Proceedings of Australasian Conference on Robotics and Automation, 3-5 Dec 2012, Victoria University of Wellington, New Zealand.

Path Planning with Maximum Expected Map Deformation

Mark Whitty, José Guivant
School of Mechanical and Manufacturing Engineering
University of New South Wales, Australia
m.whitty, j.guivant@unsw.edu.au

Abstract

Path planning for mobile robots during map construction requires the consideration of not only traversal costs but also belief state deformation, both of which have been studied independently. We therefore introduce a method for planning safe paths over cost maps given the Maximum Expected Deformation (MED) over the entire belief state. A process akin to convoluting the cost map is applied, whereby the cost map is dilated by a kernel whose size varies as a function of the MED at each point in the map. The dilated cost map is then used to generate the optimal policy that is guaranteed to be safe in the presence of map deformation. A traditional EKF-SLAM process was simulated and the resulting belief state used to both calculate the feasibility of a given path and the optimal policy for minimising traversal cost.

1 Introduction and Literature Review

The changing belief state during a stochastic mapping process such as Simultaneous Localisation And Mapping (SLAM) poses challenges for planning paths for autonomous mobile robots. Online SLAM approaches and faster versions of full SLAM approaches generally maintain or output a complete map and pose estimate on a regular basis. Translating this updated map into a plan similar to a grid world is trivial. Any subsequent updates to the map can cause deformation at any location, and hence the entire plan needs to be updated on a regular basis. The paths are required to be consistent with the belief at all times, or collision is likely to occur in which case they are considered to be unsafe.

Path planning in the presence of map uncertainty has followed two main themes in the existing literature; (1) using absolute map position uncertainty and (2) using information gain for efficient exploration of free space while avoiding obstacles.

In the first theme, Guibas *et al.* [2009] proposed an efficient method for planning on roadmaps with bounded absolute position uncertainty. Here, line segments were represented by their endpoints which were modelled by probability distributions with finite support. It included an efficient method of checking for collision free paths between the endpoints by discretising the edges and maintaining a bound on the probability for each configuration.

Censi *et al.* [2008] cited an extensive list of papers considering planning with uncertainty and went on to show that planning in the space of poses$\times$covariances reduced the problem to a deterministic search. They considered two problems - minimising execution time, while remaining localised; and minimising the final uncertainty, with a limit on the total time allowed. The uncertainty minimisation could be seen as a problem of optimal planning for exploration during SLAM, however the computation time was highlighted as a major drawback of this approach. Gonzalez and Stentz [2009] improved this approach by using linear landmarks and, through the use of a binning function for reducing the dimensionality of the search space, was able to demonstrate faster planning times, although still significantly slower than real-time.

In the second theme, Bourgault *et al.* [2002] and Makarenko *et al.* [2002] introduced metrics for guiding exploration by combining the expected quality of localisation, traversal distance and map information gain. The total utility of each point was linearly weighted between these three individual utilities. Their work was based on Fujimura [1996], who discussed the use of multiple objective optimisation for calculating a plurality of robot paths, akin to running Dijkstra's algorithm with multiple objectives. Rocha *et al.* [2005] implemented a multiple robot exploration strategy that explicitly represented uncertainty through the map's entropy. Thrun [1998] proposed a hybrid metric-topological planning construct based on binary occupancy grids during SLAM, but the occupancy grid was assumed to be deterministic for the purposes of planning.

Bryson and Sukkarieh [2008] used observability analysis to show that only the relative positions between features and the robot could ever be observed by traditional SLAM, and never the absolute positions of landmarks or the vehicle. This work was then extended to evaluate the information gain of a set of possible paths for an unmanned aerial vehicle. The information gain was estimated based on a regular distribution of expected landmarks which were distributed according to an a priori feature density. Several decision rules were imposed to select the path so as to maintain the localisation uncertainty below a threshold.

Given a path to a goal point, Huang *et al.* [2004] calculated the time optimal control policy which minimised the time taken to traverse the path while the localisation and map uncertainty was kept below a prescribed level. This was rather restrictive in that no spatial flexibility was allowed, as opposed to the work of active SLAM, where no such constraint was placed on the path. Fang *et al.* [2005] did preliminary work with bearing only SLAM and used a two stage planning process. Firstly, planning with the objective of improving the map quality while respecting a localisation accuracy bound and secondly to obtain the time-optimal path while respecting a localisation accuracy bound and a mapping accuracy bound. The two stage process was specifically designed for bearing only localisation, where substantial movement was necessary to reduce the initial map and localisation uncertainty to a useful level.

Leung *et al.* [2008] then combined frontier based exploration [Yamauchi *et al.*, 1998] with Model Predictive Control (MPC) and used an attractor point which was chosen according to the mode of operation - improve localisation, improve map or explore. Huang and Gupta [2008] introduced RRT style planning during SLAM, but assumed the map was not modified between particle filter updates, and that localisation was achievable within the map. Very recently, the idea of using RRT on a series of charts which parametrise the manifold which may deform during loop closure was presented by Jaillet and Porta [2011].

Approaches considered in the first theme handle map uncertainty only by using the absolute position uncertainty of the pose and/or the map and are therefore inefficient when large scale rigid-body translation and rotation of the belief state occurs, as is common during loop closure in SLAM. This is also true of the majority of approaches to planning in the second theme, but they are focussed on map quality maximisation rather than guaranteeing safety during planning. They also generally use occupancy grids with a binary classification of space — free space having zero cost and obstacles of constant (often infinite) cost — thereby ignoring traversal costs that may be provided independently of the localisation and mapping process.

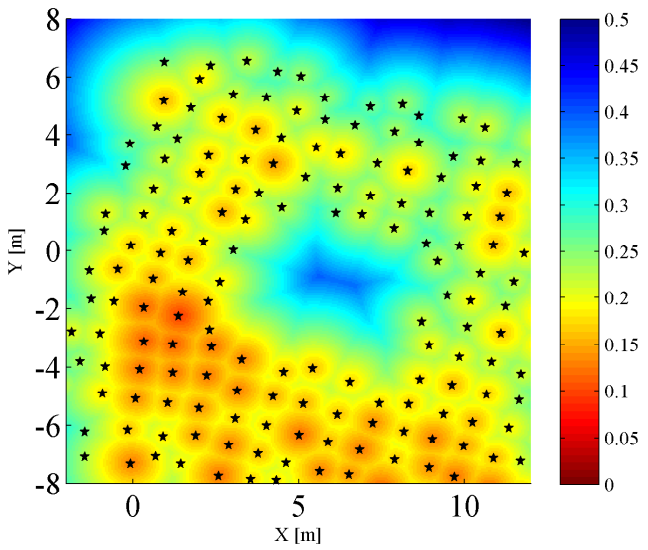

Figure 1: Maximum expected deformation of the belief state during SLAM based on sample point spacing and deformation metric values at neighbouring sample points. Sample points are shown as black stars. Higher values indicate higher expected deformation.

For mobile robots, emphasis has frequently been placed on increasing the efficiency of planners for autonomous operation. Here we present an alternative approach — to generate plans that are guaranteed to be safe before they are actually traversed, allowing for as much map deformation as possible — with the aim of vastly reducing the amount of replanning necessary. The objective of this paper is therefore to present a method for planning safe paths over cost maps given the Maximum Expected Deformation (MED) over the entire belief state.

1.1 Outline

In Section 2 we briefly outline how the MED may be obtained. Section 3 the presents the contribution of this paper - namely a method for using both cost maps and the MED together to generate safe paths. The results from a SLAM simulation where this method has been applied are analysed in Section 4 before the conclusions are drawn in the final section.

2 Maximum Expected Deformation

In our previous work, we introduced several deformation metrics [Whitty and Guivant, 2009] and showed how they could be used to plan paths efficiently during large-scale belief state deformation [Whitty and Guivant, 2011]. The deformation metrics were based on the analysis of pairwise distances between sample points in

the belief state. Thus, they only measured deformation at discrete points and even with stochastic information were only able to estimate the map uncertainty at those points.

However, it can be shown that from stochastic information about the locations of these discrete sample points, a probabilistic upper bound on the the rigidity of the belief state can be inferred. Not only is this true at discrete points, but it can be interpolated and extrapolated across the entire continuous belief state. This bound, which we term the Maximum Expected Deformation (MED), can be determined everywhere such that with further evolution of the belief state, local deformation will always be below that bound. The main requirement of this bound is that the density of landmarks is regular, in a similar fashion to the results from Bryson and Sukkarieh [2008]. Further details will be given in a forthcoming paper.

Figure 1 shows an example of how the MED varies across the belief state. Clusters of red indicate that landmarks are highly correlated and thus unlikely to deform locally, as at (1, -3). Hence in this area future deformation of the belief state is unlikely and therefore confidence in the positions of obstacles in this area is high. In the centre of the belief state, (6, 0), the local map uncertainty is inferred to be high and therefore the positions of any known obstacles in that area are likely to change and interfere with a planned path in that region. To guarantee that a planned path is safe, or does not obtain a higher traversal cost than initially calculated, the plan and the MED are linked in the novel approach presented below.

3 Relating Planning and Deformation

Given the traversal cost over the belief state, there are many planning algorithms [Jaillet *et al.*, 2010] which can generate paths between any two points. We focus on a simple example, spatially discretising the traversal costs into a rectangular grid, which we term a cost map, over which Dijkstra's algorithm is run to calculate the lowest cost path from any point to a goal point. The optimal policy, i.e. a mapping from state to action, is also defined by Dijkstra's algorithm, giving the optimal action to take from any state to reach the goal point.

Many methods for generating cost maps exist, and the focus of this paper is to manage deformation together with a cost map for planning, so from henceforth we define the costmap in a straightforward fashion. Simply, each cell is assigned a scalar cost value, for example the traversability of that cell or the risk of collision if that cell is reached. In this paper, the traversability is extracted from a single instance of a SLAM belief state for illustrative purposes.

If the cost can only take a binary value, we term this

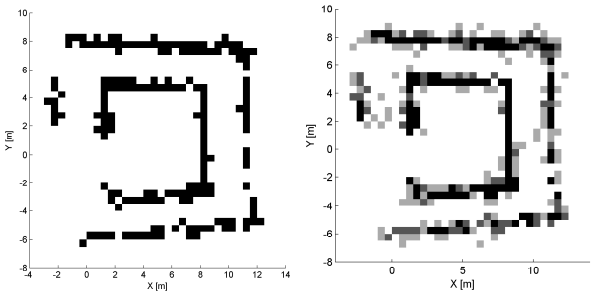

(a) A binary cost map, with obstacles in black.

(b) A non-binary cost map. Darker tones indicate higher total cost.

Figure 2: Cost maps on a regular grid.

a binary classification of space, as shown in Figure 2a, whereas a non-binary classification of space allows each cell to take any positive real value, as in Figure 2b. In addition, transitions between cells are assigned small positive values to prevent unnecessarily long paths. Therefore the concept of optimality is defined over the action set consisting of eight neighbours and seeks to minimise the total traversal cost from the summation of the cell costs and transition costs.

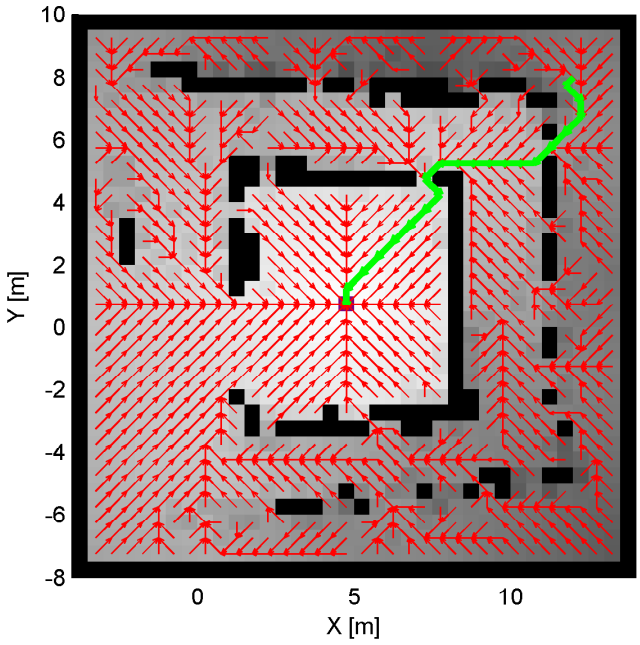

Figure 3: Global optimal policy, shown by red arrows, to reach destination at centre for the cost map shown in Figure 2b. A sample path is overlaid in green. Note in particular that the policy is optimal given the cost map and the action set. The cell colour indicates the total traversal cost to reach the destination, the summation of cell and transition costs.

Figure 3 then shows the resulting optimal policy to reach the destination at the centre by red lines, with the green line highlighting one complete path from source to destination. The greyscale cell colours are in proportion to the total traversal cost to reach the destination cell.

3.1 Convolution of Cost Map

Now, given the maximum expected deformation over the subsection of the belief state that is covered by the cost map, we convolute the cost map according to this deformation. The convolution is applied in an isotropic manner, where higher cost regions are expanded in all directions by a distance given by the magnitude of the deformation. This convoluted cost map combines the cost of traversal with the uncertainty derived from a stochastic mapping process.

The convolution acts as a grayscale morphological dilation, to borrow from image processing terminology, narrowing corridors and increasing the safety barrier around obstacles. This concept has frequently been applied in path planning to reduce the robot planning problem from considering the shape of the robot to one involving just a point in the map. However the key difference is in the use of a different kernel in different locations, reflecting the different magnitudes of expected deformation. It is important to observe that the dilation is not applied uniformly across the map, but varies with the expected deformation at every point. The concept of spatially adaptive morphological operations was expounded by Debayle and Pinoli [2005], although in the context of image processing. Given that the deformation is calculated on a discrete grid and indeed, dilation can only be applied to a discrete grid, we discretise the magnitude of deformation according to the grid resolution. Maximum expected deformation that is more than half a cell width is dilated by the full cell width, while deformation of more than 1.5 cell widths is dilated by two full cell widths etc. The dilation operator applies isotropic dilation as far as the discrete nature of the kernel allows.

Given a grid G of $U \times V$ cells, we define the neighbourhood of each cell, $N(i,j)$, as the set of neighbouring cells into which the contents of the current cell could deform. Given a 2D cost map, $c_1(i,j)$, also defined over the same grid, we convolute the cost map according to

$$c_{conv}(i,j) = \max_{\forall (u,v) \in \Phi(i,j)} (c_1(u,v)) \quad (1)$$
$$\Phi(i,j) = \{(u,v) : (i,j) \in N(u,v)\}.$$

The value of the convoluted cost, $c_{conv}(i,j)$ is taken as the maximum over all cells whose deformation can reach the point (i,j). In the isotropic case, the neighbourhood is defined as the set of all cells within a radius of R, i.e. $N(i,j) = \{(i+\Delta i, j+\Delta j) : \Delta i^2 + \Delta j^2 \leq R^2\}$.

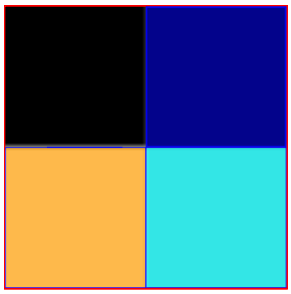

(a) Maximum expected deformation. (darker is higher)

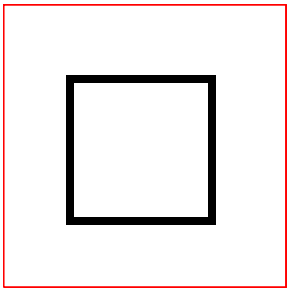

(b) Original cost map.

(c) Section of deformation less than 0.5 grid cells (black).

(d) Section of deformation between 0.5 and 1.5 grid cells (black).

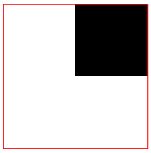

(e) Section of deformation between 1.5 and 2.5 grid cells (black).

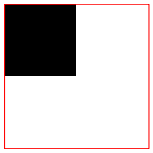

(f) Section of deformation between 2.5 and 3.5 grid cells (black).

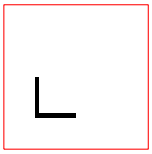

(g) Undilated section of cost map.

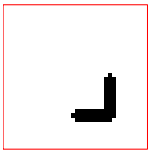

(h) Section dilated by 1 cell.

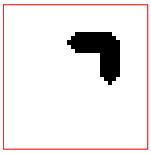

(i) Section dilated by 2 cells.

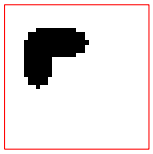

(j) Section dilated by 3 cells.

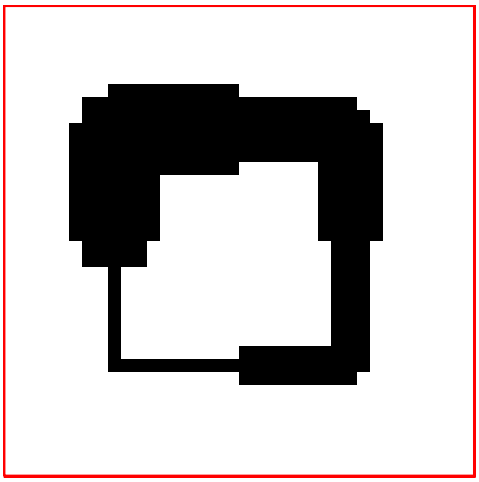

(k) Resulting cost map, convoluted by varying magnitudes according to maximum expected deformation.

Figure 4: The maximum expected deformation (a) provides the magnitude of dilation for every cell in the grid containing the cost map (f). The dilation magnitudes are discretised, in this case according to the number of cells the maximum expected deformation corresponds to, as shown in (c - e). For each magnitude, the corresponding sections of the cost map are dilated by that magnitude, as shown in (b - h). The convoluted cost map is calculated as the maximum cost in each cell from each of these dilated sections.

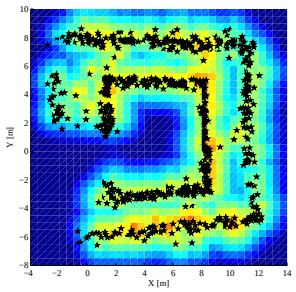

(a) Maximum expected deformation. (darker is higher)

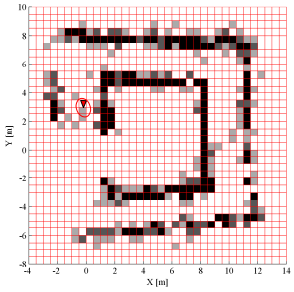

(b) Original cost map.

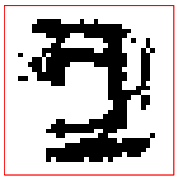

(c) Deformation less than 0.5 grid cells (black).

(d) Deformation between 0.5 and 1.5 grid cells (black).

(e) Deformation between 1.5 and 2.5 grid cells (black).

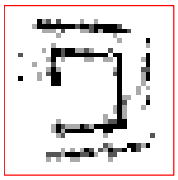

(f) Undilated section of cost map.

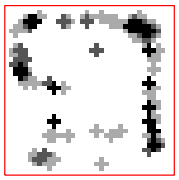

(g) Cost map sections dilated by 1 cell.

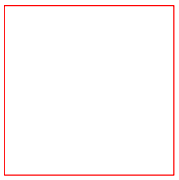

(h) Cost map sections dilated by 2 cells (blank).

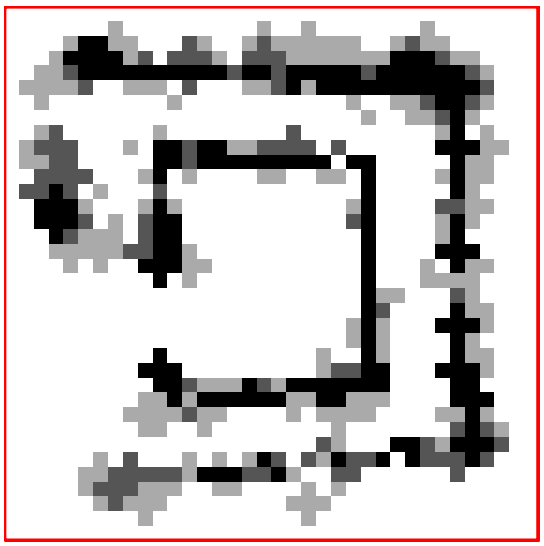

(i) Resulting cost map, convoluted by varying degrees according to maximum expected deformation. Compare with (c).

Figure 5: An example formatted similarly to Figure 4, here using a more realistic cost map and maximum expected deformation obtained by simulation and described in Section 4. The maximum expected deformation (a) provides the magnitude of dilation for every cell in the grid containing the cost map (b).

Since the maximum expected deformation is generally smooth, Φ could be replaced by N in Equation (1), giving

$$c_{conv}(i,j) = \max_{\forall(u,v)\in N(i,j)} (c_1(u,v)). \quad (2)$$

This form of the equation is slightly more intuitive, but assumes there are no large discontinuities in the maximum expected deformation.

Figure 4 illustrates the process on a simple cost map (b), using deformation magnitudes corresponding to 0, 1, 2 and 3 cells repectively as shown in (a), with orange being 0 and black being 3. Once the discrete dilation magnitude has been calculated for every cell in the cost map, the cost map is partitioned according to these magnitudes using a bit mask. The partitions are shown in (c - f) with the active section highlighted in black. The dilation operator with the appropriately sized kernel is applied to each section, giving (g - j). Finally, the resulting dilated sections are recombined by taking the maximum cost in each cell from each dilated section, giving the convoluted cost map shown in (k).

4 SLAM Simulation

Figure 5 illustrates the process on a cost map generated during a SLAM simulation. The robot traversed the simulated corridors in an anticlockwise direction. The sample points in this example correspond to the landmarks in a traditional EKF-SLAM formulation. At each sample point the current local deformation and its covariance are calculated. An upper bound on the expected deformation, called the Maximum Expected Deformation (MED), is then calculated and propagated over the full map. In (d) the cost map has been simply calculated by tallying the number of sample points in each grid cell, for want of better data to represent obstacles in the environment. The sample points and maximum expected deformation are plotted in (a), with red (lighter) representing low expected deformation –logically near the landmarks –and blue representing high expected deformation. Note that the MED is independent of global translations and rotations, so unlike the absolute position uncertainty, does not appear to increase linearly before loop closure.

The sections where the maximum expected deformation is less than half a cell are shown in black in (c). The majority of the map has a maximum expected deformation that requires one cell of dilation, given by the black area in (d). A very small component at the edge of the map requires dilation by two cells. The dilation operator is applied to each section, giving (f - h), where (f) shows the undilated section of the cost map for completeness. The section of the map dilated by two cells contained only cells of zero cost, so it did not contribute to the convoluted cost map, hence (h) is empty. Finally,

Figure 6: Dilating the entire cost map from Figure 5b by the maximum expected deformation in the map greatly overestimates the cost of traversal. The dark areas of high cost indicate the traversing the "corridors" is impossible.

the resulting dilated sections are recombined by taking the maximum cost in each cell from each dilated section, giving the convoluted cost map shown in (i).

If we had instead done a simple dilation by the largest deformation expected across the map, the resulting cost map would correspond to that shown in Figure 6. Clearly, traversal of the interior of the "corridors" is indicated as being impossible and this approach grossly exaggerates the traversal costs.

Careful observation of Figure 5 also shows that the cost map need not be limited to a binary classification of space as used in Figure 4, but can contain a range of discrete costs across the cells. A simple example of discrete costs is the case of a wheeled robot, where regions of mud may be of higher cost than concrete, but the principle can be applied to a wide range of environment properties. If we interpret the cost map values as being a direct measure of the consequences of traversing that cell (or "risk"), for example higher costs may imply a higher probability of collision with an obstacle, then by convoluting the cost map we are estimating the risk while inherently considering belief state uncertainty.

Consequently, given an acceptable limit on the level of risk, we have two alternatives: 1) Calculate the feasibility of a given path; or 2) Calculate the optimal (lowest cost) path that satisfies the risk constraint.

1) The first takes a given path, typically from the robot's current position to a goal point, and examines the risk at every state along that path. Integrating the risk over the path gives the feasibility of being able to safely traverse the path. If at any point along that path the individual risk exceeds a threshold, the path is marked as infeasible.

Figure 7 compares the paths generated from the non-convoluted and convoluted cost maps. Since the grayscale shade gives the total cost from each cell to reach the centre of the cost map, it can be immediately seen that the total cost of the path from the convoluted map is higher, since the cell at the start position (12, 8) is notably darker. As we are treating the cost map values as a risk, clearly there is more risk in moving from the start point to the goal once the expected deformation is considered.

2) The second treats all cost cells over a threshold (the risk threshold) as being impassable and runs Dijkstra's algorithm to generate a policy without traversing these cells. In Figure 7 these cells are drawn in black. The remaining costs are used to build the policy as usual, and since Dijkstra's algorithm gives the optimal solution, the paths generated will contain the route with the lowest risk, which is the optimal behaviour for taking safe action. Both paths shown in this figure are therefore optimal in terms of risk, however the second one is guaranteed to be safe providing the MED is correct, hence it is more robust. As long as the cells with risks over a threshold are removed, an optimality criterion from any other data source can be used to plan appropriate paths.

For increased efficiency, in the first case the convolution need only be applied near to the path to be examined. In the second case, the planning risk can also be used to adjust the frequency of calculating the MED, contributing to the overall efficiency of the system. Since changes in the expected deformation occur local to the robot, the efficiency of this approach could be improved by only performing the convolution local to the robot in the single robot case. If however, the cost map was altered in regions not proximate to the robot, or a multiple robot system was in use, further care would be needed.

5 Conclusion

We suggest that this is the first approach which explicitly relates the expected belief state deformation during SLAM with traversal costs. In Section 1 we outlined several alternative approaches to planning with uncertainty, but these either dealt with spaces that are classified as obstacles, free or unknown, ignoring traversal cost; or did not consider the type of deformation experienced during SLAM. For example Huang and Gupta [2008] calculated the collision likelihood at each node while planning a path using RRT, but this was based on localisation uncertainty and a binary occupancy grid, so no measure of the consequences or cost of collision was considered.

The method presented in this paper not only explicitly uses the expected deformation at every point in the map, but combines this with traversal costs which may be pro-

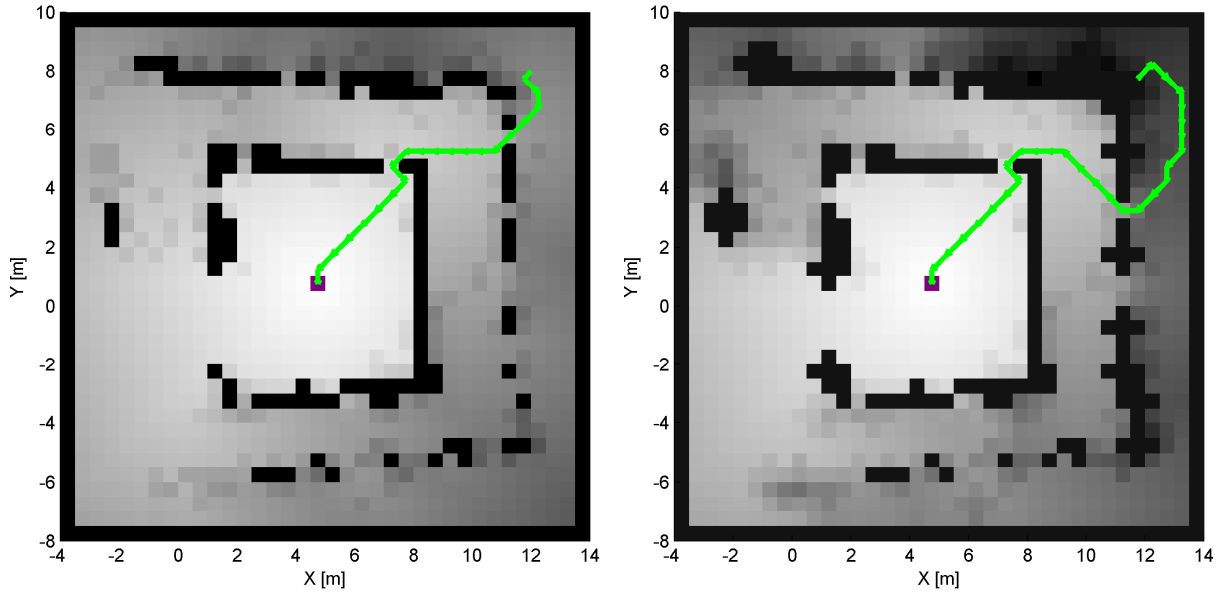

(a) Path generated on the non-convoluted cost map.

(b) Path generated on the convoluted cost map.

Figure 7: Comparing paths on the convoluted and non-convoluted cost maps. Obstacles are shown in black and other shades of gray show the total cost to reach the centre.

vided by sensors independent of those used for SLAM. Therefore, the method is able to generate guaranteed safe paths during the mapping process.

The directionality of expected deformation would also be valuable to investigate in future work, so non-isotropic convolution could be applied to reduce the estimated risk in some areas. This paper has also assumed that a global planning algorithm is feasible, and while appropriate for demonstrating its relationship with deformation metrics, integration with a hierarchical planning framework (such as [Whitty and Guivant, 2011]) is necessary to improve its scalability.

References

[Bourgault *et al.*, 2002] F. Bourgault, A. Makarenko, S. Williams, B. Grocholsky, and H. Durrant-Whyte. Information based adaptive robotic exploration. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 1, pages 540–545, Lausanne, Switzerland, 2002.

[Bryson and Sukkarieh, 2008] M. Bryson and S. Sukkarieh. Observability analysis and active control for airborne SLAM. *IEEE Transactions on Aerospace and Electronic Systems*, 44(1):261–280, 2008.

[Censi *et al.*, 2008] A. Censi, D. Calisi, A. De Luca, and G. Oriolo. A Bayesian framework for optimal motion planning with uncertainty. In *Proceedings of 2008 IEEE International Conference on Robotics and Automation*, pages 1798–1805, Pasadena, California, 2008.

[Debayle and Pinoli, 2005] J. Debayle and J.-C. Pinoli. Spatially adaptive morphological image filtering using intrinsic structuring elements. *Image Analysis & Stereology*, 24(3):145–158, November 2005.

[Fang *et al.*, 2005] G. Fang, G. Dissanayake, N. Kwok, and S. Huang. Near minimum time path planning for bearing-only localisation and mapping. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 850–855, Edmonton, Alberta, Canada, 2005.

[Fujimura, 1996] K. Fujimura. Path planning with multiple objectives. *Robotics & Automation Magazine, IEEE*, 3(1):33–38, 1996.

[Gonzalez and Stentz, 2009] J. P. Gonzalez and A. Stentz. Using linear landmarks for path planning with uncertainty in outdoor environments. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1203–1210, St Louis, Missouri, October 2009.

[Guibas *et al.*, 2009] L. J. Guibas, D. Hsu, H. Kurniawati, and E. Rehman. Bounded uncertainty roadmaps for path planning. In B. Siciliano, O. Khatib, F. Groen, G. S. Chirikjian, H. Choset, M. Morales, and T. Murphey, editors, *Algorithmic Foundation of Robotics VIII*, volume 57, pages 199–

215. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009.

[Huang and Gupta, 2008] Y. Huang and K. Gupta. RRT-SLAM for motion planning with motion and map uncertainty for robot exploration. In *Proceedings of 2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1077–1082, Nice, France, 2008.

[Huang *et al.*, 2004] S. Huang, Z. Wang, and G. Dissanayake. Time optimal robot motion control in simultaneous localization and map building (SLAM) problem. In *Proceedings of 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 3, pages 3110–3115, 2004.

[Jaillet and Porta, 2011] L. Jaillet and J. M. Porta. Path planning with loop closure constraints using an atlas-based RRT. In *Proceedings of 15th International Symposium on Robotics Research*, Flagstaff, AZ, August 2011.

[Jaillet *et al.*, 2010] L. Jaillet, J. Cortes, and T. Simeon. Sampling-based path planning on configuration-space costmaps. *IEEE Transactions on Robotics*, 26(4):635–646, 2010.

[Leung *et al.*, 2008] C. Leung, S. Huang, and G. Dissanayake. Active SLAM in structured environments. In *Proceedings of 2008 IEEE International Conference on Robotics and Automation*, pages 1898–1903, Pasadena, California, May 2008.

[Makarenko *et al.*, 2002] A. Makarenko, S. Williams, F. Bourgault, and H. Durrant-Whyte. An experiment in integrated exploration. In *Proceedings of 2002 IEEE/RSJ International Conference on Intelligent Robots and Systems.*, volume 1, pages 534–539, Lausanne, Switzerland, 2002.

[Rocha *et al.*, 2005] R. Rocha, J. Dias, and A. Carvalho. Cooperative multi-robot systems: A study of vision-based 3-d mapping using information theory. *Robotics and Autonomous Systems*, 53(3-4):282–311, December 2005.

[Thrun, 1998] S. Thrun. Learning metric-topological maps for indoor mobile robot navigation. *Artificial Intelligence*, 99(1):21–71, February 1998.

[Whitty and Guivant, 2009] M. Whitty and J. Guivant. Efficient path planning in deformable maps. In *IEEE/RSJ International Conference on Intelligent Robots and Systems.*, pages 5401–5406, St Louis, Missouri, 2009.

[Whitty and Guivant, 2011] M. Whitty and J. Guivant. Efficient global path planning during dense map deformation. In *Proceedings of the 2011 IEEE International Conference on Robotics and Automation*, pages 4943–4949, Shanghai, China, May 2011.

[Yamauchi *et al.*, 1998] B. Yamauchi, A. Schultz, and W. Adams. Mobile robot exploration and map-building with continuous localization. In *Proceedings of the 1998 IEEE International Conference on Robotics and Automation*, volume 4, pages 3715–3720 vol.4, Leuven, Belgium, 1998.