

Resilient Navigation through Probabilistic Modality Reconfiguration

Thierry Peynot, Robert Fitch, Rowan McAllister and Alen Alempijevic

Abstract This paper proposes an approach to achieve resilient navigation for indoor mobile robots. Resilient navigation seeks to mitigate the impact of control, localisation, or map errors on the safety of the platform while enforcing the robot’s ability to achieve its goal. We show that resilience to *unpredictable* errors can be achieved by combining the benefits of independent and complementary algorithmic approaches to navigation, or *modalities*, each tuned to a particular type of environment or situation. In this paper, the modalities comprise a path planning method and a reactive motion strategy. While the robot navigates, a Hidden Markov Model continually estimates the most appropriate modality based on two types of information: context (information known *a priori*) and monitoring (evaluating unpredictable aspects of the current situation). The robot then uses the recommended modality, switching between one and another dynamically. Experimental validation with a SegwayRMP-based platform in an office environment shows that our approach enables failure mitigation while maintaining the safety of the platform. The robot is shown to reach its goal in the presence of: 1) unpredicted control errors, 2) unexpected map errors and 3) a large injected localisation fault.

1 Introduction

Motion planning and control of a mobile robot necessarily involves multiple sources of uncertain information such as control uncertainty, localisation uncertainty, and mapping errors. Current research seeks to address these sources of uncertainty by

Thierry Peynot, Robert Fitch, Rowan McAllister
Australian Centre for Field Robotics, The University of Sydney, NSW 2006, Australia, e-mail: {tpeynot, rfitch, r.mcallister}@acfr.usyd.edu.au

Alen Alempijevic
Mechatronics and Intelligent Systems Group, University of Technology Sydney, NSW 2007, Australia, e-mail: Alen.Alempijevic@uts.edu.au

modelling them in the context of the planning problem [1, 5]. However, problems arising during execution of a plan are not always predictable (and hence able to be modelled). For example, it is difficult to predict localisation errors ahead of time, or to anticipate which map locations actually contain large errors. We are interested in mitigating the impact of such unpredictable errors on robot performance and safety. We introduce the term *resilience* to refer to this goal. Resilient navigation seeks to mitigate the impact of control, localisation, and map errors on the safety of the platform while enforcing the robot’s ability to achieve its goal.

In this paper, we study resilient navigation in the context of indoor mobile robots. We believe resilience is best achieved by combining the benefits of multiple independent algorithmic approaches, or *modalities*, each tuned to a particular type of environment or situation. The idea is to develop a set of modalities that covers the range of possible situations, and then to reconfigure the system dynamically in response to unpredicted errors. The key challenges are: 1) how to choose a suitable set of modalities, 2) how to represent information that describes the robot’s context, and 3) how to decide which modality is most appropriate at any given time. Because we are dealing with uncertain information, these challenges require solutions in probabilistic form.

Our approach in choosing a set of modalities is to include a motion planning strategy that requires global information, and a reactive strategy that requires only local information. These two modalities are complementary. If the navigation goal is within the field of view (FOV) of the robot, a reactive obstacle avoidance approach (e.g. [10, 8]) can be feasible. However, reactive approaches have known limitations. They can become trapped in dead-ends or U-shape obstacles, and it is difficult to obtain smooth trajectories. If the goal is located outside of the robot’s FOV, the recommended strategy is to use a motion planning algorithm that reasons more globally, especially if some prior knowledge of the environment is available. In addition, smoother and more efficient paths can be obtained (see Fig. 1). However, in cluttered environments, such a strategy can only be effective if sufficiently accurate map and global localisation are available. In addition, the control of the platform needs to be robust and precise enough to follow the planned trajectory.

An alternative is to combine the two strategies to obtain a *hybrid* system [4]. Typically, a motion planning algorithm computes a global plan, generating a list of waypoints along the computed trajectory which are passed to a reactive motion method. A drawback of these hybrid techniques is that even if the motion planner can produce smooth trajectories (or trajectories respecting some pre-defined constraints), the execution of such types of trajectories cannot be enforced. Another inconvenience is that events that provoke failure of one of the components will often provoke failure of the combination, whereas this can be mitigated by using the appropriate method at the right time. A comparison of the different strategies discussed in this paper is shown in Table 1.

Instead, we propose a modality-switching algorithm based on a *hidden-markov model* (HMM) that considers *context* and *monitoring* information. If the system is aware that path execution cannot safely handle a difficult situation such as a narrow doorway, it is appropriate to switch to a reactive strategy. This situation can be

Table 1 Comparison of navigation strategies

Strategy:	Planning	Reactive	Hybrid	Our approach
Robust to dead-ends	✓	✗	✓	✓
Robust to dynamic obstacles	✗	✓	✓	✓
Robust to errors in localisation or map	✗	✓	✓	✓
Optimised paths (when possible)	✓	✗	✗	✓

evaluated using the localisation of the robot in a map, and detecting the presence of this narrow passage. However, reasoning only on this context information will not be sufficient to handle situations where the error/uncertainty of global localisation is high, where elements of the map have moved, or where a dynamic obstacle has appeared. Fast local replanning integrating map updates can partially address this problem but is computationally expensive and can lead to instabilities in control. Therefore, we propose to choose a modality based on context information *and* monitoring information (such as proximity to obstacles observed from laser data).

We evaluate our approach through hardware experiments with an indoor mobile robot in an office environment. We show that failures can be mitigated in challenging situations while maintaining the safety and liveness of the platform. The situations we consider include: control errors, localisation errors, map errors (unexpected obstacles), and presence of an “aggressive” human dynamic obstacle.




Fig. 1 Trajectory obtained using a planner (a) and a *hybrid* approach (b). Obstacles in the map are in black. Circles represent the radius of the robot. Approximate size of the area shown: $6.5m \times 7.5m$.

2 Related Work

Previous work has considered multi-modal systems for the navigation of an indoor mobile robot. [11] proposed the Robel system: a robot controller that learns different ways of combining sensory-motor functions to achieve a navigation task. Robel uses a Markov decision process (MDP) to provide a policy. However, MDPs are generally computationally expensive and policies often have to be computed off-line or at low frequencies. Our system was designed to be efficient enough for the robot to be reactive: modality switching can happen quickly when needed. Besides, an MDP requires the states to be fully observable. [15] proposed a system based on a partially observable MDP (POMDP) that can be used to detect, diagnose and recover from faults. However, the policy is computed off-line and the robot does not have a real alternative navigation modality when the path planning strategy fails. Our approach does not require explicit detection and identification of specific faults such as a localisation error; it focuses on mitigating failures that could occur in consequence, finding alternatives to obtain robustness while maintaining safety.

Motion planning and obstacle avoidance are well-studied problems in the literature. See [6] for a comprehensive review up to 2005. More recently, researchers have sought to address motion planning under uncertainty in control [1], localisation [5] or sensing and environment map [9]. However, these studies typically require the ability to *predict* possible errors, as they need to model the uncertainty in the context of the planning problem. In this paper, we are interested in mitigating the impact of *unpredictable* errors.

3 Probabilistic Modality Reconfiguration

The approach we propose is a probabilistic framework for an indoor robot endowed with two main navigation modalities: 1) a global planner (*PLAN*), and 2) a reactive motion approach (*REACT*). In addition, a *STOP* modality is included for emergency and safety. This method builds on our previous work for an outdoor mobile robot with modalities appropriate to flat terrain and rough terrain [13].

Our approach is to estimate the likelihood of each modality being most suitable using an HMM. The HMM is appropriate since states are not directly observable and it provides a time integration that prevents jitter (too frequent modality changes, see Fig. 5). Crucially, the probabilistic approach allows the system to handle uncertainty in the different sources of information.

The goal of the HMM is to provide a *modality recommendation*. The HMM is constructed such that the number of states is equal to the number of available modalities. Fig. 2 provides a graphical representation of our three-state HMM, where each state x_k corresponds to the proposition: “modality m_k is the appropriate modality to apply at this point in time.”

Two categories of information are input to the HMM: 1) *context* information is global environmental knowledge known *a priori*, and 2) *monitoring* information is

online execution knowledge of the observed immediate situation. The framework is designed as a Markov *conditional estimation* system [2]. It estimates the conditional state $x_{k,t}$ at time t , knowing context observation until time t , $O_{1:t}$, and online monitoring information $M_{1:t}$. If the robot is endowed with N different modalities, the probability that m_k is the appropriate modality to apply at time t can be written, $\forall k \in \llbracket 1, N \rrbracket$,

$$P(x_{k,t} | O_{1:t}, M_{1:t}) \propto P(O_t | x_{k,t}) \sum_{i=1}^N P(x_{k,t} | x_{i,t-1}, M_t) P(x_{i,t-1} | O_{1:t-1}, M_{1:t-1}) \quad (1)$$

where $P(O_t | x_{k,t})$ is an observation probability (computed using the *context* information), and $P(x_{k,t} | x_{i,t-1}, M_t)$ is the conditional probability of transition from state x_i to state x_k , knowing the *monitoring* data M_t at time t . The following sub-sections describe more specifically the different modalities of the robot used in this paper and the nature of the context and monitoring information.




Fig. 2 Graphical representation of the 3-state HMM

3.1 Context

The context information relates to the distance d from the robot boundaries to the closest obstacles as seen on an *a priori* global map. This information is used to predict the likely modality at a given map location. We determined experimentally that the *PLAN* modality is likely to fail in situations where the robot is too close to obstacles, i.e. closer than a *security distance* $d_s = 0.15m$, equal to half the radius of the robot. Therefore, intuitively, the *a priori* recommendation based on context information is to use *PLAN* in areas sufficiently clear from obstacles ($d > d_s$), *REACT* in areas that are close to an obstacle on the global map ($d \leq d_s$), and *STOP* in places immediately proximal to obstacles ($d < d_c$, *critical distance*).

d (the observation O_t) is calculated online using the current localisation of the robot in the map. To integrate this observation in the system (HMM), probability density functions (pdf) are used to take into account uncertainties. The main sources of uncertainties are the (x, y) localisation of the robot in the map and the location of

the obstacles in the map itself. The map is an occupancy grid that was built using the laser of the robot assuming perfect localisation. Therefore, the map uncertainty can be expressed as $\sigma_{map} = \sigma_{laser}$, the standard deviation of the range measurements of the laser scanner. $\sigma_{laser} = 0.03 \text{ m}$ for the Hokuyo laser on our robot.

The uncertainty in the *a priori* map is independent of the uncertainty of the current localisation, as the map was built beforehand, using a different localisation. Therefore, the standard deviation on the observation of d can be expressed as the sum of the uncertainties: $\sigma = \sigma_{map} + \sigma_{loc}$, where σ_{loc} represents the localisation uncertainty provided by the algorithm mentioned in Sec. 4.

3.1.1 Modality STOP

We define the distribution of $p(O_t|x_2)$, or $p(d|STOP)$, as an inverse sigmoid centred on the critical distance d_c (see Fig. 3 in red):

$$P(d|STOP) = 1 - \frac{1 - \alpha}{1 + e^{-(d-d_c)/\sigma}} \quad (2)$$

where σ partly defines the curvature of the sigmoid. σ (similar to the standard deviation of a Gaussian) corresponds to the uncertainty in d , and $d_c = 0$ is the critical distance.

The distribution $p(d|STOP)$ represents the likelihood that observation d is made, knowing that the robot should stop. The inverse sigmoid accounts for the uncertainty in the observation and in the knowledge of this threshold value. The limit of this sigmoid, when d tends to infinity, is superior to zero (see Fig. 3). This accounts for the fact that the map does not capture all information in the world, in particular dynamic obstacles. The value of this limit represents the chance of having to stop the robot while infinitely away from map obstacles, i.e. the chance of having a dynamic object appearing within less than d_c of the robot, α . It is crucial to account for the possibility of this event sufficiently so that the system maintains a chance of capturing it [7]. Thus, this value is set to a value higher than the actual probability of occurrence as would be determined statistically. In our implementation we set $\alpha = 0.05$.

3.1.2 Modality REACT

The distribution of $p(O_t|x_1)$, or $p(d|REACT)$, is defined as a sigmoid centred on d_c (see Fig. 3 in blue):

$$P(d|REACT) = \frac{1 - (\alpha + \gamma)}{1 + e^{-(d-d_c)/\sigma}} \quad (3)$$

where $\sigma = \sigma_{map} + \sigma_{loc}$, as defined earlier. To guarantee safety, the main restriction for this modality is that it cannot be used too close to obstacles ($d < d_c$), hence the sigmoid.




Fig. 3 Probability density functions for the context information (shown before normalisation, with $\sigma_{loc} = 0$). These are functions of d , representing the distance to the closest obstacles to the robot, as seen in the map.

There are two secondary restrictions. One consideration is the chance of a dynamic object appearing within a distance d_c to the robot, i.e. $\alpha = 0.05$. The other is the *a priori* chance of failure of *REACT* in general, even in an open map (recall that this modality is subject to local minima). This chance of failure highly depends on the environment, which we capture with the probability: $\gamma = 0.20$. Considering the events represented by α and γ as independent, the limit of the sigmoid $p(d|REACT)$ when d tends to infinity is set to $1 - (\alpha + \gamma) = 0.75$.

3.1.3 Modality PLAN

The distribution of $p(O_t|x_0)$, or $p(d|PLAN)$, is also defined as a sigmoid, centred on the security distance d_s (see Fig. 3 in green):

$$P(d|PLAN) = \frac{1 - \beta}{1 + e^{-(d - d_s)/\sigma}} \quad (4)$$

Note that once again the limit of the sigmoid $p(O_t|m_0)$ when d tends to infinity is lower than 1. This accounts for the chance of having dynamic objects appearing within d_s of the robot bounds. We consider the prior probability of this event to be $\beta = 0.10$ ($\beta > \alpha$), therefore the limit of the sigmoid distribution is $1 - \beta = 0.9$. For high values of d it is important to set the chance of success of *PLAN* higher than *REACT* (if the goal is far, it is known that *PLAN* is more likely to succeed), i.e. $\beta < \alpha + \gamma$. Finally, note that these distributions need to be normalised before integration in the HMM.

3.2 Monitoring

Contrary to context information, the purpose of monitoring information is to check the actual “appropriateness” of the current situation, with regard to the possible modalities, using data only observable during execution. The online monitoring uses δ , the distance from the robot bounds to the closest obstacle detected in laser measurements. Recall that the online monitoring contributes to the computation of the transition probabilities of the HMM. If the robot gets too close to obstacles seen in current laser scans while operating in *PLAN*, it should switch to *REACT*. In this way, if global localisation is temporarily inaccurate, or if obstacle points are in a different location than on the (static) global map, this situation can be handled by *REACT*, contrary to *PLAN*.

More specifically, the intuitive rules of transitions (given here without uncertainty, for convenience) are the following. The corresponding transition probabilities used in the HMM are given in parenthesis, in both full (e.g. $P(x_2|x_1, \delta)$) and equivalent reduced form (e.g. $p_{1,2}$). First, let us consider the output transitions of x_0 (i.e. *PLAN*).

- The transition $P(x_1|x_0, \delta) = p_{0,1}$ (*PLAN* to *REACT*) is likely if $d_c < \delta < d_s$, i.e. an obstacle is detected by the laser in the *intermediate* proximity of the robot.
- The transition $P(x_2|x_0, \delta) = p_{0,2}$ (*PLAN* to *STOP*) is likely if $\delta < d_c$, i.e. an obstacle is detected by the laser in the *immediate* proximity of the robot.
- $P(x_0|x_0, \delta) = p_{0,0}$ (*PLAN* to *PLAN*) is likely if $\delta > d_s$, i.e. the robot is clear from obstacles.

The other transitions can be defined similarly, using the same short notations as above: $p_{1,0} = p_{2,0} = p_{0,0}$, $p_{1,1} = p_{2,1} = p_{0,1}$, $p_{1,2} = p_{2,2} = p_{0,2}$.

Because of the uncertainty in δ (the laser measurements), these rules are defined probabilistically using sigmoid distributions similar to those defined in Sec. 3.1 and shown in Fig. 3. In this case the main source of uncertainty is the relative inaccuracy of the laser measurements, therefore the σ of the sigmoids is: $\sigma = \sigma_{laser}$. The output transition probabilities from each state are normalised, as their sum must equal 1.

4 Implementation

Our experimental platform consists of the Segway RMP100 base with onboard PCs and various sensors, including a Hokuyo UTM-30LX laser range-finder and encoders in the mobile base for odometry [12]. Localisation is computed using the Monte Carlo Localisation (MCL) algorithm [14]. The robot’s belief is represented by a set of weighted hypotheses which approximate the posterior under a common Bayesian formulation of the localisation problem. We update this distribution using data from odometry, the laser range-finder, and a predefined map of the environment.

The test area is an office environment occupied by over 25 people and consisting mainly of student workstations and fixed and movable furniture. This area is thus well-suited for evaluating real-world applicability.

4.1 Available Modalities

Modality m_0 is *PLAN*. We implemented the well-known Latombe Grid-Search algorithm [3] for nonholonomic planning, customised to find paths with minimum change in curvature. Although the name may seem to imply a discrete search space, the algorithm does use continuous coordinates. A detailed summary can be found in [6]. The planner is complete with respect to the resolution of its given proximity grid and time interval of the path set [3]. Because this proof is not constructive, we do not have a method for determining parameter values analytically. We hand-tuned them empirically and found a reasonable grid resolution of $0.2m \times 0.2m \times \frac{\pi}{8} rad$ and path set time intervals of 2 or 4 seconds. Our path set has angular velocities chosen from $\{-\frac{\pi}{4}, -\frac{\pi}{8}, -\frac{\pi}{16}, -\frac{\pi}{32}, 0, \frac{\pi}{32}, \frac{\pi}{16}, \frac{\pi}{8}, \frac{\pi}{4}\}$ and linear velocities from $\{0.2m/s, 0.1m/s\}$. Our priority queue uses a cost function that combines minimum distance to goal with minimum change in curvature.

Modality m_1 is *REACT*. This is a reactive collision avoidance method that avoids sensed obstacles. We implemented a potential field method derived from a model of human navigation [8]. This method directly controls angular acceleration and produces smooth paths. We chose this method because the robot operates in an office-like environment amenable to human-like paths. Because the laser cannot scan all 360° around the robot, the perception data that *REACT* uses is a local fusion of laser scans. Odometry is used for localisation in order to avoid the influence of errors in global localisation.

Modality m_2 is *STOP*. This is the safety modality; it applies when the robot has come too close to an obstacle and the only reasonable option is to halt. If *STOP* was provoked by monitoring information, the robot can only resume when the obstacle responsible for the stop is dynamic and has moved away. To account for this, our system continues to evaluate the HMM recommendation even though the robot is stationary.

4.2 Modality Switching

PLAN is the default starting modality, as it has the highest prior probability. When switching from *PLAN* to *REACT*, a goal waypoint must be chosen. We initially choose the next waypoint on the last path computed by *PLAN*. However, because of localisation or map errors, this waypoint may intersect an obstacle. In this case, the next waypoint of the plan that is confirmed as clear from obstacles becomes the new

goal for *REACT*. It is then preferable to switch back to *PLAN* quickly to avoid the risk of *REACT* falling into local minima.

5 Experiments

The experimental validation in this section illustrates the resilience of our probabilistic reconfiguration approach, which allows the mitigation of unpredictable failures. Examples of causes of such failures are: errors of the controller while executing a planned path, errors in the map (i.e. presence of objects that could not be integrated in the map early enough) and large localisation uncertainty or error. We also compare our method to simple threshold rules for modality switching. Results were obtained using the platform described in Sec. 4.

The illustrations show the estimated robot trajectory during each test, using coloured points to represent the modality used at the time. The selected modality corresponds to the highest output probability $P(x_{k,t}|O_{1:t}, M_{1:t})$ at each time step t . The HMM and the modality selection were updated at $10Hz$.

5.1 Modality Reconfiguration in Static Environment

5.1.1 Unpredicted Control Error

The experiment in Fig. 4(a) illustrates how our framework allowed the system to maintain the robot’s safety in the presence of unpredicted errors of the controller during execution of the planned trajectory. The robot started executing a planned path similar to the one in Fig. 1(a), which was successful using *PLAN* only. However, at the (expected) end of the turn around the first corner, the controller “overshot”, risking the safety of the platform. This event was detected by our system, which switched to *REACT* to recover. When safe, the robot returned to the *PLAN* modality to complete its mission.

5.1.2 Going Through a Narrow Doorway

We also tested the robot’s ability to follow a corridor and then pass through a $0.85m$ -wide doorway (the robot’s diameter is $0.6m$). As the corridor is reasonably large (about $1.7m$ in average), the robot first used *PLAN* and only switched to *REACT* to negotiate the doorway passage.

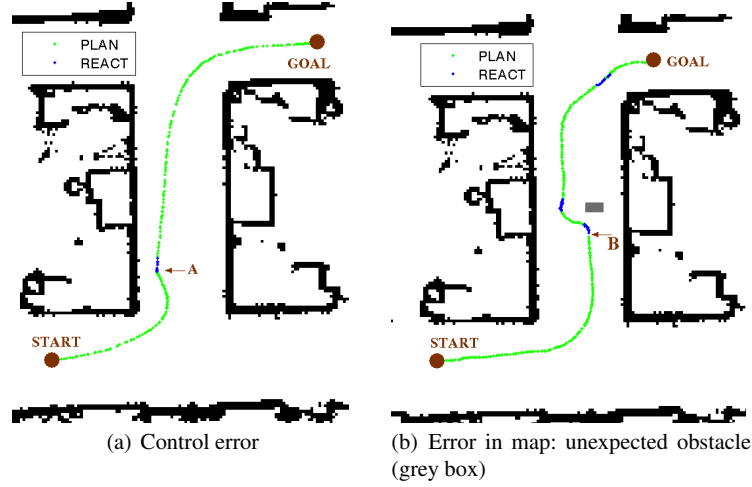


Fig. 4 Examples of robot trajectory executed with modality switching. Known obstacles in the map are shown in black, while the colour points show the (estimated) positions of the centre of the robot. Green means the recommended modality is *PLAN*, while blue means the recommended modality is *REACT*. Approximate size of the area shown: $6.5m \times 7.5m$.

5.2 Unexpected Map Error and Comparison to Simple Threshold Rules

Fig. 4(b) shows another example of modality switching to negotiate an unexpected situation safely: an unpredictable large error in the map. This situation is caused by the presence of an unexpected obstacle. This simulates a map error. In order to avoid the box, the robot switches to *REACT*, then returns to *PLAN* once the situation is safe and the map is more consistent with the current observation. A likely collision was thus avoided.

Fig. 5 illustrates a similar test using a recommendation based on simple logical rules comparing d and δ to “hard” thresholds equal to d_c and d_s . It can be seen that such strategy can provoke frequent undesirable modality switches, contrary to the HMM of our approach.

5.3 Presence of Dynamic Obstacles

We validated that the robot is resilient to the presence of highly dynamic obstacles. In the test shown in Fig. 6(a), a pedestrian coming from the top left of the scene walked quickly towards the robot. On approach, the robot first switched to *REACT* and then tried to evade (event C). Once the human had left the vicinity, the robot




Fig. 5 Comparison with a simple threshold strategy. The robot encounters an unexpected obstacle in (a). (b) shows the chosen modality over time for the 10s surrounding event A in (a). Modality switching with fixed thresholds results in unacceptable oscillation (top) compared to our method (bottom).

could resume its mission. A similar situation occurred again later in the test, with an even more sudden appearance of the human in the FOV of the robot. This event was again safely handled by the robot. This test shows that, although the robot can nominally execute optimised trajectories, it can also safely react to dynamic obstacles, comparably to a pure reactive motion strategy.

5.4 Injected Localisation Fault

In this experiment, a significant localisation fault was artificially created by introducing a sudden and unexpected offset of $1m$ to the output of the localisation estimator. Fig. 6(b) shows the clear offset to the right between the estimated position and the reality. However, the robot was still able to safely achieve its mission by switching to *REACT* when appropriate.

The context information is shown in Fig. 7 and the corresponding partial probabilities of modality recommendation are shown in Fig. 8. The moment of the localisation fault injection is clearly visible at $t = 35s$ (event D) on both figures. Fig. 7 indicates that with context information alone the robot should definitely *STOP*. If the robot could rely only on its current global localisation estimate and map, it would not have been safe to continue, since according to its estimated position the robot appears to be *on* the location of known obstacles in the map (see Fig. 6(b)). However, our system recommended a better alternative: a prudent switch to *REACT*. The

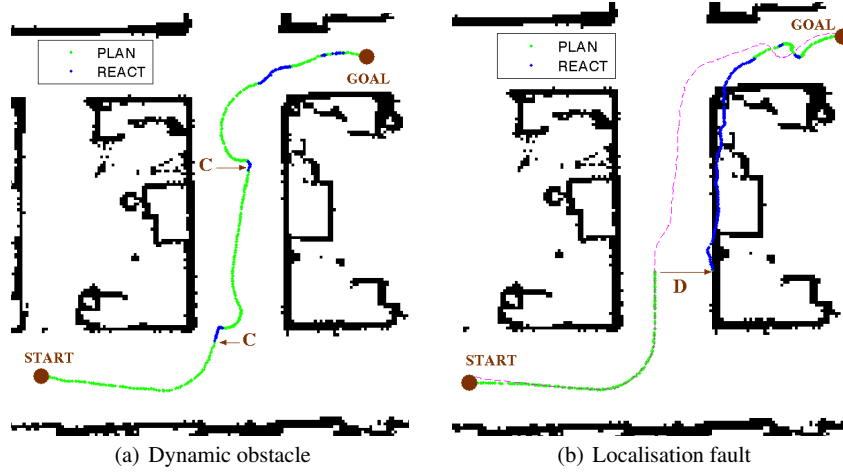


Fig. 6 (a): Presence of highly dynamic obstacles (events C). (b): Modality switching with injected localisation fault (sudden offset to the right of 1m, see D). As a result, the estimated positions of the robot appear to be *on* the right wall (blue line). The magenta dashed line shows the reference localisation.




Fig. 7 Partial probabilities for context information, corresponding to the test in Fig. 6(b).




Fig. 8 Partial probabilities for all three states (i.e. modality recommendation), corresponding to the test in Fig. 6(b).

localisation algorithm then progressively corrected its estimation up to a point when the situation became comfortable enough again to use *PLAN* to finish the mission. This shows our system to be resilient to large unpredicted localisation or map errors.

6 Conclusion

We have shown that a multiple modality strategy for resilient navigation, based on a probabilistic framework, can be applied to an indoor mobile robot to combine the advantages of navigation modalities while maintaining the safety of the platform. In our implementation, the robot is able to plan and execute smooth paths (minimising change in curvature) when possible, while being very reactive when needed. The system was applied online and shown to be reliable and robust in the presence of map errors and large localisation uncertainties or offsets. The concept is demonstrated with one choice of planning and reactive modality, however, these planning and reactive methods are easily interchangeable.

Future work will exploit the monitoring and context information for diagnosis and recovery of particular components of the system. Currently, both types of information are only exploited to compute a modality recommendation. However, we saw in Sec. 5.4 that the discrepancy between context and monitoring indicates a likely error in the map or global localisation. This could be used to actively repair these components, while the robot uses a reactive modality.

7 Acknowledgments

This work was supported in part by the Australian Centre for Field Robotics (ACFR) and the NSW State Government. This material is based on research sponsored by the Air Force Research Laboratory, under agreement number FA2386-10-1-4153. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon.

References

1. Alterovitz, R., Simeon, T., Goldberg, K.: The stochastic motion roadmap: A sampling framework for planning with markov motion uncertainty. In: *Robotics: Science and Systems II* (2007)
2. Arnaud, E., Memin, E., Cernuschi-Frias, B.: Conditional filters for image sequence based tracking - application to point tracking. *IEEE Transactions on Image Processing* **14**(1), 63–79 (2005)
3. Barraquand, J., Latombe, J.: Nonholonomic multibody mobile robots: Controllability and motion planning in the presence of obstacles. *Algorithmica* **10**, 121–155 (1993)
4. Brock, O., Khatib, O.: High-speed navigation using the global dynamic window approach. In: *Proc. of IEEE International Conference on Robotics and Automation* (1999)
5. Bryand, A., Roy, N.: Rapidly-exploring random belief trees for motion planning under uncertainty. In: *Proc. of IEEE International Conference on Robotics and Automation* (2011)
6. Choset, H., Lynch, K.M., Hutchinson, S., Kantor, G.A., Burgard, W., Kavraki, L.E., Thrun, S.: *Principles of Robot Motion: Theory, Algorithms, and Implementations*. MIT Press, Cambridge, MA (2005)
7. Dearden, R., Clancy, D.: Particle filters for real-time fault detection in planetary rovers. In: *12th International Workshop on Principles of Diagnosis* (2002)
8. Huang, W.H., Fajen, B.R., Fink, J.R., Warren, W.H.: Visual navigation and obstacle avoidance using a steering potential function. *Robotics and Autonomous Systems* **54**(4), 288–299 (2006)
9. Kurniawati, H., Bandyopadhyay, T., Patrikalakis, N.M.: Global motion planning under uncertain motion, sensing, and environment map. In: *Robotics: Science and Systems VII* (2011)
10. Minguez, J., Montano, L.: Nearness diagram navigation (ND): Collision avoidance in troublesome scenarios. *IEEE Transactions on Robotics and Automation* **20**(1), 45–59 (2004)
11. Morisset, B., Ghallab, M.: Learning how to combine sensory-motor functions into a robust behavior. *Artificial Intelligence* **172**, 392–412 (2008)
12. N. Kirchner et al.: RobotAssist - a Platform for Human Robot Interaction Research. In: *Proc. of ARAA Australasian Conference on Robotics and Automation* (2010)
13. Peynot, T., Lacroix, S.: Selection and monitoring of navigation modes for an autonomous rover. In: O. Khatib, V. Kumar, D. Rus (eds.) *Experimental Robotics: The 10th International Symposium on Experimental Robotics*, vol. 39, pp. 121–130. Springer-Verlag, Berlin (2008)
14. Thrun, S., Fox, D., Burgard, W., Dellaert, F.: Robust monte carlo localization for mobile robots. *Artificial Intelligence* **128**(1–2), 99–141 (2000)
15. V. Verma R. Simmons, J.F.: Probabilistic models for monitoring and fault diagnosis. In: *2nd IARP and IEEE/RAS Joint Workshop on Technical Challenges for Dependable Robots in Human* (2002)