

Autonomous Single Camera Exploration

Teresa Vidal-Calleja¹, Alberto Sanfeliu¹ and Juan Andrade-Cetto²

¹Institut de Robòtica i Informàtica Industrial, UPC-CSIC. Llorens Artigas 4-6, Barcelona 08028, Spain.

²Centre de Visió per Computador, UAB. Edifici O, Campus UAB, Bellaterra 08193, Spain.

Abstract— In this paper we present an active exploration strategy for a mobile robot navigating in 3D. The aim is to control a moving robot that autonomously builds a visual feature map while at the same time optimises its localisation in this map. The technique chooses the most appropriate commands maximising the information gain between prior states and measurements, while performing 6DOF bearing-only SLAM at video rate. Maximising the mutual information helps the vehicle avoid ill-conditioned measurements appropriate to bearing-only SLAM. To validate the approach, extensive simulations over rugged terrain have been performed. Moreover, experimental results are shown for the technique being tested with a synchro-drive mobile robot platform.

Index Terms—Bearing only SLAM, Exploration.

I. INTRODUCTION

Autonomous exploration is a process in which an observer can interact with its surroundings by moving about and collecting information in order to learn about the environment [24]. Within the mobile robotics context, much attention has been paid to the second part of this process, collecting information that is. The technique is known as simultaneous localisation and mapping (SLAM in short) [19], and over the past 20 years, SLAM systems have matured from producing indoor planar range-based maps [9–11], to 2D outdoor maps [21, 23], to 3D outdoor maps [3], to trinocular and stereovision based maps [6, 14], and more recently, to monocular (bearing-only) mapping [7, 8].

Less attention has been paid however to the first part of the problem, that of actively exploring while mapping. Noteworthy, and given the probabilistic nature of the Bayesian approach to the solution of the SLAM problem, entropy reduction has recently gained popularity as a map building strategy for driving a robot during a SLAM session in order to minimise uncertainty [1, 9, 15].

Given the real-time characteristics of the visual SLAM system we use, fast and efficient action evaluation is of utmost importance. Fortunately enough, the elements needed to validate the quality of actions with respect to

entropy reduction are readily available from the SLAM priors [4], and, by making enough implementation adaptations, we are able to evaluate in real time the value of a limited number of actions. The technique has already been tested for an unconstrained moving camera [22], and this communication presents the natural step forward, evaluating the technique during constrained motion.

Action evaluation with respect to information gain has already been implemented for other SLAM systems, but little to no effort has been expended on the real-time constraint. One such approach makes use of Rao-Blackwellized particle filters [20]. When using particle filters for exploration, only a very limited number of actions can be evaluated due to the complexity in computing the expected information gain. The main bottleneck is the generation of the expected measurements each action sequence would produce, which is generated by a ray-casting operation in the map of each particle. In contrast, measurement predictions in a feature-based EKF implementation can be computed much faster, having only one map posterior per action to evaluate, instead of the many a particle filter requires. Moreover, in [20] the cost of choosing a given action is subtracted from the expected information gain with a user selected weighting factor. A more theoretically sound approach is presented in this work, in which the cost of performing a given action is inherently taken into account when evaluating the entropy for the set of possible priors.

Sim has also addressed decision making for the robot exploration problem, as an optimisation problem for a restricted hand-crafted set of exploratory policies [16], as a sequential decision making problem (POMDPs) [18], and by updating an information surface in a SEIF implementation [17]. These contributions, however, only test the strategies for very small planar point-based simulated environments, and remain to be tested in real-world applications. In order to avoid local maxima, the approach presented in [17] explicitly avoids loop closing by discarding repeated poses during trajectory search. Our previous experience with real-time vision-based SLAM has shown us however, that short loop closing is essential for consistent bearings-only mapping. Moreover, [24] suggests that a gradient strategy for uncertainty reduction would not fal-

This research is supported by the Spanish Ministry of Education and Science by an FPI scholarship to TVC, under project DPI 2004-5414 to AS, and by project TIC 2003-09291 and the EU PACO-PLUS project FP6-2004-IST-4-27657 to JAC.

ter on top of a local maximum. The reason being that the information surface being ascended is continuously changing as new data are added. Maximally informative posteriors come from locations with large variance, and when measurements iterate over the same states, the prediction variance will be reduced to the level of sensor noise, flattening the information surface with the effect of “pushing” the robot away from that location. Consequently, in this work, we choose to concentrate in a greedy real-time steepest descent approach to entropy reduction for a monocular SLAM system, rather than on planning for large sequences of actions.

Other approaches include, for example, a multirobot stereo-vision occupancy grid-based SLAM system [13], with best single-step look ahead chosen on the basis of overall map entropy reduction. In such a discrete representation of the map posterior, overall map entropy is computed as the sum of individual entropies for each grid cell. Bryson *et al.* on the other hand, present simulated results of the effect different vehicle actions have with respect to the entropic mutual information gain [2]. The analysis is performed for a 6DOF aerial vehicle equipped with one camera and an inertial sensor, for which landmark range, azimuth, and elevation readings are simulated, and data association is known.

In this paper, we have opted for a strategy that chooses those actions that maximise the mutual information between states and measurements. Notice that maximising an information criterion might result in uncertain actions being chosen, since their reduction of uncertainty once a measurement has taken place would be larger. Other reported approaches maximise present to future posterior entropy differences instead. With our chosen strategy overall entropy decay may happen at a lower pace, at the expense of actually choosing exploratory actions instead of homeostatic ones.

The rest of the paper is distributed as follows. Section II presents a brief overview of the vision-based bearing only SLAM system we use. Section III is devoted to a discussion from first principles on the value of expected measurements in reducing overall state entropy. This gives rise to the actual action selection policy used, which is described in detail in Section IV. Section V presents an evaluation of our exploration strategy for a 3D simulated environment; and Section VI contains actual experimental results. Finally, some concluding remarks are given in Section VII.

II. EKF 6DOF BEARING-ONLY SLAM

The SLAM problem is usually formulated as the probabilistic estimation of a multivariate state, containing the pose of a moving platform $\mathbf{x}_{v,k}$ (be it a robot, a wearable

device, an UAV, etc), as well as all learned feature location estimates \mathbf{y} . The objective is to compute the posterior $p(\mathbf{x}_{v,k}, \mathbf{y} | U^k, Z^k)$ conditioned on the history of motion commands U^k and feature measurements Z^k ; that with a Kalman filter is approximated as a Gaussian distribution with mean $\mathbf{x}_{k|k}$ and covariance $\mathbf{P}_{k|k}$.

A. Unconstrained Camera Motion

Considering initially that our sensor is a camera, and that it is free to move in any direction in $\mathbb{R}^3 \times SO(3)$, we adopt the same smooth unconstrained constant velocity motion model as in [22],

$$\mathbf{x}_{v,k+1|k} = \begin{bmatrix} \mathbf{p}_{k+1|k} \\ \mathbf{q}_{k+1|k} \\ \mathbf{v}_{k+1|k} \\ \boldsymbol{\omega}_{k+1|k} \end{bmatrix} = \begin{bmatrix} \mathbf{p}_{k|k} + (\mathbf{v}_{k|k} + \mathbf{a}_k \Delta t) \Delta t \\ \mathcal{Q} \mathbf{q}_{k|k} \\ \mathbf{v}_{k|k} + \mathbf{a}_k \Delta t \\ \boldsymbol{\omega}_{k|k} + \boldsymbol{\alpha}_k \Delta t \end{bmatrix}. \quad (1)$$

Suffice to say that $\mathbf{p} = [x, y, z]^\top$ and $\mathbf{q} = [q_0, q_1, q_2, q_3]^\top$ denote the camera pose (three states for position and four for orientation using a unit norm quaternion representation), and $\mathbf{v} = [v_x, v_y, v_z]^\top$ and $\boldsymbol{\omega} = [\omega_x, \omega_y, \omega_z]^\top$ denote the linear and angular velocities, respectively, corrupted by zero mean normally distributed linear and angular accelerations $\mathbf{a} = [a_x, a_y, a_z]^\top$, and $\boldsymbol{\alpha} = [\alpha_x, \alpha_y, \alpha_z]^\top$. The quaternion transition matrix is

$$\mathcal{Q} = \cos\left(\frac{\Delta t \|\boldsymbol{\Omega}\|}{2}\right) \mathbf{I} + \frac{2}{\|\boldsymbol{\Omega}\|} \sin\left(\frac{\Delta t \|\boldsymbol{\Omega}\|}{2}\right) \boldsymbol{\Omega}_\times$$

with $\boldsymbol{\Omega} = [0, \omega_x, \omega_y, \omega_z]^\top$ the angular velocity vector expressed in quaternion form, and $\boldsymbol{\Omega}_\times$ its skew-symmetric matrix representation.

B. Constrained Camera Motion

It is assumed, however, that such camera is attached to a mobile robot navigating in a 3D terrain. The mobile robot is controlled by linear and angular velocities $\mathbf{u} = [\mathbf{v}_r, \omega_r]^\top$ which are tangent to the terrain surface. In simulating the robot motion taking into account surface contact at all times, we can substitute the previous motion prediction model with a constrained model for the continuous transition of the optic centre of the camera

$$\begin{bmatrix} \mathbf{p}_{k+1|k} \\ \boldsymbol{\theta}_{k+1|k} \end{bmatrix} = \begin{bmatrix} \mathbf{p}_{k|k} \\ \boldsymbol{\theta}_{k|k} \end{bmatrix} + \boldsymbol{\Gamma} \mathbf{u}_k \Delta t, \quad (2)$$

where

$$\boldsymbol{\Gamma} = \begin{bmatrix} -\cos \psi \sin \theta & -l \cos \psi \cos \theta \cos \phi \\ \sin \psi & -l \cos \psi \cos \theta \sin \phi \\ \cos \psi \cos \theta & -l \cos \psi \sin \theta \\ \cos \phi & -\sin \psi \cos \phi \\ \sin \phi & \cos \psi \cos \phi \\ 0 & -\sin \psi \end{bmatrix},$$

$\boldsymbol{\theta} = [\psi, \theta, \phi]^\top$ is a *yaw, pitch, roll* representation of \mathbf{q} , and l is the distance between the axle centre of the mobile robot and the camera optic centre.

C. Measurement Model

Our 6DOF Single Camera SLAM system extracts salient point features from images, building a map of their 3D coordinates. Image projection priors are estimated with a full perspective wide angle camera

$$\begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} u_0 - u_c/\sqrt{d} \\ v_0 - v_c/\sqrt{d} \end{bmatrix}$$

where the position of a 3D map point is first transformed into the camera frame $\mathbf{y}_i^c = \mathcal{R}(\mathbf{y}_i - \mathbf{p})$, with \mathcal{R} the rotation matrix equivalent of \mathbf{q} , and $u_c = fk_u x^c/z^c$, $v_c = fk_v y^c/z^c$. The radial distortion term is $d = 1 + K_d(u_c^2 + v_c^2)$, and the intrinsic calibration of the camera is known — focal distance f , principal point (u_0, v_0) , pixel densities k_u and k_v , and radial distortion parameter K_d .

These priors are then compared against actual measurements using a nearest neighbour test within a 3σ elliptical search region inside the innovation covariance \mathbf{S}_i for each image estimate (see [22] for details).

New features are initialised using the approach presented by Davison in [7].

III. INFORMATION GAIN

The exploration strategy proposed in this paper is aimed specifically at maximising the mutual information between the state and consequent measurement priors, both resulting from an action in the form of a motion command. Different commands give rise to better or worse priors (in an entropic sense), and we want to select, from a limited test set, the one that produces the most expected reduction in entropy for the entire state, once the consequent measurement has taken place.

The mutual information for these two continuous probability density functions is defined as [4]

$$I(X; Z) = E \left[\log \frac{p(\mathbf{x}|\mathbf{z})}{p(\mathbf{x})} \right]. \quad (3)$$

For our Gaussian Multivariate case, the prior distribution is simply $p(\mathbf{x}) = N(\mathbf{x}_{k+1|k}, \mathbf{P}_{k+1|k})$, whereas, the conditional is given by the Kalman posterior $p(\mathbf{x}|\mathbf{z}) = N(\mathbf{x}_{k+1|k+1}, \mathbf{P}_{k+1|k+1})$ with the updates

$$\mathbf{x}_{k+1|k+1} = \mathbf{x}_{k+1|k} + \mathbf{P}_{k+1|k} \mathbf{H}^\top \mathbf{S}^{-1} (\mathbf{z}_{k+1} - \mathbf{z}_{k+1|k}) \quad (4)$$

$$\mathbf{P}_{k+1|k+1} = \mathbf{P}_{k+1|k} - \mathbf{P}_{k+1|k} \mathbf{H}^\top \mathbf{S}^{-1} \mathbf{H} \mathbf{P}_{k+1|k} \quad (5)$$

Substituting Eqs. 4 and 5 in Eq. 3, and taking the expectation, the Mutual Information between our state and measurement priors evaluates to the difference between prior and posterior state entropies

$$I(X; Z) = \frac{1}{2} (\log |\mathbf{P}_{k+1|k}| - \log |\mathbf{P}_{k+1|k+1}|).$$

In other words, maximising the mutual information between the state and measurement priors we end up choosing the motion command that most reduces the uncertainty in the state due to the knowledge of the consequent measurement as a result of a particular action.

IV. CONTROL STRATEGY

In this section we present the guidance strategy for our mobile robot performing SLAM with a single wide-angle camera. The control scheme is based on computing the instant robot accelerations that maximise mutual state and measurement information gain. Actions in the form of impulse accelerations guarantee smooth platform velocity change. The chosen command is then integrated to produce the input velocity that is sent to the robot. Given the real-time limitations of our system, only a limited number of actions can be evaluated at each step. These are the discrete set from Table I.

TABLE I
ACTION SET

Action	Linear Acceleration	Angular Acceleration
0	0	0
1	0	$-\dot{\omega}_r$
2	0	$\dot{\omega}_r$
3	$-\dot{v}_r$	0
4	\dot{v}_r	0
5	$-\dot{v}_r$	$-\dot{\omega}_r$
6	\dot{v}_r	$\dot{\omega}_r$

To compare the actions, the motion model from Eq. (2) is used to predict the prior mean $\mathbf{x}_{k+1|k}$ for each instant acceleration in the set, propagating the covariances by computing the corresponding Jacobians. Map features priors are also used to simulate the expected observations using the camera measurement model and the state prior.

The posterior covariance is then computed taking into account only known features inside the camera field of view.

At each time step we compute, in turn, the mutual information for one action in the set, using the prior and posterior covariance matrices. That is, for every linear and angular instant acceleration combination. Every 15th cycle, once all possible actions have been evaluated for a lapse of at least 8 cycles, the action that maximises the mutual information is chosen, and a new velocity input is sent to the system.

It is assumed a fixed number of expected features will be found within a 3D unexplored room. During the action selection process, the unknown features are taken into account in the covariance matrix initialized with large uncertainty.

V. SIMULATIONS

Extensive simulations have been performed using the constrained motion model for the mobile robot from Eq. (2), navigating in uneven 3D terrain, and using a full perspective wide angle camera model as sensing device. Unfortunately, this model is too restrictive due to the planar approximation of the terrain when computing priors. For this reason, it is only used to transform from linear and angular robot velocities to Cartesian velocities. The actual estimation of the robot pose and velocities is performed however, with the unconstrained motion model.

The aim is to choose impulsive acceleration commands for the mobile robot in order to explore the whole room while trying to reduce most the uncertainty. Accelerations are applied only every 15th step, and in between action decision, null acceleration is set, i.e. constant velocity behaviour is chosen until a new action is decided.

The control action is chosen from the discrete set of instant linear and angular accelerations shown in the Table I. The values for \dot{v}_r and $\dot{\omega}_r$ that produced the results shown in this section are $0.5m/s^2$ and $0.3rad/s^2$ respectively. The simulated environment shown contains 25 unknown features and 6 known features uniformly distributed in the room. Our simulated wheeled mobile robot is navigating over a 3D sinusoidal surface.

Figure 1(a) shows the trajectory followed by the vehicle and the initialised features with their uncertainty plotted as 2σ level hyperellipsoids. The expected covariance matrix is extended with the unknown feature uncertainties with diagonal values of $5m^2$ each to avoid homeostasis. Entropy reduction is computed using the extended covariance. The instant at which new features are added to the state are shown in Figure 1(b). Moreover, state estimation errors are shown in Figure 2 for the camera pose. Notice how when the terrain abruptly changes, the estimated velocities become underestimated in the direction the ter-

rain changed. Thus in simulating vehicle motion, a more elaborate model taking into account surface discontinuities ought to be considered for very rough terrains.

The selected actions reduce the camera pose and velocity uncertainty first, tracking features with low uncertainty. After that, the variance for unvisited features with large uncertainty is reduced as new features are added. Interestingly enough, the system autonomously explores by repeatedly choosing a negative linear acceleration. The effect is to augment the camera field of view with the consequent inclusion of new feature in the model, but still maintaining known features in sight, thus keeping the vehicle well localised at all times. In contrast to our previous experiments reported with a free-moving hand-held camera [22], it is more difficult in this constrained motion setting to actively perform short loop closure orthogonal to the field of view. The reason being that the robot cannot achieve saccadic motions in the way a free-moving camera can.

At this point we can argue how the same tracking (unconstrained constant velocity 6DOF motion model) and action selection strategies (maximising the mutual information between states and measurements) is capable of choosing different exploratory manouvers depending on the characteristics of the platform: short loop closing for a 6DOF free-moving camera, and backwards linear motion increasing the field of view for a mobile robot.

VI. EXPERIMENTS

Our main concern was to test the strategy during real-time vision-based SLAM execution. This Section is devoted to a discussion on such results. The experiments were conducted on the mobile platform shown in Fig. VII, with a wide-angle camera rigidly attach to the robot body, and for which an updated version of the single camera SLAM system reported in [5] was setup.

Within a room, the robot starts approximately at rest with some known object in view to act as a starting point and provide a metric scale to the proceedings. The robot moves, translating and rotating constrained by the 3D terrain, such that various parts of the unknown environment come into view. The aim is to estimate and control the 6DOF camera pose continuously, promptly and reliably during arbitrarily long periods of movement. This will involve accurately mapping (estimating the locations of) a sparse set of features in the environment.

The whole process is running at $15fps$. Since our mutual information measure requires evaluating the determinant of the full covariance matrix (enlarged with the unvisited features) at each iteration, single motion predic-

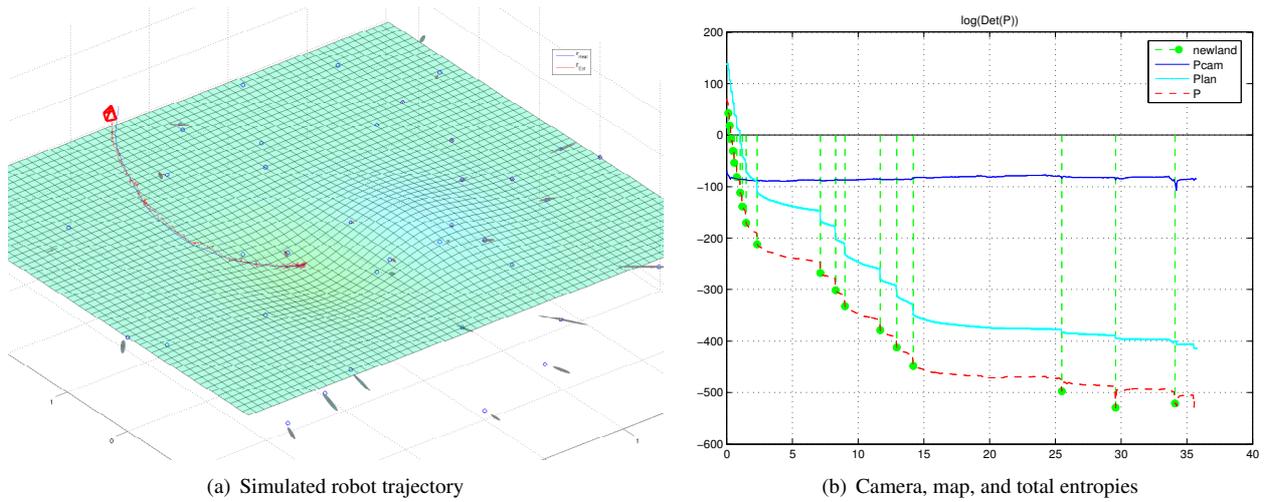


Fig. 1. Simulation of a mobile robot actively exploring a room. The mutual information maximisation strategy produces a nearly linear motion tangent to the surface. The vehicle starts at the shown terrain depression and proceeds backwards slightly rotating to increase map coverage. (\mathbf{r}_{Real} and \mathbf{r}_{Est} are the real and estimated vehicle trajectories, the label `newland` and the green dots and dotted vertical lines represent the value of entropy at the instant when new landmarks are initialised. `Pcam`, `Plan`, and `P` indicate the robot, map, and overall entropies.)

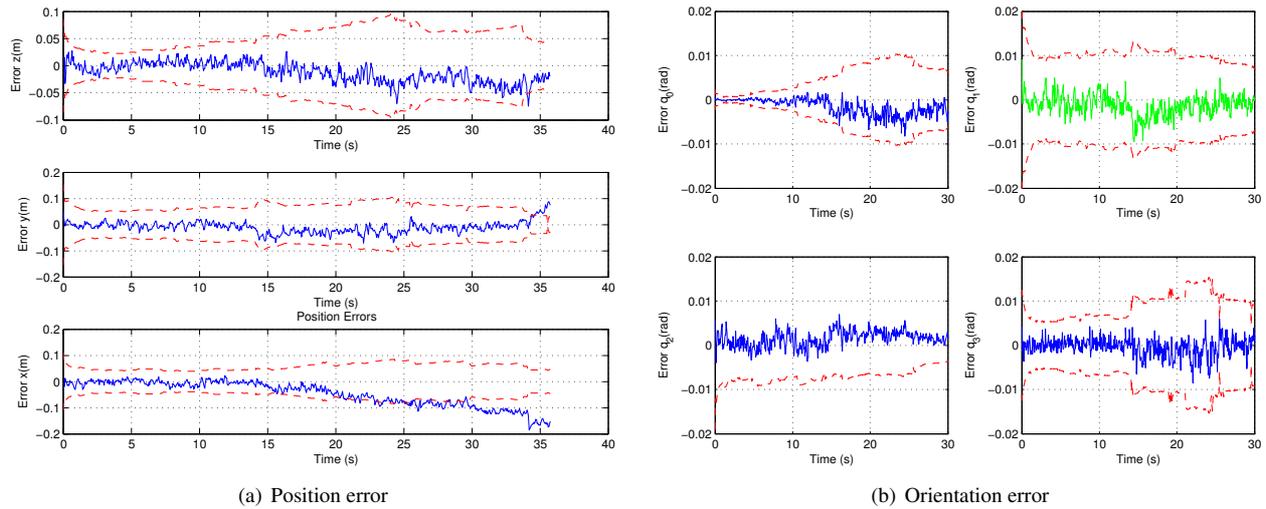


Fig. 2. Estimation errors for camera position and orientation and their corresponding 2σ variance bounds. Position errors are plotted as x, y, and z distances to the real camera location in meters, and orientation errors are plotted as quaternions.

tions are evaluated one frame at a time. It is only every 15th frame in the sequence that all mutual information measures are compared, and the best action is sent to the mobile robot. For the experiments, the acceleration magnitudes were set to $\dot{v}_r = 0.1m/s^2$ and $\dot{\omega}_r = 5deg/s^2$. When computing posteriors, these are all predicted for the duration that would take them to the end of the 15th frame, each action in turn being evaluated for a slightly shorter period of time. The motivation is that we want to be able to test actions in the basis of their effect at the very same point in time (at the end of the 15th frame). In order to evade any bias related to the time spent in evaluating the effect of actions, these are randomly ordered at each iter-

ation.

As with the simulated setting, the robot navigates in uneven terrain as shown in Figures VII and 4(a). In the plot, the estimated path (blue continuous line) is shown in 3D, as opposed to the vehicle odometry which is restricted to the Z-X plane. The orientation angle from Figure 4(b) indicates the vehicle orientation with respect to the world axis Y (orthogonal to the white sheet of paper placed in front of the robot, which serves as global reference consistent to the world XZ plane). Estimation in this case is similar to the measure provided by the encoders.

As in the simulated case, our mutual information-based action selection strategy for this constrained motion case



Fig. 3. The mobile robot platform used in the experiments.

autonomously explores the room driving the vehicle back and forth, but mostly backwards, enlarging the field of view by pulling away from the initial view.

Figure 4(b) gives account of the actions sent to the robot, and shows as most frequent actions iterations between positive and negative linear acceleration. The feature map and camera pose are updated and displayed in real-time in the graphical user interface. Figure 5 shows a sequence of frames from the same experiment, that show the robot driving away from the start known features.

VII. CONCLUSION

This paper has presented an autonomous information-driven exploration strategy for a wheeled mobile robot equipped with a single wide angle camera and navigating in uneven terrains. The approach is based in choosing the action that maximises the information gain between state and measurement priors. Simulation and experimental results consistently show a behaviour in which the robot pulls back from an initial configuration, by having the camera search for more features whilst reducing its own pose uncertainty.

The reported camera trajectories are simple because a) the robot is commanded by acceleration impulses that tend to drive the robot through smooth velocity changes, and b) the real-time constraints of the implementation allow only for the evaluation of a very limited set of possible actions. The computational complexity in computing entropy does not permit large maps, in that case submapping will be a good solution.

It is worth noting that no high-level task-dependent path planning is being performed whatsoever. The exploratory actions are chosen purely in the context of entropy minimisation. We foresee that planning under uncertainty while mapping requires moving ahead from the approach presented in this paper involving local action se-

lection, to longer term planning including task description. One approach to the problem we seek to explore is by planning in partially observable continuous domains via value iteration over POMDPs [12].

ACKNOWLEDGEMENTS

We are grateful to Andrew Davison for discussions and software collaboration.

BIBLIOGRAPHY

- [1] F. Bourgault, A. Makarenko, S. Williams, and B. Grocholsky. Information based adaptive robotic exploration. In *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, Lausanne, Oct. 2002.
- [2] M. Bryson and S. Sukkarieh. An information-theoretic approach to autonomous navigation and guidance of an uninhabited aerial vehicle in unknown environments. In *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, Edmonton, Aug. 2005.
- [3] D.M. Cole and P.M. Newman. 3D SLAM in outdoor environments. In *Proc. IEEE Int. Conf. Robot. Automat.*, Orlando, May 2006. To appear.
- [4] A. Davison. Active search for real-time vision. In *Proc. IEEE Int. Conf. Comput. Vision*, pages 66–73, Beijing, Oct. 2005.
- [5] A. Davison, Y. Gonzalez-Cid, and N. Kita. Real-time 3d SLAM with wide - angle vision. In *Proc. 5th IFAC/EURON Sym. Intell. Auton. Vehicles*, Lisbon, Jul. 2004.
- [6] A. J. Davison and D. W. Murray. Simultaneous localisation and map-building using active vision. *IEEE Trans. Pattern Anal. Machine Intell.*, 24(7):865–880, Jul. 2002.
- [7] A.J. Davison, W.W. Mayol, and D.W. Murray. Real-time localisation and mapping with wearable active vision. In *Proc. IEEE Int. Sym. Mixed and Augmented Reality*, Tokyo, Oct. 2003.
- [8] R.M. Eustice, H. Singh, and J.J. Leonard. Exactly sparse delayed-state filters. In *Proc. IEEE Int. Conf. Robot. Automat.*, pages 2428–2435, Barcelona, Apr. 2005.
- [9] H.J.S. Feder, J.J. Leonard, and C.M. Smith. Adaptive mobile robot navigation and mapping. *Int. J. Robot. Res.*, 18:650–668, 1999.
- [10] U. Frese, P. Larsson, and T. Duckett. A multigrid algorithm for simultaneous localization and mapping. *IEEE Trans. Robot.*, 21(2):1–12, 2005.
- [11] F. Lu and E. Milios. Globally consistent range scan alignment for environment mapping. *Auton. Robot.*, 4(4):333–349, 1997.
- [12] J.M. Porta, M.T.J. Spaan, and N. Vlassis. Robot planning in partially observable continuous domains. In *Robotics: Science and Systems I*, Cambridge, Jun. 2005.
- [13] R. Rocha, J. Dias, and A. Carvalho. Cooperative multi-robot systems: A study of vision-based 3-D mapping using information theory. *Robot. Auton. Syst.*, 53:282–311, 2005.
- [14] S. Se, D. Lowe, and J. Little. Mobile robot localization and mapping with uncertainty using scale-invariant visual landmarks. *Int. J. Robot. Res.*, 21(8):735–758, Aug. 2002.
- [15] R. Sim. Stable exploration for bearings-only SLAM. In *Proc. IEEE Int. Conf. Robot. Automat.*, pages 2422–2427, Barcelona, Apr. 2005.
- [16] R. Sim, G. Dudek, and N. Roy. Online control policy optimization for minimizing map uncertainty during exploration. In *Proc. IEEE Int. Conf. Robot. Automat.*, New Orleans, Apr. 2004.
- [17] R. Sim and N. Roy. Active exploration planning for SLAM using extended information filters. In *Proc. 20th Conf. Uncertainty in AI*, 2004.
- [18] R. Sim and N. Roy. Global A-optimal robot exploration in SLAM. In *Proc. IEEE Int. Conf. Robot. Automat.*, pages 673–678, Barcelona, Apr. 2005.
- [19] R. C. Smith and P. Cheeseman. On the representation and estimation of spatial uncertainty. *Int. J. Robot. Res.*, 5(4):56–68, 1986.
- [20] C. Stachniss, G. Grisetti, and W. Burgard. Information gain-based

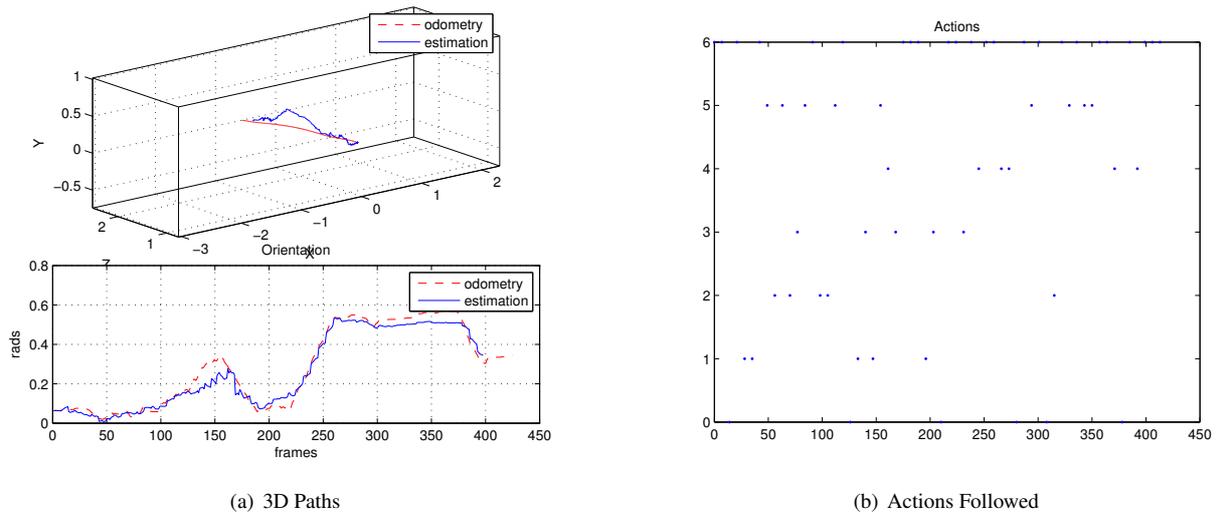


Fig. 4. Trajectories and Actions during the experiment.

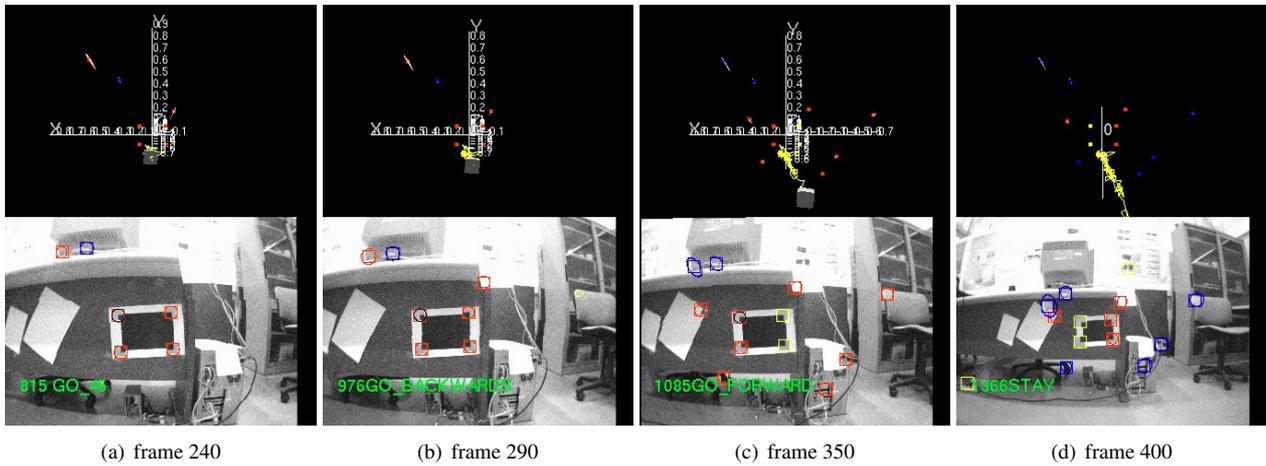


Fig. 5. Snapshots of the Graphical User Interface during autonomous exploration

- exploration using Rao-Blackwellized particle filters. In *Robotics: Science and Systems I*, pages 65–72, Cambridge, Jun. 2005.
- [21] S. Thrun, Y. Liu, D. Koller, A. Y. Ng, Z. Ghahramani, and H. Durrant-Whyte. Simultaneous localization and mapping with sparse extended information filters. *Int. J. Robot. Res.*, 23(7-8):693–716, Jul. 2004.
- [22] T. Vidal-Calleja, A.D. Davison, J. Andrade-Cetto, and D.W. Murray. Active control for single camera SLAM. In *Proc. IEEE Int. Conf. Robot. Automat.*, pages 1930–1936, Orlando, May 2006.

- [23] C-C. Wang, C. Thorpe, and S. Thrun. Online simultaneous localization and mapping with detection and tracking of moving objects: theory and results from a ground vehicle in crowded urban areas. In *Proc. IEEE Int. Conf. Robot. Automat.*, volume 1, pages 842–849, Taipei, Sep. 2003.
- [24] P. Whaite and F.P. Ferrie. Autonomous exploration: Driven by uncertainty. *IEEE Trans. Pattern Anal. Machine Intell.*, 19(3):193–205, Mar. 1997.