Learning Safe Navigation in Uncertain Environments

Maria João Rendas and Stefan Rolfes*

Laboratoire d'Informatique, Signaux et Systèmes de Sophia Antipolis (I3S)
2000 route des Lucioles Bat. Algorithmes/Euclide
F-06410 Biot, FRANCE

Abstract. We present a strategy for the navigation of mobile robots in open scarce uncertain environments. Our work assumes that the robot has no access to external informations about its absolute position, and that the environment consists of a set of objects, whose separation is large compared to the range of the perception sensors. Our approach is based on: *(i)* maximizing the probability of successful transition between any pairs of known environmental components, by requiring, if necessary, that the robot uses other objects as intermediate landmarks, optimally combining periods of deliberative motions and perceptual behaviors, and *(ii)* guaranteeing, by controlling the extension of purposive robot's motions, and resorting if necessary to active search for perceptual information, that a reliable path to the homing region always exits. The paper presents experimental results that clearly demonstrate the increase of robot's autonomy resulting from our navigation strategy.

1 Introduction

In this paper, we present a navigation strategy for mobile robots operating in *open space regions, were a set of isolated objects are present, the distance between objects being large compared to the range of the robot's perception sensors* (hypothesis H_1). We further assume that the robot has *no access to absolute positioning information and no a-priori knowledge of its environment* (hypothesis H_2). This set of assumptions is representative of the operating conditions that can be found, for instance, in underwater robotics where:

H_1: distinctive perceptual features are scarce, while the range of perception sensors (sonar, vision) is small;
H_2: external positioning information, such as GPS, is not available inside the water column, leading to global positioning based on acoustic references, whose effective range is limited by the height of the water column; a priori information, when available, is extremely unreliable.

Combined, H_1 and H_2 impose the need for *important periods during which odometry is the sole means to localize the robot*, i.e., the areas inside which perception-based navigation is possible are disconnected, and represent a small fraction of the entire workspace. The largest risk associated to platforms operating in this kind of unbounded environments is to **get lost**, being unable to come back to their designated homing region (were, for instance, they must be picked-up or recharge their batteries). The control of the platform must then be addressed as a multi-criteria optimization problem, which attempts to *minimize risks*, while at the same time *minimizing consumed resources* (energy, distance, or time of passage).

With this motivation, we propose in this paper a navigation control strategy which is dictated by the necessity of **guaranteeing the recovery of the platform**, and which is based on three fundamental tools: *(i)* the possibility of surveying the **current loss risk**; *(ii)* the ability of assigning a **risk** and *(iii)* associating an **utility** to future actions.

Our approach to the problem is grounded on the following observations: *(a)* uncertainty about the robot's state increases with the degree of maneuvering required by the trajectory, *(b)* no external positioning information being available, the only possibility to minimize the uncertainty at the end of deliberate motions is to minimize it at their beginning. The first observation, *(a)*, resulted in restricting nominal deliberate motions to straight line trajectories,

* Email: {rendas,rolfes}@i3s.unice.fr

while *(b)* lead to the definition of a discrete set of points in the configuration space of the robot, which we designate by "gateways," and which are associated to the detected objects: before starting a deliberate motion out of an object, the robot uses its acquired knowledge and perceptual sensors to "dock" at a gateway point, which is chosen as the configuration that offers the best possible "initial conditions" for the subsequent period of odometric navigation.

During deliberate motions (not supported by perceptual data), the robot continuously supervises its security, by **determining the probability that it is able to return to its homing region**. If this probability falls below a prescribed level, the deliberate motion is aborted, and the robot is directed to one of the known objects, from where it is sure (up to a given success probability) that it will be able to reach in a safe manner the homing area.

To travel between known objects, our control strategy leads the robot, through an alternation of deliberate and perception based periods, along the path that presents the **smaller predicted risk**, i.e., for which the probability of reaching the designated target object is larger. We point out that this criterion automatically yields a compromise between the deliberate and behaviorist approaches to path planning. If the positioning errors during purposive displacements are large, a path is found that relies heavily on the possibility of periodic repositioning using environmental knowledge, increasing the duration of perception-based maneuvers. Otherwise, the length of straight line segments is maximized, and a path close to the shortest path between initial and final positions found, as our experimental results show.

To access target points inside unexplored regions, the robot is directed to the point in the known portion of the workspace that offers the **best expected utility** for reaching the target point. If necessary, it will execute, based on that point, an active exploration of its neighboring region, in search for new perceptual features that can safely support its progression to the goal. This active exploratory behavior leads to an effective extension of the autonomy range of the robot (defined as the size of the region that it can access without getting lost) and is necessary for safe operation in scarce environments.

2 Related Work

The navigation architecture presented in the paper differs from previously published works in several ways.

Open scarce environments

Most of the literature concerning the navigation of mobile robots using perceptual information has been developed for dense structured environments such as office spaces [1, 2, 3, 4]. They rely heavily on a priori assumptions about the environment (for instance, [3] assumes scenarios composed of hallways, crossings and doors), which they use to constraint the set of possible robot's actions (e.g. [2], where the only possible robot's actions are STOP, MOVE FORWARD or TURN RIGHT or LEFT by 90^o). The question central to our approach, of loosing the robot, is not relevant in these restricted scenarios.

World representation

We use an approximate uncertain high-level description of the morphologic characteristics of the environment, which allows us, using a detailed statistical model of the uncertainty affecting actuators and sensor data, to predict the confidence on future actions of the robot. Most previous published works fall under two categories:

- Occupancy grid approaches, e.g. [10], which yield probabilistic low-level metric descriptions of the workspace, requiring unbounded memory and heavy computations for action planing. Our approach is also based on a metric description of the operating region. However, the use of a high-level geometric description in our work allows the use of analytical tools for prediction of the performance of actions and limits the amount of memory required for environment representation.
- Topological maps, e.g. [2]. We gather in this category all representations that tightly couple the representation of the state of the robot and the description of the environment. These approaches are suitable only for perceptually dense environments, where the robot can be constantly guided by its perception, lacking the degree of detail necessarily to analyze the exploration of new regions.

The originality of our work on environment modeling, when compared to previous approaches, stems also from the myopic characteristics of the perception sensors, which preclude constant position correction by successive frame alignment, like in [6, 7, 8, 9].

As it will be discussed briefly in a later section, our work does impose the spatial coherency of the over-all representation, a problem that has been addressed by other researchers, e.g. [8]. The major difference between our methods and previous references is the fact that we do not require memorization of all previous data, as some statistical-based approaches do, e.g. [9], nor the explicit propagation of cross-covariances between map elements, as, for instance, in [7]. Our algorithm relies on analytical approximate models for the evolution of position uncertainty along the robot's trajectory to recover this information when needed, and requires less computations and much less memory that the approaches mentioned.

Planning under uncertainty

Some authors have addressed the problem of planning the future trajectory of the robot, to achieve, under uncertainty, a given target state with maximum probability see e.g. [4, 5]. The major novelty of our work comes from the notion of "gateway points," to be precisely defined later. Contrary to [5], where the existence of separate regions of homogeneous perfect positioning is postulated, or to [4] which relies on a priori assumptions about the workspace to constraint the set of possible robot's actions, our definition is perfectly general, and strongly grounded on the geometric characteristics of the perceived objects, and on the uncertainty affecting control and perception of the robot.

Active exploration

A major element in our approach is the ability of providing the robot with an active exploration attitude. Most previous works on navigation under uncertainty consider a purely passive attitude of the robot, and address only the problem of using, at its best, the acquired perpetual data. As is has been discussed in the introduction, and is demonstrated by the examples shown later, the active behavior proposed here is a necessary condition for safe goal achievement when working in rarefied environments. In [10] and [11] the problem of actively controlling the robot in order to minimize uncertainty is addressed, with the goal of decreasing the uncertainty associated either to the robot's state or to the world description. In [10], the authors trigger active relocalisation procedures whenever the robot's uncertainty falls above a pre-defined threshold, selecting the action that will result in a largest decrease of the entropy of the robot's state. Two major distinctions with respect to our work can be singled out: *(i)* application of this control strategy is not motivated by, and does not guarantee, extension of the effective robot's autonomy range, *(ii)* this approach imposes a spatially homogeneous upper bound on the entropy of the robot's state, not taking into account the morphological properties of the robot's vicinity (our approach is more efficient, establishing a bound that takes these into consideration). Moreover, the problem of detecting dead-end states (where progression of the robot could imply its loss) and of actively exploring new regions is not addressed by neither of these authors.

3 Assumptions and Control Strategy

In this section we describe our control strategy, addressing in particular the issues of path planing and of triggering active exploration, i.e., switching from nominal mission execution to an active search for new perceptual information that can safely sustain the continuation of the robot's mission.

Built-in features

We assume that the robot can execute the following *low-level maneuvers*:

1) advance in straight line from its current configuration,
2) make a turn (possibly in place), directing itself at a given point;

and the following *perception-based behaviors*:

3) detect and acquire an internal representation of perceptual features,

4) dock into a given configuration defined with respect to a known object,
5) recognize a previously seen object, and reposition itself with respect to it.

In the experiments presented in this paper, we used a small mini-robot (diameter of 5 cm) with two driving wheels, equipped of myopic distance sensors (maximum range of about 10 cm) and odometric counters. Using the distance sensors, the robot can detect objects, and incrementally make a description of their shape, by executing a trajectory parallel to the object, details can be found in [13]. It can subsequently use this acquired representation to place itself at a given configuration, defined relative to the objet, [14]. An algorithm for shape recognition, [15], allows the robot to recognize objects that have previously been mapped and reposition itself with respect to them.

Mission

With generality, we consider that the robot's mission is defined by a sequence of target points, defined relative to its initial position and orientation, which it must reach sequentially. The homing region (the region were it must be able to return) contains the initial position of the robot, and an object is located there (serving as a "lighthouse" of limited range). This assumption is not restrictive, since, if the robot can be placed at a given position, then we should also be able to place there an object. Information about this object is either given to the robot, or autonomously acquired as the first step of its mission.

Internal representation

Throughout its mission, the robot builds an internal representation of all detected objects, that we designate by $\mathcal{E}$. We assume that $\mathcal{E}$ is a collection of M objects $\mathcal{E} = \{O_1, \ldots, O_M\}$. Each object O_m is characterized by: its position and orientation, described by the vector p_m, and its morphological properties (shape and, eventually, other attributes). In our work, we assign a probability distribution to p_m and use fuzzy sets to describe object's shape [13, 15]. Based on the morphological properties of object O_m we identify the set of gateway points $\{p_{mj}\}_{j=1}^{M}$ at which the uncertainty at the beginning of the deliberate motion in the direction of object O_j is minimized. Besides each object's position and attributes, we also assume known a measure of confidence of the relative positions and orientations of all pairs of objects. In our work, this relative uncertainty is described by conditional distributions $P(p_m|p_n)$. Note that these conditional distributions are inferred using statistical analysis and require only memorization of the robot's trajectory during acquisition of $\mathcal{E}$, i.e., on the order by which the object's were detected. Each object defines a local frame, inside of which the robot's uncertainty is fundamentally dependent on the object's perceptual features. Moving from one frame to another is governed by the conditional distributions mentioned above.

Planing graph

Let $\mathcal{G} = (\mathcal{N}, \mathcal{A})$ be a graph, whose nodes are the complete set of gateway points $\mathcal{N} = \{p_{ij}\}_{i,j=1}^{M}$, where we make $p_{ii} = p_i$, the object's positions and orientations defined previously. We call $\mathcal{G}$ the planning graph. Two sets of arcs $\mathcal{A} = \mathcal{A}_u \bigcup \mathcal{A}_d$ are defined between a subset of elements of $\mathcal{N}$. $\mathcal{A}_u$ is a set of *undirected* arcs $\{\beta_{ij}\}_{i=1,\ldots,M; j\neq i}$, linking node p_{ii} to the gateways from O_i to all other objects O_j (the notion of gateway points will be precised later, and their relevant property is that they are the configurations in the local frame attached to the departing object that maximize the probability of success of the action "move from O_i to O_j"). All these nodes have unit weight: $b(i,j) = 1$. These arcs correspond to executing a given sequence of perceptual behaviors, moving inside frame i, with reference to object O_i. The other set of *directed* arcs, $\alpha_{ij} \in \mathcal{A}_d$, corresponds to the allowed set of purposive motions, and link each gateway point p_{ij} in object O_i to the node p_{jj}. The weights of these arcs, $a(i,j)$, express the **degree of confidence** on passing directly from one object to another [16]. Figure 1 shows an example of $\mathcal{G}$ for an environment with 3 objects.

Going from one object to another

We now explain our control strategy, considering first, for ease of explanation, the case were the robot is located in the vicinity of one object, and must reach another mapped object.

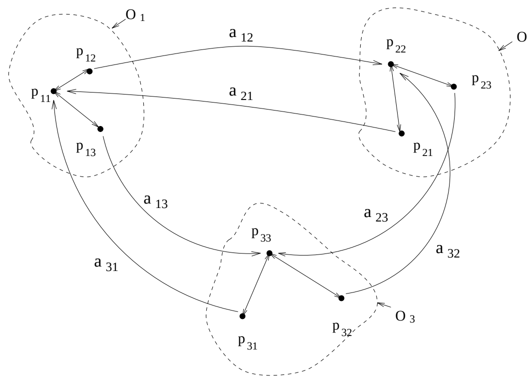

Fig. 1.: Planing graph

With the underlying constraint of trajectories that are a sequence of straight line segments linking detected objects, the planning problem is equivalent to searching for a minimum weight trajectory in $\mathcal{G}$ linking node p_{ii} to node p_{jj}. Let $t_{ij} = [i, t_{ij}^1, \ldots, t_{ij}^k, j]$ be one of the possible paths in $\mathcal{G}$ linking nodes p_{ii} to node p_{jj}. According to the probability algebra, the risk associated to this path is

$$C(t_{ij}) = a(i, t_{ij}^1) \cdot a(t_{ij}^1, t_{ij}^2) \cdots a(t_{ij}^k, j),$$

where, as we said before, $a(i,j)$ is the **degree of confidence** on passing directly from O_i to O_j. The determination of these probabilities is based on the prediction of the performance of all the basic actions involved in getting from one object to another (precision of docking, accumulated error over purposive motions), and is based on tools from parametric estimation theory and on sample path properties of Markov processes [17, 18]. Once $\mathcal{G}$ is defined, any standard minimum length path algorithm can be used to find the optimal trajectory between any two pairs of objects.

Exploring new regions

To handle situations where an unknown region of the workspace must be reached, the control strategy can no longer be based on the arc probabilities introduced above, since no object node is associated with the target regions. Our approach is based on the definition, for each target point T_n, of a "virtual node" V_n that has no shape description associated to it (and that has, thus, an empty set of gateway points). For each object node O_i, we define a new set of points e_{in} in the configuration space of the robot in the corresponding local frame, that we designate by "exploration gateways", and create new arcs linking O_i to each "virtual object" V_n. Each such arc is labeled by an **utility** measure $u(i,n)$ reflecting its usefulness as a starting point for active exploration in the direction of T_n. This measure of utility, with values in the unit interval, is mainly dictated by the probability of reaching the target point (up to a distance equal to the sensor's range) from the corresponding "exploration gateway."

When trying to attain a virtual object V_n, the robot searches for the exploratory gateway e_{in} which offers the largest utility $u(i,n)$. It directs itself to that node (using the safest route indicated by $\mathcal{G}$), and docks at e_{in}, in the direction of the virtual object V_n, initiating an exploration of the unknown region in the direction of T_n. We do not address here the problem of defining the actual exploration strategy that is used at each "exploration gateway," assuming just the following: *(i)* during the active exploration the robot will never loose the possibility of returning to the "exploration gateway," i.e., it moves within a bounded autonomy region determined by the morphological characteristics of O_i; *(ii)* if an object exists inside this autonomy region, the robot will find it, i.e., the search procedure is complete; and *(iii)* the value of $u(i,n)$ is continuously updated during the search, monotonically decreasing with the duration of the search procedure, and becoming zero if no object is found within the range of safe autonomy from e_{in}.

Assumption *(i)* above is verified by continuously **supervising the robot's risk level**, determining the probability that it is able, from its current position, to return to one of the known objects. If this probability falls below a prescribed level (close to one), the current motion is aborted, and the robot directed to one of the objects previously found, from where it initiates a new exploration. To compute this risk, the robot must be able to evaluate

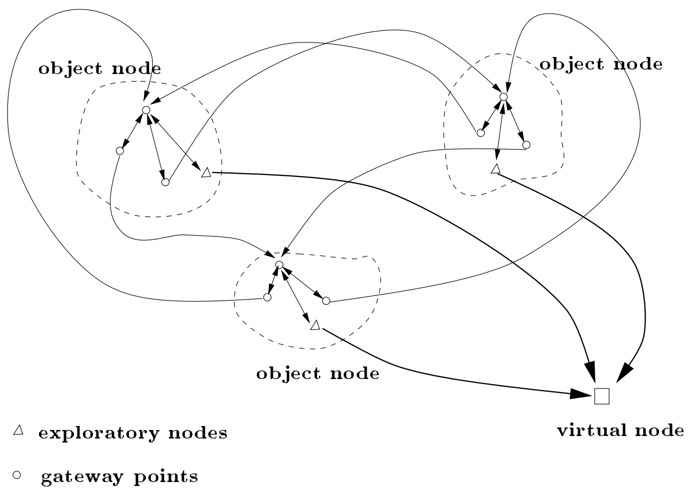

Fig. 2.: augmented graph, showing one virtual node and the exploratory gateways leading to it.

the probability of attaining any of the objects already in its map from its current position. This probability is computed using analytical tools that provide upper bounds for the evolution of the uncertainty during odometric navigation periods, [16, 17].

When executing an exploration in the direction of a target point T_n, three end conditions may happen:

1) the robot can directly reach the target point from the exploratory gateway,
2) a new object is detected within the autonomy range of e_{in};
3) $u(i, n)$ becomes zero.

In case *(i)* the robot updates $\mathcal{G}$, integrating in its environment description the morphological characteristics of the target region, which becomes then a normal object node of the graph. All exploratory gateways to T_n are deleted, and new gateways points to and from the new object, and the corresponding arc labels, defined. In case *(ii)*, $\mathcal{G}$ is augmented by the new object (along with corresponding normal and exploratory gateways) and a new exploratory procedure is repeated from the gateway with largest utility value. In case *(iii)*, an attempt is made to find an alternative exploratory gateway in object O_i, e_{in}, and the complete procedure re-started from the current most useful exploratory gateway leading to target T_n. If all remaining exploratory nodes have zero utility, the mission is aborted, declaring unreachability of the target point.

We remark that this control strategy necessarily leads to environment representations for which the corresponding planing graph is connected, i.e., it imposes that at least one safe path exists between any two objects.

4 Uncertainty modeling and prediction

In this section we briefly address the problem of computing the performance indexes used by our architecture and the definition of the gateway points.

Performance predicition

The probability of reaching a given object from a given uncertain initial configuration is formally related to a first passage time of the positioning error process. Previous studies [16] have shown that one component of the error process during linear motions is well approximated by a Brownian motion. The determination of the statistics of first passage times of multivariate processes is a difficult problem, for which the majority of the analytical results concern asymptotic behavior. Working with a one-dimensional projection of the error process, we use available results on scalar Markov processes, [12], to determine a lower bound on the probability of reaching a desired target object.

Let $p = (x, \theta)$ be the initial robot's configuration in the vicinity of object O_i, and $\Sigma(x, \theta)$ its uncertainty. The vehicle is launched in a nominal uniform linear motion from this position along direction θ, directed at the center of the target object. Making a change of variables,

$$\begin{bmatrix} y' \\ x' \end{bmatrix} = R_\theta \begin{bmatrix} x - x(0) \\ y - y(0) \end{bmatrix}, \qquad R_\theta = \begin{bmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{bmatrix},$$

such that y' coincides with the direction of the linear motion, and the initial position is placed at the origin of the new coordinate system, we use the following lower bound on the probability of reaching R:

$$Pr(\text{reaching } R \,|\, p) \geq Pr\left(\forall_{t \in [0, T_R]},\ x'(t) \in [-\ell_R, \ell_R]\right), \tag{1}$$

where T_R is the time required to reach the center of the target with constant nominal velocity v_0, with a given probability $\gamma \simeq 1$, and $2\ell_r$ is the spatial extention of the target in the direction orthogonal to the nominal direction of motion.

Note that the previous expression is a lower bound on the probability of attaining R, since other sample paths of the process that reach the target region may exist, crossing the boundary $|x'(t) = \ell_R$ a pair number of times. This bound can be computed using the statistical characterization of the error of the modified variable $x'(\cdot)$ at the final time T_R for which closed form expression are known, see [14].

Gateways

Determination of the gateway points p_{ij} is highly dependent on the particular perception sensors used by the robot, and on its holonomic constraints. We formally define the gateway points as the points that *allow a relative positioning of the robot with respect to the departure object that yield a lower risk of not achieving the target object*, as given by the lower bound presented above. To compute this bound we must know the local initial uncertainty, when the robot *docks* into a gateway, using perception.

We assume that the estimates of the position and orientation of the robot with respect to an object O are obtained by maximizing the likelihood of the observed perceptual and proprioceptive data z given the robot's final pose p and the known object characteristics,[2]

$$\hat{p} = arg\max_p f(z \,|\, p, O),$$

and approximate the covariance matrix of the estimation error by its Cramer-Rao lower bound.

This unconstrained approach, although optimal, leads to a non-linear multivariate optimization problem, which is far too complex to be solved in real-time. Considering a restricted set of docking maneuvers for robots with two driving wheels and polygonal objects, we were able to identify a small discrete set of candidate gateway points (located near the corners of the objects). These results can be found in [17] and will be published elsewhere in the near future.

5 Experimental results

In this section we illustrate our control strategy, showing experiments realized with the mini-robot Khepera, a small circular robot ($\Phi = 5.2$ cm), with two driving wheels, and equipped of a belt of 8 infra-red sensors, that were calibrated to be used as distance sensors.

The robot starts from the neighborhood of the small trapezoidal object O_1 shown in the lower left corner of the left plot in Figure 3, its mission consisting in attaining, by that order, points T_1 and T_2 (marked with $*$ in the plot). Note that the size of the rectangular region displayed is about $2{\times}2$ meters, much larger than the robot's size, and the range of its sensors (width of the rectangular regions in the left plot).

The left plot in figure 3 illustrates the active exploration of the work space, that leads to learning a safe route to the (unmapped) target points. The plot on the right concerns planing the motion between objects already mapped.

[2] $f(x|y)$ denotes the conditional probability density of x given y.

The robot starts by observing and mapping the object close to its initial position, and defines the exploratory gateways for the two targets. It positions itself in the exploratory gateway e_{11}, and moves from there in the direction of T_1 (along the path indicated by an arrow). During the linear motion, it is stopped by the risk surveillance procedure, and redirected again to e_{11}, from where it starts, after repositioning with respect to O_1, the exploration of the unknown region in new directions, alternating on each side of the direct route to T_1. These explorations are still aborted two more times, until the robot detects a new object, O_2, which it observes and maps. New gateways points (between O_1 and O_2) and exploratory gateways (from O_2 to the target points) are defined. The robot then selects the exploratory gateway e_{21}, which offers the largest utility. The same exploratory procedure is launched from that node, which, this time, leads directly to the detection of the object placed at the target point T_1. Note that since O_2 is larger than O_1, the range of autonomy out of the new object is larger than out of O_1.

The robot includes this new object, O_3, in its world description, deletes all exploratory gateways leading to T_1, defines gateway points to and from the new object O_3 and new exploratory gateways from O_3 to T_2. It then selects the exploratory gateway to T_2 with largest utility, and starts the exploratory behavior, reaching T_2 in the first movement out of O_3, and incorporating the description of the object placed there, O_4, in its environment representation. The rectangular regions in between objects the left plot of Figure 3 represent the areas sensed by the robot's sensors during the entire exploration.

The robot must now return to its homing region, i.e., to object O_1. Figure 4 illustrates the selected route in $\mathcal{G}$ at this point. The planned trajectories are represented by the sequence of straight linear motions that compose the planned path. Each linear segment is indicated by an angular cone, with vertex at the departing gateway point, and with a long central line showing the nominal direction of the linear motion.

The robot selects the path that leads it from O_4 to O_1 through O_2, which is the trajectory that offers the smallest risk. Note that this trajectory does not correspond to executing in the reverse direction the same path as during its progression to O_4. The relatively small distance between O_4 and O_2, and the large apparent size of O_2 seen from O_4, compared to the small attraction region of O_3, make this route preferable, in spite of the larger uncertainty in the relative position of these two objects (which the robot only knows "through" O_3). The plot at the right-hand side of Figure 3 displays the estimated trajectory during the return from O_4 to O_1.

The robot docks at gateway point p_{42} and progresses in the direction of O_2. When it reaches this object, it confirms that the object perceived is indeed O_2 by comparing the observed shape – in this case two linear sides of the polygonal object – with the memorized representation of O_2 [15]. Subsequently, it docks at p_{21}, executes a deliberate straight linear motion in the direction of O_1, checks that it reached the desired target and stops.

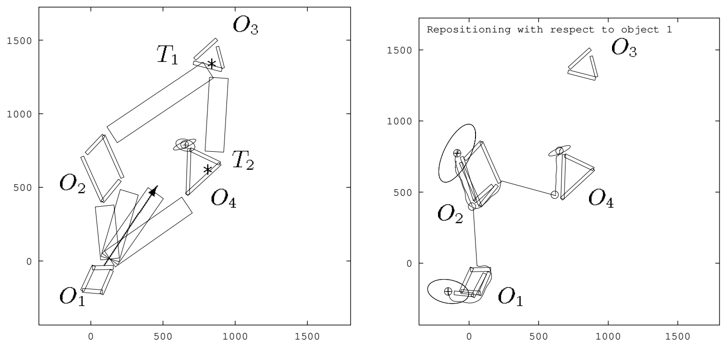

Fig. 3.: Mission with autonomy control, active exploration, and minimal risk routing.

Figure 5 illustrates the execution of the same mission without autonomy control. The two objects placed at T_1 and T_2 (O_3 and O_4) are detected, but represented at wrong positions and/or orientations, leading to robot's loss when it tries to return directly to O_1. This can be seen by the fact that the robot's estimated path goes through O_1, revealing an error at least as large as the object's size.

Figure 6 illustrates the safety of the learned representation for navigation between objects, showing motions between the two most distant objects, O_1 and O_3. The plots at the bottom of this Figure show the planned path,

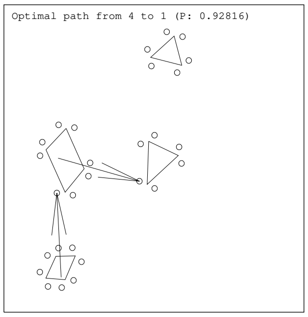

Fig. 4.: Routing graph for returning to O_1 from O_4.

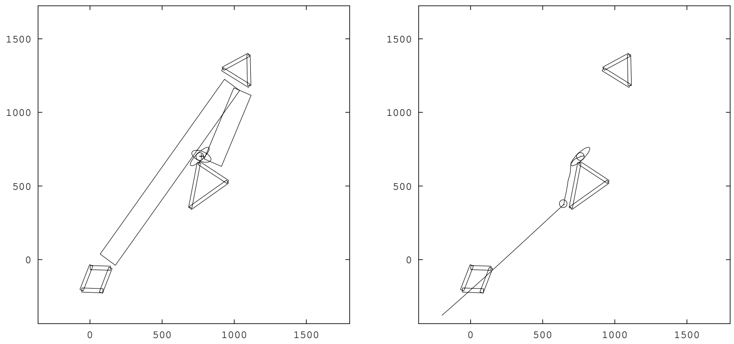

Fig. 5.: Executing the mission of Fig. 3 without autonomy control.

and the upper diagrams the position of the robot after it recognized the target object, and acquired its position in the local frame defined by it.

6 Conclusions

The paper addresses the control of platforms working in extended spatial regions where isolated objects are place at distances which may be very large compared to the perception range of the robot's sensors. We presented a navigation control strategy that allows the robot to switch between nominal mission and an exploratory behavior, searching for new perceptual information that can safely sustain its progression. This architecture is based on a statistical model of the uncertainty affecting the robot's state and its knowledge of the operating environment – which is incrementally acquired during mission evolution. The exploratory behavior is triggered whenever the probability of being able to a designated recovery region is smaller than a threshold.

References

1. Sven Koening, Reid Simmons. Unsupervised Learning of Probabilistic Models for Robot Navigation. *IEEE Int. Conf. on Rob. and Autom.* (ICRA'96).
2. Hagit Shatkay, Leslie Pack Kaelbling. Learning Topological Maps with Weak Local Odometric Information. Proceedings of the Int. Joint Conf. on Artificial Intell., IJCAI, 1997.
3. Sebastian Thrun. Learning Maps for Indoor Mobile Robot Navigation. *Artificial Intelligence*, to appear.
4. Sven Koenig, Richard Goodwin, Reid Simmons Robot Navigation with Markov Models: A Framework for Path Planing and Learning with Limited Computational Resources. Lecture Notes in Artificial Intelligence, Volume 1093: Reasoning with Uncertainty in Robotics; L. Dorst, M. van Lambalgen, and R. Voorbraak (Eds.); pages 322-337; 1996.
5. A. Lazanas, J.C. Latombe. Motion Planning with Uncertainty: A Landmark Approach. Artificial Intelligence, 76(1-2): 285-317, 1995.

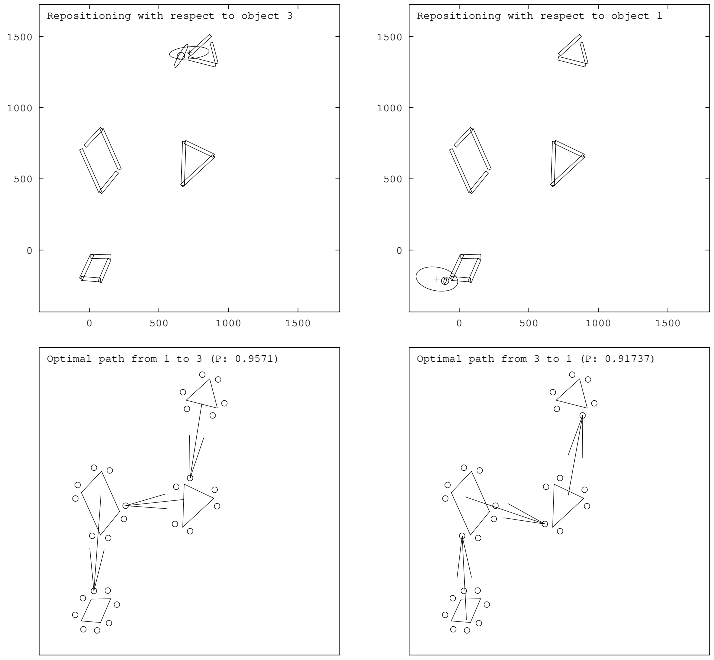

Fig. 6.: Navigating between known objects: $O_1 \rightarrow O_3$ (left) and $O_3 \rightarrow O_1$ (right).

6. Sean Engelson, Drew McDermott. Error Correction in Mobile Robot Map Learning. *IEEE Int. Conf. on Rob. and Autom.* (ICRA'92), pp: 2555-2560, Nice, France, May 1992.
7. J.A. Castellanos, J.M. Martinez, J. Neira, J.D. Tardós. Simultaneous Map Building and Localization for Mobile Robots: A Multisensor Fusion Approach. *IEEE Int. Conf. on Rob. and Autom.* (ICRA'98), pp: 1244-1249, Louvain, Belgium, May 1998.
8. F. Lu, E. Milios. Globally Consistent Range Scan Alignment for Environment Mapping. *Autonomous Robots*, 4, 333-349, 1997.
9. Sebastian Thrun, Dieter Fox, Wolfram Burgard. Probabilistic Mapping of An Environment By a Mobile Robot. *IEEE Int. Conf. on Rob. and Autom.* (ICRA'98), pp: 1546-1551, Louvain, Belgium, May 1998.
10. Michael Beetz, Wolfram Burgard, Dieter Fox, Armin Cremers. Integrating Active Localization into High-Level Robot Control Systems. Robotics and Automation Systems (special issue SIRS'96), to appear.
11. D. Fox, W. Burgard, and S. Thrun. Active MArkov Localization for Mobile Robots. accepted for publication in: Robotics and Autonomous Systems (special issue Euromicro'97), to appear. %
12. Leo Breiman. Probability. Classics in Applied Mathematics Series, SIAM, Philadelphia, 1992.
13. M. João Rendas and Stefan Rolfes. Propagating Multiple Hypothesis for Unknown Environment Learning. Int. Symp. on Intelligent Robotic Systems (SIRS'96), Lisbon, Portugal, July 1996.
14. S. Rolfes and M. João Rendas. A Two-Folded Approach for Positioning in Uncertain Environments. tech. report, Lab. Inf. Syt. de Sophia Antipolis (I3S), 98-#. February 1998.
15. S. Rolfes and M. João Rendas. Shape Recognition: a Fuzzy Approach,
16. M. João Rendas and Isabel Lourtie. Prediction of Positioning Uncertainty Accumulated Over Linear Motion Using Dead-Reckoning. CNRS, Laboratoire I3S (96-42), September 1996.
17. M. João Rendas and Isabel Lourtie. *Navigation and Mapping.* Deliverable D3, Project ESPRIT/LTR 20185 - NARVAL, February 1997.
18. M. João Rendas, José Santos-Victor, Jean-Yves Tigli, Isabel Lourtie, Luc Pronzato, Marie-Claude Thomas, Alexandre Bernardino, Miguel Ribo, " Uncertainty Modeling and Perceptual Guiding for Safe Operation in Unknown Environments," accepted for publication in the International Journal on Intelligent Control Systems, March 1997. *IEEE Int. Conf. on Rob. and Autom.* (ICRA'98), 1998.

This article was processed using the TEX macro package with SIRS98 style