Integration of 3D Lines and Points in 6DoF Visual SLAM by Uncertain Projective Geometry

Daniele Marzorati † Matteo Matteucci* Davide Migliore* Domenico G. Sorrenti†

*Dept. Electronics and Information, Politecnico di Milano, Milano, Italy

{matteucci,migliore}@elet.polimi.it

†Dept. Informatica, Sistemistica e Comunicazione, Università di Milano-Bicocca, Milano, Italy

{marzorati,sorrenti}@disco.unimib.it

***Abstract*— In this paper we face the issue of fusing 3D data from different sensors in a seamless way, using the unifying framework of uncertain projective geometry. Within this framework it is possible to describe, combine, and estimate various types of geometric elements (2D and 3D points, 2D and 3D lines, and 3D planes) taking their uncertainty into account. By means of uncertain projective geometry, it is possible to derive simple bilinear expressions to represent join and intersection operators using only three matrices as operator. In particular, we are interested in using 3D information coming from different (logical) vision sensors observing the same scene, to improve map accuracy. The experimental section shows that it is possible to improve both mapping accuracy and pose estimation while performing SLAM with a mobile robot, by integrating sensor information coming from trinocular feature-based vision and correlation based stereo.**

***Index Terms*— Sensor Fusion, 6DoF Hierarchical SLAM, Uncertain Projective Geometry, Computer Vision.**

I. Introduction

Simultaneous Localization and Mapping, SLAM hereafter, is a well-known problem in mobile robotics since many years [6], [14], [20]. A very relevant aspect in SLAM concerns the representation of the entries in the world model and the management of their uncertainty; improper uncertainty management induces errors in robot localization and world mapping, which therefore suffers of geometric inconsistencies. These prevent practical use of mobile robotics technology whenever an a priori and reliable map is not available.

Many robot activities requires a full 3D knowledge of the observed environment features; a few examples are: motion, which is constrained by table legs and steps; cleaning, which has to be performed also under tables and chairs while fire extinguishers, hanging off-walls, has to be avoided; books to be moved, which are on top of tables, etc. (see Figure 1). Most of these items are not perceivable with the ubiquitous 2D laser range finders (LRF). It is therefore relevant to map the full 3D robot workspace, but this has been not so common up to now for SLAM systems. Most of the works dealing with 3D data bases on 3D LRFs (e.g., [21]); these devices provide clouds of 3D points, and this makes difficult pursuing other robot tasks like, e.g. the semantic classification of places [18], which are required for a real indoor service robot.

In this paper, our main objective is to provide a general framework for 3D sensor fusion, for vision based SLAM, that takes into account uncertainty in projective geometry and provides a mean for seamless integration of several information sources, e.g. 3D line segments, 3D planes, clouds of 3D points, etc. By looking at vision as a main source of information we naturally come across the issue of sensor fusion, as it is possible to build several *logical sensors* (e.g, line segments, corners, affine-covariant regions, etc.) on top of the same physical device [7, 8]. Each sensor provides a noticeably different level of robustness and accuracy. It is therefore of uttermost importance to be able to integrate, i.e., to associate and fuse, data provided by different logical sensors to get the most for precision and robustness.

Even though we are here proposing to use just the geometric information provided by 3D vision systems, we think that the full richness of the output of vision systems is necessary for other tasks. Moreover, as pointed out by C. Angle, an indoor robot with realistic sale expectations cannot base on a costly sensing suite; a, perhaps provoking, estimate of 10US$ cost for a complex robot sensor was reported in [1]. 3D vision-based sensing includes the capabilities obtainable from costly (3D or 2D) LRFs. We therefore think we need to deal with robotics tasks like SLAM with 3D data from vision. Nowadays, stereo vision is quite reliable and cheap both for the cost and for the power consumption (not a secondary issue for autonomous robots) and integrated devices are already available on the market [5, 17].

II. Uncertain Projective Geometry

Uncertain projective geometry is a framework used to represent the geometric entities and their relationship intro-

Fig. 1. Robots without 3D perception cannot clean under the table (left) neither avoid bumping into the open window (right).

duced by Heuel [9, 10]. This framework is able to describe, combine, and estimate various types of geometric elements (2D and 3D points, 2D and 3D lines and 3D planes) taking their uncertainty into account. These elements are represented using homogeneous vectors, allowing to derive simple bilinear expressions to represent join and intersection operators. This is obtained using only three matrices (construction matrix): $\mathbf{S}(\cdot)$ (for 2D points and 2D lines), $\mathbf{O}(\cdot)$ (for 3D lines) and $\mathbf{\Pi}(\cdot)$ (for 3D points and 3D planes). To get a line from two 2D points we can use the operator:

$$\mathbf{l} = \mathbf{x} \wedge \mathbf{y} = \mathbf{S}(\mathbf{x})\mathbf{y} \tag{1}$$

$$\mathbf{S}(\mathbf{x}) = \frac{\partial \mathbf{x} \wedge \partial \mathbf{y}}{\partial \mathbf{y}} = \begin{pmatrix} 0 & -x_3 & x_2 \\ x_3 & 0 & -x_1 \\ -x_2 & x_1 & 0 \end{pmatrix}, \tag{2}$$

the same hold to join two 3D points into a 3D line:

$$\mathbf{L} = \mathbf{X} \wedge \mathbf{Y} = \mathbf{\Pi}(\mathbf{X})\mathbf{Y}, \tag{3}$$

$$\mathbf{\Pi}(\mathbf{X}) = \frac{\partial \mathbf{X} \wedge \partial \mathbf{Y}}{\partial \mathbf{Y}} = \begin{pmatrix} W_1 & 0 & 0 & X_1 \\ 0 & W_1 & 0 & -Y_1 \\ 0 & 0 & W_1 & -Z_1 \\ 0 & -Z_1 & Y_1 & 0 \\ Z_1 & 0 & -X_1 & 0 \\ -Y_1 & X_1 & 0 & 0 \end{pmatrix}. \tag{4}$$

Again we can join a 3D point with a 3D line into a 3D plane:

$$\mathbf{A} = \mathbf{X} \wedge \mathbf{L} = \mathbf{O}(\mathbf{L})\mathbf{X}, \tag{5}$$

$$\mathbf{O}(\mathbf{L}) = \frac{\partial \mathbf{X} \wedge \partial \mathbf{L}}{\partial \mathbf{X}} = \begin{pmatrix} 0 & L_3 & -L_2 & -L_4 \\ -L_3 & 0 & L_1 & -L_5 \\ L_2 & -L_1 & 0 & -L_6 \\ L_4 & L_5 & L_6 & 0 \end{pmatrix}. \tag{6}$$

These construction matrices are useful tools to derive new geometric entities from other ones, e.g., a 3D line from two 3D points, a 3D point from the intersection of two 3D lines, etc.; at the same time, being bilinear equations these operators represent themself the Jacobian of the transformation which is used for the uncertainty propagation in the construction process.

Finally, these matrices can be used to express various geometric relations between pair of elements: incidence, identity, parallelism and orthogonality. Using these relations we can generate probabilistic tests to establish relationships between entities and formulate a simple estimation process, for fitting an unknown entity β to a set of observations $\tilde{\mathbf{y}}$ constrained by a set of relationship $\mathbf{w}(\tilde{\mathbf{y}}, \beta) = \{w(\tilde{\mathbf{y}}, \beta)\}$.

Suppose we have a set of observations, described by equation:

$$\tilde{\mathbf{y}}_\mathbf{i} = \mathbf{y}_\mathbf{i} + \mathbf{e}_\mathbf{i}, \tag{7}$$

where $\mathbf{e}_\mathbf{i} \sim N(\mathbf{0}, \mathbf{Q})$, to estimate the unknown entity it is possible to use an iterative algorithm in two steps:

1) Estimate the unknown entity using the relationship between the unknown and the observations $\mathbf{w}(\tilde{\mathbf{y}}, \beta) = \mathbf{0}$ and the homogeneus constraint $\mathbf{h}(\beta) = \mathbf{0}$. This can be obtained minimizing

$$\begin{aligned} &\mathbf{\Theta}(\tilde{\mathbf{y}}, \beta, \lambda, \mu) = \\ &= \tfrac{\mathbf{1}}{\mathbf{2}}(\mathbf{y} - \tilde{\mathbf{y}})^\mathbf{T}\mathbf{Q}_\mathbf{y}^{-\mathbf{1}}(\mathbf{y} - \tilde{\mathbf{y}}) + \lambda^\mathbf{T}\mathbf{w}(\tilde{\mathbf{y}}, \beta) + \mu^\mathbf{T}\mathbf{h}(\beta), \end{aligned} \tag{8}$$

where λ and μ are Lagrangian multipliers, respectively for the relationships among the entities and the homogenity constraints.

2) Re-evaluate the constraints on the observations by another iterative process, updating the observations by the use of the relationship, the homogeneus constraint and the new entity estimated. This secon step is done to propagate the new information to the entities in the relationshps as well.

Being our operators bilinear, we can estimate a new entity z, from two entities x and y, with a simple matrix multiplication:

$$z = f(x, y) = U(x)y = V(y)x, \tag{9}$$

where $U(x)$ and $V(y)$ are, at the same time, the bilinear operators and the Jacobian of the x and y entity respectively.

Assuming the entities to be uncertain, the pairs (x, Σ_{xx}), and (y, Σ_{yy}), and, possibly, the covariances Σ_{xy} between x and y are required for computing the error propagation:

$$\begin{aligned} &(z, \Sigma_{zz}) = \\ &\left(U(x)y, [V(y), U(x)] \begin{pmatrix} \Sigma_{xx} & \Sigma_{xy} \\ \Sigma_{xy} & \Sigma_{yy} \end{pmatrix} \begin{bmatrix} V^T(y) \\ U^T(x) \end{bmatrix} \right). \end{aligned} \tag{10}$$

In case of independence between x and y one obtains:

$$(z, \Sigma_{zz}) = \left(U(x)y, U(x)\Sigma_{yy}U^T(x) + V(y)\Sigma_{xx}V^T(y)\right) \tag{11}$$

To check the geometric relationship between two geometric entities it is possible to use a statistical test on the distance vector d defined through the previous bilinear equation. In particular a relation can be assumed to hold if the hypothesis

$$H0 : d = U(x)y = V(y)x = 0, \tag{12}$$

cannot be rejected. Notice that the hypothesis H_0 can be rejected with a significance level of α if

$$T = d^T\Sigma_{dd}^{-1}d > \varepsilon_H = \chi^2_{1-\alpha;n}. \tag{13}$$

The covariance matrix Σ_{dd} of d is given by first order error propagation as

$$\Sigma_{dd} = U(x)\Sigma_{yy}U^T(x) + V(y)\Sigma_{xx}V^T(y).$$

In general Σ_{dd} may be singular, if d is a n x 1 vector, r is is the degree of freedom of the relation R and $r < n$. The singularity causes a problem, as we have to invert the covariance matrix, but, at least for projective relationshps, it can be guaranteed that the rank of Σ_{dd} is not less than r (see Heuel [9, 10]).

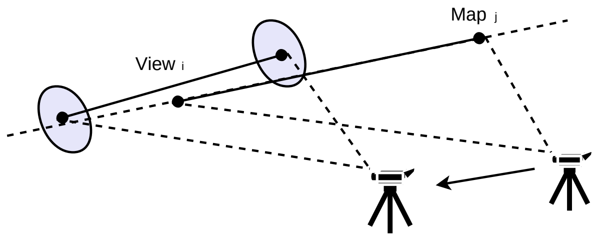

Fig. 2. The “moving window” problem.

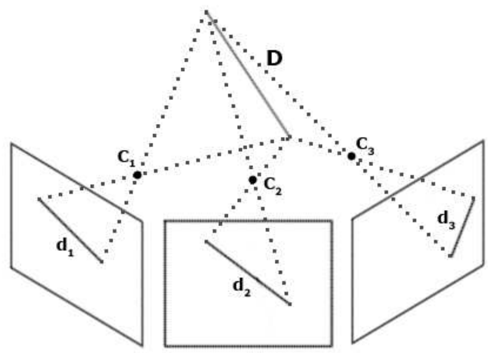

Fig. 3. 3D segment-based reconstruction for a trinocular stereoscopic system.

III. 6DoF Visual SLAM

Our interest in the framework described in the previous section comes from the issue of integrating 3D points, sensed by a stereo camera like the one in [12], with the 3D segments used in a previously developed algorithm for 6DoF hierarchical SLAM, sensed basing on trinocular stereo vision. The algorithm uses hierarchical map decomposition, uncertainty modeling for trinocular 3D data, and 6DoF pose representation.

In [16] we discussed in details some algorithms for data association and the importance of using a proper criterion to match features in the view with features in the map. Usually the point-to-point distance is considered as an appropriate criterion for single segment matching and much of the effort is devoted in finding a good association strategy for dealing with the exponential complexity of finding the best match for the whole view. In that paper we showed how a better criterion for 3D segment matching results in a better data association almost independently from the algorithm for interpretation tree traversal (i.e., data association algorithm).

The approach we proposed is based on a multi-criteria evaluation, for associating segments in the view with map segments. The reason for discarding the point-to-point criterion is mainly due to the problem of the moving-field-of-view in the sensing system, which turns in a moving window on the world feature(s), see Figure 2. More precisely, the segment extrema are induced by the reduced field of view and are not always related to real extrema in the world; when the sensing system moves it senses new extrema, which could result in new segments, at each step; this can easily become a problem for the classical point-to-point distance[1].

By using the uncertain projective geometry framework we are able to extend the original system by integrating the 3D segments coming from the feature-based trinocular stereo with the 3D points detected by the correlation-based stereo camera. Our idea is to improve the original SLAM algorithm integrating segments and points into 3D segments (by using the math introduced in Section II), before introducing this information in the EKF-based state filter. In this way we can reuse the original filter, since we still base on segments to perform the SLAM. Moreover, being uncertain projective geometry a probabilistic framework, we can take into account the different uncertainties in percepts so to have a consistent estimate of the measurement uncertainties in the filter. In the following, we describe the details of the representation used for the sensed data and the mechanism used for their integration.

[1]Our proposal is of interest also for 2D-3DoF SLAM systems which groups 2D data points into 2D lines, because this moving-field-of-view issue applies there too.

A. Segment-based stereo vision

This perception channel is a widespread and well known system, which reconstructs the scene in terms of 3D segments. In order to give out such data the system has to deal with segments since the very first (image) processing step, this is the intended meaning of the term segment-based 3D reconstruction system here used. The segments are represented by the 3D coordinates of their extrema. This choice is in agreement with the intrinsically 3D nature of indoor environments.

Our system bases on the trinocular approach [2]. As depicted in Figure 3, the processing starts with looking for 2D segments in the images, and then for correspondences between the different images. The last step is the computation of the parameters of the 3D segment, represented by the 3D coordinates of the endpoints. In Figure 3 $\mathbf{D}$ is the 3D scene segment, $\mathbf{C}_i$ and $\mathbf{d}_i$ are respectively the projection center and the projection of $\mathbf{D}$ on image i.

In the uncertain projective geometry framework, a 3D line is represented by a 6-coordinates vector $\mathbf{L}$ in a Plücker form:

$$\mathbf{L} = \begin{pmatrix} L_1 & L_2 & L_3 & L_4 & L_5 & L_6 \end{pmatrix}^T . \tag{14}$$

The end-points $(\mathbf{X_1}, \mathbf{X_2})$ are computed projecting lines for each 2D end-point and collecting the intersections with the estimated 3D line. The nearest intersections will be the end-points of the 3D segment represented by the touple:

$$\mathbf{Segm} = \langle \mathbf{L}, \mathbf{X_1}, \mathbf{X_2} \rangle \tag{15}$$

where:

$$\mathbf{X_i} = \begin{pmatrix} wX & wY & wZ & w \end{pmatrix}^T \tag{16}$$

are the segment extrema with covariance matrix:

$$\mathbf{\Sigma_i} = \begin{pmatrix} \sigma_{XX} & \sigma_{XY} & \sigma_{XZ} & 0 \\ \sigma_{YX} & \sigma_{YY} & \sigma_{YZ} & 0 \\ \sigma_{ZX} & \sigma_{ZY} & \sigma_{ZZ} & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} . \tag{17}$$

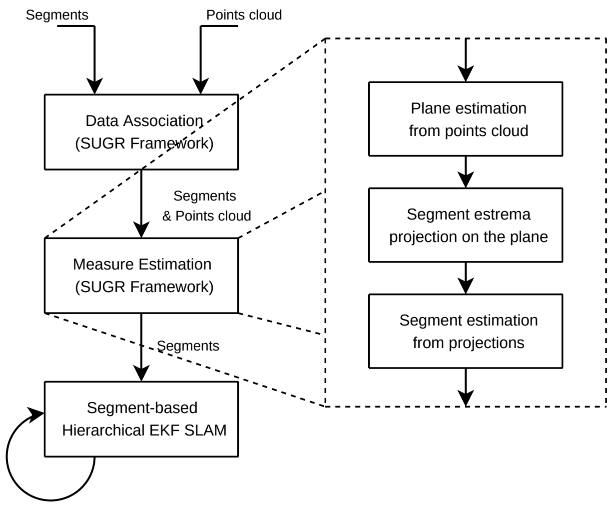

Fig. 4. The sensor fusion mechanism

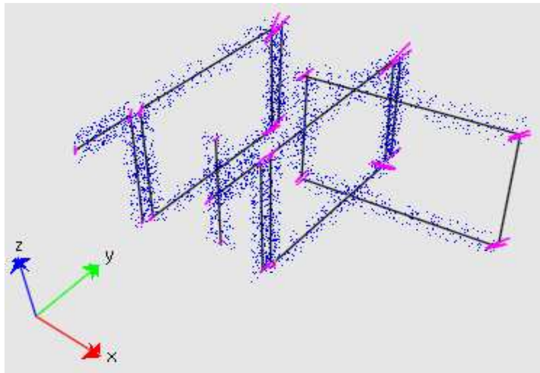

Fig. 5. Segments and points simulated for a single view. Notice, at the edges of segments, the uncertainty ellipses as computed by the real trinocular system

This system has been implemented to provide the 3D extrema as mean and covariance, i.e., approximating the non-linear transformation by a Jacobian-based uncertainty propagation of the Gaussian noise in pixel detection and projection parameters. In other words, cameras calibration provides also the covariance matrix of the parameters, so that 3D extrema can be given out by the system altogether with a covariance matrix, to represent the measurement uncertainty as a first order, i.e., normal, probability distribution.

Such systems date a long time ago and are quite widespread in the computer vision and robotics communities. Our implementation differs from the original only in the use of the Fast Line Finder algorithm [11], in the polygonal approximation phase.

B. Correlation-based stereo vision

The correlation-based system computes matches between local areas of the two images by evaluating the similarity of the regions. Each small area in the first image is correlated with other areas in the second image. The maximum correlation value, for each pixel, is computed and a disparity image is generated, which permits to obtain 3D points.

This method produces quite dense results (for each pixel in the first image we have a point in the 3D space, if the corresponding pixel in the second image is found). In this way we obtain a very large number of 3D points, for each activation, which are represented by their three coordinates in the 3D space (see eq.16 and 17).

C. Sensor fusion mechanism

Sensor fusion is performed before each segment measure is passed to the SLAM algorithm thus we say it is made “out of the filter”. Following the flow chart in Figure 4, we perform hypothesis test to assign a sets of points to each segment by using the Statistically Uncertain Geometric Reasoning (SUGR) framework of Section II. In this data association phase, we identify three sets of points, for each segment, using three hypothesis tests: two tests, one for each extrema, are devoted to checking if there are points coinciding with the segment extrema; the third test aims at searching points incident to the line passing through the two extrema. It is important to notice that the framework proposed allows us to estimate, in a simple way, the 3D line passing trough the two extrema, with its uncertainty. It is therefore possible to identify points incident to the line by taking into account also the uncertainty in line estimation.

Having performed such tests, we are now able to integrate each set of points with the corresponding segment updating both the position and the uncertainty of its extrema. Also this activity is performed “outside the filter” and aims at generating a new measure for the perceived segments. To have a more robust segment estimation, we decided to perform the integration by following the three steps procedure outlined herefter.

This is mainly due to the presence of points that satisfy the test because of their large uncertainty; they usually belong to the plane incident with the segment, but not necessary to the segment itself. For each segment we:

1) estimate a 3D plane incident to each point in the subset matching the line incidence hypothesis test
2) estimate the new extrema of the segment, estimating the two points incident to the plane and equal to the old extrema, i.e., the point sets that passed the first two tests
3) estimate the new segment by using these two projected extrema

It now is possible to pass the new segments estimated in such a way as new improved measures for the EKF in the segment-based Hierachical SLAM algorithm described in details in [15, 16].

IV. Experimental Validation

In this section we show the capabilities of the framework presented in this paper, using a simulator for both the feature-based and the correlation-based stereo cameras. Given a map of segments and the robot trajectory, we simulated the image formation process on the two devices, as well as the uncertain measurements of the world, as depicted in Figure 5.

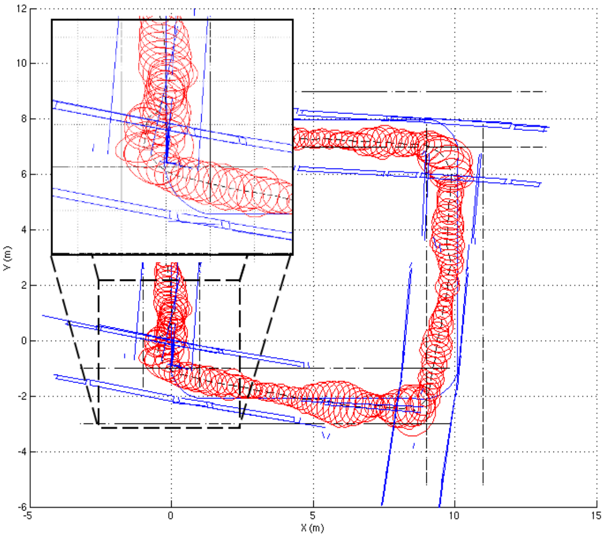

(a) Comparison with ground truth for trinocular

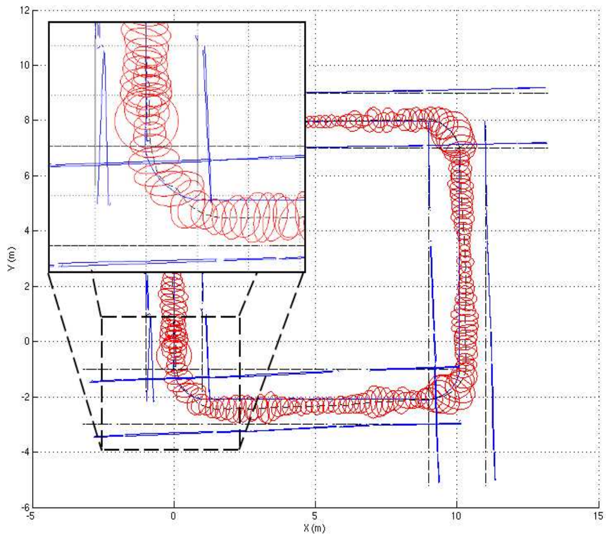

(b) Comparison with ground truth for sensor fusion

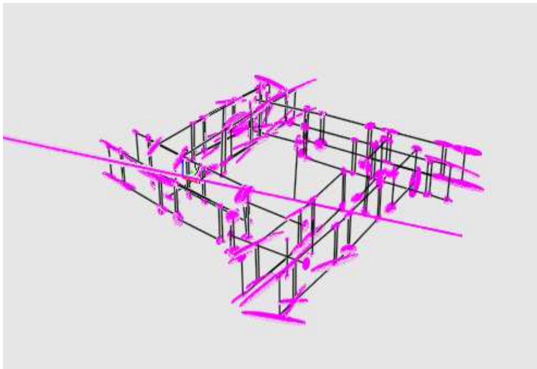

(c) 3D map from trinocular

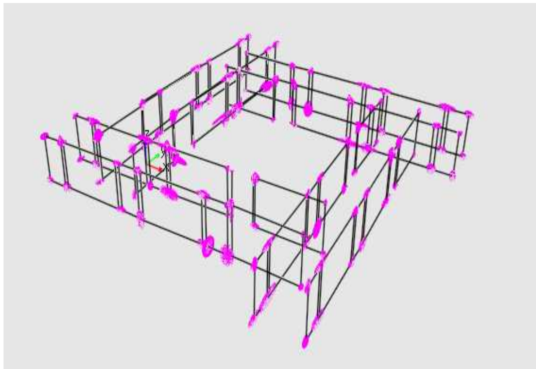

(d) 3D map from sensor fusion

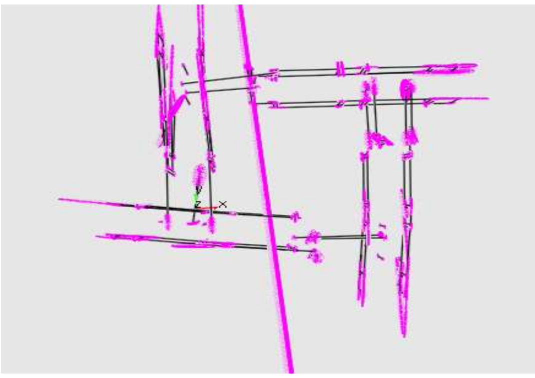

(e) Top view 3D map from trinocular

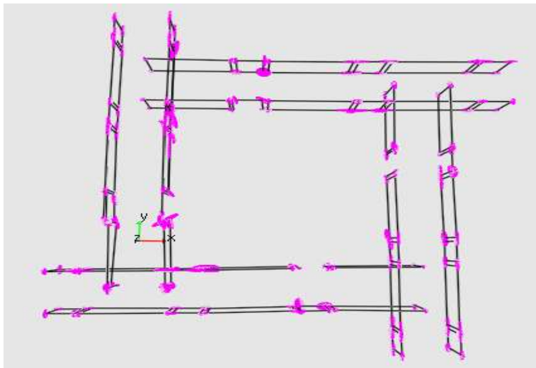

(f) Top view 3D map from sensor fusion

Fig. 6. Comparison between the same SLAM system, which takes in input different data: on the left, the input is the trinocular data; on the right, the input is the data fused by using the uncertain geometry framework proposed.

We have two reasons for using a simulated environment to test the proposed method. First, with a simulation environment we have access to the ground truth and we can perform numerical comparison about the consistence of the EKF. Second, we are not able to use the dataset collected for the experimental work in [15] and [16] because we did not have the SVS (correlation-based) camera at that time.

The two algorithm, i.e., the one with and the one without sensor fusion, have been compared on the very same data. In Figure 6 it is possible to compare the results of the approach proposed in this paper with our previous work. The plots refer to a circular trajectory of 40m and the SLAM algorithm has been stopped before loop closure. In Figures 6(a) and 6(b), we plotted the $\pm 3\sigma$ confidence ellipse around the estimated robot pose. It can be noticed how the estimate of robot pose with the trinocular based system eventually becomes inconsistent while using sensor fusion helps in maintaining it consistent.

Figures 6(c), 6(d), 6(e) and 6(f) allow mapping comparison. Notice that by using trinocular data and SVS points we are able to reduce the uncertainty in the segment extrema (ellipses in the plots) and to improve the overall accuracy.

V. Conclusion and Future Works

This paper presents a 6DoF SLAM algorithm that exploits the uncertain projective geometry framework to perform sensor fusion of 3D data coming from different sensors. The framework allows to represent 3D segments and 3D points with their associate uncertainty. Although the algorithm presented in this paper deal with vision data, it is straightforward to include other source of information in the framework. For instance, laser scans could be integrated either using single measures as points or extracting 3D lines or 3D planes, when dealing with 3D laser scans.

We faced sensor fusion by exploiting uncertain projective geometry, outside the SLAM EKF state filter, to provide a new “virtual” sensor. This had the purpose of reducing measurement errors and improving the results obtained with the SLAM algorithm. The described framework could be used within the SLAM algorithm as well, for data association, trough hypothesis test, and filter update. We are actually working on the latter by adopting a Smoothing and Mapping approach [4] to obtain independence between the features and thus independent fusion procedures for each of them.

In the last years many researchers have demonstrated that it is possible to use pictorial features from a monocular camera to perform SLAM as well. Notable examples are the works of Lacroix et al.[13], Davison et al.[3] and Lowe et al.[19]. The key idea, in this kind of approaches, is to use interest points as features, trying to fulfill the data association task by the use of a point descriptors or correlation based methods. Usually, to have significant landmarks, these descriptors are chosen to be invariant to image scale, rotation, and partially invariant (i.e. robust) to changing viewpoints, and illumination. From a SLAM point of view this is very useful not only for the data association but also because it allows the robot to identify each possible loop closure.

We are presently working on the integration of these pictorial feature in the framework proposed in this paper. By manipulating uncertain geometric entities, we are able to treat interest points as delimited portions of uncertain planes, enriched by an appearance information represented by the image patch located around the point on the image plane. Performing the data association among frames, we could estimate the most probable plane where the feature lies and, at the same time, update the appearance information depending on its position relative to the camera orientation.

Acknowledgment

This work has partially been supported by the European Commission, Sixth Framework Programme, Information Society Technologies: Contract Number FP6-045144 (RAWSEEDS).

References

[1] Colin Angle. Invited talk at the annual euron meeting. author is CEO of IS-Robotics; Euron: www.euron.org, March 2004.
[2] N. Ayache and F. Lustman. Trinocular stereo vision for robotics. *IEEE Trans. on PAMI*, 12(1), 1991.
[3] A. Davison. Real-Time Simultaneous Localisation and Mapping with a Single Camera. In *Proc. of ICCV*, pages 1403–1410, October 2003.
[4] F. Dellaert and M. Kaess. Square Root SAM: Simultaneous localization and mapping via square root information smoothing. *Intl. J. of Robotics Research*, 25(12), Dec 2006.
[5] Videre Design. `http://www.videredesign.com/`.
[6] H. F. Durrant-Whyte. *Integration, coordination and control of multi-sensor robot systems*. Kluwer Academic, 1987.
[7] T. Henderson, C. Hansen, and B. Bhanu. The specification of distributed sensing and control. *J. of Robotic Systems*, 2:387–396, 1985.
[8] T.C. Henderson and E. Shilcrat. *Logical sensor systems*. Ablex Publishing Corp. Norwood, NJ, USA, 1995.
[9] S. Heuel. Points, lines and planes and their optimal estimation. (2191):92–99, September 2001.
[10] S. Heuel. *Uncertain Projective Geometry: Statistical Reasoning for Polyhedral Object Reconstruction*. Springer, 2004.
[11] P. Kahn, L. Kitchen, and E. M. Riseman. A fast line finder for vision-guided robot navigation. *IEEE Trans. on PAMI*, 12(11), Nov 1990.
[12] Kurt Konolige. Small vision system. hardware and implementation. In *ISRR*, Hayama, Japan, 1997.
[13] T. Lemaire, S. Lacroix, and J. Sola. A practical 3d bearing-only slam algorithm. In *Proc. of IROS*, 2005.
[14] F. Lu and E. Milios. Globally consistent range scan alignment for environment mapping. *Autonomous Robots*, 4:333–349, 1997.
[15] D. Marzorati, D. G. Sorrenti, and M. Matteucci. 3d-6dof hierarchical slam with trinocular vision. In *Proc. of ECMR*, 2005.
[16] D. Marzorati, D. G. Sorrenti, and M. Matteucci. Multi-criteria data association in 3d-6dof hierarchical slam with 3d segments. In *Proc. of ASER06*, 2006.
[17] Point-gray. `http://www.ptgrey.com/products/bumblebee2/`.
[18] A. Rottmann, O. Martínez Mozos, C. Stachniss, and W. Burgard. Place classification of indoor environments with mobile robots using boosting. In *Proc. of AAAI*, 2005.
[19] S. Se, D. Lowe, and J. Little. Global localization using distinctive visual features. In *Proc. of IROS*, Lausanne, Switzerland, October 2002.
[20] J. D. Tardós and J. A. Castellanos. *Mobile robot localization and map building: a multisensor fusion approach*. Kluwer Academic, 1999.
[21] O. Wulf, B. Wagner, and M. Khalaf-Allah. Using 3d data for monte carlo localization in complex indoor environments. In *Proc. of ECMR*, 2005.