 2951–2973. doi:10.5194/nhess-14-2951-2014.

Rochoux MC, Emery C, Ricci S, Cuenot B, Trouvé A (2015) Towards predictive data-driven simulations of wildfire spread – Part II: Ensemble Kalman Filter for the state estimation of a front-tracking simulator of wildfire spread. *Natural Hazards and Earth System Sciences* **14**, 2951–2973. doi:10.5194/nhess-14-2951-2014.

Rossi L, Molinier T, Akhloufi M, Tison Y, Pieri a (2010) A 3D vision system for the measurement of the rate of spread and the height of fire fronts. *Measurement Science and*

Technology **21**, 105501. doi:10.1088/0957-0233/21/10/105501.

Rossi L, Molinier T, Pieri A, Akhloufi M, Tison Y, Bosseur F (2011) Measurement of the geometric characteristics of a fire front by stereovision techniques on field experiments. *Measurement Science and Technology* **22**, 125504. doi:10.1088/0957-0233/22/12/125504.

Rossi L, Molinier T, Akhloufi M, Pieri A, Tison Y (2013) Advanced stereovision system for fire spreading study. *Fire Safety Journal* **60**, 64–72. doi:10.1016/j.firesaf.2012.10.015.

Rudz S, Chetehouna K, Séro-Guillaume O, Pastor E, Planas E (2009a) Comparison of two methods for estimating fire positions and the rate of spread of linear flame fronts. *Measurement Science and Technology* **20**, 115501. doi:10.1088/0957-0233/20/11/115501.

Rudz S, Chetehouna K, Séro-Guillaume O, Pastor E, Planas E (2009b) Comparison of two methods for estimating fire positions and the rate of spread of linear flame fronts. *Measurement Science and Technology* **20**, 115501. doi:10.1088/0957-0233/20/11/115501.

Rudz S, Chetehouna K, Hafiane A, Laurent H, Séro-Guillaume O (2013) Investigation of a novel image segmentation method dedicated to forest fire applications. *Measurement Science and Technology* **24**, 75403. doi:10.1088/0957-0233/24/7/075403.

Shakesby RA (2011) Post-wildfire soil erosion in the Mediterranean: Review and future research directions. *Earth-Science Reviews* **105**, 71–100. doi:10.1016/j.earscirev.2011.01.001.

Sobel IE (1970) Camera Models and Machine Perception. Stanford University.

Stow DA, Riggan PJ, Storey EJ, Coulter LL (2014) Measuring fire spread rates from repeat pass airborne thermal infrared imagery. *Remote Sensing Letters* **5**, 803–812. doi:10.1080/2150704X.2014.967882.

Sullivan AL (2009a) Wildland surface fire spread modelling, 1990–2007. 1: Physical and quasi-physical models. *International Journal of Wildland Fire* **18**, 349–368. doi:10.1071/wf06143.

Sullivan AL (2009b) Wildland surface fire spread modelling, 1990–2007. 2: Empirical and quasi-empirical models. *International Journal of Wildland Fire* **18**, 369–386. doi:10.1071/wf06142.

Sullivan AL (2009c) Wildland surface fire spread modelling, 1990–2007. 3: Simulation and mathematical analogue models. *International Journal of Wildland Fire* **18**, 387–403. doi:10.1071/wf06144.

Toulouse T, Rossi L, Celik T, Akhloufi M (2015) Automatic fire pixel detection using image processing: a comparative analysis of rule-based and machine learning-based methods. *Signal, Image and Video Processing* 1–8. doi:10.1007/s11760-015-0789-x.

Veraverbeke S, Sedano F, Hook SJ, Randerson JT, Jin Y, Rogers BM (2014) Mapping the daily progression of large wildland fires using MODIS active fire data. *International Journal of Wildland Fire* **23**, 655. doi:10.1071/WF13015.

Zajkowski TJ, Dickinson MB, Hiers JK, Holley W, Williams BW, Paxton A, Martinez O, Walker GW (2016) Evaluation and use of remotely piloted aircraft systems for operations and research - RxCADRE 2012. *International Journal of Wildland Fire* **25**, 114–128. doi:10.1071/WF14176.

Automated tracking of active wildland fires through aerial infrared imaging and unsupervised edge detectors

Table 1. Input parameters for the unsupervised edge detection algorithm. Values used for all study cases.

Parameter	Value
Gaussian filter standard deviation (σ)	*4*
Hysteresis threshold candidates	$c = \{(c1, c2) \mid c1 < c2\}$ such that $c1 \in [0.2, 0.7] \wedge (\exists k \in \mathbb{Z})[c1 = 0.05 \cdot k]$ $c2 \in [0.5, 0.9] \wedge (\exists k \in \mathbb{Z})[c2 = 0.05 \cdot k]$

Table 2. Most relevant properties of the selected experimental footages.

Footage no.	Experiment	Size of the experimental area	Fuel complex	Recording position
1	Laboratory combustion table	4.5 m^2	Straw	Fixed elevated (non-nadir) position
2	Small-scale field burning	136 m^2	Aleppo pine saplings	Fixed elevated (non-nadir) position
3	Large-scale field experiment	4 ha	Mallee shrubland	Hovering helicopter
4	Large-scale field experiment	9 ha	Heath shrubland	Hovering helicopter
5	Large-scale field experiment	25 ha	Mallee woodland	Hovering helicopter

Table 3. Main technical properties of the TIR camera used.

Commercial name	AGEMA Thermovison 570-Pro, FLIR Systems
Detector type	Focal Plane Array (FPA)
Spectral range (μm)	7.5 – 13
Spatial resolution (pixels)	240 x 320
Field of view (deg.)	24 x 18
Temporal resolution (fps)	5

Table 4. Benchmarking of the proposed algorithm when optimum thresholds were approximated by the average corresponding to the first 50 frames.

Footage no.	Average FoM		Average Mean Distance (m)		Average Baddeley Distance (m)	
	Updated threshold	Constant threshold	Updated threshold	Constant threshold	Updated threshold	Constant threshold
1	0.7415	0.7091	0.0188	0.0231	0.0181	0.0221
2	0.6227	0.5648	0.3383	0.4800	0.3439	0.4747
3	0.6783	0.6783	4.0247	4.0247	3.8089	3.8089
4	0.6774	0.6805	3.7701	3.7175	3.7364	3.6931
5	0.6529	0.6538	10.016	9.9782	9.9634	9.9281
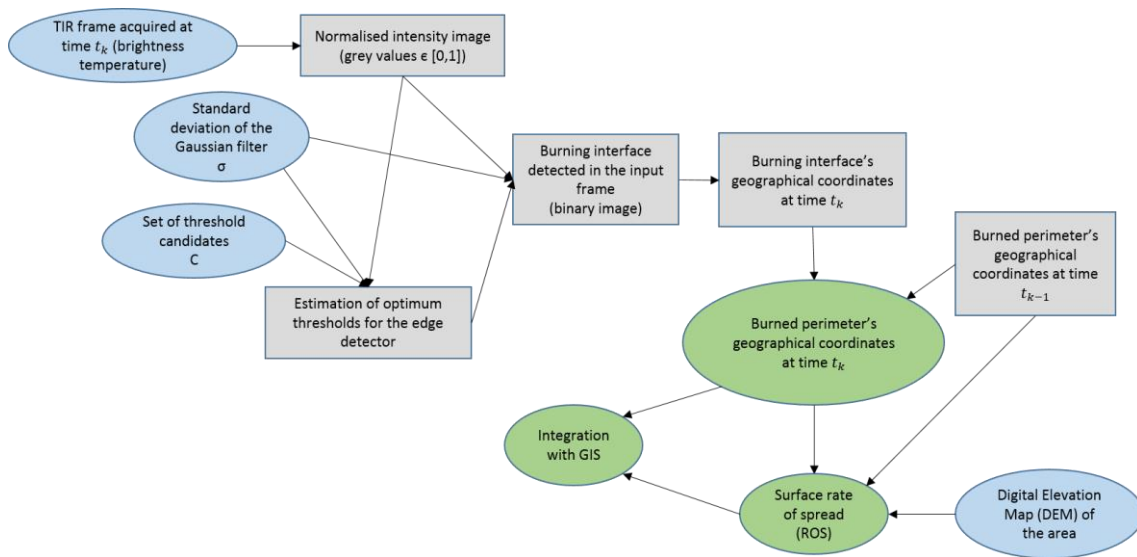
Figure 1. Block diagram of the complete algorithm. Blue ellipses represent inputs, green ellipses represent outputs and grey rectangles, intermediate computation steps.

Automated tracking of active wildland fires through aerial infrared imaging and unsupervised edge detectors
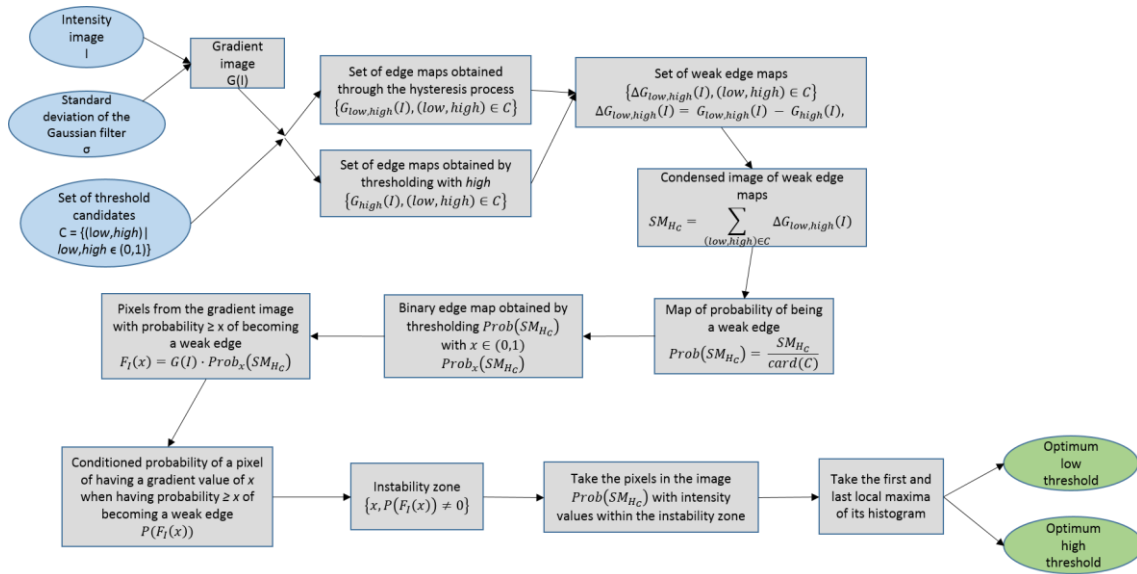
Figure 2. Block diagram of the optimum thresholds' estimation for the edge detector. Blue ellipses represent inputs, green ellipses represent outputs and grey rectangles, computation steps.
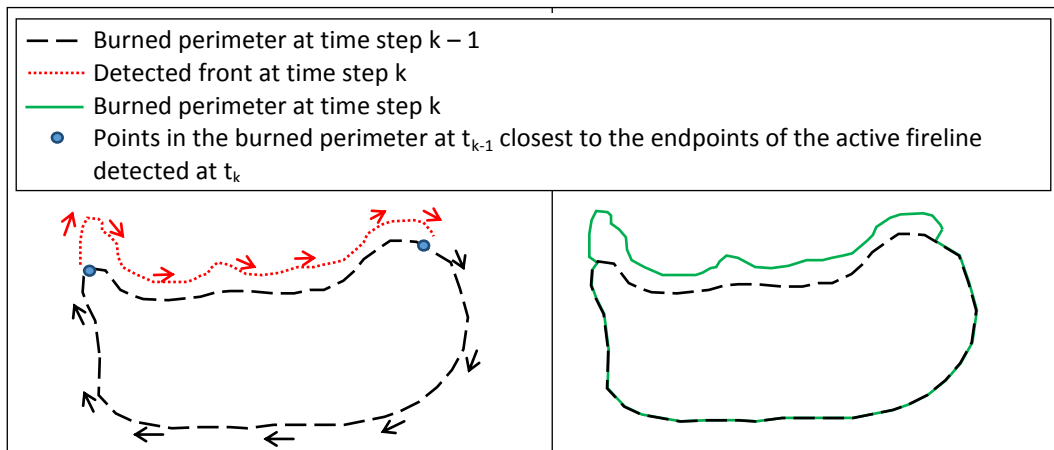
Figure 3. Sketch of the fire perimeter updating process. Newly burned area is added to the previously computed fire scar.
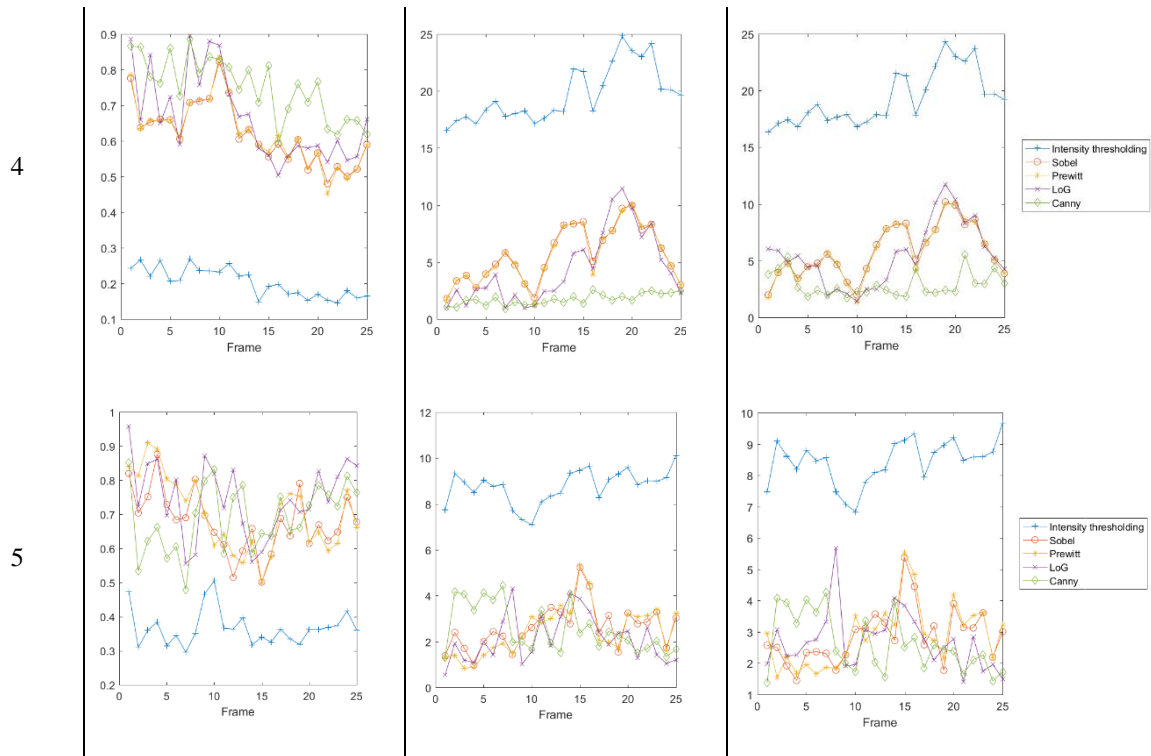
Footage no.	TIR spectrum	Visible spectrum
1	1.5 m	
2	17 m	
3	200 m	
4	300 m	
5	500 m	
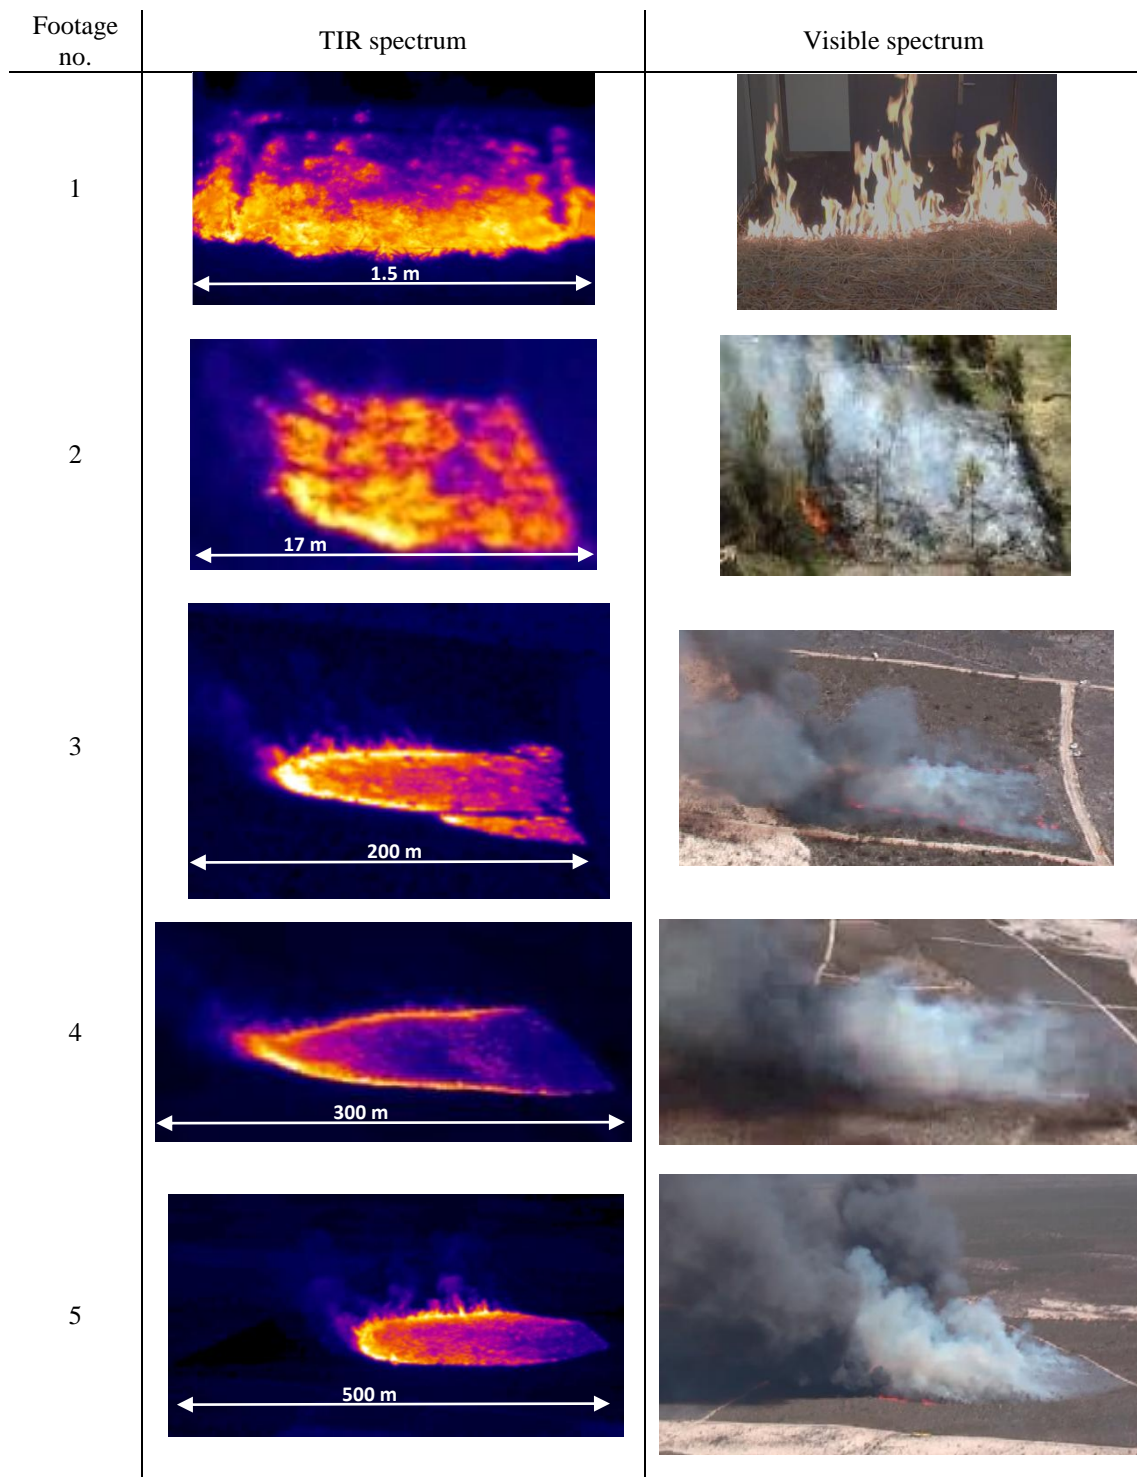
Figure 4. Sample frames of the analysed TIR footages (left column). Visible images are also displayed for comparison. Notice the little utility of visible images in a real-scale operational scenario, mainly due to the presence of smoke.

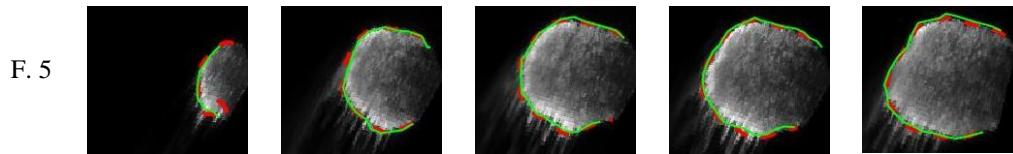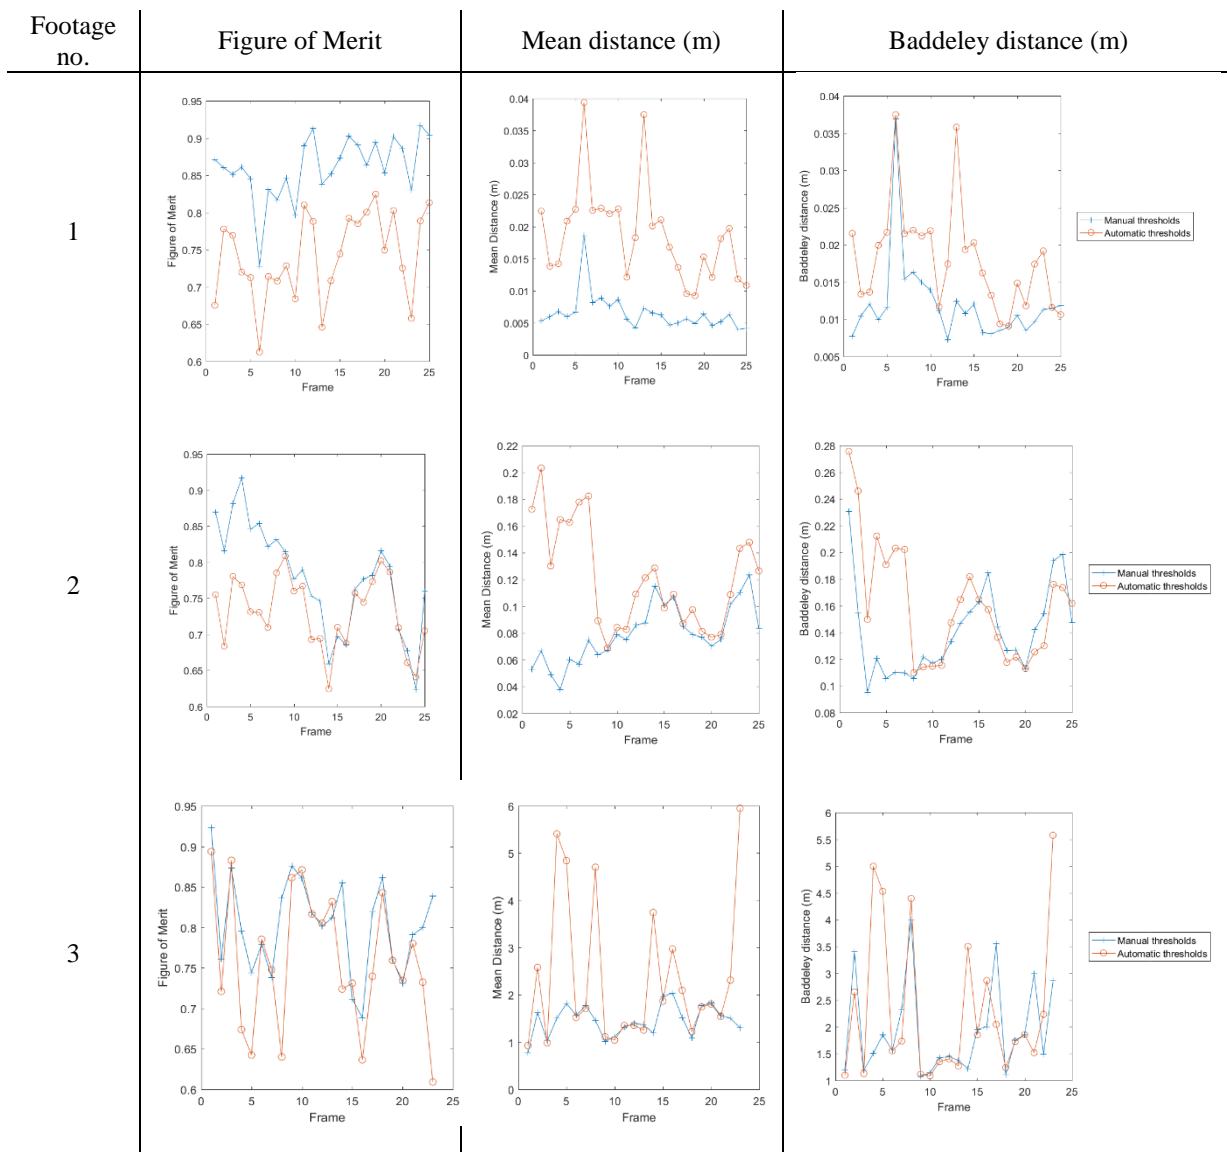Automated tracking of active wildland fires through aerial infrared imaging and unsupervised edge detectors

Footage no.	Original TIR frame	Global intensity thresholding	Sobel gradient operator	Prewitt gradient operator	Laplacian of a Gaussian (LoG)	Canny detector
1						
2						
3						
4						
5						

Figure 5. Qualitative results of the application of the different edge detection methods to some representative frames from all studied footages.
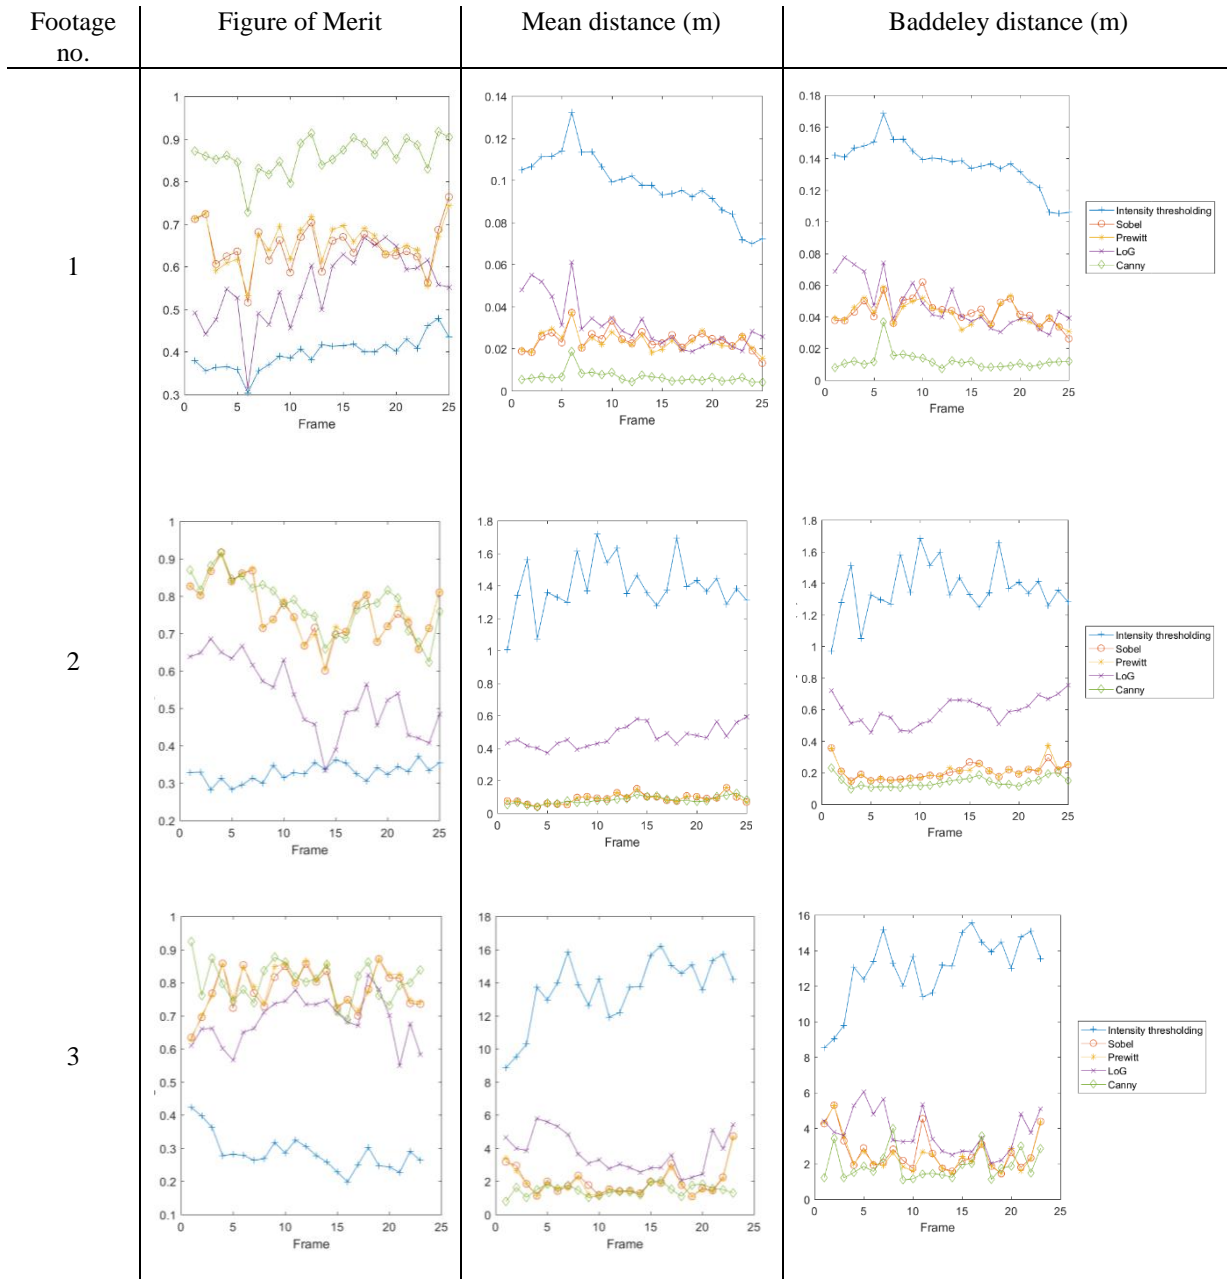
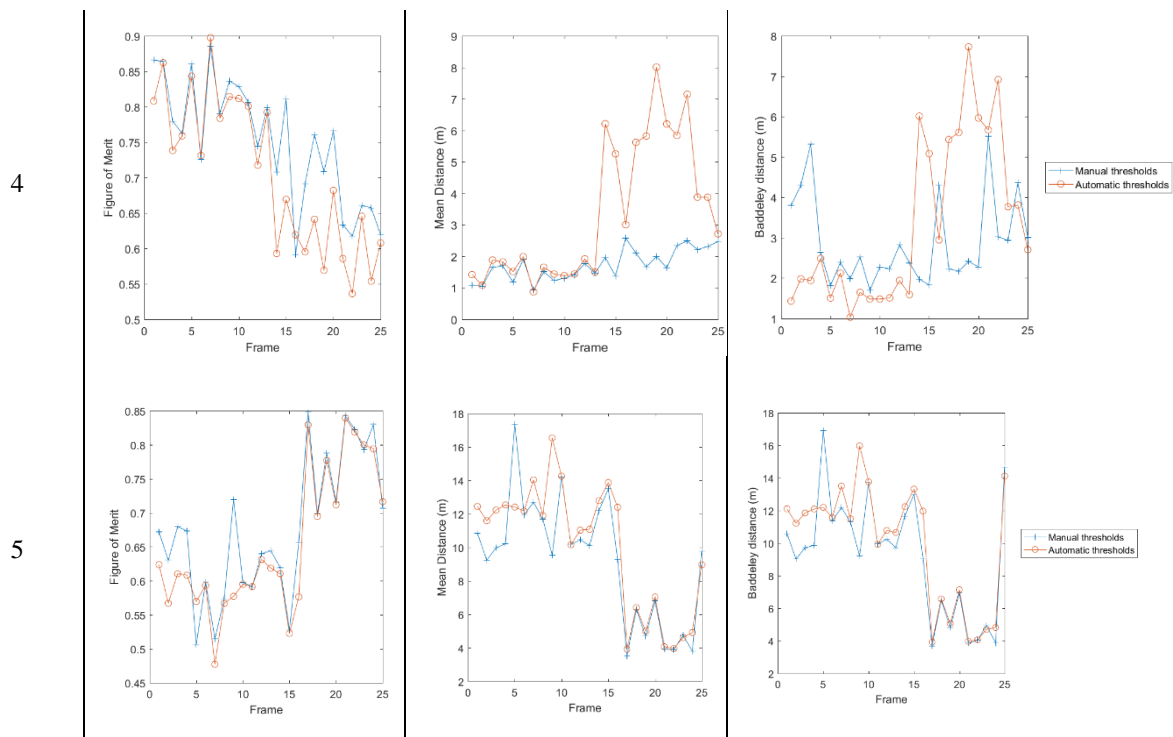| Footage no.	Figure of Merit	Mean distance (m)	Baddeley distance (m)
1			
2			
3			

Figure 6. Similarity indices, for 25 frames in each footage, between ground truth isochrones and the output of different algorithms. Best matching is obtained when FoM values approach 1 and distance-based indices approach 0. Edge detectors' parameters were optimised manually in all cases.
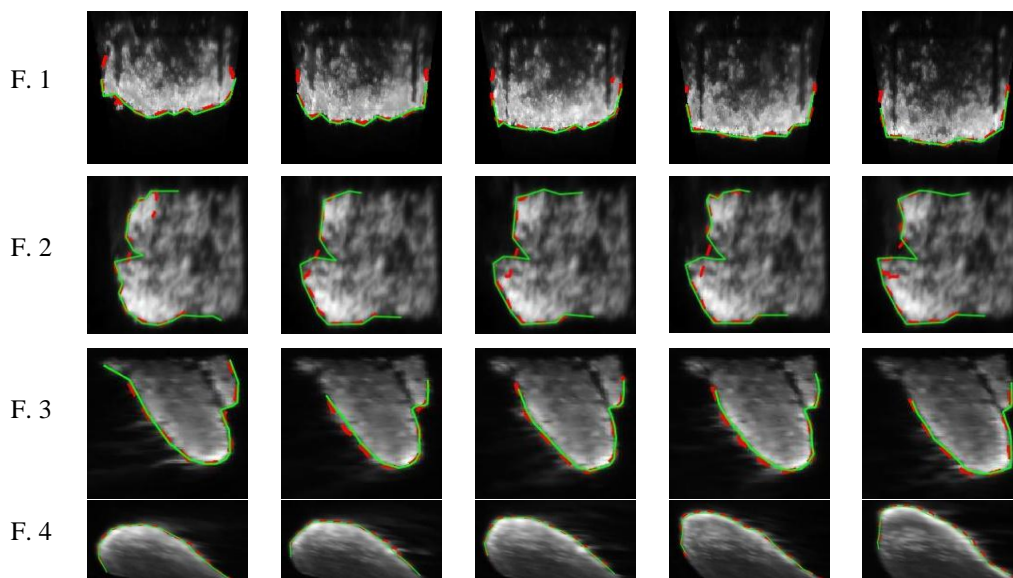
F. 5

Figure 7. Sample fire fronts detected by the automatic algorithm (red dashed line), superimposed on original georeferenced TIR images and compared to the corresponding ground truth, i.e. manually delineated fronts (solid green line). *F.* stands for *footage*.

Footage no.	Figure of Merit	Mean distance (m)	Baddeley distance (m)
1			
2			
3			

Figure 8. Similarity indices, for 25 frames in each footage, between the ground truth isochrones and the output of the Canny algorithm with manually optimised or automatically estimated parameters.
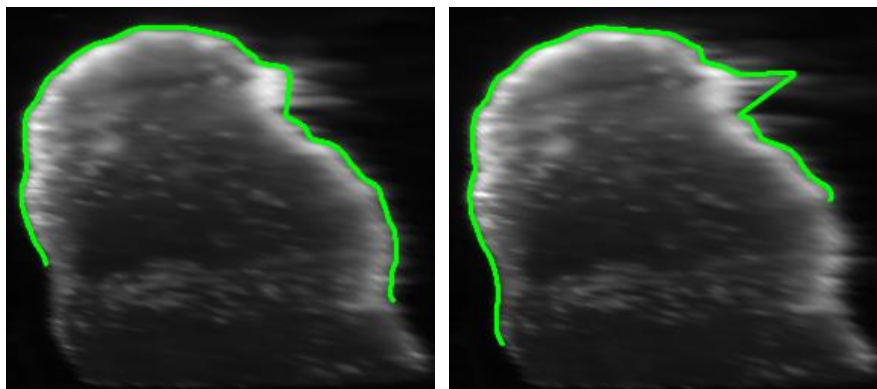

Figure 9. Example of successful (left) and unsuccessful (right) behaviour of the algorithm in the presence of flames. Frames belong to footage no. 4 and are spaced 2 seconds. Automatically detected fire fronts are displayed in green.

Automated tracking of active wildland fires through aerial infrared imaging and unsupervised edge detectors
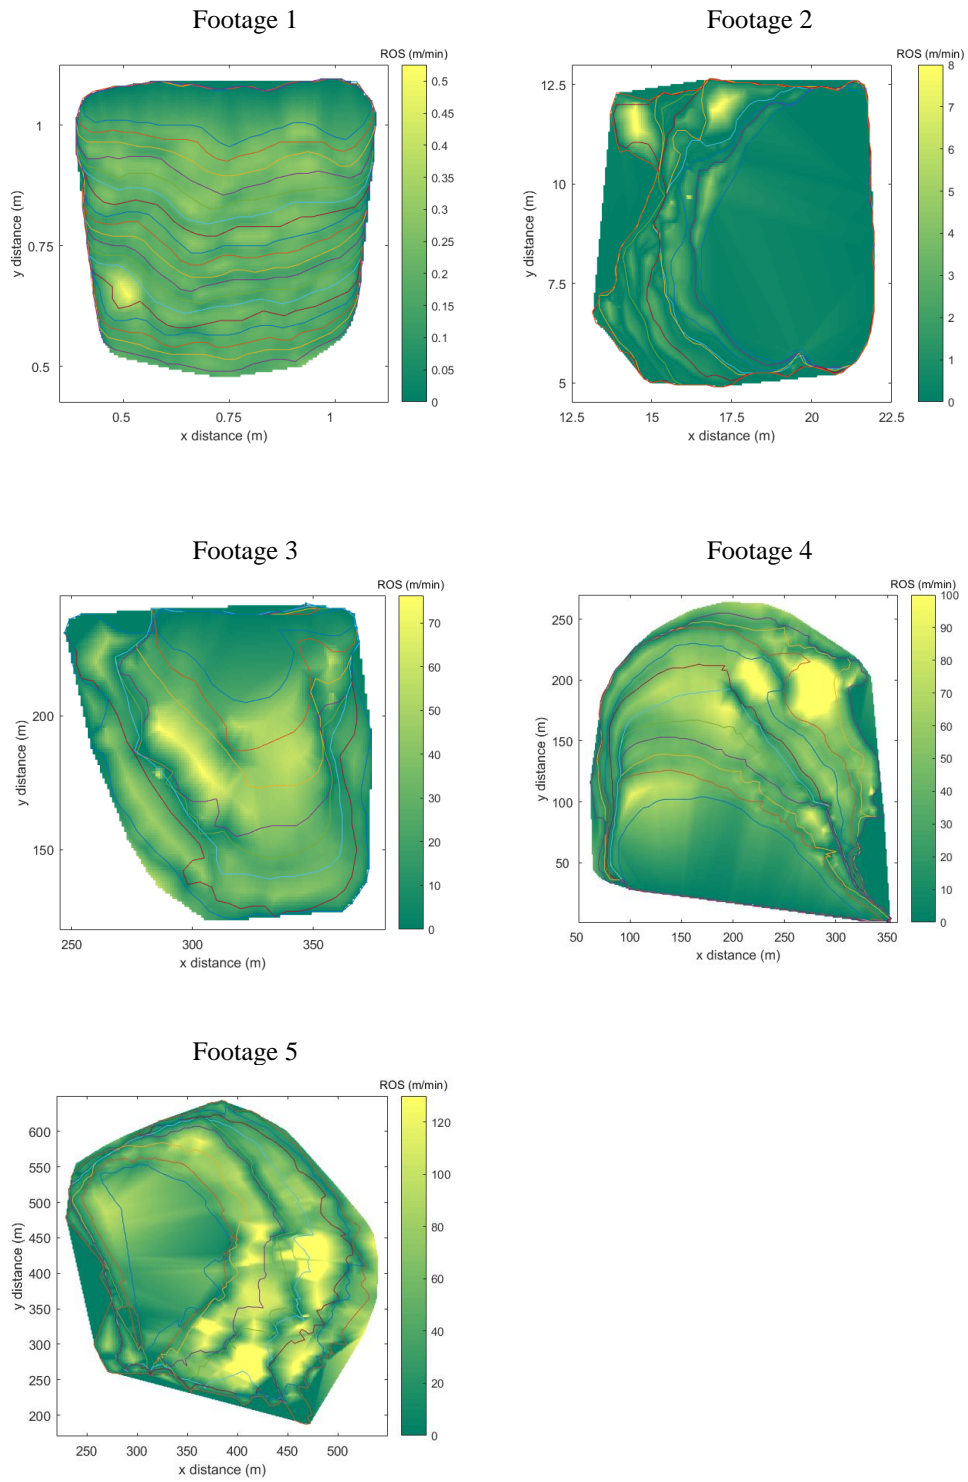

Figure 10. ROS fields obtained automatically for the 5 studied footages.

Automated tracking of active wildland fires through aerial infrared imaging and unsupervised edge detectors
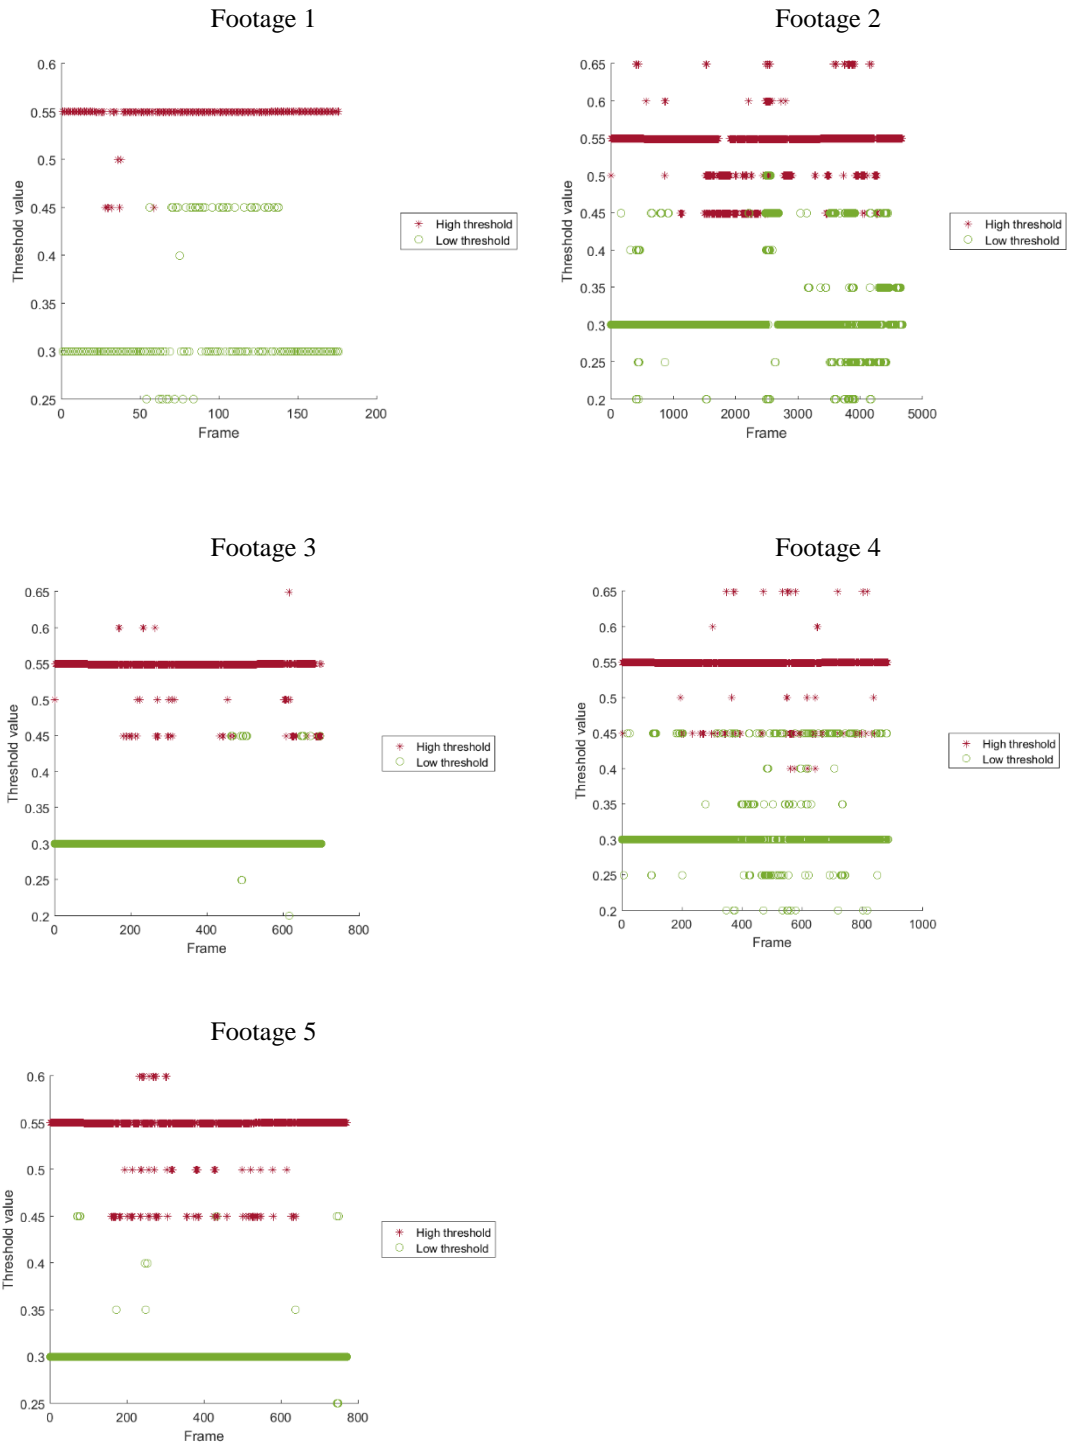

Figure 11. Optimum thresholds automatically estimated by the unsupervised edge detector.