Artificial Vision in Extreme Environments for Snowcat Tracks Detection

Alberto Broggi, *Associate Member, IEEE,* and Alessandra Fascioli*, Member, IEEE*

***Abstract*—This paper describes the image processing techniques designed to localize the tracks of snowcats for the automation of transportation of goods and people during the Italian scientific missions in Antarctica. The final goal is to enable a snowcat to automatically follow the preceding one in a train-like fashion. A camera is used to acquire images of the scene; the image sequence is analyzed by a computer vision system which identifies the tracks and produces a high level description of the scene. This result is then forwarded to a further software module in charge of the control of the snowcat movement. A further optional representation, in which markers highlighting the tracks are superimposed onto the acquired image, is transmitted to a human supervisor located off board. This system has been tested in the Italian test site and was under testing in the South Pole during the early 2002 Italian scientific mission. The paper also briefly describes an alternative solution based on an evolutionary approach.**

***Index Terms*—Autonomous vehicles in extreme conditions, extreme robots, machine vision.**

I. Introduction

HIS PAPER presents the artificial-vision algorithms developed to autonomously drive a platoon of snowcats. This research is a part of the Ente per le Nuove Tecnologie, l'Energia e l'Ambiente, Italy (ENEA) Surface Antarctic Robot (SAR) project, aimed at facilitating the transportation of people and goods in the Antarctica region during the Italian scientific missions in the South Pole. The logo of the Italian missions is shown in Fig. 1.

Researchers and equipment arrive from Italy via New Zealand on a ship which stops in the harbor of the "Baia Terra Nova" permanent Italian base (74° 41′ 42″ south, 164° 07′), shown in Fig. 2. Part of the goods and people must reach the "Dome Concordia" inland Italian base or the "Dumont d'Urville" French base, which are about 1 200 and 2 300 km away from the harbor, respectively. This distance can be covered by helicopter, but goods generally travel by a platoon of snowcats, whose average speed is 10–12 km/h.

The RAS Project has been launched in order to ease the long and stressful procedure of people and equipment transportation in the unfriendly South Pole environment; the main goal of this project is to partially automate this procedure.

Manuscript received August 31, 2001; revised March 15, 2002. This work was supported in part by the Ente per le Nuove Tecnologie, l'Energia e l'Ambiente, Italy (ENEA), within the RAS Project, and in part by the Italian National Research Council (CNR) in the framework of the MADESS2 Project. The Guest Editor for this paper was K. Ikeuchi.

The authors are with the Dipartimento di Ingegneria dell'Informazione, Università di Parma, 43100 Parma, Italy (e-mail: broggi@ce.unipr.it; fascal@ce.unipr.it).

Publisher Item Identifier 10.1109/TITS.2002.802936.

Fig. 1. Logo of the Italian research programs in Antarctica.

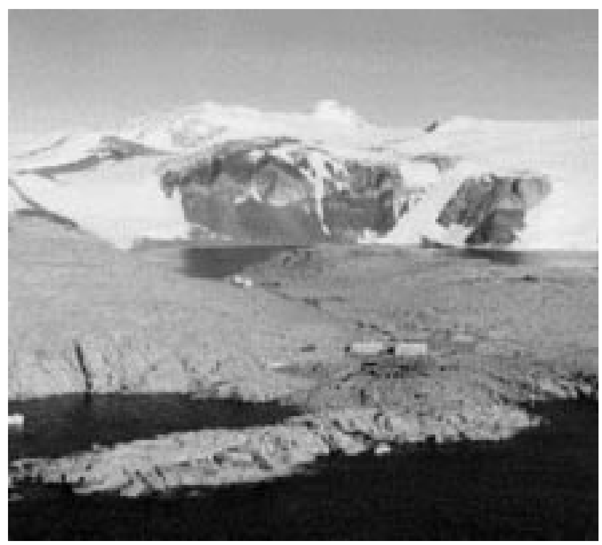

Fig. 2. Aerial view of the Italian base in Terra Nova, Antarctica.

In the first implementation, the leading vehicle will be manually driven by an expert driver; in a second step, as shown in Fig. 3, it will be equipped with a camera and a TV antenna which will broadcast live images to the base and will allow remote driving.

All the remaining vehicles of the platoon will follow automatically in a train-like fashion. Moreover, since cracks in the ice can put both the driver and the snowcat itself in serious danger , it is imperative that the following vehicles follow the same precise path defined by the first vehicle. Since even small drifts from the original driving path defined by the human driver can be extremely dangerous, an extremely precise detection of the tracks left by the previous vehicle, a correct measurement of their position, and a smooth control of the actuators must be carefully designed, tested, and evaluated.

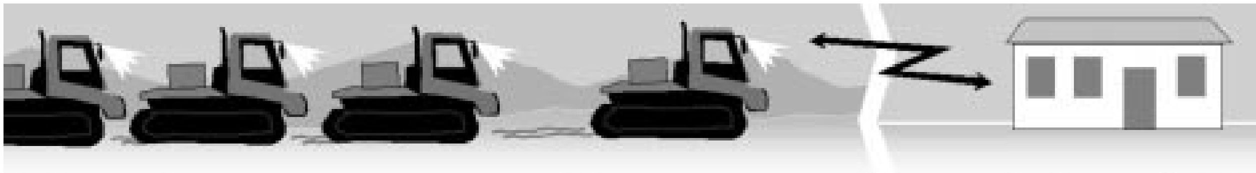

Fig. 3. Platoon of snowcat vehicles; the first will be either manually driven by an expert driver or remotely driven from the Italian base; all the others will follow automatically, using visual information only.

Fig. 4. Prototype vehicle during a test in the Italian test site.

Fig. 5. Sensing capabilities of the platoon leading vehicle.

Fig. 6. Inside view of the vehicle during a test; the stereo cameras are visible on the left-hand side.

Fig. 7. External view of the stereo cameras.

A preliminary test phase showed that the most promising sensor that should be able to deliver sufficiently precise measurements is a vision sensor (camera). Many other devices have been considered [1]–[3], even active ones, since the specific working site would not present any problem due to interference or to environmental pollution [4]. However, vision seems the sensing capability that may deliver the highest performance in terms of precision of the localization. Data are currently acquired from a monocular camera installed inside the driving cabin, but a stereo pair has been integrated to allow stereo image processing in the future. Fig. 4 shows the prototype of the vehicle follower, while Fig. 5 shows all the sensing capabilities of the platoon leading vehicle. Besides cameras, a global positioning system (GPS) will be also used for self-localization, but due to the impossibility to use differential GPS (DGPS) in the area, their effectiveness will need to be confirmed.

Fig. 6 shows the internal equipment, while Fig. 7 shows an external view of the stereo cameras. The calibration of the cameras is done thanks to a set of markers at known distances on a planar surface; Fig. 8 shows the test site on the Italian Alps and the grid used for calibration.

Due to the extreme conditions of the working environment, where temperatures can reach $-80\,^\circ$C, the terrain is completely covered by snow or ice, strong sunlight and reflections may be present, and no specific ground references are available nor assumptions can be made on the terrain slope, this application is extremely challenging and presents many additional problems with respect to the driving of unmanned vehicles on traditional (un)structured roads [5].

For this reason, an extremely careful analysis and design of the processing techniques is mandatory.

Several approaches have been considered due to the low visibility of white tracks on a white background, and specific filters have been developed in order to cope with the typical problems of this environment. The high problem complexity is slightly reduced by the low speed of the vehicle, which permits focus on the localization of tracks in a reduced close area only.

Moreover, in the automatic driving of road vehicles [6] a special emphasis is generally given to the exploitation of *a priori* knowledge in order both to speed up the computation and

Fig. 8. Image of the test site on the Italian Alps; a patch of snow with markers is used to calibrate the cameras.

make the detection robust. In this case, only a little knowledge about the environmental conditions can be exploited: generally no other vehicle or building is seen by the camera, and the only markings on the ice are due to the preceding vehicle. On the other hand, no assumptions can be made with respect to a possible flatness of the area ahead of the vehicle, nor to a given range of illumination of the scene. In other words, hilly conditions must be considered as well and, therefore, the camera orientation generally used in road environments (low toward the road ahead) cannot be replicated here. Besides the acquisition of a large amount of insignificant data during driving in flat areas, the framing of a large portion of the sky can raise another important problem: since in the working site the sun may be very low on the horizon, no specific camera orientation can overcome the problem of direct sunlight into the vision system. This is an extremely difficult issue that must be carefully considered in the development of vision algorithms.

The approach used to solve this artificial-vision problem was suggested by the experience of the research group within the automotive field [7]. Within the ARGO Project, aimed at the development of an intelligent vehicle able to drive autonomously in real-traffic conditions and on real roads, many vision algorithms were developed and implemented. Unfortunately, the extremely different environmental conditions did not allow us to use the same algorithms, but the experience made on lane detection was of basic importance to design the system used on the snowcat. In particular, a simple approach was preferred with respect to more sophisticated ones because of its easy implementation on the simple processing engine available on the snowcat. Also, the requirement for a fast response was a key issue that suggested the development of a simple but efficient algorithm.

This paper is organized as follows. Section II discusses the characteristics of the working environment which make the application particularly challenging, Section III describes the details of the vision algorithm, Section IV presents some results, Section V illustrates a different approach to the problem which is under development, and finally Section VI presents future project developments.

II. Environmental Characteristics

The environmental characteristics of the Antarctica region are very challenging and the automatic driving of a vehicle in these conditions is extremely different from traditional highway or urban applications.

The main differences are due to the coverage of the driving area with snow or ice, and the localization of other vehicles' tracks on different kinds of snow or ice requires specific algorithms able to adapt to different scenarios.

As shown in Fig. 9(a) and (b), the tracks' characteristics can vary considerably: due to different sun positions, in the first image the tracks are darker than the background, while in the second image the tracks are brighter than the surrounding area. Besides a weak brightness gradient, another invariant that could be exploited is the brightness variance, or—in other words—the texture. Unfortunately, due to the high brightness of the environment, the snow texture provides very weak information. As can be seen in all the images of Fig. 9 the difference of texture between tracks and background is generally small.

In some cases the shadow of the vehicle itself or of the mountains are captured by the camera [see Fig. 9(c)]. Due to the very high contrast of shadows on snow, it is impossible to detect weak brightness gradients (the tracks) in the region inside the shadow, which, therefore, must be eliminated from the analysis. In particular, it is necessary to remove the high brightness gradient generated by shadows, and keep and enhance the weak tracks' edges.

As mentioned, strong sun or light reflections can cause the appearance of reflections patterns in the image, as shown in Fig. 9(d) and (e). This disturbing effect is also caused by the inevitable presence of small icy particles on the windshield in the region in front of the camera and scratches on the windshield itself.

No assumptions on terrain slope can be made: in this application domain, no a priori knowledge on the flatness of the region in front of the vehicle can be used to simplify the localization algorithm. As can be clearly seen from Fig. 9(e) and (f)—acquired with only a few seconds of distance—the slope can change abruptly, making it difficult even to define an area of interest in the image.

Furthermore, the change in terrain slope can also affect the camera orientation with respect to the sun and, thus, can modify the quantity of light acquired from the sensor. This is also visible in Fig. 9(e) and (f), in which in the former—due to strong sunlight—the snow brightness is lower than in the latter.

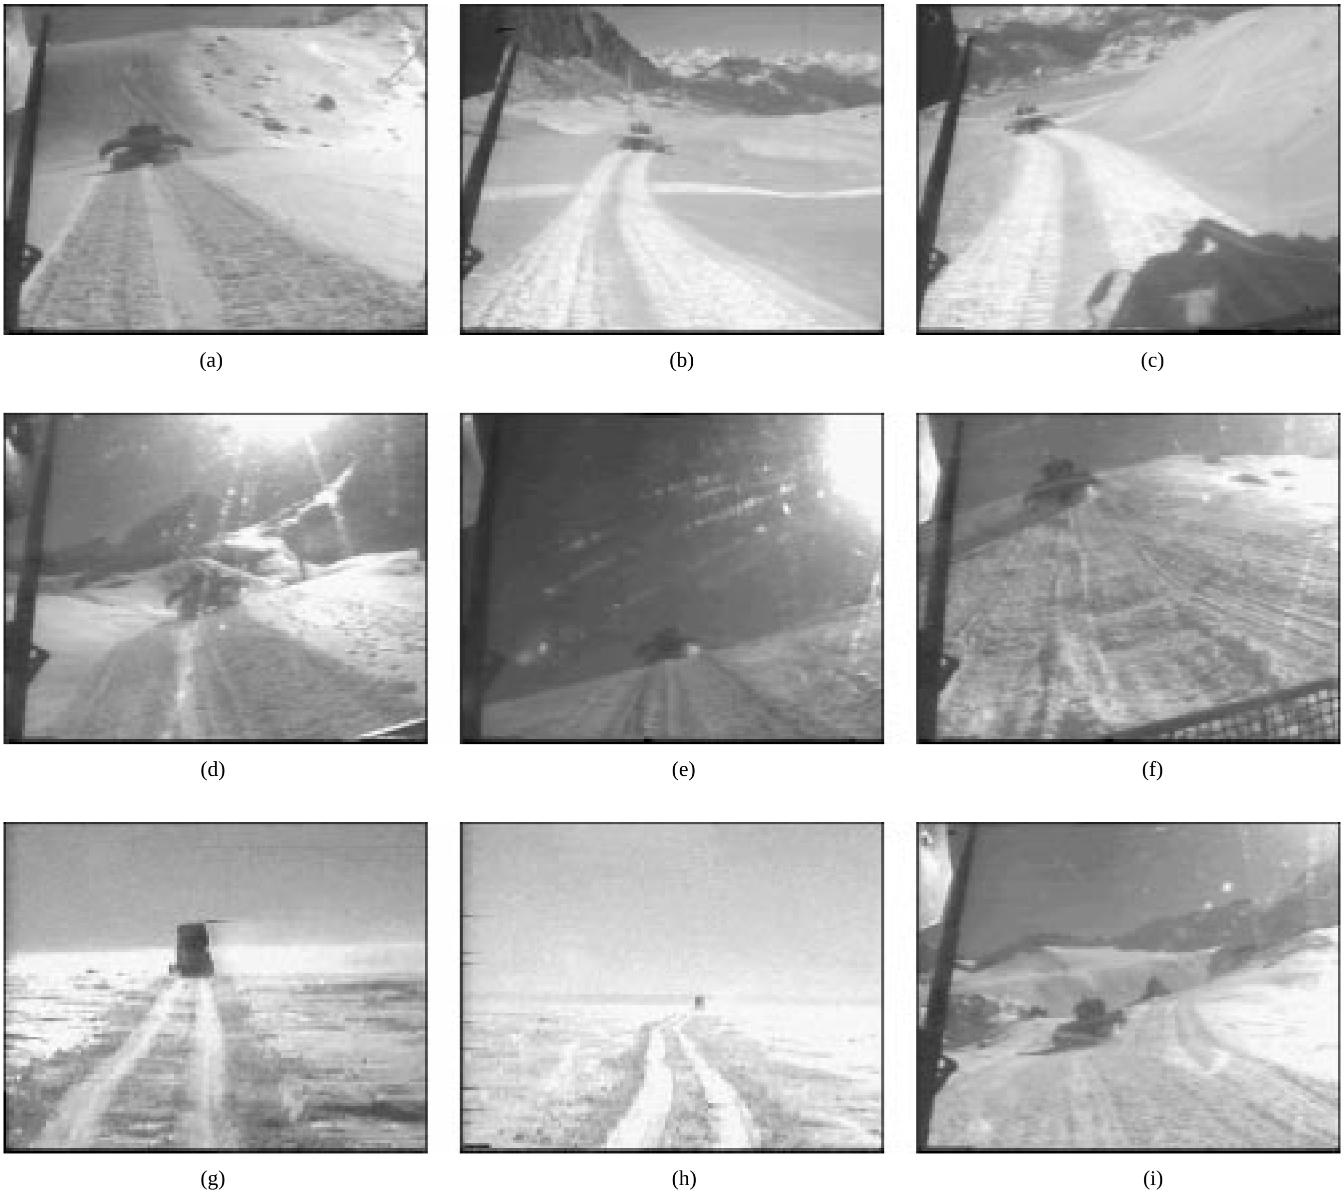

Fig. 9. Images from both the Antarctica region and the Italian Alps.

The specific traveling conditions may also affect the tracks shape and appearance: in case the ahead vehicle is towing a sledge, the tracks will appear as two compact and uniform stripes surrounded by background with a higher brightness variance, as in Fig. 9(g) and (h). On the contrary, the tracks texture when no sledge is used is characterized by a higher brightness variance than the background (see all the other images of Fig. 9).

Finally, no ground references at all can be exploited, as shown by Fig. 9(g) and (h).

In this first version of the system, problems of divergence from previous tracks as visible in Fig. 9(i) are not considered. Furthermore, in some of the sequences acquired for the first tests, the snowcat was equipped with a shovel—visible in the bottom of Fig. 9(f)—and a windshield wiper is present in almost all images. Both objects are filtered out through a specific filter, as discussed in the following section.

III. Tracks Detection

This section presents the description of the processing steps for tracks detection; Fig. 10 sketches the corresponding block diagram.

In order to reduce the complexity of the detection of snowcat tracks in a snowy environment, some assumptions are taken. In the first place, thanks to the low speed of the vehicle, the localization of the tracks in a *nearby area* suffices for the automatic driving of the vehicle. Secondly, focusing on a close region ahead of the vehicle, the nearest portion of the tracks is supposed to be straight and their position is assumed to be slowly varying from frame to frame.

Therefore, for each track border a specific area of interest is defined and analyzed (see Fig. 11). In these two regions, edges are extracted by means of a classical gradient-based approach (Sobel operator), followed by thresholding on both the edges'

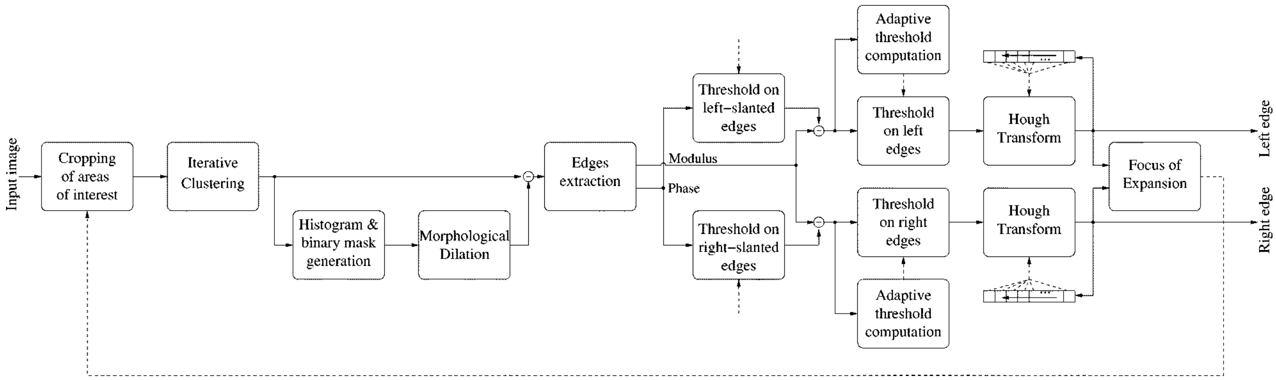

Fig. 10. Block diagram of the complete processing.

(a)

(b)

(c)

Fig. 11. (a) Original image. (b) Area of interest for left border. (c) Area of interest for right border.

phase and modulus. The application of a phase threshold is suggested by perspective considerations: due to the perspective effect the track's border appear as oblique lines. The threshold on the modulus is aimed at extracting the sharpest edges which should be due to the preceding snowcat's track. In order to reduce the sensitivity to both noise and the threshold value itself, a preliminary clustering is applied as well as a specific filter to mask the presence of shadows and/or dark objects. The resulting edge images (right and left border) present a sequence of disconnected regions of high gradient in correspondence of the snow blocks moved by the snowcat. They are then used to recover a linear approximation of the tracks position by means of the Hough transform. The steps of this algorithm are detailed in the following.

Since the contrast between the track and the snowy or icy ground is generally low, the edge extraction phase can benefit from a preliminary clustering step. An iterative procedure proposed in [8] has been used, which is able to enhance also weak and isolated intensity discontinuities. It repeatedly substitutes each pixel's brightness with a weighted average computed over its neighborhood. The definition of the weights comprises a function of the neighborhood which enhances sharp edges and preserves smooth edges, while averaging uniform areas. Fig. 12 shows the result of seven iterations of a 3×3 filter applied to Fig. 11(a).

Fig. 12. Result of the iterative clustering. The procedure is actually applied to the areas of interest only; the whole clusterized image is here presented for displaying purposes.

Subsequently, possible dark objects, such as shadows or framed parts of the internal or external snowcat equipment (windshield wiper, shovel, etc.), have to be detected since they disturb the high-level track recognition phase. An histogram of the pixels intensity is computed in the union of the two search regions, in order to devise a brightness threshold which allows to discriminate between the snow, which is generally bright, and shadows or other dark objects. In this way, sharp edges deriving from dark objects can be masked out, while leaving smooth edges generally representing the tracks position. Fig. 13(a) shows the gray-level histogram computed in the search areas of the clusterized image: the contribution of dark objects [a small shadow in the right bottom corner and part of the windshield wiper on the left of Fig. 11(a)] can be clearly distinguished from the bright ground. The result of the threshold (a binary

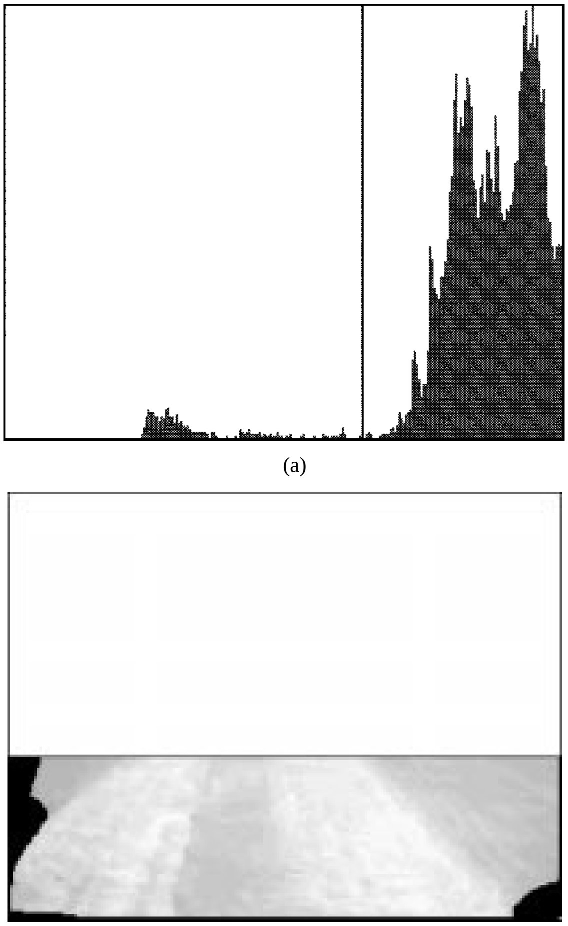

(a)

(b)

Fig. 13. (a) Histogram of gray-level values. (b) Result of masking.

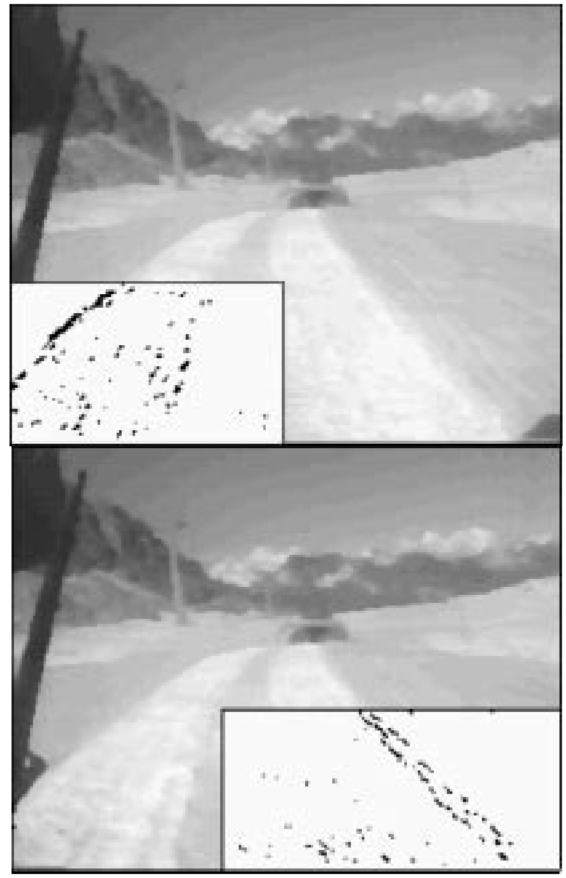

Fig. 14. Edges extracted in the two different areas of interest.

image to be later used as a mask) is then dilated with a 5×5 morphological structuring element [9] to enlarge the masked areas. In Fig. 13(b) the result of masking is presented.

Once dark objects are masked out, edges are extracted in the union of the two search regions by means of a Sobel operator. To separately detect the two tracks' borders, the gradient based filtering is followed by thresholding the edges' phase so to extract edges belonging to forward slanting oblique borders in the left area of interest, and edges belonging to backward slanting oblique borders in the right area of interest. Such filtering has been designed to rely on edges' direction only—and not on a complete 360° phase—in order to work both with tracks brighter and darker than the ground [see, for example, Fig. 9(a) and (b)]. Oblique edges point are then thresholded with respect to their modulus by means of an adaptive threshold which maintains constant from frame to frame the density[1] of surviving edges in the search area. The value of this parameter was experimentally computed from the analysis of several different sequences. In this way, the process is adjusted to the different illumination conditions and variable contrast between tracks and ground: a constant density of edges is obtained by lowering the threshold when the luminance difference is low and raising the threshold when the contrast is high. The threshold value is easily determined from a cumulative histogram of the gray-level intensity values. Fig. 14 shows the edges extracted from Fig. 13(b) in the two different areas of interest.

The Hough transform is then applied to localize the straight line that best fits the edge points of each track border. When selecting the line which gains the highest score in the parameter space, a region centered on the average position of the track in the previous few frames is considered, in order to exploit the strong temporal/spatial correlation. In order to take into account that the tracks edges are generally not perfectly aligned, a local average is performed in the parameter space. This operation smoothes the score of isolated peaks, while preserving the strength of cluster of close peaks.

Once two lines approximating the nearest portion of the track borders have been selected (see Fig. 15), the focus of expansion (FOE) is determined by computing their intersection. The position of the FOE is compared to the previous few ones: if it is too distant from previous results or if it exits from a specific area whose size and position have been determined from the analysis of several sequences (see Fig. 16), the current result is discarded.

Moreover, the two search areas are dynamically resized: their height and width are adapted using the average of the FOE position in a few previous images; the FOE position encodes information on the terrain slope and the relative orientation between the vehicle and terrain (see Fig. 17).

A. Different Use of the Hough Transform

In order to extract the edges of objects from an image, the application of a gradient based operator which gives a measure of

[1]The density of edges is defined as the number of edges divided by the size of the search area, which is dynamically varied as explained in the following.

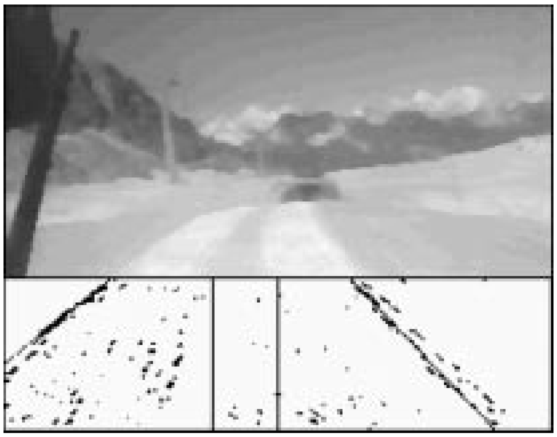

Fig. 15. Straight lines that approximate the track borders.

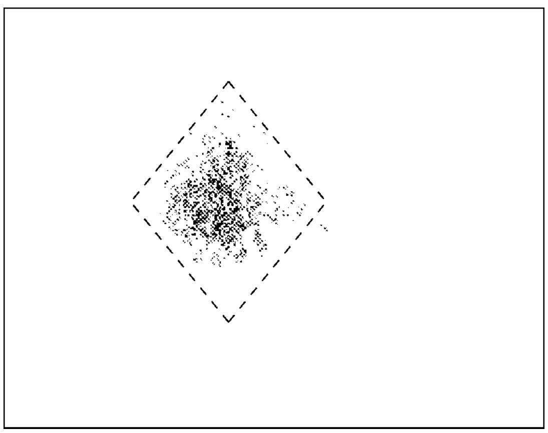

Fig. 16. Image representing the recurrence of the FOE position: the darker the point, the higher the frequency of occurrence in the considered sequences; the dashed bounding box represents the area used to validate the final result.

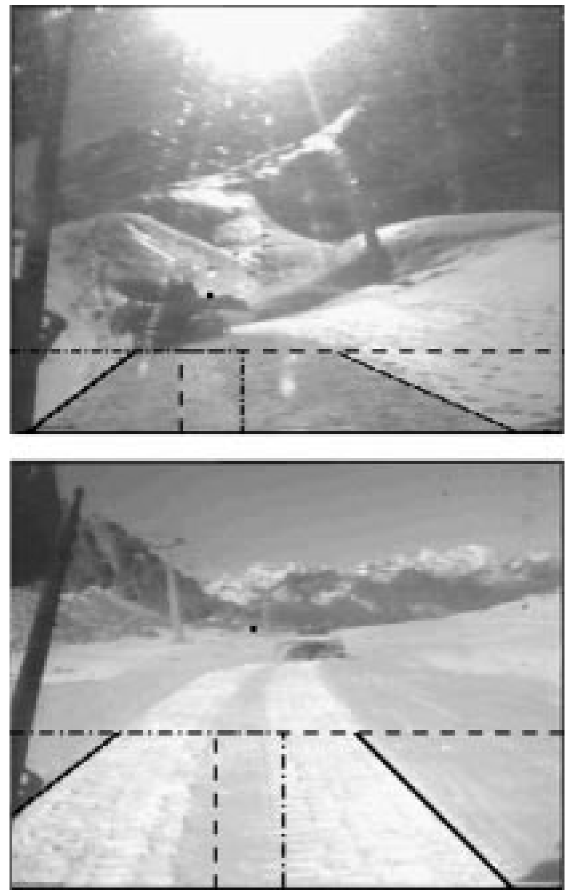

Fig. 17. Different search areas are considered depending on the FOE position, displayed with a black dot.

the brightness discontinuity, such as the Sobel operator, has to be followed by a thresholding operation. It selects the pixels in which the gradient is sufficiently high for the point to be considered an edge.

Unfortunately, following the application of the filter which masks out dark objects, the portion of the non masked image exclusively includes bright snowy soil, therefore, the pixels brightness is not distributed along the whole gray-level scale, but is rather concentrated around the highest gray-level values [see Fig. 13(a)]. Therefore, the luminance differences between each pixel and its neighbors measured by the gradient operator are very small. Thus, the application of a threshold to extract edges in such conditions is very critical.

For this reason a different use of the Hough transform (sketched in Fig. 18) has been considered to avoid the need for a thresholding step which is particularly critical when the gray values are very concentrated. In this procedure the edge points contribute to the score of each line passing through them with their gray value, rather than with a fixed vote. In this way, sharp edges have a higher weight than smooth edges in the selection of the best line, and nevertheless it is not necessary to fix a threshold above which the edge is considered sufficiently strong to be able to vote for the lines.

In Fig. 19(a) the right border of the track is not clearly marked and the classical fixed vote approach fails to detect it. On the other hand, the new variable-weight Hough transform is able to recognize it, as shown in Fig. 19(b). This new alternative procedure is also being experimented on other sequences and its actual effectiveness is under evaluation.

IV. Discussion of Results

Fig. 20 shows some results of snowcat track detection in different conditions.

Fig. 20(a)–(d), (g), and (h) present the result on the corresponding images of Fig. 9. The remaining images of Fig. 9 represent very critical situations where the track cannot be distinguished from background with the current algorithm. In such cases, the system is anyway able to realize and signal its failure in the detection of the track by evaluating the quality of the selected border lines in terms of number of encompassed edge points, position and orientation, and the position of the resulting FOE.

Conversely, Fig. 20(e)–(f) and (i) illustrates situations where the detection is successful even if noisy or critical conditions such as shadows, sun reflections, unknown terrain slope, and dark objects are present.

Generally, the tracks feature a quasilinear behavior in the region of interest. Anyway, it can happen that when approaching a curve—particularly when facing a hill—the tracks begin to deviate from this assumption and the Hough transform result is not precise. For this reason, a simple extension to the Hough

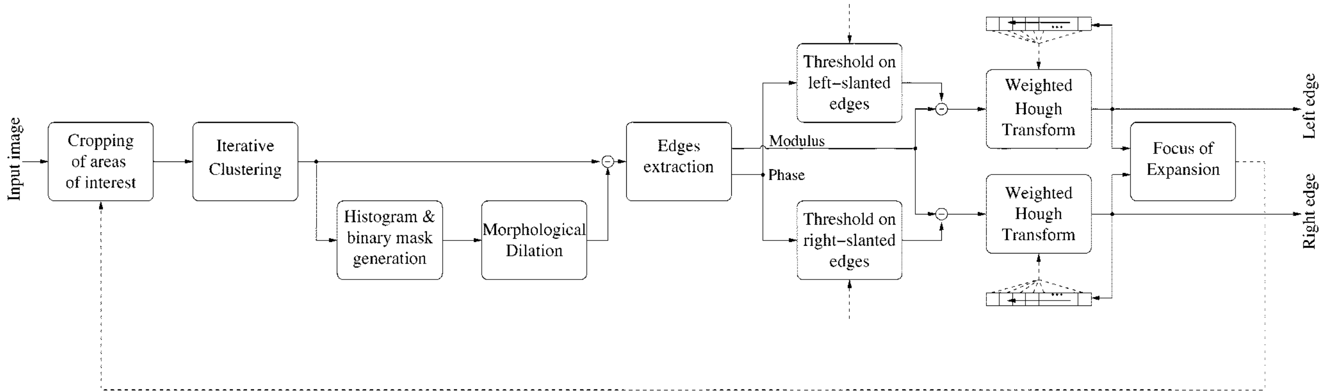

Fig. 18. Block diagram of the processing for the weighted-Hough transform.

(a)

(b)

Fig. 19. (a) Result of the application of the fixed vote Hough transform. (b) Result of the application of the variable-weight Hough transform.

transform is being developed. Instead of looking for a single maximum, when the peak is formed by a cluster of adjacent pixels of similar value some representatives are selected. In this way, the final result is not a single line, but polylines formed by more than one segment.

In any case, in order to address this problem a totally different new approach is being developed which provides a free form chain of pixels as a result. It is based on a stochastic method, where independent agents move on the image trying to walk on high-intensity edge points. The cooperation between the agents assures that the knowledge of the position of high intensity pixels will be shared, and that noncompletely random choices will be made. The following section briefly describes this new approach.

V. Solution With a Different Approach

An evolutionary approach based on the ant colony optimization (ACO) method has been applied to the same problem of snowcat tracks detection. The initial stage of this algorithm is the same: a low-level processing is performed aimed at extracting the image edges, after masking noisy areas to eliminate disturbing objects such as the windshield wiper, shovel, shadows, etc.

Starting from the same edge image, the problem of recognizing the tracks is here converted into that of tracing a consistent line from the lower border of the image up to the line of the horizon that is the shortest and passes through as many edge points as possible. The recognition can be accomplished by means of an algorithm based on the ACO approach. The ACO [10] is a distributed meta-heuristic for hard combinatorial optimization problems. The idea of the methodology is drawn from the behavior of real ant colonies searching for food.

The edge image under analysis functions as a two dimensional map featuring the territory where the artificial ants move. The ants of the first exploring groups have no knowledge of the territory and choose their path randomly, with the heuristic of preferring edges to homogeneous regions. When they reach the top line of a given search area, a pheromone trail is dropped: the shorter the path, the stronger the pheromone trail left by a given ant. The following sets of ants observe the pheromone and trade their random decision with the experience of precedent ants. As more cycles are completed, ants pay more attention to the accumulated knowledge (pheromone) than to the simple heuristic of following the edges. Thus, the whole colony quickly finds a short way to reach the destination line by passing through as many edges as possible. Fig. 21 shows two examples of the tracks recognized by the ants as the best path.

For its intrinsic evolutionary nature, the algorithm is characterized by a set of parameters which influence the solution. As an example Fig. 22 presents the results of two executions of the algorithm driven by different parameters. Indeed the

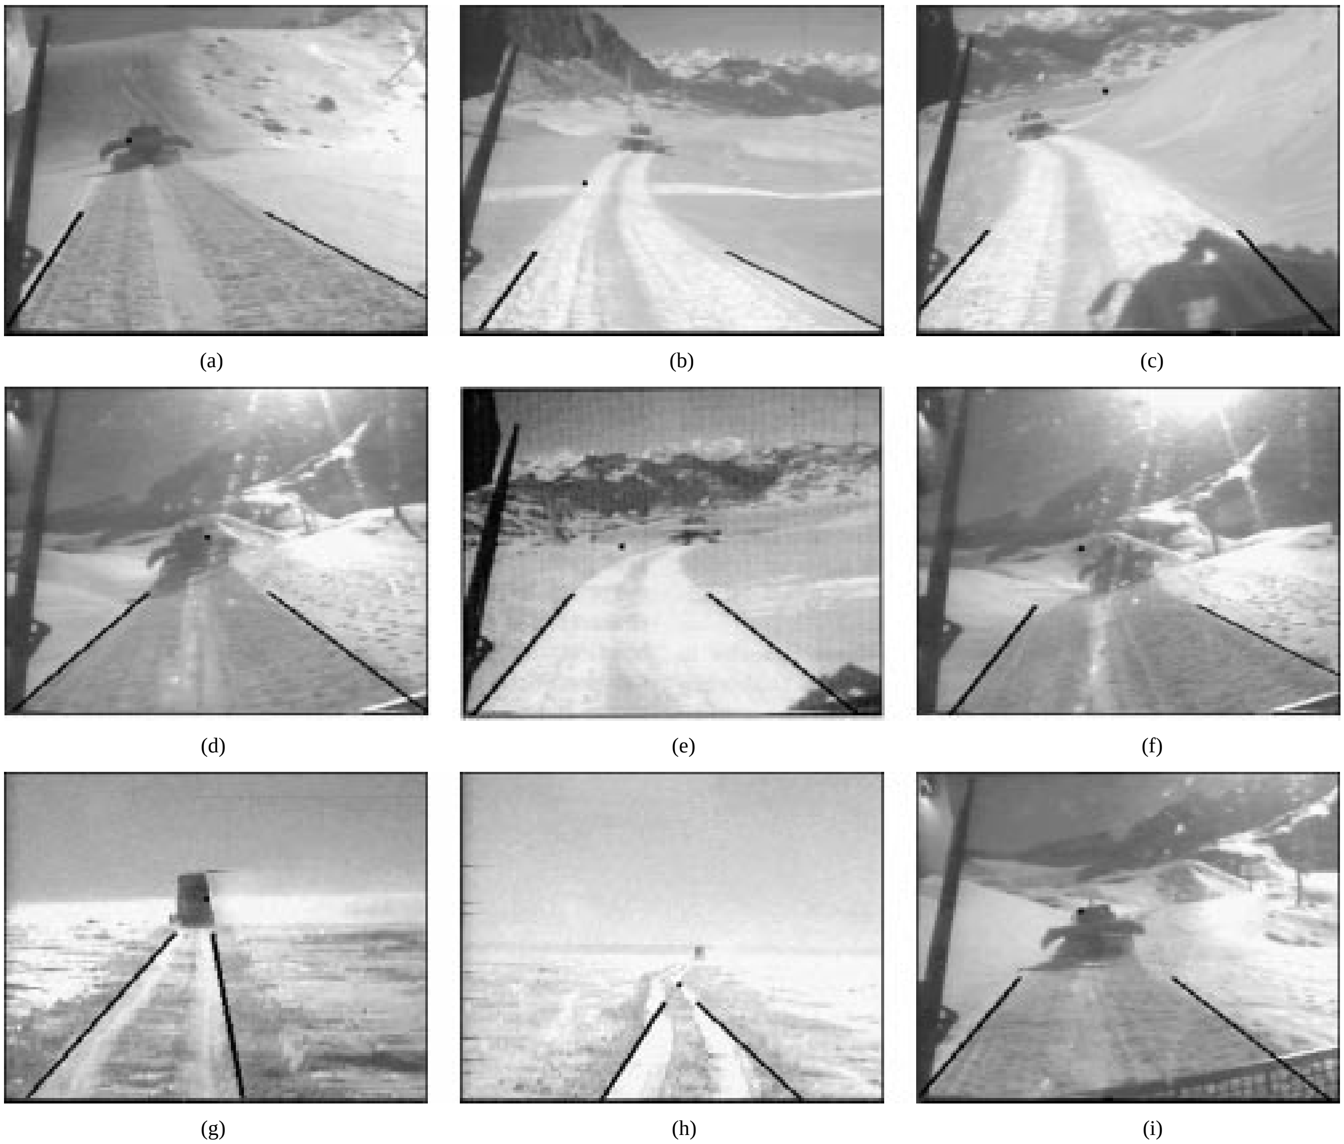

Fig. 20. Results of snowcat track detection in different conditions.

tuning of such parameters represents a critical aspect of this kind of approach.

From the results of a set of preliminary tests this approach appears to be a promising solution deserving further investigation, which is currently under development and will be described in a subsequent paper.

VI. Conclusion and Future Research

The setup and the algorithms discussed in this paper have already been tested on board of the prototype vehicle available to all the research units in the Italian test site. Further tests will be performed on prerecorded images coming from both the Italian test site and the South Pole (the latter have been acquired on previous missions to Antarctica).

In the last months the prototype vehicle was shipped to the South Pole so that during the current scientific mission to Antarctica (from January 2002 to March 2002), the vehicle can be tested on-site and in real environmental conditions. The possibility to test the vision system under illumination conditions that are unique on earth will be of basic importance to adapt the system with respect to possible unexpected problems.

Besides extremely low temperatures, one of the most important issue that must be carefully tested is the behavior of the vision system during stormy conditions. Due to different weather, it is impossible to replicate or emulate antarctic conditions on the Italian test site, and it is, therefore, imperative that exhaustive tests are performed on-site. The tests will take at least two months, during which not only the vision system but also other equipment such as speedometer, radar, lidar, GPS, and TV broadcasting system will be tested.

During the test period, the research unit will continue working on the artificial-vision project and will address two other important issues:

- precise computation of the terrain slope in front of the vehicle;
- possible use of ground penetrating radar (GPR) to detect snow or ice thickness and density.

Currently, slope detection is obtained as an indirect result of track detection, since the position of the FOE gives a rough in-

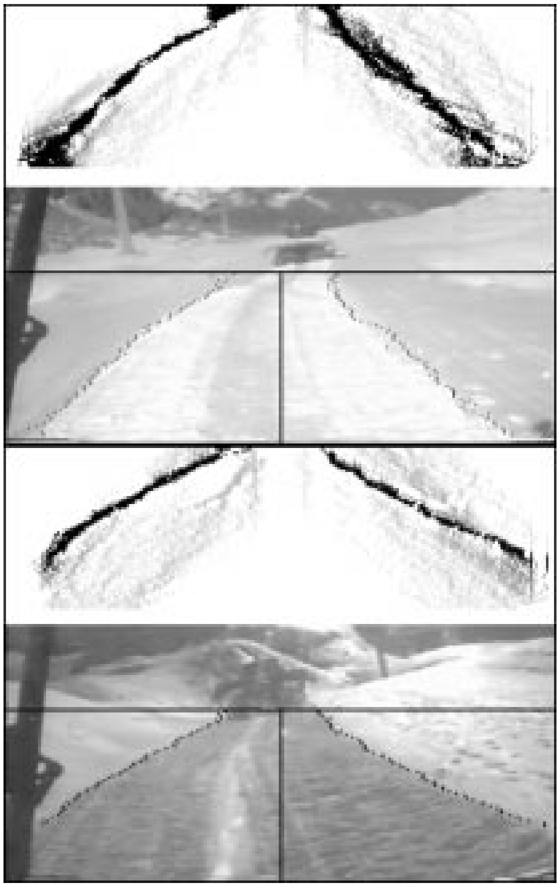

Fig. 21. Examples of the best path selected by the ants; the upper part of each image encodes the number of ants passed on each pixel.

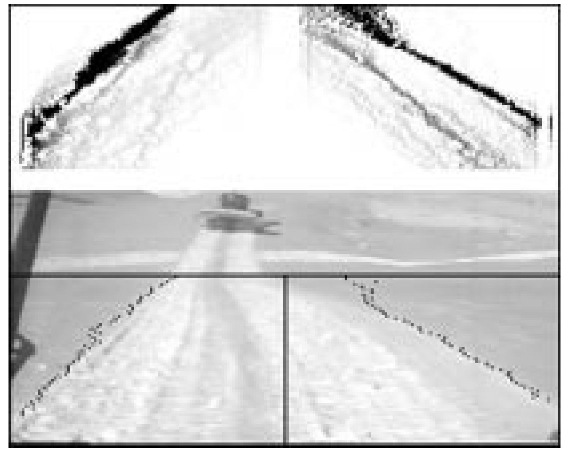

(a)

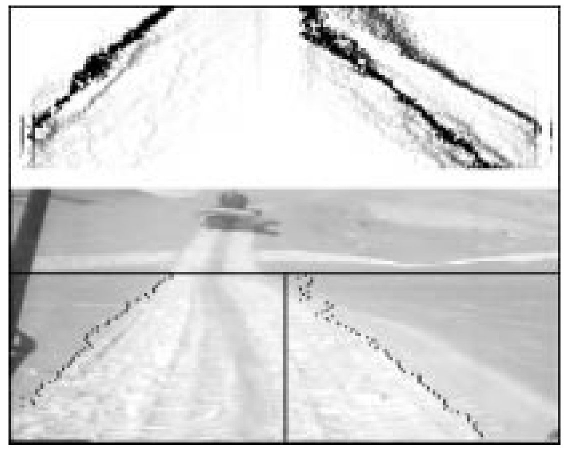

(b)

Fig. 22. Results of two executions of the algorithm with different set of parameters. (a) Wrong right track is selected. (b) Correct right track is selected.

dication about the orientation of the tracks with respect to the vehicle. The FOE horizontal position, together with the measurement of the snowcat pitch and roll, can be used to assess the orientation of the path, and, thus, used to steer the vehicle. On the other hand, the FOE vertical position can be used to obtain an indication about the slope of the terrain ahead.

Since the driving of a snowcat vehicle is extremely complex, even manually, an extremely precise measurement of terrain slope becomes mandatory. Therefore, one of the next research stages will be the use of both cameras—as a stereo pair—to get an accurate three-dimensional (3-D) reconstruction in a limited portion of the terrain ahead.

Another future research stage is based on the use of a GPR, they provide high resolution geophysical imaging of the shallow subsurface. This ability to investigate and map the subsurface using continuous 2-D profiles is advantageous to many engineering and environmental applications, but in our case mainly to safeguard the antarctic robot from terrain dangers and detect subsurface objects. Moreover, the radar data can be used to characterize snow/ice/bedrock stratigraphy. GPR position measurements may also enable the mapping of the location of hazards and interesting subsurface objects and features.

Another problem that does not seem to be easily solvable is the presence of light reflections in the acquired images. Since the camera must be kept inside the driving cabin, due to the very low temperatures of the Antarctic region, images are acquired through the windshield. This causes reflection problems when the windshield presents scratches and when snow or small ice particles stick on the windshield itself.

These new research directions will be dealt with in the months to follow, together with the new evolutionary approach that, although requiring a complex parameter tuning stage, seems to deliver high qualitative performance.

ACKNOWLEDGMENT

The authors would like to thank Dr. C. Moriconi of ENEA, Italy, and his institution for their support of this research, and for the use of the photographs of the Italian base in Terra Nova, Antarctica. Gratitude goes also to the staff of the ski resort of Passo Tonale (Italian Alps) for their help during the tests.

REFERENCES

[1] N. Vandapel, S. Moorehead, W. R. L. Whittaker, R. Chatila, and R. Murrieta-Cid, "Preliminary results on the use of stereo color cameras and laser sensors in antarctica," presented at the Int. Symp. Experimental Robotics, Sydney, Australia, Mar. 1999.

[2] A. Foessel, S. Chheda, and D. Apostolopoulos, "Short-range millimeter-wave radar perception in a polar environment," presented at the Proc. Field and Service Robotics Conf., Pittsburgh, PA, Aug. 1999.

[3] S. Moorehead, R. Simmons, D. Apostolopoulos, and W. R. L. Whittaker, "Autonomous navigation field results of a planetary analog robot in antarctica," presented at the Int. Symp. Artificial Intelligence, Robotics and Automation in Space, Noordwijk, The Netherlands, June 1999.

[4] A. Broggi, M. Bertozzi, and A. Fascioli, "Requirements for visual perception of automotive environments," in *Multimedia Video-Based Surveillance Systems: Requirements, Issues and Solutions, Kluwer International Series in Engineering and Computer Science*, G. L. Foresti, P. Mähönen, and C. S. Regazzoni, Eds. Boston, MA: Kluwer, Oct. 2000, vol. 573, ch. 1.3.

[5] A. Broggi and E. D. Dickmanns, "Applications of computer vision to intelligent vehicles," *Image Vis. Comput. J.*, vol. 18, pp. 365–366, Apr. 2000.

[6] M. Bertozzi, A. Broggi, and A. Fascioli, "Vision based intelligent vehicles: State of the art and perspectives," *J. Robot. Auton. Syst.*, vol. 32, pp. 1–16, June 2000.

[7] A. Broggi, M. Bertozzi, G. Conte, and A. Fascioli, "ARGO prototype vehicle," in *Intelligent Vehicle Technologies*, L. Vlacic, F. Harashima, and M. Parent, Eds, London, U.K.: Butterworth-Heinemann, June 2001, ch. 14.

[8] A. Broggi, "Parallel and local feature extraction: A real-time approach to road boundary detection," *IEEE Trans. Image Processing*, vol. 4, pp. 217–223, Feb. 1995.

[9] J. Serra, *Image Analysis and Mathematical Morphology*. New York: Academic, 1982.

[10] A. Colorni, M. Dorigo, and V. Maniezzo, "Distributed optimization by ant colonies," in *Proc. Eur. Conf. Artificial Life*. Cambridge, MA, 1991, pp. 134–142.

Alberto Broggi (S'90–A'96) received the Dr.Eng. (Master's) degree in electronic engineering and the Ph.D. degree in information technology from the Università di Parma, Parma, Italy, in 1990 and 1994, respectively.

From 1994 to 1998, he was a Full Researcher in the Dipartimento di Ingegneria dell'Informazione, Università di Parma. From 1998 to 2001, he was an Associate Professor of artificial intelligence in the Dipartimento di Informatica e Sistemistica, Università di Pavia, Pavia, Italy. Since 2001, he has been a Professor of computer science at the University of Parma. He is the Coordinator of the ARGO project, with the aim of designing, developing, and testing the ARGO autonomous prototype vehicle, equipped with special active safety features and enhanced driving capabilities. He is the author of more than 120 refereed publications in international journals, book chapters, and conference proceedings. He is actively involved in the organization of scientific events, is on the editorial boards and program committees of many international journals and conferences, and has been invited to act as Guest Editor of journals and magazines theme issues on topics related to intelligent vehicles, computer vision application, and computer architectures for real-time image processing. His research interests include real-time computer vision approaches for the navigation of unmanned vehicles, and the development of low-cost computer systems to be used on autonomous agents.

Alessandra Fascioli (S'97–M'99) received the Dr.Eng. (Master's) degree in electronic engineering and the Ph.D. degree in information technology from the Università di Parma, Parma, Italy, in 1996 and 2000, respectively.

Currently, she is a Temporary Researcher at the Università di Parma. Her research interests focus on real-time computer vision and computer architectures for automatic vehicle guidance and image processing techniques based on the mathematical morphology computational model.

Ms. Fascioli is member of the IEEE Computer Society, Italian Association for Artificial Intelligence (AI*IA), and International Association for Pattern Recognition (IAPR). From 1996 to 1999 she chaired the local IEEE student branch at the Università di Parma.