Where is the rat? Tracking in low contrast thermographic images

Guillaume-Alexandre Bilodeau[1], Ramla Ghali[1], Sébastien Desgent[2]
Pierre Langlois[1], Rana Farah[1], Pier-Luc St-Onge[1], Sandra Duss[2], and Lionel Carmant[2]
[1]École Polytechnique de Montréal
P.O. Box 6079, Station Centre-ville Montréal (Québec), Canada, H3C 3A7
[2]Pediatry, Sainte-Justine Hospital
3175, Côte Ste-Catherine, Montréal, (Québec), Canada, H3T 1C5
guillaume-alexandre.bilodeau@polymtl.ca, ramla.ghali@polymtl.ca, sdesgent@gmail.com
pierre.langlois@polymtl.ca, rana.farah@polymtl.ca, pier-luc.st-onge@polymtl.ca
duss.sandra@gmail.com, lionel.carmant@umontreal.ca

Abstract

This paper presents a method to track an animal in low-contrast thermographic images in order to obtain its body temperature. This work was done in the context of the study of atypical febrile seizures. To solve this tracking problem, we propose a method based on morphological operations on the area to track using regions resulting from consecutive frame differences. A Gaussian model is then used to classify tracked area pixels into animal and background pixels to further remove outliers. The temperature of the animal is taken as the mean of the tracked area. Experimental results show that we obtain, in general, temperature estimation within 1^{o}C from ground-truth for videos as long as 16000 frames.

1. Introduction

Atypical febrile seizures (prolonged, lateralized, or repetitive seizures with fever) have important clinical implications because of their recognized association with the mesial temporal lobe epilepsy syndrome (MTLE), which is the most common of intractable epilepsies. Retrospective studies suggest that between 30% and 60% of patients with MTLE have histories of atypical febrile seizures occurring early in childhood ([5]). Despite these findings, the events that lead to atypical febrile seizures remain poorly understood, and the suggested risk factors for this condition, which include neurologic and/or perinatal abnormalities, early-life stress and a low threshold temperature at the onset of convulsions, have yet to be validated with animal models.

Thus, the development of an animal model of atypical hyperthermic seizures is a critical step in understanding the progression from febrile seizures to MTLE. In that trend and during the past ten years, researchers have developed animal models based on a two-hit hypothesis where an early-life insult, such as the combination of an artificial cortical lesion (e.g. microgyria) one day after birth (P1) or the induction of chronic early-life stress during the first nine days of life (e.g. daily cortisol administration) followed by prolonged hyperthermic febrile seizures occurring later at P10 in early postnatal life, lead to predisposition to temporal lobe epilepsy in adult male rats ([10]). In that model, the P10 animals are subjected to a hyperthermia experiment where their core temperature is increased to 45-50^{o}C to provoke a generalized convulsion (GC) ([9]). Previous experiments, using a rectal probe during hyperthermia, have shown that the threshold temperature and latency of GCs were significantly lower in insulted pups than in controls [10]. However, the classical use of thermometers of different types such as internal rectal electrodes, implantable probes or external infrared paw or tail sensor clips can be quite stressful for the animals and sometimes cause incontrollable variability in data acquisition that may partly hide the true effects of a treatment. Therefore, the development of non-invasive ways to measure the animal core temperature generates tremendous interest in this field of research. Using real-time thermographic images represents a promising advantage in that context by eliminating stressful probes for physiological measurements in the subjects and by getting more accurate estimates of possible differences between experimental conditions or treatments. Thus, in an ideal scenario, it is very important to be able to monitor accurately the temperature of the subject regularly without disturbing it during the experiment. In general, an error of about 1^{o}C is acceptable in this case.

An efficient way to obtain a large amount of data for this

experiment is to use a thermographic camera that performs temperature measurement at 30 Hz. Every pixel of an image obtained with a thermographic camera corresponds to a temperature measurement. By selecting in each frame the pixels that correspond to the body of a subject, it is then possible to get continuous temperature measurement of the animal. The challenge is to select pixels that are on the body of the subject. The goal of this work is thus to select pixels on the body of a subject at every f frames in the context of hyperthermia experiments mimicking fever induced atypical febrile seizures.

Problem definition: For the hyperthermia experiment, a single subject is placed in a Plexiglas box approximately 0.027 m^3 (1 $cu - ft$) in volume. The initial temperature inside the box is approximately 20 ^{o}C. A flow of hot air is forced into the box, progressively heating its contents to a temperature of approximately 50 ^{o}C over a period of 8 to 12 minutes. A thermographic camera's lens is placed through a hole pierced in the box cover, giving an overhead view of the subject. At the beginning of the experiment, the subject can be easily segmented from the background with simple thresholding, as shown in Figure 1(a). As the temperature inside the box increases, the rat's temperature becomes very similar to the temperature of the box' floor, as shown in Figure 1(b). Then, a temperature inversion occurs and the rat's temperature is lower than the temperature of the box' floor, as shown in Figure 1(c and d). In these two latter cases, tracking the rat is difficult because of low image contrast (the contrast is around 20 gray levels, or 2.5 ^{o}C). Furthermore, the rat sometimes urinates, which creates low temperature spots as shown in Figure 1(d).

The combination of low contrast, temperature inversion, and continuous temperature variation of the subject and of the background preclude the use of simple thresholding or classification by training. We postulate that tracking is required. Once the rat's position is estimated, its temperature can be measured, which is the ultimate requirement for this experiment.

Related works: Some researchers have studied automatic monitoring of the temperature of subjects. Notably, in the work of Sunderam et al. [11], thermographic images of the faces of six patients were acquired every hour and during seizure events as indicated by real-time EEG analysis. Thermal images were filtered manually to remove images where occlusion occurred. Since the face was in the middle of the image, the temperature measured was the maximum in the center region and there were no tracking requirements. Tracking was performed in the works of Torabi et al. [12] and Bilodeau et al. [1, 2] to measure the temperature of a moving subject. In Bilodeau et al. [1], a mean-shift tracker and a model based on the histogram of initial body temperature were used for tracking a shaved patch on the body of a rat. Torabi et al. [12] and Bilodeau et al. [2] instead used a particle filter tracker that models the shaved patch using edges around the patch and the mean intensity. Although these methods performed well for their intended body temperature measurement application, they are not applicable to our problem because of the low contrast between the body and its environment. Everything looks the same and there are no strong edges.

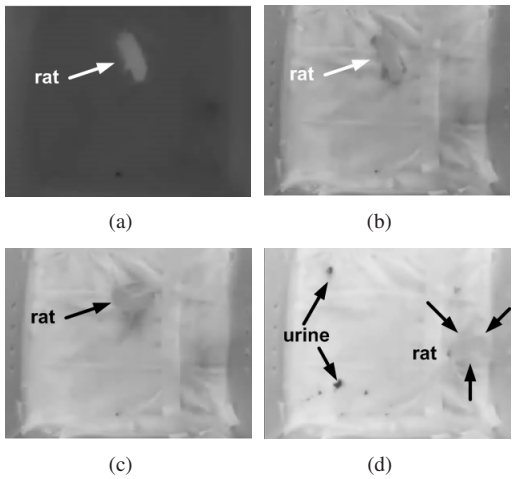

Figure 1. Animal to track at different frames during the experiments. a) At the beginning of the experiment (animal hotter than background), b) close to the moment of temperature inversion, c) after temperature inversion (animal colder than background), d) by the end of the experiment (just before the occurrence of the GC and lost of the righting reflex.

If we now consider the literature on low-contrast tracking, again, there is not a large body of works. Most of these works are interested in tracking small low-contrast objects ([3, 4, 6, 8]) in static background scenes. The main task is to distinguish small objects from camera noise. In Davies et al. ([3, 4]), objects are detected by using Daubechies wavelet filters over a sequence of frames to distinguish noise from small objects. Detected regions are validated using a Kalman filter. Padfield et al. [8] used a similar approach for detecting cells. In Hsieh et al. [6], frame differencing is used to locate object, and then noise is subtracted by learning a Gaussian noise model and neighbourhood encoding. The center of the detected region is used as a region of interest (ROI), and watershed segmentation is used to refine the ROI contours. Wang et al. [14] started also from frame differencing with thresholding. Detected objects are validated with a particle filter, but here, it is only point-based tracking with no concern with the object's shape. Underwater videos have also low-contrast [13]. In that case, background subtraction was used followed by pixel saliency measurements. Such measures are not discriminative enough because of the dynamic nature of our background that makes many regions salient. Our proposed method follows similar principles to the work of Hsieh et

al. [6], but instead, frame differencing is used to reshape and find the boundary of the object of interest.

Contributions: To solve this tracking problem, we propose a method based on morphological operation on the area to track. Our intuition to solve this problem is based on the following observation. The detection of the subject's body in a still image is difficult for a human observer (see Figure 1). However, the human observer can easily detect the subject if two frames are shown in which the subject moves even in a low contrast situation. Thus, our method is based on the modification of an initial area using morphological operations based on the motion of the subject. In experiments, we show that this approach allows the tracking of the subject throughout 10-12 minutes videos.

The paper is structured as follows. Section 2 presents our proposed method. Section 3 presents validating experiments and section 4 concludes the paper.

2. Methodology

Our proposed method was designed based on the following hypotheses:

- If no motion is detected, the temperature should be measured in the same area as for the previous frame;
- If there is motion in an area that intersects with the previous subject area, the area corresponding to the subject should be updated;
- If there is motion in a large area that does not intersect with the previous subject area, that new area should correspond to the subject area after a large displacement if its temperature is about the same (smooth temperature change constraint).

The requirement of smooth temperature change is necessary because the tracking environment might change drastically as a result of rat urination, which produces sudden apparition of large colder regions (see Figure 1d). Figure 2 illustrates the region updating process that allows subject tracking.

2.1. Initial subject area detection

At the beginning of the experiment, the subject's body is hotter than its environment. Simple thresholding may thus be applied. The initial subject area is the largest connected component A^1 such that

$$\forall(x,y) \in A^1, I^1(x,y) > \tau, \tag{1}$$

where $I^1(x,y)$ is the intensity of a pixel of the initial frame of the experiment video at coordinates x and y, and τ is a threshold. A^1 is the binary mask that represents the subject's body at frame 1. Note that the thermographic camera outputs are values between 0 and 255 representing temperatures based on the selected measurement range.

2.2. Motion detection and filters

Once the initial subject area is found, our proposed tracking method aims essentially at updating the subject area by adding and removing motion pixels given intersection and temperature smoothness criteria. To enforce the temperature smoothness criteria, background and urine pixels are removed to keep only the pixels corresponding to the animal.

Motion of the subject is evaluated by using frame differencing and thresholding to remove noise. Motion pixels I_M^t are obtained such that

$$\forall(x,y) \in I_M^t, \left|I^t(x,y) - I^{t-1}(x,y)\right| > \epsilon, \tag{2}$$

where ϵ is a threshold to remove noisy pixel that do not correspond to any motion and to remove small changes due to the heating of the environment. This value is preselected and kept constant for all frames. It is estimated based on the camera noise, which is around 2 or 3 intensity levels.

I_M^t will contain outlier motion areas that correspond in general to urine, which is significantly colder. To enforce the temperature smoothness criteria, these pixels are removed using statistics on the temperature in the current frame. These outlier pixels I_O^t are pixels such that

$$\forall(x,y) \in I_O^t, \left|I_M^t - \mu_I^t\right| > 3 * \sigma_I^t, \tag{3}$$

where μ_I^t is the mean and σ_I^t is the standard deviation of the intensities in the current frame. Both I_M^t and I_O^t are binarized resulting in B_M^t and B_O^t, and the filtered motion binary mask B_{FM}^t is

$$B_{FM}^t = B_M^t \cap \neg B_O^t. \tag{4}$$

To ensure that all outliers are removed, prior to applying Eq. 4, B_O^t is dilated using

$$B_O^t = B_O^t \oplus S, \tag{5}$$

where S is a 3×3 structuring element (a 3×3 matrix of ones). Finally, closing is used to fill the holes in B_{FM}^t with

$$B_{FM}^t = B_{FM}^t \bullet S. \tag{6}$$

2.3. New region selection and updating

The binary mask A^t that represents the subject's body at frame t is updated based on the number of motion pixels η in B_{FM}^t:

- if $|B_{FM}^t| < \eta$

$$A^t = A^{t-1}\, XOR\, B_{FM}^t \tag{7}$$

This means that motion pixels are added at positions outside A^{t-1} and pixels are removed at positions inside A^{t-1}. If there is no motion $A^t = A^{t-1}$.

(a) (b) (c)

Figure 2. Subject's body area modification via motion. The pixels corresponding to background were made darker to show the body area clearly. (a) initial body area, (b) consecutive frame difference (areas with motion), and (c) new body area by updating with motion pixels.

- else

$$A^t = B^t_{FM} \tag{8}$$

2.4. Temperature measurement and final tracking area

As already mentioned, to estimate the temperature of the rat in each frame of the video, we have to find an adequate process to distinguish the pixels of the background from those in the foreground. At this point, we only have a very coarse classification. At the outset, we thought of applying a classification algorithm such as the KNN and the Naive Bayes algorithms. However, these algorithms are not effective since we do not have a well-defined training set to classify the data appropriately in one of three classes (animal, background or urine), because the temperature of each class changes dynamically during the experiment. Furthermore, the background class has a large number of pixels compared to the other two classes. It extends to as much as 97% of the data. This makes the probability of classifying an instance larger in this class. An instance of the animal or urine class is more likely to be assigned to the background class and subsequently be misclassified.

Thus, we studied the distribution of data in the background and animal classes, and found an appropriate method of classification to remove outliers from animal data. The temperature distribution of the background class is approximately Gaussian (see Figure 3(a)). To a greater extent, as verification, we used the Q-Q plot graphical tool (Quantile to Quantile plot) [15]. A Q-Q plot is a graphical method to compare an unknown probability distribution with a normal distribution. In this case, the unknown probability distribution is the distribution of the background. In such a plot, a normal distribution is represented by a straight line. So, if a distribution is normal, the plot will be close to linear and not linear in the opposite case. Figure 3(b) shows the Q-Q plot of the class background for a typical frame. From this figure, the linearity of the points suggests that the distribution of the background class is normal. The normal distribution assumption holds for most frames. We are thus using a classifier based on this assumption.

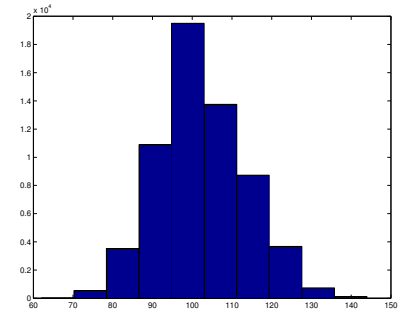

(a)

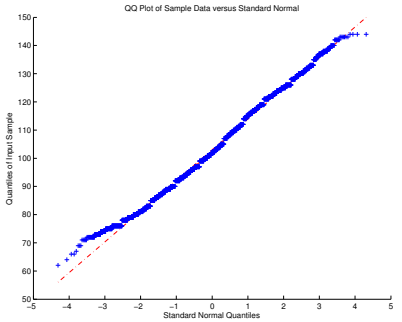

(b)

Figure 3. Analysis of background pixels. (a) Typical distribution, (b) Q-Q plot of the typical distribution.

For a normal distribution, the probability of dispersion of data around the mean (μ) is known as a function of standard deviation (σ). It satisfies the following property that is often referred as the empirical rule or as the three-sigma rule [7]: about 68% of the set fall within one standard deviation of the mean (between $\mu - \sigma$ and $\mu + \sigma$), about 95% of the values lie within two standard deviations of the mean, and finally nearly all values (99.7%) lie within three standard deviations of the mean.

Based on this rule, we reconstructed the animal class for each frame of the video as follows:

1. take a pixel $I^t(x, y)$ of class animal in A^t;

2. compare $|I^t(x,y) - \mu_{bg}|$ with σ_{bg}, where μ_{bg} and σ_{bg} are the estimated parameters of the background distribution for the current frame. $I^t(x,y)$ is added to a new animal area A^t_{TM} if it satisfies the condition defined by

$$\left|I^t(x,y) - \mu_{bg}\right| > \alpha * \sigma_{bg}, \tag{9}$$

where α is a factor to adjust the threshold. Otherwise, it is rejected.

3. Finally, the temperature of the rat is calculated by averaging the new animal area A^t_{TM} with

$$T^t = T_{min} + \left((mean(I^t(x,y) \in A^t_{TM})/255) * (T_{max} - T_{min})\right), \tag{10}$$

where T_{min} and T_{max} are the minimum and maximum temperatures measured by the thermographic camera, respectively, corresponding to pixel intensities of 0 and 255. The scale was set to linear on the thermographic camera.

From the result of this classification, the final subject area mask A^t used for tracking in the next frame is after closing

$$A^t = A^t_{TM} \bullet S. \tag{11}$$

3. Experimental validations

To validate our proposed method, we used eight thermographic videos of hyperthermia experiments. As stated in the introduction, in these experiments, P10 (10 days old) animals were individually placed in a Plexiglas box ($30 \times 30 \times 30cm$ with 16 small holes) in which a hair-dryer produced an airflow to increase their core temperature to 45-50^{o}C to provoke a generalized convulsion. To acquire the videos, a hole was pierced in the cover of the box to insert the camera lens. Thus, the animal was filmed from the top. The videos have a resolution of 320 $\times$ 240 and are grayscale. All videos are taken at about 29.4 frames per second. The thermographic camera (FLIR Thermovision A40M) was set with a linear scale between 20 and 53 ^{o}C. We tested our proposed method using three values of α (0.7, 1.0 and 1.3, see Eq. 9). Table 1 gives the specifications of each video.

To evaluate the tracking method, we used the RMS error on the temperature as in previous works ([12, 2]). A partial ground-truth was generated by selecting one out of every 200 frames over the whole video sequence and by manually indicating the animal area. The temperature value was calculated as in Eq. 10. This gives a set of F ground-truth temperatures T_{GT}. Recall that our objective is to automatically select pixels corresponding to the animal to obtain an estimate of the temperature of its body. Thus, we do not do comparison with actual body temperature measured with another instrument. By its specification, the thermographic camera has shown to be precise to 0.08^{o}C. The temperature measurements by the tracking algorithm for these frames were then compared with the ground-truth. The evaluation metric is the root mean square error defined as

$$T_{rms} = \sqrt{\frac{1}{F}\sum_{i=1}^{F}(T^i - T^i_{GT})^2}. \tag{12}$$

We selected this metric on temperature estimation error instead of using a metric on animal location accuracy, because the exact location of the animal is not required as long as the tracked area intersects to some extent with the actual rat position. That is, if the tracked area intersects with the animal area, our classifier will select only the intersecting pixel for temperature estimation. As such, our method does not depend on the exact location of the animal. Furthermore, the location of the animal cannot be reduced to a single point (x,y) and is instead a 2D distribution of points in the image. The temperature measurement error allows us to verify if there is intersection (small temperature difference) or not (large temperature difference) for each tested frame.

Table 2 gives the RMS errors found for the three values of α and for the temperature obtained only based on the center pixel of A^t. First, using the center pixel to obtain the temperature is not a good strategy because A^t does not

Table 1. Description of experiments

Video name	Duration (s)	Number of frames
Rat1	352	10222
Rat8	287	8310
Ratb	523	15174
Ratc	577	16752
Ratd	339	9831
Ratf	380	11181
Raton171ms	524	15741
Raton174mc	406	12178

Table 2. RMS temperator errors in ^{o}C obtained for the different videos

Video name	Center pixel	$\alpha = 0.7$	$\alpha = 1$	$\alpha = 1.3$
Rat1	3.97	0.68	0.68	0.80
Rat8	3.97	0.67	0.67	0.79
Ratb	7.25	0.78	0.87	1.12
Ratc	4.09	1.62	1.85	2.38
Ratd	1.39	1.43	2.42	2.96
Ratf	3.19	0.90	1.04	1.33
Raton171ms	5.07	4.38	4.64	4.95
Raton174mc	2.62	0.93	1.12	1.39

always have a convex shape. Furthermore, the temperature across the body of the animal is not uniform and thus a mean temperature needs to be computed. This shows that classification to obtain many subject's body pixels is required prior to temperature computation with Eq. 10. For all videos, we notice that with $\alpha = 0.7$, we get the smallest RMS error. This is explained by the fact that fewer pixels were rejected in A^t, because we reject only values that are close to the mean of the background class (with $\alpha = 0.7$, only 47.6% of data belongs to the background class). This means that the mean of the animal class and of the background class are nearby (low contrast, around 15-20 gray levels). Thus, a smaller α just removes pixels that obviously belong to the background, while a larger α starts to reject also pixels from the animal class. By removing many pixels from the animal class, we end up with a large proportion of outliers that are both far from the mean of the background and far from the mean of the animal.

In addition, we notice that in 5 of 8 videos (Rat1, Rat8, Ratb, Ratf, Raton174mc), RMS errors are smaller or near 1^oC (or 2% error), which is acceptable to separate trends as stated in the introduction. Note also that the temperature over the whole rat's body is not uniform. Results are not as good for videos of Ratc, Ratd, and Raton171ms. For those videos, the RMS error is between 2% and 5% (1.4 to 4.4^oC), which is high in comparison to the other videos. In fact, those videos contain some tracking errors caused by the rapid movement of the rat in conjunction with the presence of urine, which splits the rat's body into two parts. In that case, the tracker may get stuck at the position of the previous frame. If the rat makes a large movement again, tracking may recover by way of Eq. 8. On the other hand, if the rat stays still, or does not move much, the tracker will stay in the wrong position, and we get temperature measurements from the background. This is what happened for Raton171ms starting from frame 3000.

4. Summary and conclusions

In this paper, we presented a method to track the body area of a rat in low contrast thermographic videos. Mathematical morphology and frame differencing are used to adapt the body area from frame to frame. Results show that most of the time, we can successfully track the body of a rat for videos a long as 16 000 frames to obtain temperature estimation within 1^oC. The results allow the correct classification between insulted and control rats in 7 out of 8 videos from hyperthermia experiments (87.5%). In case of tracking failure, the tracking errors are caused by urine in conjunction with the motion of the rat that splits it and disconnects the rat area. Large movement often allows track recovery. Such situation of tracking failure may be detected by enforcing the temperature smoothness constraints more strictly, which would allow the detection of the track loss.

References

[1] G.-A. Bilodeau, M. Levesque, J. Langlois, P. Lema, and L. Carmant. Thermographic body temperature measurement using a mean-shift tracker. In *The International Conference on Bio-inspired Systems and Signal Processing (BIOSIGNALS 2009)*, pages 18 –24, January 2009. 56

[2] G.-A. Bilodeau, A. Torabi, M. Levesque, C. Ouellet, J. P. Langlois, P. Lema, and L. Carmant. Body temperature estimation of a moving subject from thermographic images. *Machine Vision and Applications*, 2010. 56, 59

[3] D. Davies, P. Palmer, and M. Mirmehdi. Detection and tracking of very small low contrast objects. In *9th British Machine Vision Conference*, pages 599–608, 1998. 56

[4] D. Davies, P. L. Palmer, and M. Mirmehdi. Robust tracker of small, fast-moving low-contrast targets. In *Proceedings of IX European Signal Processing Conference, Volume III*, pages 1545–1548. Typorama Editions, September 1998. 56

[5] J. Engel. Intractable epilepsy: Definition and neurobiology. *Epilepsia*, 42:1528 – 1167, 2001. 55

[6] F.-Y. Hsieh, C.-C. Han, N.-S. Wu, T. C. Chuang, and K.-C. Fan. A novel approach to the detection of small objects with low contrast. *Signal Processing*, 86(1):71 – 83, 2006. 56, 57

[7] D. Montgomery and G. Runger. *Applied statistics and probability for engineers*. John Wiley and Sons, inc., 2003. 58

[8] D. Padfield, J. Rittscher, and B. Roysam. Defocus and low cnr detection for cell tracking applications. In *Microscopic Image Analysis with Applications in Biology*, pages 1–6, September 2008. 56

[9] M. Scantlebury, S. Gibbs, B. Foadjo, P. Lema, C. Psarropoulou, and L. Carmant. Febrile seizures in the predisposed brain: a new model of temporal lobe epilepsy. *Ann Neurol.*, 58(1):41 – 49, 2005. 55

[10] M. Scantlebury, P. Ouellet, C. Psarropoulou, and L. Carmant. Freeze lesion-induced focal cortical dysplasia predisposes to atypical hyperthermic seizures in the immature rat. *Epilepsia*, 45(6):592 – 600, 2004. 55

[11] S. Sunderam and I. Osorio. Mesial temporal lobe seizures may activate thermoregulatory mechanisms in humans: an infrared study of facial temperature. *Epilepsy and Behavior*, 49(4):399–406, August 2003. 56

[12] A. Torabi, G.-A. Bilodeau, M. Levesque, J. Langlois, P. Lema, and L. Carmant. Measuring an animal body temperature in thermographic video using particle filter tracking. In *Advances in Visual Computing*, volume 5358 of *Lecture Notes in Computer Science*, pages 1081–1091. Springer Berlin / Heidelberg, 2008. 56, 59

[13] D. Walther, D. Edgington, and C. Koch. Detection and tracking of objects in underwater video. In *Computer Vision and Pattern Recognition, 2004. CVPR 2004. Proceedings of the 2004 IEEE Computer Society Conference on*, volume 1, pages I–544 – I–549 Vol.1, june 2004. 56

[14] X. Wang, G. Zhang, and Y. Chu. A robust approach to the detection and tracking of small targets with low contrast. In *VLSI Design and Video Tech., 2005. Proceedings of 2005 IEEE Int. Workshop on*, pages 296 – 299, May 2005. 56

[15] M. Wilk and R. Gnanadesiken. Probability plotting methods for the analysis of data. *Biometric*, 55(1), 1968. 58