Practical Computer Vision Techniques for Interactive Robotic Heads

Oscar Deniz[2], Javier Lorenzo[1], Modesto Castrillon[1], Luis Anton[1], Mario Hernandez[1] and Gloria Bueno[2]
[1]*Universidad de Las Palmas de Gran Canaria,*
[2]*Universidad de Castilla-La Mancha*
Spain

1. Introduction

In the last years there has been a surge in interest in a topic called social robotics. As used here, social robotics does not relate to groups of robots that try to complete tasks together. For a group of robots, communication is simple, they can use whatever complex binary protocol to "socialize" with their partners. For us, the adjective social refers to humans. In principle, the implications of this are much wider than the case of groups of robots. Socializing with humans is definitely much harder, not least because robots and humans do not share a common language nor perceive the world (and hence each other) in the same way. Many researchers working on this topic use other names like human-robot interaction or perceptual user interfaces. However, as pointed out in (Fong et al., 2003) we have to distinguish between conventional human-robot interaction (such as that used in teleoperation scenarios or in friendly user interfaces) and socially interactive robots. In these, the common underlying assumption is that humans prefer to interact with robots in the same way that they interact with other people.

Human-robot interaction crucially depends on the perceptual abilities of the robot. Ideal interaction sessions would make use of non-invasive perception techniques, like hands-free voice recognition or computer vision. Computer vision is no doubt the most useful modality. Its non-invasiveness is the most important advantage.

In this paper, a number of computer vision techniques for human-robot interaction are described. All of them have been used in a prototype social robot called CASIMIRO (Figure 1), an animal-like head that stands on a table and has the goal of interacting with people, see (Deniz, 2006) for details.

2. Omnidirectional vision

Most social robots built use two types of cameras: a wide field of view camera (around 70 deg), and a foveal camera. Recently, interest in omnidirectional vision has increased in robotics. Omnidirectional vision allows to capture images that span 360º. Four main techniques are being used to achieve this:

- Cameras with fish-eye lenses, which have a very short focal length (called dioptric systems.

Source: Advances in Robotics, Automation and Control, Book edited by: Jesús Arámburo and Antonio Ramírez Treviño,
ISBN 78-953-7619-16-9, pp. 472, October 2008, I-Tech, Vienna, Austria

Open Access Database www.i-techonline.com

- Cameras with curved (convex) mirrors mounted in front of a standard lens (called catadioptric systems). This is the most common variant.
- Sets of cameras mounted in a ring or sphere configuration. The information flow is time consuming.
- An ordinary camera that rotates around an axis and takes a sequence of images that span 360º. Mechanical looseness can appear.

Fig. 1. CASIMIRO.

The omnidirectional camera shown in Figure 2 is a catadioptric setup that gives the robot a 180 deg field of view, which is similar to that of humans. The camera is to be placed in front of the robot. The device is made up of a low-cost USB webcam, construction parts and a curved metallic surface looking upwards, in this case a kitchen ladle. As for the software, the first step is to discard part of the image, as we want to watch only the frontal zone, covering 180 degrees from side to side. Thus, the input image is masked in order to use only the upper half of an ellipse, which is the shape of the mirror as seen from the position of the camera.

A background model is obtained as the mean value of a number of frames taken when no person is present in the room. After that, the subtracted input images are thresholded and the close operator is applied. From the obtained image, connected components are localized and their area is estimated. Also, for each connected component, the Euclidean distance from the nearest point of the component to the centre of the ellipse is estimated, as well as the angle of the centre of mass of the component with respect to the centre of the ellipse and its largest axis. Note that, as we are using an ellipse instead of a circle, the nearness measure

obtained (the Euclidean distance) is not constant for a fixed real range to the camera, though it works well as an approximation. The robot uses this estimate to keep an appropriate interaction distance. The background model M is updated with each input frame:

$$M(k+1) = M(k) + U(k) \cdot [I(k) - M(k)] \tag{1}$$

, where I is the input frame and U is the updating function:

$$U(k) = \exp(-b \cdot D(k)) \tag{2}$$

$$D(k) = a \cdot D(k-1) + (1-a) \cdot |I(k) - I(k-1)| \tag{3}$$

a (between 0 and 1) and b control the adaptation rate. Note that M, U and D are images, the x and y variables have been omitted for simplicity. For large values of a and b the model adaptation is slow. In that case, new background objects take longer to enter the model. For small values of a and b, adaptation is faster, which can make animated objects enter the model.

Fig. 2. Omnidirectional camera.

The method described up to this point still has a drawback. Inanimate objects should be considered background as soon as possible. However, as we are working at a pixel level, if we set the a and b parameters too low we run the risk of considering static parts of animate objects as background too. This problem can be alleviated by processing the image D. For each foreground blob, its values in D are examined. The maximum value is found, and all the blob values in D are set to that level. Let the foreground blobs at time step k be represented as:

$$B_i = \{x_{ij}, y_{ij}\} ; i = 1, .., NB ; j = 1, .., N_i \quad (4)$$

There are NB blobs, each one with N_i pixels. Then, after (3) the following is applied:

$$m_i = \max_{j=1,..,N_i} D(x_{ij}, y_{ij}, k) ; i = 1, .., NB \quad (5)$$

$$D(x_{ij}, y_{ij}, k) = m_i ; i = 1, .., NB ; j = 1, .., N_i \quad (6)$$

With this procedure the blob only enters the background model when all its pixels remain static. The blob does not enter the background model if at least one of its pixels has been changing.

3. Face detection

Omnidirectional vision allows the robot to detect people in the scene, just to make the neck turn towards them (or somehow focus its attention). When the neck turns, there is no guarantee that omnidirectional vision has detected a person, it can be a coat stand, a wheelchair, etc. A face detection module should be used to detect people (and possibly facial features). Facial detection commonly uses skin-color as the most important feature. Color can be used to detect skin zones, though there is always the problem that some objects like furniture appear as skin, producing many false positives. Figure 3 shows how this problem affects detection in the ENCARA facial detector (M. Castrillon-Santana et al., 2005), which (besides other additional cues) uses normalized red and green color components for skin detection.

Fig. 3. Skin color detection. Note that wooden furniture is a distractor for facial detection. Both the bounding box and the best-fit ellipse are rather inaccurate (left).

In order to alleviate this problem, stereo information is very useful to discard objects that are far from the robot, i.e. in the background. Stereo cameras are nowadays becoming cheaper and faster. A depth map is computed from the pair of images taken by a stereo camera situated under the nose of the robot. The depth map is efficiently computed with an included optimized algorithm and library. The map is thresholded and an AND operation is performed between this map and the image that the facial detector uses. Fusion of color and depth was also used in (Darrell et al., 1998; Moreno et al., 2001; Grange et al., 2002). The results are shown in Figure 4. Note that most of the undesired wood colored zones are filtered out.

Fig. 4. Skin color detection using depth information.

4. Person recognition

In (Schulte et al., 1999) three characteristics are suggested as critical to the success of robots that must exhibit spontaneous interaction in public settings. One of them is the fact that the robot should have the capability to adapt its human interaction parameters based on the outcome of past interactions so that it can continue to demonstrate open-ended behaviour. CASIMIRO is intended to interact with people. Humans will be the most important "object" in its environment. Data associated to humans (gathered throughout the interaction) should be stored in memory, so that the robot could take advantage of previous experiences when interacting with them. Breazeal (Breazeal, 2002) argues that to establish and maintain relationships with people, a sociable robot must be able to identify the people it already knows as well as add new people to its growing set of known acquaintances. In turn, this capacity will be part of the robot's autobiographical memory.

In order to make this person memory possible, gathered data should be unambiguously associated to the correct person. Facial recognition would be the perfect approach. However, the experience of the author with face recognition is somewhat negative: face recognition still does not work well in unrestricted scenarios. Recognition rates fall as more time passes since the training samples were taken. Illumination, pose and expression variations normally reduce recognition rates dramatically.

Colour histograms of (part of) the person's body could also be used as a recognition technique. Colour histograms are simple to calculate and manage and they are relatively robust. The price to pay is the limitation that data in memory will make sense for only one day (at the most). Colour histograms of a person's body were used for short-term identification people in (Maxwell, 2003; Kahn, 1996; Maxwell et al., 1999) and also for people tracking (Krumm et al., 2000; Collins & Dennis, 2000).

CASIMIRO achieves person identity maintenance by using colour histograms in conjunction with a simple person tracking algorithm. Tracking is done in 1D, for the interesting position is the angle of the person with respect to the robot.

The implemented tracking algorithm is very simple. Each person is represented as a single point in two sets of horizontal positions (positions range from 0 to 180) at times t - 1 and t. The association of points between the two sets is obtained as that which minimizes the total sum of distances between points of the two sets. This minimization involves a factorial search, though it is practical for the number of people that will be expected to interact with the robot.

Ties can appear, for example in the case of crossings, see the example of Figure 5. These ties are broken by selecting the association with lowest variance of distances, 1 with A and 2 with B in the case of the example. This always selects non-crossings.

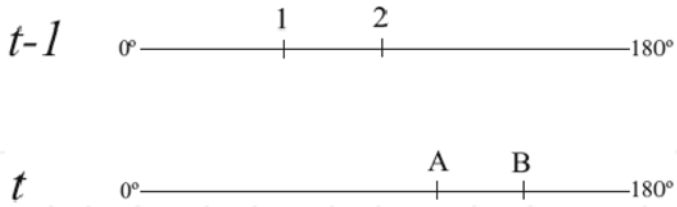

Fig. 5. Tie in sum of distances. The sum of distances |1 - A| + |2 - B| is equal to |1 - B| + |2 - A|. Without further information, we cannot know if the two individuals have crossed or not.

Crossings are detected by considering that, in a crossing, there is always a fusion and a separation of person blobs. Person blobs are detected by the omnidirectional vision system (see above). Fusions and separations are detected as follows:

- A blob fusion is detected when the number of blobs in the whole omnidirectional image decreases by one at the same time that one of the blobs increases its area significantly.
- A blob separation is detected when the number of blobs in the image increases by one at the same time that a fusioned blob decreases its area significantly.

The only way to know if a there is a crossing is by maintaining some sort of description of the blobs before and after the fusion. Histograms of U and V colour components are maintained for each blob. The Y component accounts for luminance and therefore it was not used. Whenever a separation is detected, the histograms of the left and right separated blobs are compared with those of the left and right blobs that were fusioned previously. Intersection (Swain and Ballard, 1991) was used to compare histograms (which must be normalized for blob size). This procedure allows to detect if there is a crossing, see Figure 6.

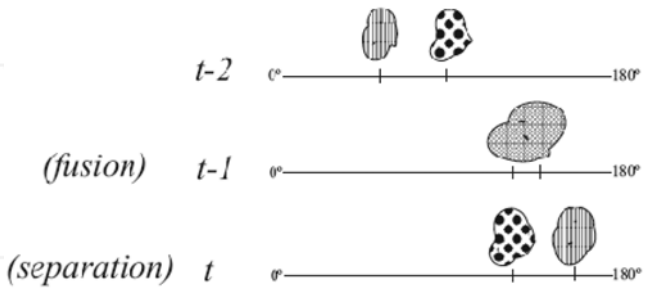

Fig. 6. Crossings can be detected by comparing blob histograms at fusion and separation events.

The histogram similarities calculated are shown in Figure 7. A crossing is detected if and only if (b + c) > (a + d). Note that in the comparison no threshold is needed, making crossing detection relatively robust.

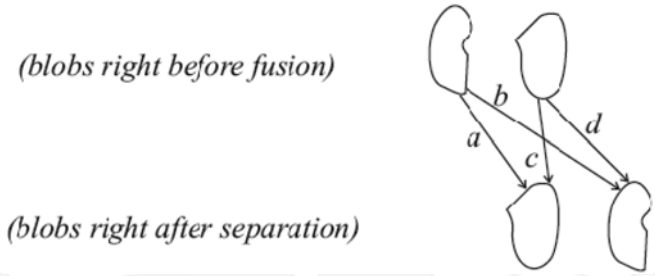

Fig. 7. Blob similarities calculated.

In order to achieve person identification, a set of Y-U histograms are stored for each person detected. The zone from which these histograms are calculated is a rectangle in the lower part of the image taken from the stereo camera placed under the nose of the robot. The rectangle is horizontally aligned with the centre of the face rectangle detected, and extends to the lower limit of the image (chest and abdomen of standing people will always occupy that lower part of the image), see Figure 8. The upper edge of the rectangle is always under the lower edge of the face rectangle detected. The width of the rectangle is proportional to the width of the face rectangle detected.

Fig. 8. Region used for person identification.

When the robot fixates on a person that the tracking system has labelled as new (the tracking system detects a new person in the scene when the number of foreground blobs increases and no blob separation is detected), it compares the histograms of the fixated individual with those of previously met individuals. This search either gives the identity of a previously seen individual or states that a new individual is in the scene. In any case the set of stored histograms for the individual is created/updated.

5. Head Nod/Shake detection

Due to the fact that practical (hands-free) voice recognition is very difficult to achieve for a robot, we decided to turn our attention to simpler (though useful) input techniques such as head gestures. Head nods and shakes are very simple in the sense that they only provide yes/no, understanding/disbelief, approval/disapproval meanings. However, their

importance must not be underestimated because of the following reasons: the meaning of head nods and shakes is almost universal, they can be detected in a relatively simple and robust way and they can be used as the minimum feedback for learning new capabilities.

The system for nod/shake detection described in (Kapoor & Picard, 2001) achieves a recognition accuracy of 78.46%, in real-time. However, the system uses complex hardware and software. An infrared sensitive camera synchronized with infrared LEDs is used to track pupils, and a HMM based pattern analyzer is used to the detect nods and shakes. The system had problems with people wearing glasses, and could have problems with earrings too. The same pupil-detection technique was used in (Davis and Vaks, 2001). That work emphasized the importance of the timing and periodicity of head nods and shakes. However, in our view that information is not robust enough to be used. In natural humanhuman interaction, head nods and shakes are sometimes very subtle. We have no problem in recognizing them because the question has been clear, and only the YES/NO answers are possible. In many cases, there is no periodicity at all, only a slight head motion. Of course, the motion could be simply a 'Look up'/'Look down'/'Look left'/'Look right', though it is not likely after the question has been made.

For our purposes, the nod/shake detector should be as fast as possible. On the other hand, we assume that the nod/shake input will be used only after the robot has asked something. Thus, the detector can produce nod/shake detections at other times, as long as it outputs right decisions when they are needed. The major problem of observing the evolution of simple characteristics like intereye position or the rectangle that fits the skin-color blob is noise. Due to the unavoidable noise, a horizontal motion (the NO) does not produce a pure horizontal displacement of the observed characteristic, because it is not being tracked. Even if it was tracked, it could drift due to lighting changes or other reasons. In practice, a horizontal motion produces a certain vertical displacement in the observed characteristic. This, given the fact that decision thresholds are set very low, can lead the system to error. The performance can be even worse if there is egomotion, like in our case (camera placed on a head with pan-tilt).

The proposed algorithm uses the pyramidal Lucas-Kanade tracking algorithm described in (Bouguet, 1999). In this case, there is tracking, and not of just one, but multiple characteristics, which increases the robustness of the system. The tracker looks first for a number of good points to track over the whole image, automatically. Those points are accentuated corners. From those points chosen by the tracker we attend only to those falling inside the rectangle that fits the skin-color blob, observing their evolution. Note that even with the LK tracker there is noise in many of the tracking points. Even in an apparently static scene there is a small motion in them.

The method is shown working in Figure 9. The LK tracker allows to indirectly control the number of tracking points. The larger the number of tracking points, the more robust (and slow) the system. The method was tested giving a recognition rate of 100% (73 out of 73, questions with alternate YES/NO responses, using the first response given by the system).

What happens if there are small camera displacements? In order to see the effect of this, linear camera displacements were simulated in the tests. In each frame, an error is added to the position of all the tracking points. If (D_x,D_y) is the average displacement of the points inside the skin-color rectangle, then the new displacement is D_x+e_x and D_y+e_y. The error, which is random and different for each frame, is bounded by $-e_{max} < e_x < e_{max}$ and $-e_{max} < e_y < e_{max}$. Note that in principle it is not possible to use a fixed threshold because the error is unknown. The error also affects to the tracking points that fall outside the rectangle.

Assuming that the objects that fall outside the rectangle are static we can eliminate the error and keep on using a fixed threshold, for $(D_x + e_x) - (F_x + e_x) \approx D_x$ and $(D_y + e_y) - (F_y + e_y) \approx D_y$. For the system to work well it is needed that the face occupies a large part of the image. A zoom lens should be used. When a simulated error of $e_{max} = 10$ pixels was introduced, the recognition rate was 95.9% (70 out of 73). In this case there is a slight error due to the fact that the components F_x and F_y are not exactly zero even if the scene outside the rectangle is static.

Fig. 9. Head nod/shake detector.

Another type of error that can appear when the camera is mounted on a mobile device like a pan-tilt unit is the horizontal axis inclination. In practice, this situation is common, especially with small inclinations. Inclinations can be a problem for deciding between a YES and a NO. In order to test this effect, an inclination error was simulated in the tests (with the correction of egomotion active). The error is a rotation of the displacement vectors **D** a certain angle α clockwise. Recognition rates were measured for different values of α, producing useful rates for small inclinations: 90% (60 out of 66) for α = 20, 83.8% (57 out of 68) for α = 40 and 9.5% (6 out of 63) for α = 50.

6. Habituation

Habituation is a filtering mechanism that has received a lot of attention in physiology and psychology. In particular, some researchers have investigated the mechanisms of habituation in animals, being one of the most known works the study of the Aplysia's gillwithdrawal reflex (Castellucci et al., 1970). When the animal's siphon is touched, its gill contracts for a few seconds. If the siphon is stimulated repeatedly, the gill-withdrawal effect tends to disappear. Crook and Hayes (Crook & Hayes, 2001) comment on a study carried out on two monkeys by Xiang and Brown who identified neurons that exhibit a habituation mechanism since their activity decreases as the stimulus is shown repeatedly. Stanley's model (Stanley, 1976) of habituation, proposed to simulate habituation data obtained from the cat spinal cord, has been widely used in the literature. This model describes the decrease efficacy y of a synapsis by the first-order differential equation:

$$\tau \frac{dy(t)}{dt} = \alpha\big(y_0 - y(t)\big) - S(t) \tag{7}$$

where y_0 is the normal, initial value of y, S(t) represents the external stimulation, τ is a time constant that governs the rate of habituation and α regulates the rate of recovery. Equation (7) ensures that the synaptic efficacy decreases when the input signal S(t) increases and returns to its maximum y0 in the absence of an input signal.

The model given by (7) can only explain short-term habituation, so Wang introduced a model to incorporate both short-term and long-term habituation using an inverse S-shaped curve (Wang, 1995):

$$\tau \frac{dy(t)}{dt} = \alpha z(t)\big(y_0 - y(t)\big) - \beta y(t) S(t) \tag{8}$$

$$\frac{dz(t)}{dt} = \gamma z(t)(z(t) - l) S(t) \tag{9}$$

where α, y_0 and γ have the same meaning than in (7), & regulates the habituation and z(t) decreases monotonically with each activation of the external stimulation S(t), and models the long-term habituation. Due to this effect of z(t) after a large number of activations, the recovery rate is slower.

Note that novelty detection is a concept related to habituation. Novelty detection is the discovery of stimuli not perceived before and so habituation serves as a novelty filter (Stiles & Ghosh, 1995). From an engineering viewpoint, perceptual user interfaces, like humanlike robots, should be endowed with a habituation mechanism. The interest is twofold. First, it would be a filtering mechanism, discarding (or minimizing the importance of) repetitive information while paying attention to new experiences. This is in part motivated by the desire to distinguish between artificial and human signals. Artificial signals are often static or repeat with a fixed frequency. We do not want our robot to pay much attention to the hands of a wall-mounted clock. Instead, it would be more interesting to detect nonrepetitive stimuli, such as a conversation or a sudden loud noise. Note that we generally consider monotonous signals as those having a fixed frequency or frequencies (which can be zero, that is, the signal does not change) but signals whose frequency changes in a periodic pattern could also be considered monotonous. Higher scales are also possible but we do not consider them in this work because they are very hard to visualize and real examples of them are not so common. Second, habituation would lead to a more human-like behaviour, as perceived by users of the interface. As an example of this, consider Kismet. Someone can catch the eye of the system while waving a hand in its visual field of view, but if the stimulus is repetitive for a long time the system can show a lack of interest in it. Many aspects of Kismet's mental architecture are directly or indirectly influenced by the detection of monotonous sensory signals: stimulation and fatigue drives and the arousal dimension of its affect space (and in turn some emotional states, like surprise, boredom or interest). If we use the model of Equation (7) we can obtain undesired effects with certain stimuli. A periodic input signal (with frequency greater than zero) can produce a response that does not exhibit habituation. This is due to the fact that the model does not account for changing stimuli, but for continuous ones. In order to include this fact in the model, we propose to use an auxiliary signal which will be zero when the stimulus is stationary or with a fixed frequency, and one otherwise, and use this signal as an input to the habituation model (7). The auxiliary signal, which basically detects monotonous stimuli, is obtained from the

spectrogram of the stimulus itself. The spectrogram is a time-frequency distribution of a signal, and it is based on the Fourier Transform with a sliding window, see Figure 10 (Holland et al., 2000).

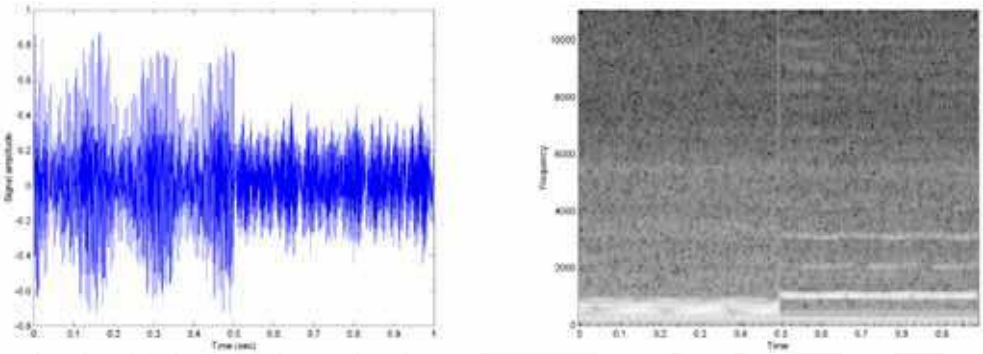

Fig. 10. Audio signal (left) and its corresponding spectrogram (right).

Spectrograms are computed from windows of the input signal. These windows, of length l, overlap by l - 1 samples. Let each spectrogram be represented as a matrix M, in which rows represent frequencies and columns represent time. We calculate the variance of each row of M, which produces a column vector v. The norm of this vector v is a measure of how monotonous the input signal is. The norm will be high when the signal is changing, and low otherwise. Thus, the auxiliary signal needed is simply the thresholded norm of v. The amplitude of the input signal affects the power content of the spectrograms, and in turn the norm of v. Thus, prior to calculating the FFT the input signal must be normalized dividing each input window by the sum of its absolute values. A value of 1 for the auxiliary signal will mean that there are changes in the input signal, while a value of 0 indicates that the input signal is monotonous. Once the auxiliary signal is available, the model (7) is used to get the desired habituation behaviour, as controlled by parameters τ and α. The auxiliary signal is then, for a given threshold T:

$$A = \begin{cases} 1 \ \ if \ \ |v| > T \\ 0 \ \ if \ \ |v| \leq T \end{cases} \tag{10}$$

With this method both static and fixed frequency stimuli can be detected. The algorithm described above was implemented to test it with different input signals. The first experiments that we present use only the first level mentioned above.

In order to gather signals from the visual domain, we recorded video containing a yellow bright stimulus (a yellow card) that was moved in a repetitive fashion, see Figure 11-a). Using simple segmentation techniques we extracted the centroid of the card on each frame (384x288) and summed the x and y pixel coordinates to form the one-dimensional signal of Figure 5.29-b). The sequence of card movements throughout the recording was: horizontal movement, random (aperiodic) movement, vertical movement and vertical movement at a different frequency than the previous one. The results appear in Figure 12. Windows of 128 samples were used, and the variance threshold was set to 1000.

The habituation mechanism described here was implemented in CASIMIRO, for signals in the visual domain only, i.e. images taken by the stereo camera (see Section 3). The difference between the current and previous frame is calculated. Then it is thresholded and filtered with Open and Close operators. Also, blobs smaller than a threshold are removed. Then the centre of mass of the resultant image is calculated. The signal that feeds the habituation algorithm is the sum of the x and y components of the centre of mass. This way when the image does not show significant changes or repetitive movements are present for a while the habituation signal grows. When it grows larger than a threshold, an inhibition signal is sent to the Attention module, which then changes its focus of attention. Neck movements produce changes in the images, though it was observed that they are not periodic, and so habituation does not grow.

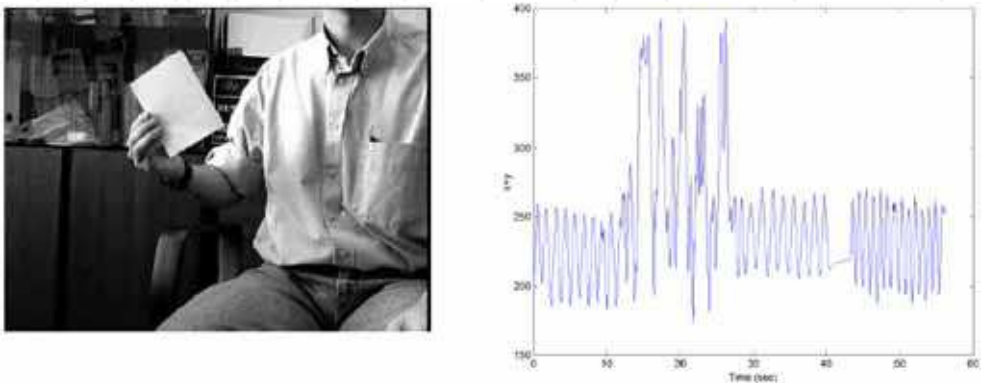

a) b)

Fig. 11. a) Video recording used for the visual habituation experiment, b) unidimensional signal extracted from it.

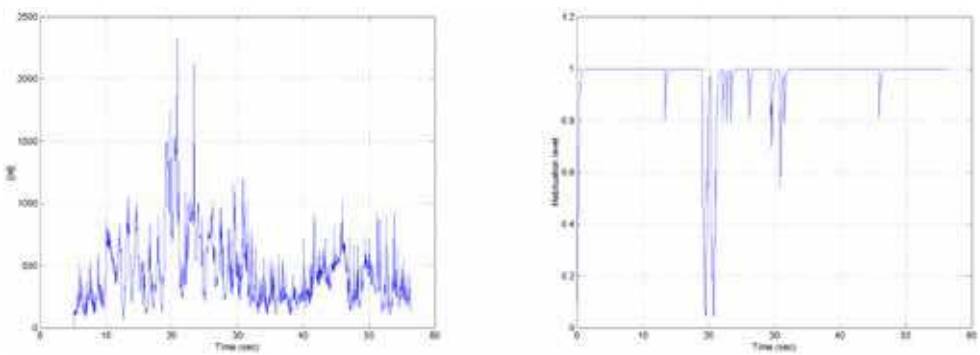

a) b)

Fig. 12. a) Evolution of the (l2) norm of the variance vector v, b) habituation level, using τ = 5, α = 1.

7. Other abilities

In this section two minor abilities of the robot are described: an owner recognition method that is simple yet powerful enough and a “hand waving” detector.

Owner recognition:

The recognition of the robot’s owner or caregiver is a potentially important feature. The recognition may not affect the creator at all, but it could impress other people present in the room. Person recognition is difficult to achieve, the identification of the robot’s owner may be more feasible, for the problem reduces to two classes (owner and not-owner). Still, using face or voice for owner recognition would lead to essentially the same lack of robustness of face and speech recognition for multiple individuals. Who is the owner of the robot? The correct answer is: the person who buys it. However, we are more interested in the role of the person who cares the robot and uses it more frequently (this person is generally the owner too). This person is the one who switches the robot on, which is usually done from a certain part of the robot or from a certain computer. That cue may be exploited to recognize the robot owner or caregiver. *Amazing Amanda*, a doll released in autumn of 2005, is able to recognize the girl that takes the mother role. Once the doll is activated, it starts asking questions. That way, the doll can "learn" the girl’s voice patterns. From that moment on, the doll is able to recognize the utterances of its "mommy". Other voices can lead Amanda to say "You don’t sound like Mommy".

Such technique may seem rather ad hoc. However, the approach finds striking examples in nature. Lorenz, one of the founders of ethology, found that, upon coming out of their eggs, geese follow and become attached to the first moving object that they encounter. He showed this by rearing the geese from hatching. From that moment on the geese would follow him. Such phenomenon, which also appears in mammals, is known as imprinting (Lorenz, 1981).

Fig. 13. The computer from where CASIMIRO is started. The interaction space is on the right.

In the case of CASIMIRO, the main computer (from which the robot is switched on) is situated behind the robot, on the same table, see Figure 13. A camera was placed on top of that computer. The owner detection module uses that camera to search for a skin coloured blob in the image. When the robot is switched on this module will detect a skin coloured blob. The camera has a wide-angle lens, and a relatively low resolution of 160x120 is used. When no blob is encountered in the image the module notifies the Attention module of that event. At that moment the owner detection module exits in order to free CPU resources. Once it has been notified by the owner detection module, the Attention module considers the owner as the first blob that "enters" the omnidirectional camera image from the left. The "owner" property is stored along with the individual in the tracking process. This simple procedure is a form of imprinting. In a sense, the robot finds its owner-caregiver in the first human it sees. It does not stores any biometric features to recognize the owner after being switched on, only its position.

Hand-waving detector:

In initial tests with the robot we saw that people tended to wave his/her hand to the robot. We decided to detect this and make the robot respond with a funny phrase. The contour of each foreground blob is obtained and compared with the contour of the same blob in the previous frame. The comparison is made using Hu moments (Hu, 1962). When the comparison yields too much difference between the contours (i.e. a threshold is exceeded) a flag is activated. When the face detection module (Section 3) detects more than one skin colour blob in the image and a certain amount of motion and the flag is activated then the Boolean perception HANDWAVING is set to TRUE.

8. Conclusion

A number of simple but useful computer vision techniques have been described, suitable for human-robot interaction. First, an omnidirectional camera setting is described that can detect people in the surroundings of the robot, giving their angular positions and a rough estimate of the distance. The device can be easily built with inexpensive components. Second, we comment on a color-based face detection technique that can alleviate skin-color false positives. Third, a simple head nod and shake detector is described, suitable for detecting affirmative/negative, approval/disapproval, understanding/disbelief head gestures. Fourth, a habituation-detection scheme has been described, allowing the robot to focus on natural signals in the environment. Finally, two minor owner recognition and hand-waving detection abilities are also described. All the techniques have been satisfactorily implemented and tested on a prototype social robot.

Future work should focus on how computer vision could aid in hands-free speech recognition, an all-important ability for an interactive robot. Progress has been already made by syncing lip motion with audio speech signals, although the processing cost is still too long for useful real-time interaction.

9. References

Bouguet, J. (1999). Pyramidal implementation of the Lucas Kanade feature tracker. *Technical report*, Intel Corporation, Microprocessor Research Labs, OpenCV documents.

Breazeal, C. L. (2002). *Designing social robots*. MIT Press, Cambridge, MA.

M. Castrillon-Santana, H. Kruppa, C. G. & Hernandez, M. (2005). Towards real-time multiresolution face/head detection. *Revista Iberoamericana de Inteligencia Artificial*, 9(27):63–72.

Castellucci, V., Pinsker, H., Kupfermann, I., & Kandel E.R. (1970). Neuronal mechanisms of habituation and dishabituation of the gill-withdrawal reflex in Aplysia. *Science*, 167:1745–1748.

Collins, G. & Dennis, L. A. (2000). System description: Embedding verification into Microsoft Excel. In *Conference on Automated Deduction*, pages 497–501.

Crook, P., & Hayes, G. (2001). A robot implementation of a biologically inspired method of novelty detection. In *Proceedings of TIMR 2001 - Towards Intelligent Mobile Robots*, Manchester, April.

Darrell, T., Gordon, G., Harville, M., & Woodfill, J. (1998). Integrated person tracking using stereo, color, and pattern detection. In *Procs. of IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 601–608, Santa Barbara, CA.

Davis, J. & Vaks, S. (2001). A perceptual user interface for recognizing head gesture acknowledgements. In *Proc. of ACM Workshop on Perceptual User Interfaces*, Orlando, Florida.

Deniz, O. (2006). *An Engineering Approach to Sociable Robots*. PhD Thesis, Department of Computer Science, Universidad de Las Palmas de Gran Canaria.

Fong, T., Nourbakhsh, I. & Dautenhahn, K. (2003). A survey of socially interactive robots. *Robotics and Autonomous Systems*, 42(3-4).

Grange, S., Casanova, E., Fong, T., & Baur, C. (2002). Vision based sensor fusion for humancomputer interaction. In *Procs. of IEEE/RSJ International Conference on Intelligent Robots and Systems*. Lausanne, Switzerland.

Holland, S., Kosel, T., Waver, R., & Sachse, W. (2000). Determination of plate source, detector separation fron one signal. *Ultrasonics*, 38:620–623.

Hu, M. (1962). Visual pattern recognition by moment invariants. *IRE Transactions on Information Theory*, vol. 8(2), 179–187.

Kahn, R. (1996). Perseus: An Extensible Vision System for Human-Machine Interaction. PhD Thesis, University of Chicago.

Kapoor, A. & Picard, R. (2001). A real-time head nod and shake detector. In *Proceedings from the Workshop on Perspective User Interfaces*.

Krumm, J., Harris, S., Meyers, B., Brumitt, B., Hale, M., & Shafer, S. (2000). Multicamera multiperson tracking for easyliving. In *3rd IEEE International Workshop on Visual Surveillance*, Dublin, Ireland.

Lorenz, K. (1981). *The Foundations of Ethology*. Springer Verlag, Heidelberg.

Maxwell, B. (2003). A real-time vision module for interactive perceptual agents. *Machine Vision and Applications*, (14):72–82.

Maxwell, B., Meeden, L., Addo, N., Brown, L., Dickson, P., Ng, J., Olshfski, S., Silk, E., & Wales, J. (1999). Alfred: the robot waiter who remembers you. In *Procs. of AAAI Workshop on Robotics*.

Moreno, F., Andrade-Cetto, J., & Sanfeliu, A. (2001). Localization of human faces fusing color segmentation and depth from stereo. In *Procs. ETFA*, pages 527–534.

Schulte, J., Rosenberg, C., & Thrun, S. (1999). Spontaneous short-term interaction with mobile robots in public places. In *Procs. of the IEEE Int. Conference on Robotics and Automation*.

Stanley, J.C. (1976). Computer simulation of a model of habituation. *Nature*, 261:146–148.

Stiles, B., & Ghosh, J. (1995). A habituation based neural network for spatio-temporal classification. In *Proceedings of the 1995 IEEE Workshop In Neural Networks for Signal Processing*, pages 135–144, Cambridge, MA, September.

Wang, D. (1995). Habituation. In M. A. Arbib, editor, The *Handbook of Brain Theory and Neural Networks*, pages 441–444. MIT Press.

Advances in Robotics, Automation and Control

Edited by Jesus Aramburo and Antonio Ramirez Trevino

ISBN 978-953-7619-16-9
Hard cover, 472 pages
Publisher InTech
Published online 01, October, 2008
Published in print edition October, 2008

The book presents an excellent overview of the recent developments in the different areas of Robotics, Automation and Control. Through its 24 chapters, this book presents topics related to control and robot design; it also introduces new mathematical tools and techniques devoted to improve the system modeling and control. An important point is the use of rational agents and heuristic techniques to cope with the computational complexity required for controlling complex systems. Through this book, we also find navigation and vision algorithms, automatic handwritten comprehension and speech recognition systems that will be included in the next generation of productive systems developed by man.

How to reference

In order to correctly reference this scholarly work, feel free to copy and paste the following:

Oscar Deniz, Javier Lorenzo, Modesto Castrillon, Luis Anton, Mario Hernandez and Gloria Bueno (2008). Practical Computer Vision Techniques for Interactive Robotic Heads, Advances in Robotics, Automation and Control, Jesus Aramburo and Antonio Ramirez Trevino (Ed.), ISBN: 978-953-7619-16-9, InTech, Available from:
http://www.intechopen.com/books/advances_in_robotics_automation_and_control/practical_computer_vision_techniques_for_interactive_robotic_heads

INTECH
open science | open minds

InTech Europe
University Campus STeP Ri
Slavka Krautzeka 83/A
51000 Rijeka, Croatia
Phone: +385 (51) 770 447
Fax: +385 (51) 686 166
www.intechopen.com

InTech China
Unit 405, Office Block, Hotel Equatorial Shanghai
No.65, Yan An Road (West), Shanghai, 200040, China
中国上海市延安西路65号上海国际贵都大饭店办公楼405单元
Phone: +86-21-62489820
Fax: +86-21-62489821