Proceedings of the SSRR '04
IEEE International Workshop on Safety, Security, and Rescue Robotics
Bonn, Germany, May 24 - 26, 2004

Semi-Autonomous Victim Search

Carolina Chang and Ángela Brando
Grupo de Inteligencia Artificial
Universidad Simón Bolívar
Caracas, Venezuela
E-mail cchang@ldc.usb.ve

Abstract

This paper describes a robotics application prototype for semi-autonomous victim search in urban environments, using a Pioneer 2-DX robot equipped with sonars, a pan-tilt-zoom CCD camera, and wireless Ethernet. The robot uses a wall-following strategy to navigate a building, and a Support Vector Machine to recognize pixels of skin-like colors in the visual input. Nearby pixels that get classified as having skin color are grouped together into blobs. The robot orients itself towards the largest color blob and approaches the object trying to keep the blob on the center of the visual field. When it has approached the object, the robot takes a picture and sends the picture via e-mail to an Incident Commander, who is also notified by cell phone text messages. The robot stays close to the suspected victim until it receives instructions from the Commander, or until the object moves away. We report the results of preliminary experiments on semi-autonomous victim search inside a simple office environment.

Keywords: *mobile robot, victim search, Support Vector Machines*

1 Introduction

Urban Search and Rescue (USAR) is becoming an important domain for robotic applications [1], [2], [3], [4]. Some of the most outstanding research has been carried out at the Perceptual Robotics Laboratory [5], [6], [7]. In 2001, the Center for Robot Assisted Search and Rescue (CRASAR) deployed several USAR robotics teams to help rescue and recover victims of the World Trade Center. Since then, CRASAR has participated in the rescue of earthquake victims worldwide. Other initiatives such as the AAAI USAR Competition and RoboCupRescue, have contributed significantly to the development of the USAR robotics field.

USAR robots are expensive because they must be able to operate in very unstructured and hazardous environments. To do so, they should be equipped with sophisticated sensors, locomotion mechanisms, and communications.

There are no USAR robots in Venezuela. The research reported in this paper was carried out using an indoor Pioneer 2-DX robot. Our robot has too many limitations as to be useful in a real USAR incident, yet we think that it can help us gain some insight about this research field. One goal of our work is to promote the development of the USAR robotics research field in our country, and to set some awareness in our Fire Departments about the benefits that robotic technologies may bring to them.

We are developing an application for teleoperated and semi-autonomous robot operation. The teleoperation mode allows the user to control the robot's movements, and the pan-tilt-zoom camera. In the semi-autonomous mode the robot explores a building looking for victims. When it finds an object of skin-like color, the robot notifies the finding to an incident Commander, via e-mail and cell-phone text messages. The Commander can chose to teleoperate the robot in order to collect more data about the victim and the environment, or can ask the robot to continue the autonomous search.

In [8] we tested the interaction between the Pioneer robot and certified firefighters during a hazardous materials incident exercise, in teleoperation mode. In this article we describe the module for semi-autonomous victim search.

2 The Robotics Application

Pioneer 2-DX is a 9Kg, 44x38x22cm differential-drive robot (see figure 1). It has a frontal ring of eight ultrasound sensors, a pan-tilt-zoom Sony CCD camera, an internal PC104 computer, frame grabber, and linux operating system. It has wireless Ethernet communications in a range of 50m by means of an access point. The robot can operate autonomously on batteries for about 4 hours.

The software prototype was developed using

Figure 1: The Pioneer2-DX Robot

Saphira 8.0 environment, and VisLib 1.8.

We have trained a Support Vector Machine (SVM) [9] to classify pixels based on their RGB values, into skin or non-skin color. Classification of a pixel x is given by:

$$\Psi \cdot x' + b \begin{cases} > 0 & x \text{ is skin} \\ otherwise & x \text{ is not skin.} \end{cases} \quad (1)$$

where Ψ and b describe the decision surface obtained after training the SVM, and x' is the input vector transformed using a polynomial kernel. The input vector has two components that represent the normalized RGB values of a pixel:

$$R_n = R/(R + G + B) \quad (2)$$

$$G_n = G/(R + G + B) \quad (3)$$

This input normalization for skin classification reduces the dimensionality of the input, and factors out variations on luminance, that do not contribute to this segmentation task [10].

After training we reduced the number of support vectors and accelerated the classification using the Burges method described in [11].

Figure 2 shows the classification of the pixels of two different images. In panels (a) and (c) the robot's visual input includes a person lying on the floor (simulating a victim). Panels (b) and (d) show the classification of the pixels using the SVM.

In the semi-autonomous mode, the robot navigates an environment by means of a wall-following strategy. This strategy is commonly chosen by certified firefighters when they have to explore unknown or hazardous environments with reduced visibility, as described in [8]. We combine the wall-following behavior with pan movements of the camera, to search for victims. Nearby pixels of skin color are grouped together into color blobs. When the robot detects a blob, it

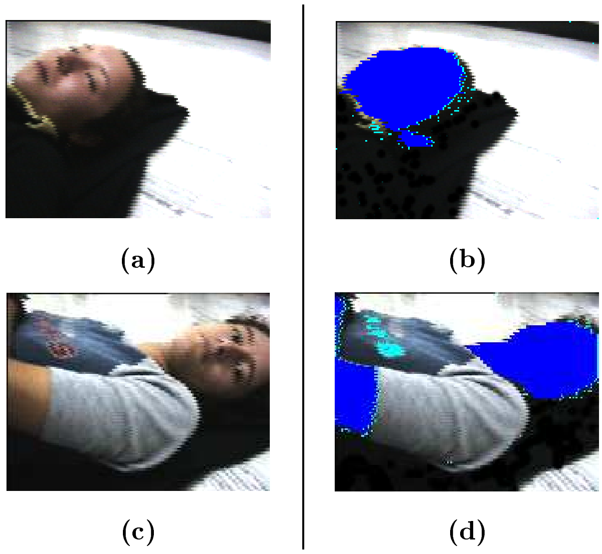

Figure 2: Skin detection using a SVM. Pixels of the images (a) and (c) that were classified as skin are shown in blue color in panels (b) and (d).

abandons the wall-following behavior and heads towards the location of the blob. If more than one blob is detected, the robot approaches the largest blob. To do so, the robot moves its body to re-center the camera, and to center the color blob in the visual image as shown in figure 3. A red square delimits the skin-color blob detected by the robot. Left panels of the figure show the location of the blobs when they are first detected by the robot. Right panels are the results of re-centering the robot's body and camera. Notice that there are two blobs in panel (c), but the blob formed by the left arm has less skin-color pixels than the other blob. Therefore,the robot re-center the largest blob, which includes the victim's head and right arm.

Then the robot heads towards the blob. It uses ultrasound sensors to stop near the object. For the time being we are not dealing with the problem of surmounting obstacles that may lie between the robot and the victim, because our indoors robot would have to move around the obstacles. However, a USAR robot may climb the obstacles or use other strategies to approach the victim.

When the robot is close to the skin-color object, it takes a a picture of the suspected victim and sends the picture via e-mail to a an incident Commander. The Commander also receives a notification of the finding via cell phone text messages. We added this feature assuming that in the future the robot would have more autonomy performing the task of victim search,

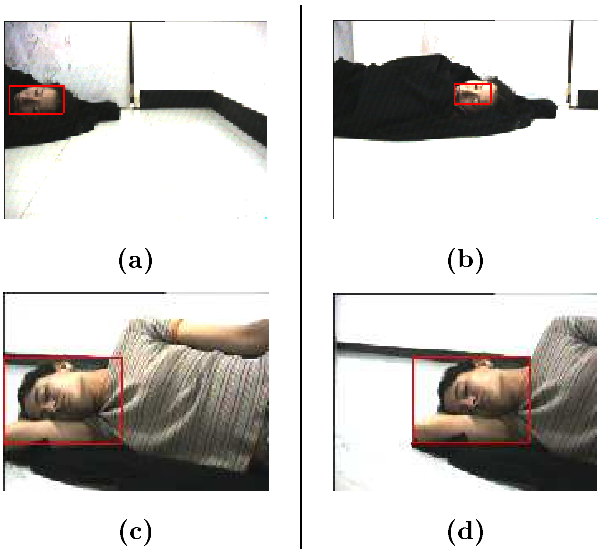

Figure 3: Detection of skin-color blobs and recentering. Skin is detected in panels (a) and (c). The robot re-centers the blobs in the image in panels (b) and (d).

and that teams of many robots would be controlled by a single Commander. In such case, the Commander would need to pay attention to the robot only when important events take place. In practice, currently there is always a human being watching remotely and on-line the visual input of our robot. In our application the robot stays nearby the detected object until it receives a command from a human operator or until the object moves away. The operator can tell the robot to ignore the object (*i.e.* it was a false alarm) or the operator can switch the robot to the teleoperation mode to collect more data about the victim condition and environment.

The robot does not decide by itself whether the object that it found is a victim or not. This is a very difficult and delicate task. We think the victim recognition should be performed by a human operator, and that the robot requires more sophisticated sensors and strategies to perform this task. For example, the victim may be covered by debris, or blood, or have no regions of skin exposed (*e.g.*, the victim could be a firefighter fully covered by an encapsulated suit). In this first prototype, we restrict the scope of our autonomous search to the detection of skin color in the visual image. However, during the task the robot can be switched to teleoperation mode at any time, so that a human operator can take control of the search for victims. In confined or hazardous environments, the use of a robot can reduce the risk for human rescuers,

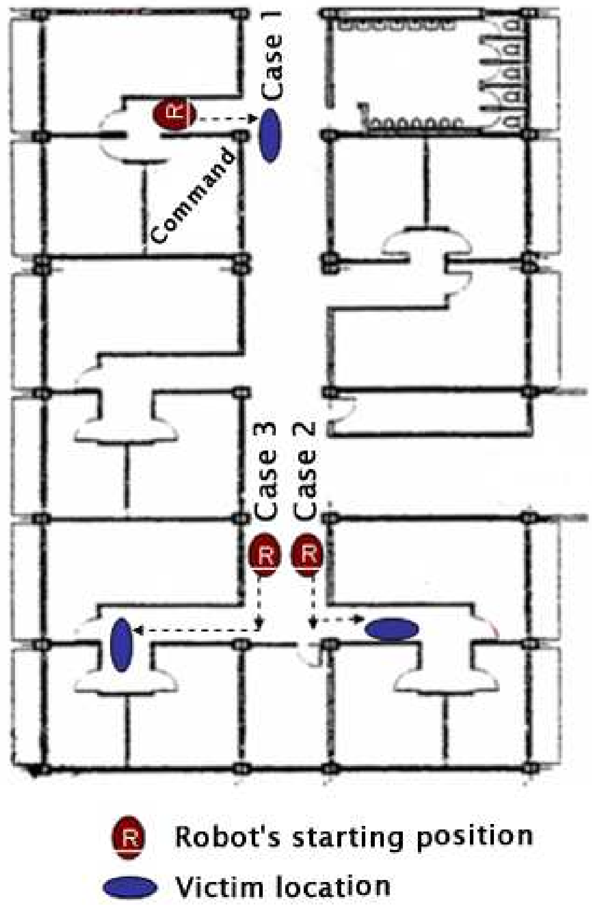

Figure 4: *Map of the building. In the experiments the robot had to detect a person lying on the floor at three different locations. See text for details.*

and help find victims, even if the robot is fully controlled by a human operator.

3 Victim Search

To train the SVM we chose 235 examples of skin color, and 235 examples of non-skin color, taken from our work-environment. The SVM was tested on 320 points (160 points for each class). We obtained a performance of 84% in the classification task. Objects such as doors and boxes were classified by the SVM as having skin color, so in the future we would need to consider other features such as size and shape of the objects, in order to improve the filtering.

We ran the experiments in our office environment. Figure 4 shows the area of the building searched by the robot. We stayed at the office labeled "Command" while the robot explored the environment. We

Table 1: Results of the Victim Detection Experiments. Cases 1, 2, 3 indicate to the location of the Victim and the robot's initial position according to the map of figure 4. Each Case was tested 10 times.

Experiment	Successes	Failures
Case 1	7	3
Case 2	9	1
Case 3	8	2

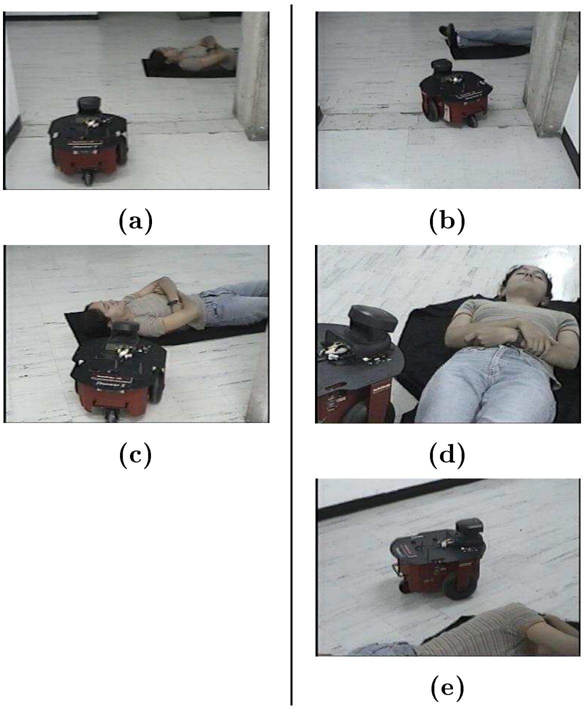

Figure 5: Victim detection in Case 1. (a) and (b) are two different victim positions. In (c) and (d) the robot was successful at detecting and approaching the victim. In (e) the robot did not detect the skin of the victim.

(a) (b)

Figure 6: A trial of the Case 3 experiment. (a) The robot enters the inner aisle where the victim is located. (b) The robot detects the victim.

tested the behaviors of victim detection and approach on three different locations. In case 1 the victim lies on the main aisle, while in cases 2 and 3, the victim lies on inner aisles, which are narrower. The victim of case 3 is more hidden in the aisle, and has most of the skin covered with a black blanket to reduce the area of the color blob detected. Panels (b) and (c) of figure 3 correspond to a Case 3 experiment.

We repeated the experiment 10 times for each case, starting from a given robot position. At each location the victim adopted different positions on different trials. Table 1 shows the results of our experiments. In all cases the robot detected the victim at least on 70% of the trials. The worst performance was obtained in case 1. The small skin-color blob that we used in case 3 did not produce a significant reduction of the robot's ability to detect and approach the blob.

Figure 5 shows some of the results of case 1. In the left panels the robot first encounters the victim's head, while in the right panels the feet of the victim are seen first, as shown in panels (a) and (b). In 7 out of the 10 trials the robot detected and approached the victim, as in panels (c) and (d). The other three times the robot did not detect the skin of the victim, as in panel (e).

Figure 6 shows a successful trial of case 3. Even though the victim is hidden inside the inner aisle (panel (a)) the robot detects the victim (panel (b)). At this task the robot succeded 8 out 10 trials.

4 Conclusion

We have described a prototype for semi-autonomous victim search. The robot can be controlled by teleoperation, or can search for victims semi-autonomously by means of color segmentation of its visual input. This application needs to be improved in order to become useful, since it would fail under many conditions, such as dim lightning, debris-covered skin, skin-like color objects, and so on.

We think that the process of recognition of human beings in collapsed structures is a task that requires sophisticated sensor such as thermal cameras, as well as great mobility and communications. Our indoors robot is not well suited to perform urban search and rescue tasks, yet this research is allowing us to gain some insight about the USAR robotics domain.

Acknowledgment

We thank Boston University Department of Cognitive And Neural Systems for donating the P2-DX robot. Prof. Robin Murphy and Jen Casper from the USF Robotics Perceptual Lab kindly lent us the tapes of their 2001 field tests. Firefighters of the *Cuerpo de Bomberos Voluntarios de la Universidad Simón Bolívar* and many friends have volunteered their time to our experiments. Dr. Paolo Gaudiano made this research possible.

References

[1] A. Kobayashi and K. Nakamura, "Rescue robot for fire hazards," in *Proceedings of the 1983 International Conference on Advanced Robotics*, 1983, pp. 91–98.

[2] J. G. Blitch, "Artificial intelligence technologies for robot assisted urban search and rescue," *Expert Systems with Applications*, vol. 11, pp. 109–124, 1996.

[3] R. Masuda, T. Oinuma, and A. Muramatsu, "Multi-sensor control system for rescue robot," in *Proceedings of the IEEE/SICE/RSJ International Conference on Multisensor Fusion and Integration for Intelligent Systems*, 1996, pp. 381–387.

[4] J. S. Jennings, G. Whelan, and W.F. Evans, "Cooperative search and rescue with a team of mobile robots," in *Proceedings of the 1997 International Conference on Advanced Robotics*, 1997, pp. 193–200.

[5] J. Casper, R. R. Murphy, and M. Micire, "Issues in intelligent robots for search and rescue," in *Proceedings of SPIE Ground Vehicle Technology II*, 2000.

[6] R. Murphy, J. Casper, J. Hyams, M. Micire, and B. Minten, "Mobility and sensing demands in usar," in *Proceedings of IECON-2000*, Nagoya, Japan, October 2000.

[7] J. Casper and R. Murphy, "Workflow study on human-robot interaction in usar," in *Proceedings of ICRA2002*, 2002, pp. 1997–2003.

[8] A. Brando and C. Chang, "Firefighter-robot interaction during a hazardous materials incident excercise," in *11th International Conference on Advanced Robotics*, Coimbra, Portugal, June 2003, pp. 658–663.

[9] V. N. Vapnick, *The Nature of Statistical Learning Theory*, Springer-Verlag, Berlin, Germany, 1995.

[10] J. Cai, A. Goshtasby, and C. Yu, "Detecting human faces in color image," in *1998 International Workshop on Multi-Media Database Management Systems.*, 1998.

[11] Christopher Burges, "Simplified support vector decision rules," in *13th International Conference on Machine Learning*, 1996.