Machine Vision and Applications (2003) 15: 92–100
Digital Object Identifier (DOI) 10.1007/s00138-003-0133-2

Machine Vision and Applications

Fusion of laser and visual data for robot motion planning and collision avoidance

Haris Baltzakis[1,2], **Antonis Argyros**[1], **Panos Trahanias**[1,2]

[1] Institute of Computer Science, Foundation for Research and Technology – Hellas (FORTH), P.O. Box 1385, Heraklion, 711 10 Crete, Greece
[2] Department of Computer Science, University of Crete, P.O.Box 1470, Heraklion, 714 09 Crete, Greece

Received: 10 May 2002 / Accepted: 18 December 2002
Published online: 7 October 2003 –

Abstract. In this paper, a method for inferring scene structure information based on both laser and visual data is proposed. Common laser scanners employed in contemporary robotic systems provide accurate range measurements, but only in 2D slices of the environment. On the other hand, vision is capable of providing dense 3D information of the environment. The proposed fusion scheme combines the accuracy of laser sensors with the broad visual fields of cameras toward extracting accurate scene structure information. Data fusion is achieved by validating 3D structure assumptions formed according to 2D range scans of the environment, through the exploitation of visual information. The proposed methodology is applied to robot motion planning and collision avoidance tasks by using a suitably modified version of the vector field histogram algorithm. Experimental results confirm the effectiveness of the proposed methodology.

Key words: Sensor fusion – Laser range scanner – Stereo vision – Collision avoidance – Motion planning

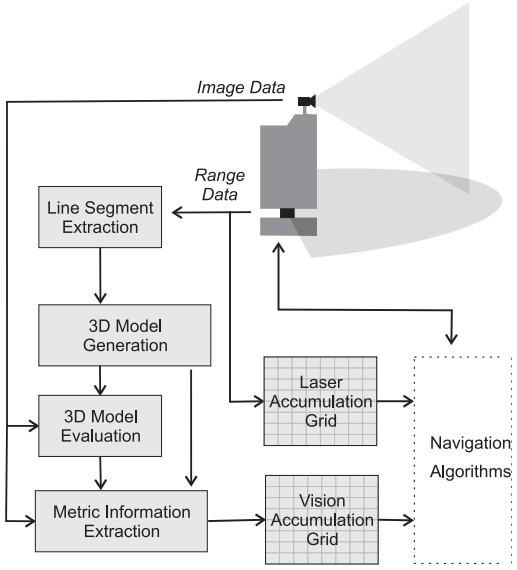

Fig. 1. Block diagram of the proposed methodology

1. Introduction

Laser scanners mounted on mobile robots have recently become very popular for various indoor robot navigation tasks. Their main advantage over vision sensors is that they are capable of providing accurate range measurements of the environment in large angular fields and at very fast rates. Although suitable for many navigation tasks [4,7,13], the quantity of information encapsulated in 2D laser scans may, in certain cases, be insufficient for more crucial and demanding robotic tasks, such as obstacle detection and collision avoidance [3,6,9,12,14]. This is because various objects common even in the simplest indoor environments, such as chairs, tables, or shelves on walls, are sometimes invisible to laser scanners and thus absent or misinterpreted in the resulting 2D profiles.

In this paper, we propose a methodology for fusing laser with visual information in order to infer accurate 3D information. Simple 3D models of the environment consisting of a flat horizontal floor surrounded by vertical planar walls are initially constructed according to 2D laser range data. Vision is then utilized in order to (a) validate the correctness of the constructed model and (b) qualitatively and quantitatively characterize inconsistencies between laser and visual data wherever such inconsistencies are detected. Scene information recovered through this procedure is directly applied to motion planning and collision avoidance through an appropriately modified version of the vector field histogram method [3].

The proposed methodology differs from classical vision-based depth estimation approaches in that visual depth information is extracted only wherever laser range information is proved to be incomplete. The increased computational efficiency meets the real-time demands for collision avoidance and facilitates its application to other crucial robotics tasks as well. Moreover, since information encapsulated in visual data serves to supplement laser range information, inherent advantages of both sensors are maintained, leading to implementations combining accuracy, efficiency, and robustness at the same time. The proposed methodology has been tested

Correspondence to: P. Trahanias (e-mail:trahania@ics.forth.gr)

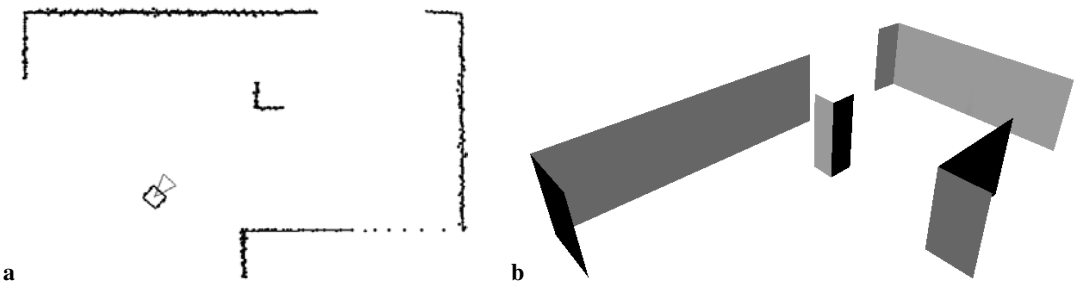

Fig. 2a,b. 3D model creation. **a** Range data and line segments. **b** Resulting 3D model

on both synthetic and real data. The results obtained are very promising and demonstrate its effectiveness.

2. Method description

Based on a single 2D range scan, a local 3D model of the robot's environment is constructed. This is based on the assumption that the environment consists of a flat horizontal floor surrounded by piecewise vertical planar walls. This assumption holds in most indoor environments. To facilitate the construction of the 3D model, range measurements are first grouped into line segments. For this task, a homebuild segmentation algorithm is used, which is based on the iterative end point fit (IEPF) algorithm [5]. Each of the resulting line segments, according to the assumptions made about the environment, corresponds to a vertical planar surface in the resulting model.

In order to evaluate the developed 3D model, a pair of images acquired by a calibrated stereo vision rig is used. Points from the first image are ray-traced to the 3D model, and 3D coordinates are estimated. Based on this information, image points are reprojected onto the frame of the second camera. If the assumed 3D model is correct, then the image constructed through reprojection should be identical to the one actually acquired by the second camera. Wherever the model is not correct, images should differ. A local correlation of image intensity values reveals regions with such inconsistencies indicating that the corresponding parts of the environment do not conform to the assumptions made when building the 3D model.

Vision is further utilized in these cases to provide additional depth information in regions where the 2D range data proved to be insufficient. For this purpose, image intensity matches are sought in the vicinity of the initial position, falsely indicated by the 3D model, along the direction of the epipolar line [8]. All intensity matches are transformed to real-world coordinates and accumulated, along with laser range information, on a local 2D occupancy grid that is utilized by the collision avoidance module to produce proper motion commands.

A block diagram of the proposed methodology is given in Fig. 1. In the next section, the process of creating local 3D models of the environment according to laser data is described. Utilization of visual data to evaluate the generated 3D models and extraction of visual metric information are discussed in Sects. 4 and 5. Finally, in Sect. 6, details regarding the application of the proposed methodology for collision avoidance are provided.

3. 3D model generation

In order to build the 3D model of the environment, range measurements (interpreted as points on 2D profiles of the robot's environment) are first grouped into line segments. For line segment extraction, a three-stage algorithm has been implemented. Range measurements are initially grouped to clusters of connected points according to their sphere-of-influence graph [10]. Clusters are then further grouped to line segments by utilizing the iterative end point fit (IEPF) algorithm [5]. Finally, line segment parameters are reestimated by a line fitting procedure.

For generating the local 3D model of the environment, an infinite horizontal plane (floor) is assumed right below the robot, at a known distance from the robot's coordinate system (the position of the range finding device). Then, line segments defined in the previous step are extended to form rectangular vertical surfaces of infinite height. More specifically, for each line segment, the plane that is perpendicular to the floor and contains the line segment is inserted into the 3D model. The coordinate system of the generated 3D model is assumed to coincide with the coordinate system of the robot.

Figure 2a shows the line segments extracted by the algorithm described above for a simple artificial environment. The extracted lines are superimposed on the simulated range measurements. A rendered view of the resulting 3D model is depicted in Fig. 2b.

4. Model evaluation

Let M be the 3D model built as described in Sect. 3, according to range data acquired at time t_1, and let I_1 be an image acquired by a camera c_1 at the same time instant.

For each image point $p_1 = (x_1, y_1)$ of I_1, the coordinates (X, Y, Z) of the corresponding 3D point P can be computed by ray-tracing it to the model M. If the assumptions made for constructing the 3D model are correct, the (X, Y, Z) coordinates found by the above procedure correspond to a real-world point on M. Let I_2 be a second image, acquired by the second camera of the stereoscopic system. Since the coordinate system of c_2 with respect to the coordinate system of M is also known, the projection $p_2 = (x_2, y_2)$ of $P = (X, Y, Z)$ on c_2 can also be computed.

The above procedure of ray-tracing points of I_1 to find 3D world coordinates and back-projecting them to I_2 leads to analytical computation of point correspondences between I_1 and I_2. If the assumptions made to form the model M are correct, corresponding image points would actually be projections of

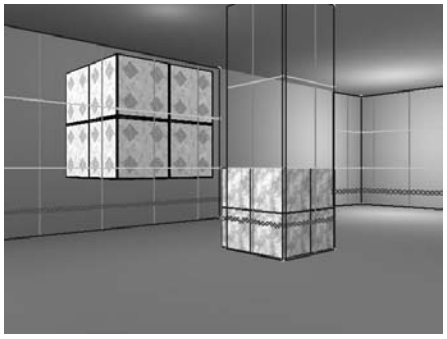

a

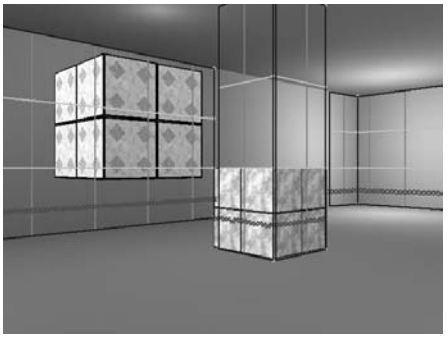

b

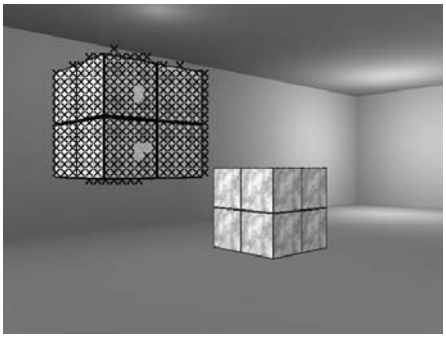

c

Fig. 3a–c. Example of the 3D model evaluation process. **a, b** Two frames of an artificial scene containing two cubes. The 3D model and the range finder points are also projected onto the images. **c** Results of the evaluation process, projected onto the second image. Areas where vision and laser data are inconsistent are marked with "x"s

the same world points and thus should share the same attributes (color, intensity values, intensity gradients, etc.). If this is not the case (e.g., image points have different attributes), then a strong indication exists that the model is locally invalid. The normalized crosscorrelation metric [5] is employed to evaluate the correctness of the calculated point correspondences. Low values of the calculated crosscorrelation correspond to regions within the images depicting parts of the environment that do not conform to the 3D model.

Figures 3a and 3b show two frames of a synthetic scene captured by a simulated stereo vision rig at the position corresponding to the range data shown in Fig. 2a. For convenience, wire-frames of the 3D model, extracted as described in Sect. 3, together with range finder measurements, are projected on the images. The scene contains two cubes; one lying at approximately 1 m above the floor (on the left part of the image) while the other (on the right) is placed directly on the floor. Figure 3c demonstrates the results of the evaluation process. Evidently, the algorithm succeeds in correctly detecting the cube on the

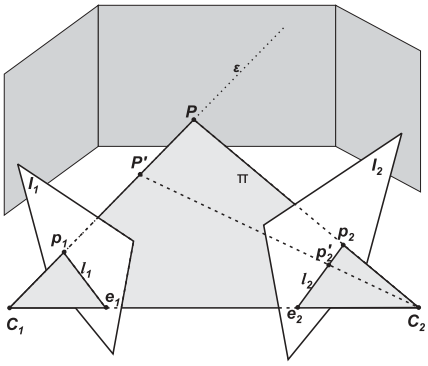

Fig. 4. Epipolar geometry for the two cameras

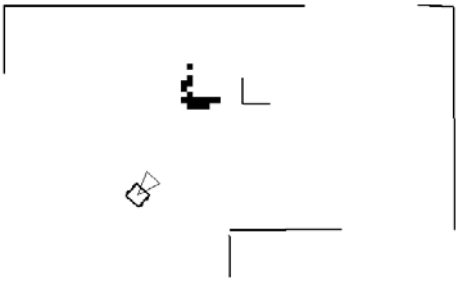

Fig. 5. Extraction of metric information by utilization of visual data

left of the image that is "invisible" to the range finder since it is floating above its scanning plane. The cube on the right of the image does not yield any unmatched areas since it lies on the floor and hence is visible to the laser scanner. It should be noted at this point that the area above the right cube should also yield inconsistencies, which are not identified due to complete lack of texture in the specific artificial scene.

5. Extraction of visual depth information

In the previous section, vision was used in order to detect inconsistencies between the 3D model that was extracted based on the 2D laser range information and the real 3D environment. Having already detected such inconsistencies, visual information can be further exploited to provide quantitative information (i.e., depth values) for the regions where laser measurements proved inadequate in modelling the real environment.

Let us assume the camera configuration depicted in Fig. 4. The 3D point P lying on the model M is projected onto the point p_1 on the left image (I_1) and onto the point p_2 on the right image (I_2). The epipolar plane π created by point P and the camera centers C_1 and C_2 intersects the image planes in lines l_1 and l_2, the epipolar lines [8].

Suppose that the image coordinates of p_1 are known and that neither the location of the 3D point P nor point p_2 in the second image is known. It has already been shown that by ray-tracing point p_1 to the 3D model we can compute the 3D coordinates of point P and by projecting the latter to the second camera plane we can estimate the corresponding point p_2 in the second image. Suppose that by locally correlating pixel intensity values at the positions of p_1 and p_2 we discover a dissimilarity and conclude that the 3D model M is inaccurate

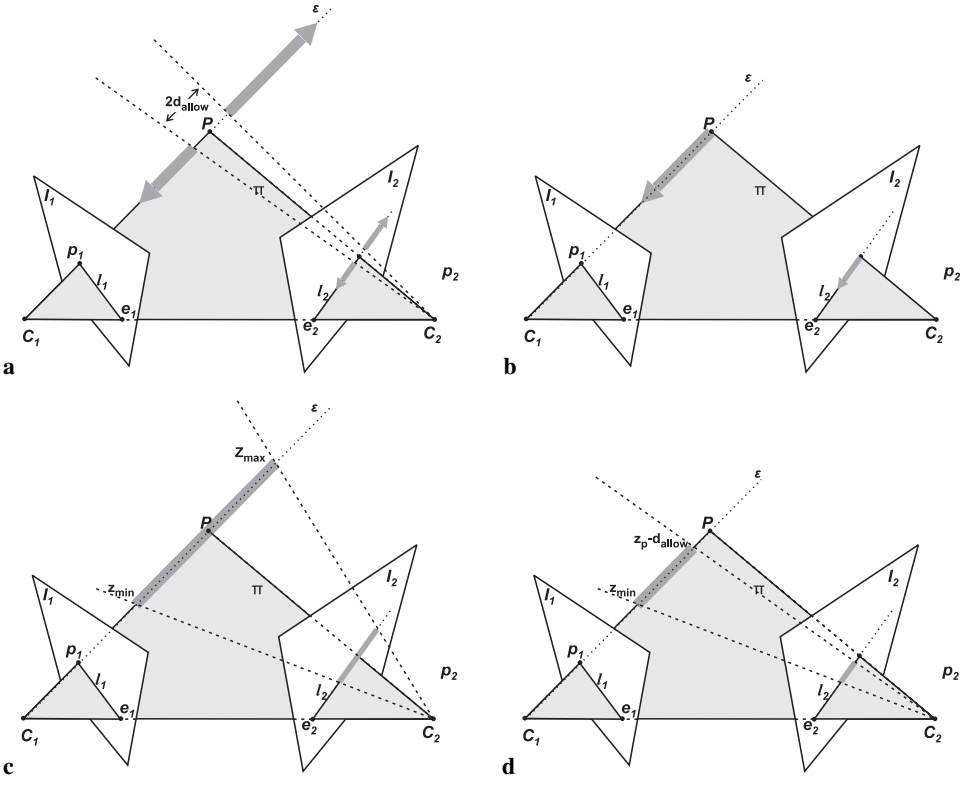

Fig. 6a–d. Visual depth computation criteria for collision avoidance. **a** $|Z_{vision} - Z_{laser}| \geq d_{allow}$, **b** $Z_{vision} \leq Z_{laser}$, **c** $Z_{min} \leq Z_{vision} \leq Z_{max}$, **d** combination of all criteria

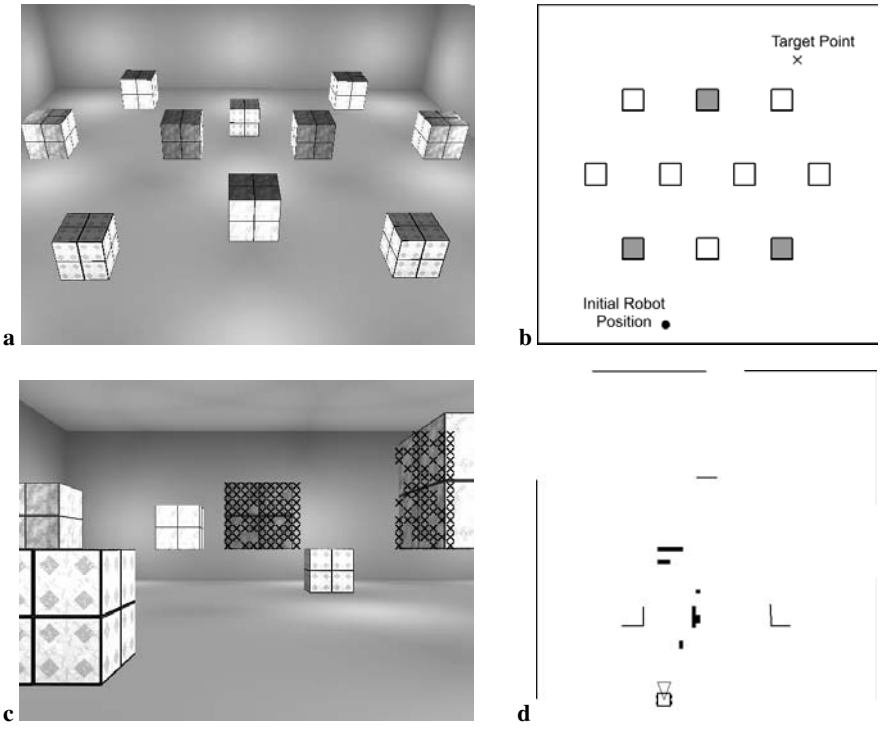

Fig. 7a–d. Visual depth extraction for collision avoidance in an artificial scene. **a, b** Rendered view and sectional plan of an artificial environment. **c, d** Visual depth extraction for collision avoidance at the initial position of the robot

and hence the 3D coordinates of point P, as implied by the model, are not correct. This raises the question of whether point P actually lies behind model M (further to the camera than assumed) or in front of it; the latter would make it a potential obstacle invisible to the range finder.

Whatever the actual depth of point P, its projection on the second camera will comply to the epipolar constraint; that is, it will lie on the epipolar line l_2. The important observation is that the shorter the actual depth of point P, the closer will be its projection p_2 to the epipole e_2. That is, if point p_1 actually corresponds to a 3D point P' closer to the first camera than point P, its projection $p_2{}'$ on the second camera will lie on the epipolar line l_2, between p_2 and the epipole e_2 (Fig. 4). If point P' lies farther from the first camera than P, its projection $p_2{}'$ will also lie on l_2 but this time in the opposite direction of e_2. If the exact location of $p_2{}'$ corresponding to p_1 were

known (known point correspondences), computation of the intersection of the line passing through points C_1 and P with the line passing though C_2 and ${p_2}'$ would yield the exact 3D location of P'.

However, accurate computation of ${p_2}'$ is not always possible due to lack of sufficient image intensity variation (texture) or simply to image noise. Relying on the assumption that for robots that move on a planar surface the projection of the obstacles on the 2D surface of motion suffices for navigation, we alleviate the problem of spurious range evidence by accumulating range estimates in a 2D occupancy grid [1, 11].

Figure 5 demonstrates the results of the procedure described above for the synthetic data set used in the previous sections. As can be easily observed, the location of the left cube shown in Fig. 3, although not visible by the range finder, is correctly identified by our method and placed in the resulting occupancy grid (Fig. 5).

6. Collision avoidance

The proposed framework for fusion of range and visual information is directly applicable to the collision avoidance task. In this section, we present a collision avoidance algorithm that utilizes range data extracted as described in the previous sections. The presented algorithm is a modified version of the vector field histogram (VFH) method [3]. VFH creates motion commands by considering the peaks and valleys of a polar obstacle density histogram. For computing the values of the polar histogram, a local occupancy grid combining both laser and visual information is calculated in real time.

6.1. Efficiency considerations

According to the procedure described in Sect. 5 for inferring depth information out of visual data, image intensity matches are sought for each pixel along its corresponding epipolar line. Elimination of depth computations so that they take place only wherever laser vision inconsistencies are detected effectively reduces the required computational time. Further reduction of computational time is feasible by proper definition of the search area along the epipolar line.

For collision avoidance, the VFH algorithm does not need to utilize all aspects of range information. For example, range information regarding obstacles that lie far from the robot (outside the active window) is not directly usable by the VFH algorithm and need not be computed. Moreover, visual range information obscured by corresponding laser range information can safely be eliminated as well. Thus a large percentage of depth computations can be safely eliminated without compromising the performance of the collision avoidance procedure.

The above observations lead to a set of criteria that dictate when extraction of visual range information is meaningful for collision avoidance:

- Criterion a: Extracted visual range differs significantly from corresponding laser range ($|Z_{vision} - Z_{laser}| \geq d_{allow}$). According to this criterion, portions of the epipolar line that correspond to 3D points lying very close to the 3D model are eliminated from the search area (Fig. 6a).
- Criterion b: Extracted visual complies with predefined, application-specific bounds ($Z_{vision} \leq Z_{laser}$). That is, the search area along the epipolar line is bound by the two points, p_2 and the epipole e_2. This criterion is graphically depicted in Fig. 6b.
- Criterion c: Extracted visual range lies neither too far nor too close to the robot ($Z_{min} \leq Z_{vision} \leq Z_{max}$). This criterion imposes a further substantial reduction of the search space by setting two boundaries along the epipolar line, corresponding to maximum and minimum depth values, Z_{min} and Z_{max}, respectively (Fig. 6c).

Each of the above criteria imposes limitations on the search area along the epipolar line. Figures 6a–c graphically depict these limitations. In Fig. 6d, the combination of all the above criteria and the resulting search area are also depicted.

6.2. Algorithm operation

Figures 7 and 8 demonstrate the application of the proposed algorithm for collision avoidance in an artificial environment containing obstacles that are both visible and invisible to the laser range scanner. A rendered view and a sectional plan of this environment are depicted in Figs. 7a and 7b. Shaded boxes in Fig. 7b indicate obstacles that are visible to the laser scanner (lying on the floor). Unshaded boxes indicate obstacles lying above the scanning plane of the laser range finder. The robot is initially positioned at the bottom of Fig. 7b, and the desired target is indicated by an "x" in the same figure. The application of the visual depth computation algorithm on the initial robot position is depicted in Figs. 7c and 7d. As can be observed, the potential obstacles that are invisible by the laser range scanner are correctly identified. For guiding the depth information extraction process, all criteria defined earlier in this section were considered.

The operation of the collision avoidance algorithm is demonstrated in Fig. 8, depicting the robot at three different time instances while following a path toward the target in the upper right corner of the scene (marked with a cross).

Figure 9 shows the occupancy maps produced by the robot after moving for a while in the artificial scene of Fig. 7. Figure 9a depicts the occupancy map constructed by utilization of laser only information, while Fig. 9b depicts a map constructed by utilization of both laser and visual information. All boxes, including those lying above the scanning plane of the laser range finder, were identified by the robot and are correctly represented on the corresponding map.

7. Results

The method proposed in this paper has been implemented and assessed on a robotic platform of our laboratory, namely, an iRobot-B21r equipped with a SICK-PLS laser range finder and a pair of cameras operating at a standard resolution of 640×480 pixels. The range finder is capable of scanning 180 degrees of the environment, with an angular resolution of one measurement per degree and a range measuring accuracy of 5 cm. An internal calibration procedure has been applied prior to testing our methodology to estimate the relative positions

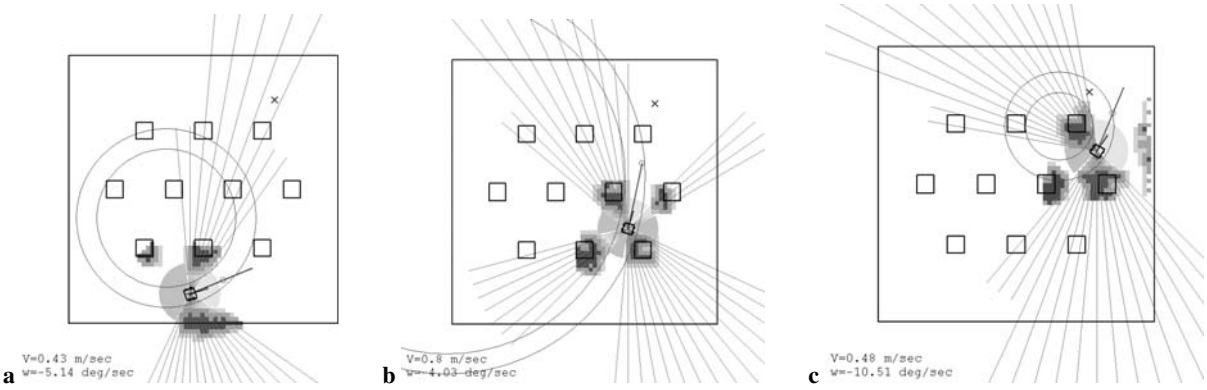

a b c

Fig. 8a–c. Operation of the motion planner in the artificial scene of Fig. 7

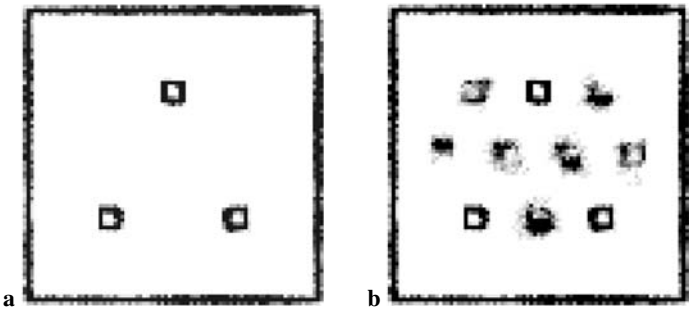

a b

Fig. 9a,b. Occupancy maps produced for the artificial scene of Fig. 8a. **a** Utilizing laser information only. **b** Utilizing both laser and visual information

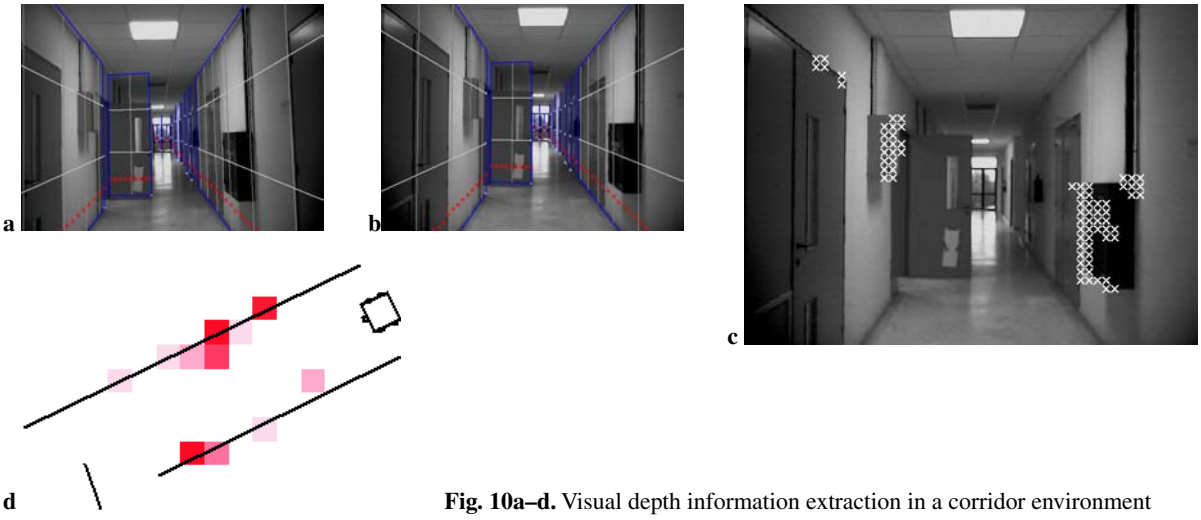

a b c d

Fig. 10a–d. Visual depth information extraction in a corridor environment

and the intrinsic parameters of all sensors. Extensive tests have been performed; in all cases, very accurate performance was observed, verifying the effectiveness and robustness of the method.

Figure 10 demonstrates the operation of the proposed method in a corridor environment. A pair of images acquired by the robot are shown in Figs. 10a and 10b. Projections of range finder data as well as of the resulting 3D model are also overlaid on the images. The results of the evaluation process, projected onto the second image, are shown in Fig. 10c. As can be verified, various structures lying on the walls of the corridor and invisible to the range finder, were correctly identified by the evaluation process. Results of the metric information extraction algorithm, in the form of an occupancy grid map, applied to these areas of range data inconsistency, are depicted in Fig. 10d. For illustration purposes, line segments used for constructing the 3D model are also overlaid.

Figure 11 demonstrates the fusion of laser and visual data in a more complex real scene. Both images used for this purpose are depicted in Figs. 11a,b. As can be observed, the scene contains obstacles that are not fully visible to the laser range finder (for example, laser beams pass below the desk to the right of images 11a,b, reflecting on the wall that is behind the desk). Results of the evaluation process are depicted in Fig. 11c, while local occupancy grids, constructed by accu-

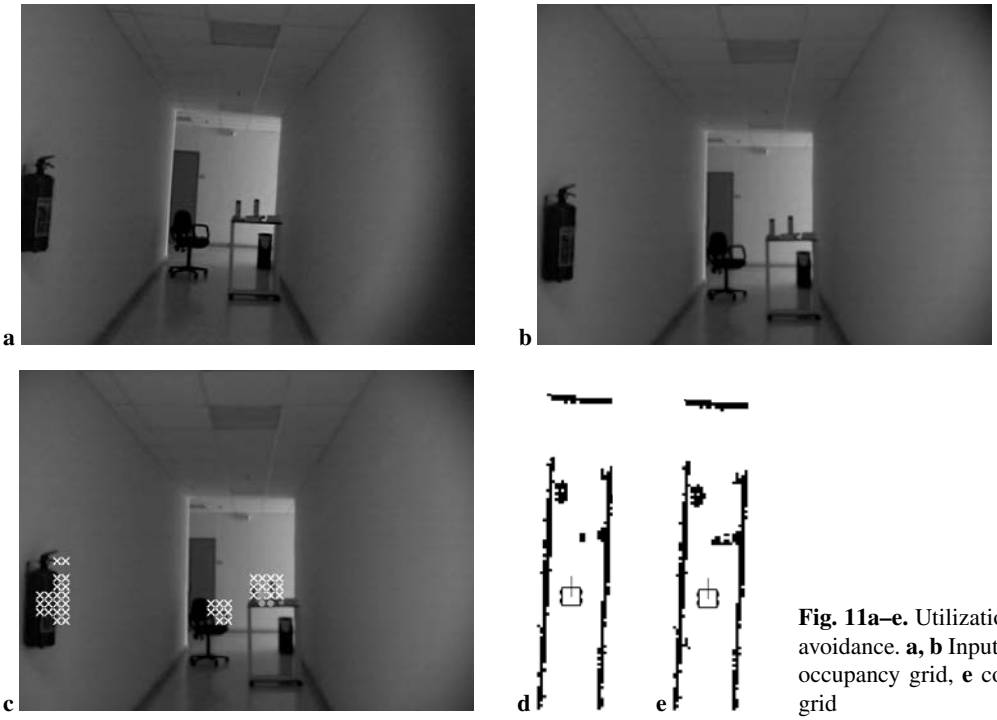

Fig. 11a–e. Utilization of visual information for collision avoidance. **a, b** Input images, **c** evaluation process, **d** laser occupancy grid, **e** combined laser and vision occupancy grid

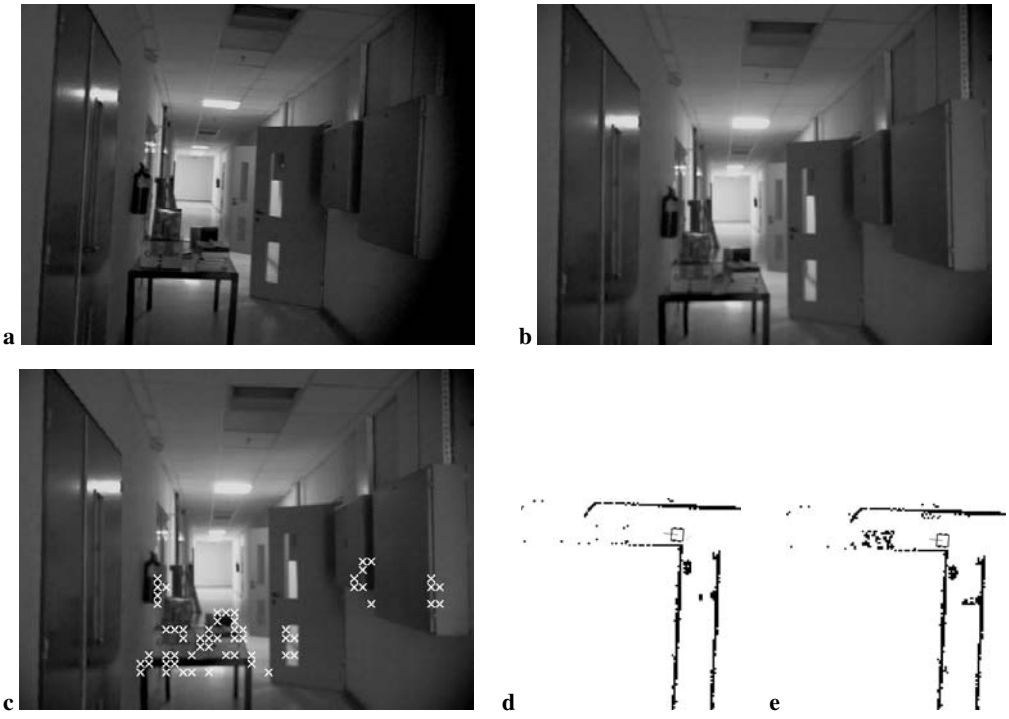

Fig. 12a–e. Utilization of visual information for collision avoidance. **a, b** Input images, **c** evaluation process, **d** laser occupancy grid, **e** combined laser and vision occupancy grid

mulating laser and laser plus visual information, are displayed in Figs. 11d and e respectively.

The operation of the proposed framework is further demonstrated in Fig. 12, depicting a desk lying on the same corridor structure outside our laboratory. As in the previous experiments, the algorithm successfully identifies the desk that is invisible to the laser range finder (only its legs are apparent in the laser scan).

Figure 13 demonstrates the application of the obstacle avoidance algorithm described in Sect. 6 for the scene of Fig. 11. The robot safely maneuvers around all obstacles in the scene in order to reach its target.

Figure 14 demonstrates the applicability of the proposed framework for other navigational tasks besides the collision avoidance task. Figures 14a,b show two occupancy grid maps constructed by the simple accumulation procedure described in Sect. 6. The first map, Fig. 14a, is constructed by accumu-

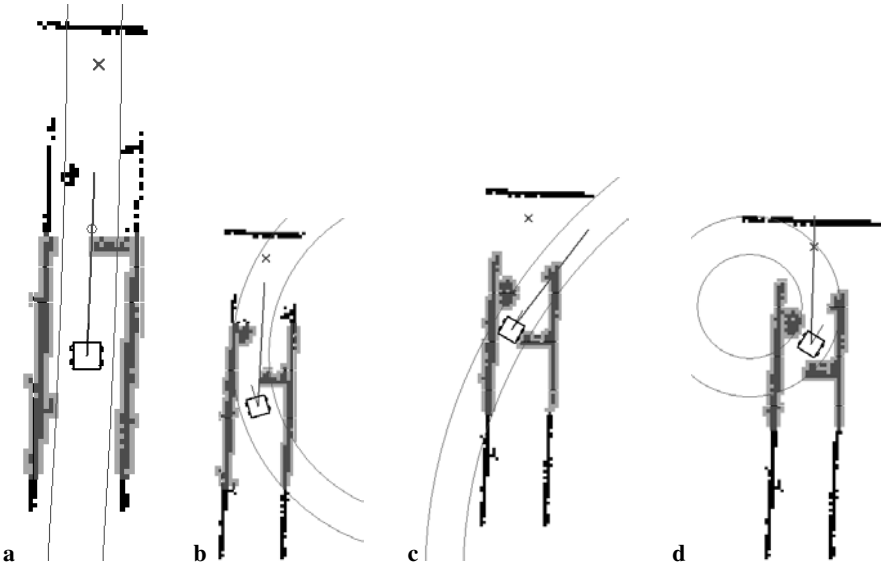

Fig. 13a–e. Collision avoidance for the scene of Fig. 11

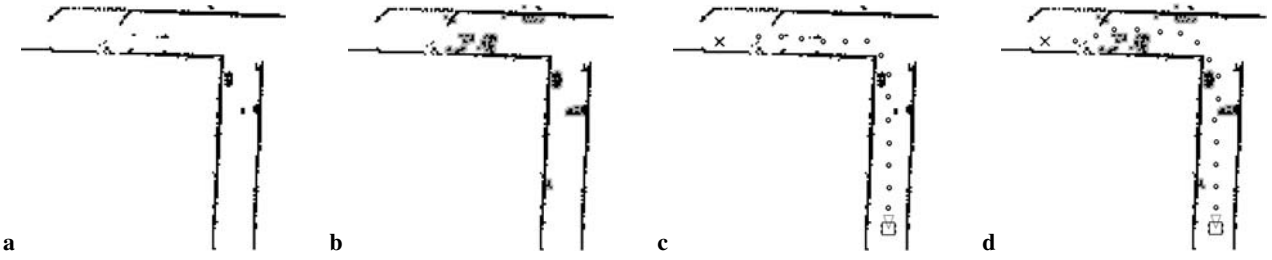

Fig. 14a–d. Global maps constructed (**a, b**) and paths planned (**c, d**) by utilization of laser information only (**a, c**) and laser and vision information (**b, d**)

lation of laser only range information, while for constructing the second map visual information was used as well. The maps were constructed by using the sensor data collected during the experiments depicted in Figs. 11 and 12. In this experiment, a global localization algorithm [2] was used to extract the required global localization information. Figures 14c,d show the application of a global path planning algorithm on the two maps depicted in Figs. 14a,b. The global path planning algorithm used for this purpose is the value iteration algorithm [13]. As expected, the path produced by using laser only information (Fig. 14c) is not correct; the algorithm suggests a path through the desk depicted in Fig. 12. On the other hand, use of visual information, as demonstrated in Fig. 14d, leads to a correct path.

8. Conclusions and future work

In this paper, a new method for fusion of range and visual data for the extraction of scene structure information has been proposed. Moreover, an application of the proposed method for real-time collision avoidance has been demonstrated. Simple 3D local models of the environment are formed based on 2D range data. Visual information is used to locally evaluate the constructed models and to detect regions that are inconsistent with visual data. In places where the evaluation process yields inconsistencies among the range data and the robot's environment, vision is utilized to provide additional metric depth information. Since pixel displacements are computed analytically by rendering image points to the model, their direct computation is not necessary. The proposed method requires two views of the environment that are acquired by a calibrated stereoscopic vision system. However, it can also be applied in cases of robots with a single camera, provided that accurate egomotion information can be computed [1].

Besides the demonstrated applicability of the proposed method for collision avoidance, the general idea presented in this paper can be used by range-based mapping and navigation algorithms to make them more accurate and robust. We believe that fusion of data provided by range and vision sensors constitutes an appropriate framework for mobile robot platforms to perform demanding navigation tasks. It is our intention to further investigate the applicability of the methodology presented in this area.

References

1. Baltzakis H, Argyros A, Trahanias P (2002) Fusion of range and visual data for the extraction of scene structure information. In: Proceedings of the international conference on pattern recognition, Quebec, 11–15 August 2002, pp 40007–40012
2. Baltzakis H, Trahanias P (2003) A hybrid framework for mobile robot localization: formulation using switching state-space models. Autonomous Robots J 15:169–191

3. Borenstein J, Koren Y (1991) The vector field histogram: fast obstacle avoidance for mobile robots. IEEE Trans Robot Autom 7(3):278–288
4. Castellanos J, Tardós J, Neira J (1996) Constraint-based mobile robot localization. In: Proceedings of the international workshop on advanced robotics and intelligence machines, University of Salford, Manchester, UK, 2–3 April 1996
5. Duda R, Hart P (1973) Pattern classification and scene analysis. Wiley-Interscience, New York
6. Fornland P (1995) Direct obstacle detection and motion from spatio-temporal derivatives. In: Proceedings of the 6th international conference on computer analysis of images and patterns, Prague, Czech Republic, 6–8 September 1995, pp 874–879
7. Gutmann JS, Konolige K (1999) Incremental mapping of large cyclic environments. In: Proceedings of the international symposium on computer intelligence in robotics and automation, Monterey, CA, 8–9 November 2000, pp 318–325
8. Kanatani K (1993) Geometric computation for machine vision. Clarendon Press, Oxford
9. Matthies L, Grandjean P (1994) Stochastic performance modeling and evaluation of obstacle detectability with imaging range sensors. IEEE Trans Robot Autom 10:783–792
10. Michael T, Quint T (1994) Sphere of influence graphs in general metric spaces. Math Comp Model 29:45–53
11. Moravec HP, Elfes A (1985) High resolution maps from wide angle sonar. In: Proceedings of the international conference on robots and automation, St. Louis, MO, 25–28 March 1985, pp 116–121
12. Stiller C, Hipp J, Rossig C, Ewald A (2000) Multisensor obstacle detection and tracking. Image Vision Comp 18(5):389–396
13. Thrun S, Buecken A, Burgard W, Fox D, Froehlinghaus T, Hennig D, Hofmann T, Krell M, Schmidt T (1996) Map learning and high-speed navigation in RHINO. Technical report IAI-TR-96-3, University of Bonn, Department of Computer Science
14. Zhang Z, Weiss R, Hanson A (1997) Obstacle detection based on qualitative and quantitative 3D reconstruction. IEEE Trans Patt Anal Mach Intell 19(1):15–26

Haris Baltzakis was born in Heraklion, Crete, Greece in 1973. He received his B.Eng. in electrical engineering from Democritus University of Thrace, Xanthi, Greece in 1997 and his M.Sc. in computer science from the University of Crete, Heraklion, Greece in 1999. He is currently a Ph.D. student in the Department of Computer Science, University of Crete, Greece. His research interests include computer vision and autonomous robot navigation.

Antonis A. Argyros is a research scientist at the Institute of Computer Science (ICS) of the Foundation for Research and Technology (FORTH) in Heraklion, Crete, Greece. He received his Ph.D. from the Department of Computer Science, University of Crete, Greece in visual motion analysis. He has been a postdoctoral fellow at the Royal Institute of Technology in Stockholm, Sweden, where he worked on vision-based, reactive robot navigation. In 1999, he joined the Computational Vision and Robotics Laboratory of ICS-FORTH, where he has been involved in many RTD projects in image analysis, computer vision, and robotics. Dr. Argyros has also served as a visiting assistant professor at the Computer Science Department of the University of Crete. His current research interests include computer vision and robotics and particularly the visual perception of motion and 3D structure, the development of robot behaviors based on visual information, and alternative visual sensors. In these research fields, he has published more than 35 papers in peer-reviewed journals and conference proceedings.

Panos Trahanias is a professor in the Department of Computer Science, University of Crete, Greece and the Institute of Computer Science, Foundation for Research and Technology – Hellas (FORTH), Greece. He received his Ph.D. in computer science from the National Technical University of Athens, Greece in 1988. He has been a research associate at the Institute of Informatics and Telecommunications, National Centre for Scientific Research "Demokritos", Athens, Greece. From 1991 to 1993 he was with the Department of Electrical and Computer Engineering, University of Toronto, Canada, as a research associate. He has participated in many RTD programs in image analysis at the University of Toronto and has been a consultant to SPAR Aerospace Ltd., Toronto. Since 1993 he has been with the University of Crete and FORTH. Currently, he is the supervisor of the Computational Vision & Robotics Laboratory at FORTH, where he is engaged in research and RTD programs in vision-based robot navigation, telematics applications, and virtual and augmented environments. Professor Trahanias has published over 75 papers in the above areas and has coordinated and participated in many funded projects.