Measuring an animal body temperature in thermographic video using particle filter tracking

Atousa Torabi[1], Guillaume-Alexandre Bilodeau[1], Maxime Levesque[2], J.M. Pierre Langlois[1], Pablo Lema[2], and Lionel Carmant[2]

[1] Department of Computer Engineering and Software Engineering,
École Polytechnique de Montréal,
P.O. Box 6079, Station Centre-ville Montréal (Québec), Canada, H3C 3A7,
{atousa.torabi, guillaume-alexandre.bilodeau, pierre.langlois}@polymtl.ca
[2] Pediatry, Sainte-Justine Hospital, 3175,
Côte Ste-Catherine,Montréal, (Québec), Canada, H3T 1C5
maxime.levesque@videotron.qc.ca, pplema@hotmail.com,
lionel.carmant@umontreal.ca

Abstract. Some studies on epilepsy have shown that seizures might change the body temperature of a patient. Furthermore, other works have shown that kainic acid, a drug used to study seizures, modify body temperature of a laboratory rat. Thus, thermographic cameras may have an important role in investigating seizures. In this paper, we present the methods we have developed to measure the temperature of a moving rat subject to seizure using a thermographic camera and image processing. To accurately measure the body temperature, a particle filter tracker has been developed and tested along with an experimental methodology. The obtained measures are compared with a ground truth. The methods are tested on a 2-hour video and it is shown that our method achieves the promising results.

1 Introduction

Neonatal seizures are convulsive events in the first 28 days of life in term infants or for premature infants within 44 completed weeks of conceptional age. Neonatal seizures are the most frequent major manifestation of neonatal neurologic disorders [1]. Population-based studies of neonatal seizures in North America report rates between 1 and 3.5 per 1000 live births. Initially thought to have little long-term consequences, we have more and more evidence that these seizures are deleterious to the developing brain [2]. Therefore, more emphasis is put on the treatment of these early life seizures. However, due to the immature connections in the neonatal brain, these seizures exhibit unusual clinical patterns, mimic normal movements and have primitive EEG patterns that are not easily recognizable. Therefore, one would be required to monitor all at-risk newborns continuously to confirm the epileptic nature of their events. At Ste-Justine Hospital, this typically represents 40 patients at any one time.

Several authors (e.g. [3], [4], [5]) have been interested in automatic detection of neonatal seizures using video recordings or electroencephalogram (EEG) pattern recognition. Preliminary data from our laboratories on an animal model of neonatal seizures suggest that body temperature monitoring could increase the automatic detection rate of significant seizure-related clinical events. The work of Sunderam et al. [6] has shown that seizures might change the body temperature of a patient. Furthermore, another work [7] has shown that kainic acid (KA), a drug used to study seizure, has impact of the body temperature of a laboratory rat. Since laboratory rats are used as an animal model to understand seizure, we are interested in monitoring continuously their temperature to understand its evolution during seizure and under KA. This paper presents our temperature measurement methodology.

There are few related works. In the work of Sunderam et al. [6], thermal 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 is the maximum in the center region and there were no tracking requirements. Other works in avian flu [8] and breast cancer [9] detection using thermography have not addressed automated continuous temperature monitoring. In our case, we are interested in continually monitoring the temperature of a rat that can move inside a perimeter, so we have to devise a more automated tracking and measuring method.

The paper is structured as follows. Section 2 presents our measurement methodology. Experimental results are presented and discussed in section 3. Section 4 concludes the paper.

2 Methodology

In this section, we first present the acquisition setup and then we present our measurement methodology.

2.1 Data acquisition

Our temperature sensor is a Thermovision A40M thermographic camera (FLIR Systems). Before acquiring animal videos, we first assessed the measurement error of the sensor for a still object in order to develop a baseline performance reference. From the manufacturer specifications, the accuracy is $\pm$ 2%, but the precision is not specified. To evaluate the camera's precision, we captured thermographic images of a wood tabletop from a fixed point of view continuously (at 27 frames/s) for approximately 30 minutes in a room at about $24°C$. The room temperature was not controlled. We selected an area of 20 $\times$ 20 pixels in the middle of the image. The camera was configured with a linear measurement range of $20°C$ to $40°C$, and pixels were quantified with 8 bits. That is, pixel values of 0 and 255 correspond to temperatures of $20°C$ and $40°C$, respectively.

This range provides a reasonable interval around the expected rat body temperature of approximately $30°C$. The interval between two adjacent pixel values is $0.078°C$. Without averaging a pixel region, the precision should be one-half of this interval, that is $0.039°C$. By averaging over a region, we may obtain a precision slightly better. The temperature of the tabletop in each frame is estimated by calculating the mean of the 10 hottest pixels in the region of interest. Assuming that the temperature changes smoothly, and because the temperature is not controlled, we computed the regression of the data using a 7^{th} order polynomial and computed the average fitting error. The average precision is the average fitting error, which is $0.021°C$ with a standard deviation of $0.026°C$. Figure 1 shows the measured temperature and the fitted polynomial. We did not validate the accuracy as we do not have the equipment to do so and in our measurements, we are only interested in the temperature variation, not in its absolute value.

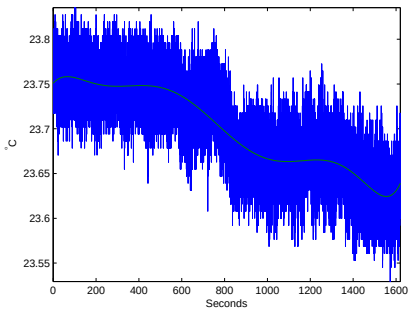

Fig. 1. Temperature measured of a table top during 1620 seconds, and fitted polynomial to evaluate the measurement error.

To acquire thermographic images during animal experiments, the rat is placed in a metal mesh cubic cage with an open top (see figure 2(a)) and the thermographic camera is angled down toward the cage. The camera is on a 525MV tripod (Manfrotto) and pointed toward the open top of the cage with an angle of approximately 20° with the vertical.

During initial experiments, we determined that the rat fur prevented precise measurements of the body temperature. Temperature measures on the whole body are not reliable as they depend on the thickness of the fur and visible area. We concluded that the rat should have an area of approximately 10 cm^2 that is shaved to measure precisely its temperature. Indeed, the head of the rat, another interesting region, which is warmer because it has less fur, is not always visible and it is occluded by a device (Neuralynx Cheetah System) to record local field potentials (LFP) inside the rat's brain. Observing the rat from the top and shaving a region on its back gives better results, as this region is almost constantly visible since the rat tends to remain on its four feet. However, using this strategy means that we have to use a tracking algorithm to follow

the shaved patch and discriminate it from the head. Figure 2(b) and 2(c) shows typical frames that must be processed.

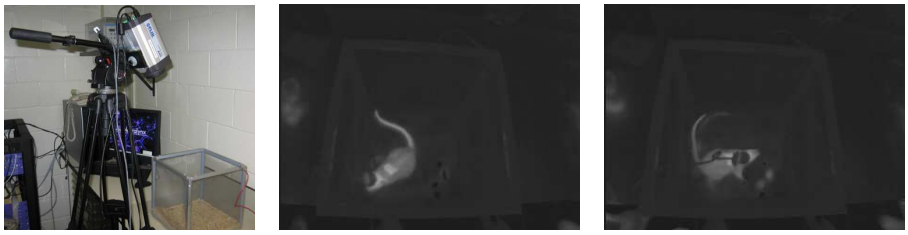

(a) (b) (c)

Fig. 2. Experimental setup. (a) Camera setup and mesh cage, (b) and (c) two frames of shaved patch to track without occlusion and with occlusion from the LFP recording device.

2.2 Particle filter tracking

For our purpose, the tracking algorithm must not lose track of the patch for the duration of the video and it must not be distracted by other hot areas like the head which is not exactly at the same temperature. We applied a particle filter tracker [10] adapted to our tracking problem. Particle filter tracking is robust to sudden movements and occlusions during long term video sequence, and thus should be appropriate for our conditions.

Our tracking algorithm is based on the following assumptions:

- the images are grayscale, with white (255) meaning hot, and black (0) meaning cold in a given range;
- the temperature of the rat is higher than its surrounding, particularly for the shaved patch and the head;
- the shaved patch can be occluded by the device that records LFP signals on the rat's head.
- the rat is always in the camera's field of view.

The particle filter is a Bayesian tracking method, which recursively in each time t approximates the state of the tracking target as a posterior distribution using a finite set of weighted samples. A sample is a prediction about the state of the tracking target.

In our work, the tracking target is modeled by an ellipse fitted to control points which lie on the edges on the perimeter of the shaved patch (figure 3). The ellipse is defined by

$$Ax^2 + Bxy + Cy^2 + F = 0; \tag{1}$$

Fig. 3. Fitted ellipse to control points on patch perimeter.

The parameters A, B, C and F are computed to optimize ellipse fitting. The particle filter state at each time step t is defined as a vector X_t of control points by

$$X_t = (Cp_x(t), Cp_y(t)) \tag{2}$$

where $Cp_x(t)$ and $Cp_y(t)$ are vectors of control points coordinates. We used the intensity and edge features as the observation model for our particle filter tracking method. For the edge feature, a dot product of unit normal vector of the ellipse at the positions of control points and the image gradient at the same positions is used as a measure. The dot product is used because the large gradient magnitude around the patch's perimeter is desirable and the gradient direction should be perpendicular to the perimeter. The gradient measure $\phi_g(s)$ for a given sample s is

$$\phi_g(s) = \frac{1}{N_{cp}} \sum_{k=1}^{N_{cp}} \left| n_k(s) \cdot g_k(s) \right|, \tag{3}$$

where $n_k(s)$ is the unit normal vector of the ellipse at the k^{th} control point, $g_k(s)$ is the image gradient at the same point, and N_{cp} is the number of control points. To facilitate the fusion of edge and intensity features, the gradient measure is normalized by subtracting the minimum gradient measure of sample set S and dividing by the gradient measure range as follow

$$\bar{\phi}_g(s) = \frac{\phi_g(s) - min_{s_j \in S}(\phi_g(s_j))}{max_{s_j \in S}(\phi_g(s_j)) - min_{s_j \in S}(\phi_g(s_j))}. \tag{4}$$

For the intensity feature, the Euclidean distance between the average intensity μ_s of the sample ellipse pixels and the average intensity μ_r of the template ellipse pixels is used as a measure. The template ellipse is the result of ellipse fitting to the chosen control points in the first frame of video sequence. The intensity measure $\phi_i(s)$ for a given sample s is

$$\phi_i(s) = \left| \mu_s - \mu_r \right|. \tag{5}$$

The intensity measure is normalized as follows

$$\bar{\phi}_i(s) = \frac{max_{s_j \in S}(\phi_i(s_j)) - \phi_i(s)}{max_{s_j \in S}(\phi_i(s_j))} \tag{6}$$

At time t, samples are selected with replacement from the sample set

$$S_{t-1} = \left\{ x_j^{t-1}, w_j^{t-1} \right\}_{j=1}^{N}, \tag{7}$$

where N is the number of samples, x_j^{t-1} is the control points coordinates of the j^{th} sample at time $t-1$, and $w_j^{t-1}$is its corresponding weight . Sample set S_{t-1} is the approximation of posterior distribution of the target state at time $t-1$. N_{s_j} samples are chosen with probability of w_j^{t-1}, which is derived for sample j by

$$w_j^{t-1} = \bar{\phi}_g(s_j^{t-1}) + \bar{\phi}_i(s_j^{t-1}). \tag{8}$$

This means that samples with high weights may be chosen several times (i.e., identical copies of one sample) and samples with small weights may not be chosen at all. In two consecutive frames, the particle filter state does not change dramatically. It is mostly a translation around ellipse major and minor axis and rotation around its center. In each time step t, samples are propagated in state space using a first order auto-regressive dynamical model defined as

$$X_t = X_{t-1} + \omega_t, \tag{9}$$

where X_t and X_{t-1} are particle filter states at time t and $t-1$ respectively, and ω_t is the stochastic part of the dynamical model. It is a multivariate Gaussian random variable and it corresponds to random translations and rotations of the ellipse. To measure temperature, the best sample is chosen by

$$s_b = argmax_{s_j \in S} \left\{ \bar{\phi}_g(s_j) + \bar{\phi}_i(s_j) \right\}. \tag{10}$$

The algorithm is the following:

1. Initialization. Manually select control points on the perimeter of the shaved patch and fit the ellipse.

   For each new frame:
2. Threshold the infrared image to keep the hottest regions (pixels with intensity more than 80). Find the gradient using a Sobel edge detector.
3. Select samples from samples set S_{t-1} based on the weights (equation 8).
4. Propagate samples using equation 9.
5. Compute the observation measurement to update new samples weights using equation 8.
6. Choose the best sample by equation 10 and calculate its temperature T_A by

$$T_A(f) = T_{min} + \left((mean(A)/255) * (T_{max} - T_{min}) \right), \tag{11}$$

where T_{min} and T_{max} are the minimum and maximum values of the temperature range selected for the camera and $mean(A)$ is the average intensity of the tracked region pixels. The area A for measuring the temperature is bounded by the ellipse fitted to the control points of the best sample.

3 Experimentation

We now present the experimentation methodology, results and discussion.

3.1 Experimentation methodology

To test our measurement method, we shot a 1hour and 57 minutes video of a Sprague-Dawley rat (Charles River Laboratories) during an experimentation using 6 mg/kg i.p. of KA (Sigma-Aldrich Canada Ltd). All experimental procedures conformed to institutional policies and guidelines (Sainte-Justine Research Center, Université de Montréal, Québec, Canada).

The camera setup was described in section 2.1. The video is 90706 frames at 12.89 frame per second with a 320 × 240 resolution and is compressed with Xvid FFDshow encoder (Quality: 100%) (http://sourceforge.net/projects/ffdshow). Our particle filter tracker is implemented in Matlab (The MathWorks) and is run on a Xeon 5150 2.66 GHz computer (Intel). For particle filter tracking, 8 control points are used for ellipse fitting. Control points are chosen with approximately the same distance from each other to cover the entire patch perimeter. Since in our experiment patch translational motion is more frequent than rotational motion, 75% of samples are specified for random translational motions and 25% of samples for random translational+rotational motions. Our tracking algorithm is tested by changing the number of samples N (45, 100 or 210 samples) and observation model (intensity or intensity+edge). The number of samples was chosen arbitrary to show their effect on tracking result and processing time. Temperature calculations were based on the 10 hottest pixels in the tracked area. T_{min} was 20 and T_{max} was 40.

Since we have a large quantity of data, to measure the performance of our tracking algorithm we used two metrics. For the first metric, we generated a partial ground truth by selecting frames at random over the whole video sequence. The four corners of the patch were selected to build a bounding polygon and the temperature value was calculated as in equation 11. This gives a set of ground truth temperatures T_{GT}. We selected F ($F = 450$) frames. 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

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

The second metric is based on the assumption that the temperature of the rat's body changes smoothly. We computed the regression of the temperatures using a 23^{rd} order polynomial (largest well-conditioned polynomial). Then, we computed the fitting error. This gives the average precision μ_m and its standard deviation σ_m. This fitting error is then compared with the fitting error obtained for a static target (a tabletop, see section 2.1) and for the ground truth. If tracking is good, we expect similar precision in the measurements.

3.2 Results and discussion

Figure 4 shows the results obtained for our test video by applying particle filter tracking using a combination of intensity and edge as observation model and 210 samples. The global decrease of the temperature between 0 and about 1800 seconds is caused by the KA. This phenomenon was previously observed [7] with a rectal thermometer at 15-30 minutes intervals. The local changes in temperature observed from 3000 seconds up to the end seem to be correlated with some seizure events (see figure 5). This conclusion is being further investigated with more experiments.

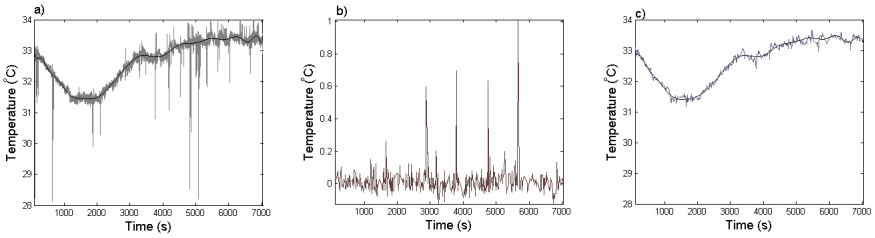

Fig. 4. Results from our tracker compared to ground truth. a) Temperature values obtained with our particle filter tracker and regression result. b) Error for the 450 ground truth points. c) Ground truth temperature values and regression result.

By comparing figure 4a) and figure 4c), one can notice that the tracking result is noisier than the ground truth. Sudden drops of temperature, more than $1°C$, are caused by tracking errors. In this experiment, occlusion by the LFP recording device caused the tracking errors around 150, 1850, 2100, 4850, 5100 seconds. In such cases, only a portion of patch is visible. During these times, particle filter tracking cannot track the patch correctly because of considerable changes in the patch shape and distraction by the rat's head. The other reason of tracking errors was the frame rate which was only 12.89 frames per second because the thermographic images were captured simultaneously with other data on the same computer. This low frame rate caused tracking errors when the rat moved too fast (i.e. wet dog shakes following some seizures) because of large displacement in the images around 3800, 4200, 4550, 5350, 6000, 6050, and 6700 seconds. Figure 4b) shows that after 3000 seconds, the tracking errors are mostly positive

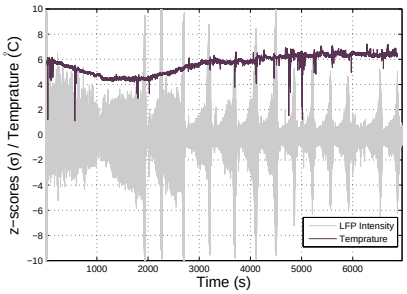

Fig. 5. Results from our tracker synchonized with LFP recordings. Temperatures were shifted by -27°C. The LFP recordings were normalized around a mean μ of 0 and a standard deviation σ of 1 (Z-scores: $z = \frac{x-\mu}{\sigma}$). Z-scores larger than $\pm 2\sigma$ correspond to seizure events.

because of the rat's quick movements which caused erroneously tracking the rat's head which is slightly warmer than the shaved patch. It is noticeable that some negative errors are not represented on this graph since the ground truth is composed of points selected at random and do not include all the tracking errors. However, it indicates that negative errors are less frequent than positive errors caused by erroneously tracking the head. Negative errors mean that tracking is neither on the head, nor on the patch.

Table 1. Average precision and root mean square error for each method and for the still object of section 2.1. N: Number of samples, Ob. model: Observation model, μ: average precision, σ: standard deviation, T_{rms}: root mean square error, FPS: frames per second.

Method	N	Ob. model	$\mu_m(\sigma_m)$ (°C)	T_{rms}(°C)	Time (s)	FPS
Still obj. (sec.2.1)	N/A	N/A	0.021(0.026)	N/A	N/A	N/A
Ground truth	N/A	N/A	0.079(0.100)	0.000	N/A	N/A
Particle filter	210	intensity	0.084(0.150)	0.104	59059	1.513
Particle filter	210	intensity+edge	0.082(0.131)	0.087	67909	1.317
Particle filter	100	intensity+edge	0.086(0.139)	0.117	32300	2.767
Particle filter	45	intensity+edge	0.093(0.188)	0.103	19989	4.472

Table 1 gives the values obtained for the metrics defined in section 3.1 and two associated processing times by changing the observation model and number of samples. Results show that by increasing the number of samples processing time increase but temperature measures change more smoothly and have average precision and standard deviation more similar to ground truth. We found out that a number of samples less than 45 is not sufficient for desirable tracking result and more than 210 does not improve tracking result but it increases the

processing time. For 100 samples, we get a larger RMS error than for 45 samples because the ground truth is only of 450 points, which is small compared to 90706 measures with the tracker. Thus, it does not cover all the errors and most errors may be located where there are no ground truth points. Results vary for different tracker runs because of the random variable of equation 9. This is particularly the case when there are few samples like 45. However, with 100 or 210 samples, the precision is better and thus the results will be more stable and should be consider a better choice than 45 samples. For small number of samples, it is the risk that tracker fails by sudden rotation and movements of the rat's body because samples are not distributed sufficiently in all directions around the tracking target. The best result that we obtained was with 210 samples using intensity+edge for observation model. Table 1 also shows that combining the intensity and edge features for observation model reduces RMS error compare to using only intensity feature. This is because modeling patch edges by an ellipse reduces tracking errors from the patch to the rat's head, which sometimes can have similar intensities, but different shapes.

It is interesting to notice that the ground truth precision is larger than the average precision obtained for a still object. At this point, we may hypothesize that is because the shaved area is deformable and its normal is not always aligned with the camera sensor's normal. Thus, the infrared radiation measured by the camera changes with the angle of the shaved area. Furthermore, as the shaved area is deformable, the skin thickness may vary regularly as it stretches depending on the rat position and attitude. Another possibility is that seizure events cause temperature changes that violate the smoothness constraints and increase the fitting error. We will test a rat in a control condition (without KA) to verify the attainable precision with a moving target. Given these results with our equipment, capture setup, and assuming smooth temperature variations, we can expect to observe phenomena that cause sudden temperature changes over a few frames larger than $0.2°C$.

4 CONCLUSION

This paper presented a methodology to measure the body temperature of a moving animal in a laboratory setting. Because of the experimental setup, uneven thickness of the fur with viewpoint and the possibility of occlusion, to measure the body temperature we needed to shave a region on rat's back. Since the head and shaved region may have different temperatures, tracking is required to measure temperature on the same body region continuously. We proposed a particle filter tracker based on shaved region intensity and shape.

Our method was tested on a 2-hour video sequence with a rat having seizures at regular intervals. Results show that our tracker achieves measurements with an RMS error less than $0.1°C$. Errors are caused by severe occlusions or by quick rat motions (wet dog shakes). Although we estimate we can observe phenomena causing changes of more than $0.2°C$, we do not obtain a precision similar to a still object. Part of this difference with camera precision is caused by the tracker,

while another part is caused by other reasons. We hypothesize that change in the orientation of the measured surface causes measurement errors, so it may not be possible to attain the precision obtained on a still object. Furthermore, in the test video, temperature changes may not be smooth and they may increase the fitting error by a polynomial.

As future works, we are interested in improving particle filter tracking by adding scaling to model motion to handle severe patch deformations because of changes in view angles. We will apply a filter to the temperature measures to remove outliers. We will also investigate the impact of changes of orientation of the measured surface. Finally, we are interested to apply this methodology in more experiments and automatically detect abnormal events based on changes in the body temperature.

References

1. Volpe, J.J.: Neonatal Seizures: Current Concepts and Revised Classification. Pediatrics **84**(3) (1989) 422–428
2. Carmant, L.: Mechanisms that might underlie progression of the epilepsies and how to potentially alter them. Advances in neurology **97** (2006) 305–314
3. Karayiannis, N., Srinivasan, S., Bhattacharya, R., Wise, M., Frost, J.D., J., Mizrahi, E.: Extraction of motion strength and motor activity signals from video recordings of neonatal seizures. Medical Imaging, IEEE Transactions on **20**(9) (Sep 2001) 965–980
4. Celka, P., Colditz, P.: A computer-aided detection of eeg seizures in infants: a singular-spectrum approach and performance comparison. Biomedical Engineering, IEEE Transactions on **49**(5) (May 2002) 455–462
5. Faul, S., Boylan, G., Connolly, S., Marnane, L., Lightbody, G.: An evaluation of automated neonatal seizure detection methods. Clinical Neurophysiology (7) (July 2005) 1533–1541
6. Sunderam, S., Osorio, I.: Mesial temporal lobe seizures may activate thermoregulatory mechanisms in humans: an infrared study of facial temperature. Epilepsy and Behavior **49**(4) (August 2003) 399–406
7. Ahlenius, S., Oprica, M., Eriksson, C., Winblad, B., Schultzberg, M.: Effects of kainic acid on rat body temperature: unmasking by dizocilpine. Neuropharmacology **43**(1) (July 2002) 28–35
8. Camenzind, M., Weder, M., Rossi, R., Kowtsch, C.: Remote sensing infrared thermography for mass-screening at airports and public events: Study to evaluate the mobile use of infrared cameras to identify persons with elevated body temperature and their use for mass screening. Technical Report 204991, EMPA Materials Science and Technology (2006)
9. Amalu, W.: Nondestructive testing of the human breast: the validity of dynamic stress testing in medical infrared breast imaging. In: Engineering in Medicine and Biology Society, 2004. IEMBS '04. 26th Annual International Conference of the IEEE. Volume 2. (2004) 1174–1177
10. Isard, M., Blake, A.: Condensation – conditional density propagation for visual tracking. International Journal of Computer Vision **29**(1) (1998) 5–28