

Path Planning for Mine Countermeasures Command and Control

Christine D. Piatko^{*a}, Carey Priebe^{**b}, Lenore Cowen^{**b}, I-Jeng Wang^{***a}, Paul McNamee^{***a}

^aThe Johns Hopkins University Applied Physics Laboratory

^bDepartment of Mathematical Sciences, The Johns Hopkins University

ABSTRACT

The future success of Navy-Marine Corps operations in the extended littoral battlespace will depend critically on organic mine countermeasure capabilities. A battlegroup commander will require tools to rapidly detect, classify, and identify mines and form a tactical picture of mined areas, so a decision can be made to punch through the minefield, avoid it, or wait for dedicated mine countermeasures forces to clear it. We introduce here a command and control framework for mine countermeasures based on probabilistic classification and multicriteria path planning under uncertainty. Data from probabilistic mine classifiers can be used by a path planning tool to generate information comparing the relative utility and risk associated with different routes through a minefield. We are developing a dynamic path planning tool that can be adjusted to manage safety versus time to objective, constructing alternate paths of varying time and risk through a minefield. It will evaluate alternative routes while searching for low risk paths. A risk management framework can be used to describe the relative values of such different factors as risk versus time to objective, giving the capability to balance path safety against other mission objectives. We are also adding visualization techniques to display mines with uncertain locations and highlight alternative routes through minefields and obstacles. Our goal is to present a commander the most useful tactical picture of mined areas for decision-making.

Keywords: mines, path planning, risk management, optimization, visualization

1. INTRODUCTION

Mine and minefield data is probabilistic due to many factors such as position uncertainty, detection probabilities, classification confidence levels, and classification errors. In the face of this uncertainty, nonetheless, a commander needs a meaningful way to measure risk in order to access the best way to cross a potential minefield, or decide that the risk to cross is too great.

Path planning is a well-studied problem in the setting where there is a single criterion to optimize. It is well known how to find the path of shortest distance, and well known how to find the path of least cost (risk). However, to optimize both criteria at once is a challenging problem. Instead of the single path quality of length/time-to-objective, we want to include other optimization criteria that describe the desirability of paths such as level of risk. In addition, it is desirable for command and control to have alternative routes presented to help make a decision. Therefore, besides a single optimal path we would like to find several choices of paths to display as alternatives. To accomplish this we consider the use of multicriteria path planning techniques as well as automatic and interactive ways to display alternative paths. By incorporating available geosituational information with the mine detections and potential routes, an integrated picture of the tactical scene can be presented.

2. PATH QUALITY CRITERIA

In order to allow reasonable trade-offs in paths we identified the following alternative path criteria that may be desirable in a mine countermeasures setting.

* Christine.Piatko@jhuapl.edu; phone 1 443 778-6584; fax 1 443 778-6904; Johns Hopkins University Applied Physics Laboratory, 11100 Johns Hopkins Road Laurel, Maryland 20723-6099;

** cep@jhu.edu and cowen@mts.jhu.edu; phone 1 410 516-7200 and 7043; fax 410 516-7459; The Johns Hopkins University, Department of Mathematical Sciences, Baltimore, Maryland 21218;

*** I-Jeng.Wang@jhuapl.edu and Paul.McNamee@jhuapl.edu; phone 1 443 778-6204 and 3816; fax 1 443 778-6904; The Johns Hopkins University Applied Physics Laboratory, 11100 Johns Hopkins Road, Laurel, Maryland 20723-6099.

2.1 Desirable path quality attributes

The most common path quality attribute is path length or time to goal, as a model of the time required to achieve a successful landing. Another important quality that we model is the probability of a successful landing, given the path risk from mines, which we describe in more detail below.

Certainly there are many other potential desirable path qualities. We may wish to have a path that does not make too many turns, since navigational constraints require minimizing course changes. We may also desire stable low-cost paths, paths for which risk increases only marginally if there is a slight deviation from an optimal path, since there may be navigation errors. Potential threats from the shore may pose additional risk that should be modeled in ways similar to mine risk.

It may also be desirable to incorporate the value of goal points to model the utility of the landing location. Currently we have developed our risk management framework and algorithms to support path length and mine risk, and it can be extended to incorporate shore threats and landing area utility as well.

2.2 Modeling point and path risk

Information about the location of mines and obstacles will always be uncertain. Sensor readings are inherently noisy and inaccurate and thus contribute to uncertainty. Furthermore, information from different sensors can have different probabilistic characterizations due to the unique characteristics of each sensor and its environment. Automated post-processing also contributes errors to the predicted locations. Appropriate probabilistic mine detection risk models can also be used for our risk management tradeoff framework. Available actuation and damage curve data for various sea mines could also be used as appropriate risk probabilities and be incorporated into our risk model.

To model uncertainties as well as classification posteriors and damage probabilities, each detected mine or obstacle location should be represented by a probability distribution rather than a single point or binary value. The corresponding shortest path problem becomes path planning under uncertainty.

An appropriate mathematically based risk model should combine different sources of risk uncertainty such as mine position uncertainty, classification uncertainty, and risk of being killed. Clearly the full treatment, modeling all relevant uncertainties and risk, requires a complex model. In this paper we use simplifying assumptions to develop an estimate of the risk associated with each path. For example, we can use a $1/r$ risk expression as a crude approximation to represent position uncertainty. Clearly, other (possibly even asymmetric) distributions are possible in a multi-sensor environment. By separating the risk module from the path planning module, we can work with a variety of probabilistic models of risk.

Here, we use a discretized model of point risk for the probability of being killed at each location. The point risk at point x (probability of being killed at point x) is $R(x) = 1 - \prod_i (1 - R_i(x))$, where $R_i(x)$ is the probability of getting killed by mine i

at point x . This $R(x)$ model incorporates spatial uncertainty in the actual location of the mine. We model this as an inverse function of distance from x to the mines, with a cutoff assumed of the range of influences of the mines. We can also scale this risk by the posterior of the classifier used to detect the mine. It can also incorporate a damage model that describes the conditional risk of location x . If a mine has been identified, actual probability of actuation and damage curves could be used.

Given this description of the risk at a point, we describe the path risk (risk of being killed following the path) as the following. For a path $P = x_1, x_2, \dots, x_n$, the risk of path P , $R(P)$, is $1 - ((1 - R(x_1))(1 - R(x_2)) \dots (1 - R(x_n)))$. By properly defining point risk, this can be transformed to a related linear function, computable using a simple sum of the point risks.

3. MULTICRITERIA PATH PLANNING

In a mine countermeasures environment, there is a fundamental trade-off between path length (based on either distance or time) and path safety. While the shortest distance path may be very risky, the commander may sometimes be willing to trade off that risk even with a high probability of failure if it could reduce the time needed to achieve a particular mission objective. Multicriteria optimization can incorporate these and other subjective criteria. Multicriteria optimization can be used to find paths that have good path quality in several criteria. Recent work in path planning for mine avoidance by autonomous

underwater vehicles (Hyland and Taylor 1993) did not consider the use of such multiple criteria, although it did incorporate some restrictions on acceptable paths such as limiting turn radius.

Multicriteria optimization problems tend to be hard. The classical approach is to develop techniques that convert the multicriteria problem into a related single criterion optimization problem. For example, if one wishes to optimize with respect to two objectives, one can form a linear combination of them, such as $\alpha \bullet \text{length} + \beta \bullet \text{risk}$, and optimize with respect to this new objective. The resulting path problem is often solvable in polynomial time. Parameters α and β give flexibility to allow trade-offs between the two different optimization criteria. This approach can be extended to allow combinations of more criteria.

A risk averse path would go totally around the minefield. If we were ignoring risk we would just take a straight-line path to the goal. Following the edges of an appropriately weighted Voronoi diagram would yield the safest path through the minefield, maintaining a maximal distance from mines. In our work, we are seeking to find intermediate solutions as well.

3.1 Path planning approximations using a grid graph

Much is known about exact algorithms and heuristics for a variety of shortest path planning problems in the simple binary setting where there is free space and obstacles that must be completely avoided. Much of the work in computational geometry has focused on obstacles represented by polygons or curves (Piatko 1993). This type of map can also be represented by a digitized binary grid, where each pixel is 0 or 1, corresponding to free space or obstacle. Using this grid, one can build a graph connecting centers of adjacent free pixels with appropriate edge lengths. The graph can be searched to locate a digitized approximation to the true shortest path. Typical methods apply the A* algorithm or Dijkstra's algorithm to find the shortest path in this grid graph (Korf 1999). The grid graph approach introduces a digitization bias, and thus does not find the exact shortest path, but obtains a reasonable approximation.

In the mine countermeasures setting this grid can be generalized to take into account the risk model for the mines described above. In this case, directed edges that connect centers of adjacent pixels are additionally labeled with the point risk of their center. Then, when path risk values are desired, instead of using distance as an optimization criterion, path weights can be computed by appropriately combining risk probabilities along candidate paths. One can then search for the "safest" or lowest risk paths as well as shortest time paths. The total risk found along the path in this way will be an approximate expression of the true integral risk along the path.

Shortest grid graph paths may not be optimal, but in many cases they can be locally improved to obey local optimality constraints (such as Snell's Law of refraction), and they can provide reasonable approximations in practice.

3.2 Linear combination approach

Our first approach is to consider the use of a linear combination of criteria. In this case, we seek to find the minimal path subject to measuring path quality using a single metric, $\alpha \bullet \text{length} + \beta \bullet \text{risk}$. This technique is similar to one used in (Mobasser 1989).

Consider a grid, with risk values r_i associated with each grid point i and lengths l_{ij} representing the estimated length/time to traverse between two adjacent gridpoints i and j . Given a starting location s and a set $z_1 \dots z_m$ of desired destinations, the paths p_k to z_k that minimize the values $O(p_k) = \alpha \sum_{ij \in p_k} l_{ij} + \beta \sum_{i \in p_k} r_i$ can be computed optimally in polynomial time using Dijkstra's algorithm for fixed α and β .

Our initial implementation is based on A* rather than Dijkstra. In this case we use the $\alpha \bullet \text{length} + \beta \bullet \text{risk}$ metric as our optimization objective for A* to find a path from a start to a goal. We can allow the user to vary α and β , or precompute "interesting" settings of α and β . The goal is to present alternative paths for different values of α and β . To precompute values we can use a fixed number of settings (for example, 10), and use $\beta = 1 - \alpha$ to display them as initial alternatives. The user could then refine from there (for example using binary search on α and β values). We are also considering generating such paths automatically using an implementation of a k shortest paths algorithm (Hadjiconstantinou and Christofides 1999). It would also be desirable to compute critical values of α and β automatically.

3.3 Constraint-based approach using dynamic graphs

An alternative approach to solving the multicriteria problem is to minimize one criterion subject to hard bounds on the others. More particularly, rather than consider the sum of risks on a path, consider the risk of a path to be instead the maximum risk R at any point along the path. This approach allows a commander to set a “maximum acceptable risk” and consider only paths that avoid regions of risk higher than the threshold. In a command and control setting, we can output several path choices, from a “short/high-risk” to a “long/low-risk” alternative by varying this risk threshold R .

Consider providing a hard “maximum acceptable risk” bound R on the allowable acceptable risk along the length of the path. We want to find the path that minimizes length/time-to-goal subject to this constraint R . Again, we use a grid graph approach that does not find the exact shortest path, but obtains a reasonable approximation. In the absence of a risk threshold we would just take a shortest path to the goal. For very small values of R , it is possible that no path would be feasible. For example we may consider there is always a background risk due to potentially undetected mines. As the risk threshold R is varied, one is considering finding paths through subgraphs of the grid graph described above (where only distance values are used to measure path quality). When R is larger than the point risk associated with any point in the region, the full grid graph is searched. As R is decreased, edges drop out of the graph because they violate the risk threshold constraint. Thus, different values of R generate different subgraphs of the full grid graph. At some point the grid graph may become disconnected and thus no acceptable path can be found.

We are exploring this constraint-based approach initially by using a variant of our A* implementation for a single fixed risk threshold R . In this case, we use minimizing length/time-to-goal as our objective function, and do not expand A* search nodes when they violate the user-specified constraint R .

This approach can be extended to allow a slider to vary this maximum acceptable risk threshold R interactively. This will dynamically present many alternative paths. We compute the different paths automatically by solving shortest path subproblems. Let E be the set of edges in the grid graph defined previously. First, we sort these edges by their risks, from least risk to most risk. Next, we consider adding edges from E in sorted order to obtain a subgraph S . We continue adding edges until we obtain a connected graph from the start to the destination(s). We then find the shortest path in this corresponding subgraph S . This is the first candidate path corresponding to the smallest possible allowable risk R . We then add the next edge, which corresponds to increasing the risk threshold R . We then compute the new shortest path. By continuing similarly we can generate many candidate paths. We may consider bucketing edges, since some paths may be very similar. They may be topologically equivalent in the way they thread through the minefield, and varying only slightly in how close they pass to various mines.

By greedily adding edges in the sorted risk (possibly in bucketed fashion), we automatically generate candidate paths for all values of R in the grid graph. Dijkstra’s algorithm can be used to recompute the shortest paths for each new subgraph S . However, recent theoretical advances in dynamic graph algorithms imply that a graphical control could be implemented which would allow such changes at interactive rates. These fast dynamic graph algorithms allow interactive changes to graphs through incremental edge insertions and deletions (King 2000). They maintain a special data structure, so that as edges are inserted or deleted from the graph properties such as the current shortest path can be updated quickly (incrementally), without resorting to recomputing answers from scratch each time. Thus, as a slider varies R , the new shortest path for that threshold can be dynamically updated and displayed at interactive rates.

3.4 Other approaches

The problem of automatically discovering various alternative paths is related to risk management for the transportation of hazardous materials (Akgun, et. Al 2000; Erkut and Verter 1998). This problem is a variation that is almost a dual to our problem, where the paths themselves might cause risk to surrounding regions. In this case, the goal is to present various alternative paths that do not endanger large population centers. The alternative paths can thus be used to allow choices based on other criteria. Typically, these problems have been solved using heuristic techniques.

We have also been investigating the potential use of k -shortest path algorithms to enumerate several path choices. There has been recent work in practical implementations to generate k simple shortest paths (Hadjiconstantinou and Christofides 1999). Ultimately, our current grid approach is seeking to estimate a weighted probabilistic Voronoi diagram of the mines. Eventually, we plan to model this by using generalized Voronoi diagrams or weighted regions for planning.

Single start point to single goal point path planning does not adequately address the command and control problem, so we are also considering segment goals. Finally, it would be useful to directly address the issue of which mines are best to clear by considering methods to iteratively find which mine or groups of mines would make the best improvements in path quality, thus indicating where it would be most valuable to expend mine-clearing resources.

4. VISUALIZATION AND INTERACTION APPROACHES

We have been developing several alternative visualization and interaction schemes together with our path-planning code in order to show the relative risks of path choices and minefield danger probabilities. We are developing this prototype interface in Java. Java is multiplatform, portable, and suitable for a distributed systems environment. Java also allows for portable interface development.

Existing decision aid tools can create geosituational displays from terrain and bathymetric databases. Such map underlays will be useful to give the context of the minefield danger areas. We do not seek to reimplement such techniques in this work.

Instead, we are concentrating on prototype interaction and visualization techniques that will be useful for interacting with a multicriteria path planning tool. We are also implementing prototype visualization techniques to display path choices and mine danger regions. The essential functionality needed includes designating start point(s) and goal point(s). We have also implemented intermediate way points. Such way points have the profound advantage that they can reduce the search time for A* through “divide and conquer,” making planning faster. We are also considering different techniques such as overlays to visualize the various choices of routes. In order to display several choices of paths, they can be displayed and overlaid at once or presented separately, for example in tabbed lists.

Other Graphical User Interface (GUI) controls are needed for the parameters of the underlying path planning algorithms, such as setting α and β factors. Sliders or text boxes can be used for these parameters as well as to select the risk constraints that can generate path choices. One may also want to keep a history of selected interesting paths.

Risk maps of the risk model grid associated with mines can be displayed using techniques common to scientific visualization, such as color scales. A grayscale version of one risk model is shown in Figure 1.

In addition to schemes above, it may be desirable to color paths along their routes to highlight the most risky points. This type of display could also prove useful for non-automated path planning to highlight riskiest areas of hand-planned routes.

5. EXAMPLE

As a concrete example to test our concepts, we have used processed COBRA ATD (Coastal Battlefield Reconnaissance and Analysis Program Advanced Technology Demonstration) contacts (Priebe, et al. 1999; Witherspoon, et al. 1995) and incorporated these contacts into our A* prototype. While this is beach minefield data, this approach can also be applied to similar sea minefield data.

Minefield detection and localization is an important problem currently receiving much attention in the engineering and scientific literature (Smith 1995; Witherspoon, et al. 1995). Witherspoon, et al. (1995) describes the operational concept for minefield reconnaissance via an unmanned aerial vehicle. Multispectral imagery of an area of interest is processed. Potential mines are located with a mine detection algorithm (Holmes, et al. (1995)). Classifiers can then be used to reduce false alarm rates.

For example, Priebe and Cowen have demonstrated the effectiveness of a generalized Wilcoxon-Mann-Whitney (GWMW) test statistic as a classifier on an existing multispectral land minefield dataset, the COBRA ATD image database (Priebe and Cowen 1998). The COBRA ATD collected this dataset and made it available through the NSWC Coastal Systems Station, Dahlgren Division, Panama City Florida. It was collected by an UAV using two side-by-side multispectral sensors, each with six different spectral bands. Feature extraction using a coefficient of variation was performed to find probable mine locations. Alternative statistics could be used for such feature extraction. The GWMW classifier showed promise in reducing the false alarm rate on this dataset (Priebe and Cowen 1998). Classifiers can also generate posterior probabilities on the detection points (Priebe, et al. 1999; Priebe, et al. 2001).

While we demonstrate our approach on this data, it can also be used on sea mine data as well as minefield (versus individual mines) data. Simulated random minefield data with non-mine bottom objects (NOMBOS) can also be used.

5.1 Linear Combination Approach

Our initial results with the COBRA ATD data demonstrate several different paths found for user-specified choices of α and β for our linear combination approach. See Figure 2. In this example we see the effect of choosing different α and β on the paths. Notice there are some digitization effects, since we are using a grid as an approximation to the optimal path.

5.2 Constraint Based Approach

We also considered the constraint-based approach using the COBRA ATD data. In Figure 3, we present some sample paths for several user-specified bounds on acceptable risk. Notice these paths are somewhat different from those found by the linear combination approach.

5.3 Issues

We are observing issues related to discretization of the search space as well as rolling length and risk into a single metric in the quality of paths found. We are exploring different discretization strategies and path quality metrics to address these issues. Faster algorithms may be possible which directly find such approximate paths. Adding directions (in this case we use 8 directions around each point) as well as finer or hexagonal grids could also be used to improve path quality.

6. CONCLUSIONS

The techniques we are developing seek to demonstrate the power of probabilistic classification, path planning and visualization for organic mine countermeasures command and control. Multicriteria path planning under uncertainty shows promise as a decision analysis tool. We have begun our investigations with readily available data as proof of concept. While these initial experiments use landmine data, our methods can also be applied to sea minefield data. We will be developing randomized data for analysis as well.

In the future we will also be developing methods to incorporate other desired risk management features, such as different relative utilities for different goal locations. Also, additional sensing as well as threat neutralization may be possible en route. Dynamic planning may be of interest to incorporate additional available sensor information or neutralization information in an on-line fashion. We will consider different techniques to support such on-line re-planning efficiently.

ACKNOWLEDGMENTS

This work is partially supported by Office of Naval Research Contract N00024-98-D-8124. Minefield data was provided by NSWC Coastal Systems Station, Dahlgren Division, Panama City, Florida, with support from United States Marine Corps Amphibious Warfare Technology. The authors thank Fernando Pineda of JHU/APL for his help developing the risk model.

REFERENCES

1. V. Akgun, E. Erkut, and R. Batta, "On finding dissimilar paths," *European Journal of Operational Research* 121, 232-246, 2000.
2. E. Erkut and V. Verter, "Modeling of Transport Risk for Hazardous Materials," *Operations Research*, Vol. 46 No. 5, Sept-Oct. 1998.
3. E. Hadjiconstantinou and N. Christofides, "An efficient implementation of an algorithm for finding k shortest paths," *Networks* 34:88-101, 1999.
4. Q. A. Holmes, C. R. Schwartz, H. J. Seldin, J. A. Wright, and L. J. Wieter, "Adaptive Multispectral CFAR Detection of Land Mines," *Proceedings of the SPIE Volume 2496: Detection Technologies for Mines and Minelike Targets*, Orlando, Florida, April 1995, pp. 421-432, 1995.
5. J. C. Hyland and F. J. Taylor. Mine avoidance techniques for underwater vehicles. *IEEE Journal of Oceanic Engineering*, 18(3):340--349, July 1993.

6. V. King, Fully Dynamic Algorithms for Maintaining All-Pairs Shortest Paths and Transitive Closure in Digraphs, Manuscript, 2000.
7. R. E. Korf, "Artificial Intelligence Search Algorithms," in Algorithms and Theory of Computation Handbook, M. J. Attallah, ed., CRC Press, 1999.
8. B. G. Mobasseri, "Path planning under uncertainty from a decision analytic perspective," In Proc. IEEE International Symposium on Intelligent Control, pp. 556--560, 1989.
9. C. D. Piatko. Geometric Bicriteria Optimal Path Problems . PhD thesis, Computer Science Department, Cornell University, 1993.
10. C. E. Priebe and L. J. Cowen. Mine detection via generalized Wilcoxon-Mann-Whitney classification. In Proceedings of the, SPIE Vol. 3392, Detection and Remediation Technologies for Mines and Mine-like Targets III, pp. 906-917, 1998.
11. C.E. Priebe, D.Q. Naiman, and L. Cope, "Importance Sampling for Spatial Scan Analysis: Computing Scan Statistic p-Values for Marked Point Processes," Computational Statistics and Data Analysis, Vol. 35, No. 4, pp. 475--485, 2001.
12. C.E. Priebe, J.-S. Pang and T. Olson, "Optimizing Sensor Fusion for Classification Performance," Proceedings of the International Conference on Imaging Science, Systems, and Technology: CISST '99, pp. 397-403, 1999.
13. D. L. Smith, "Detection Technologies for Mines and Minelike Targets," Proceedings of the SPIE Volume 2496: Detection Technologies for Mines and Minelike Targets, Orlando, Florida, April 1995, pp. 404-408, 1995.
14. Witherspoon, N.H., Holloway, J.H., Davis, K.S., Miller, R.W. and Dubey, A.C., "The Coastal Battlefield Reconnaissance and Analysis (COBRA) Program for Minefield Detection," Proceedings of the SPIE Volume 2496: Detection Technologies for Mines and Minelike Targets, Orlando, Florida, April 1995, pp. 500-508, 1995.




Figure 1. An example (grayscale) visualization of a risk field.




Figure 2. Sample paths with linear combination of criteria. Ground truth is known for this data. True mines are dark circles, false alarms are white circles. Black outlines indicate those detections with posterior ≥ 0.5 , grey outlines those with posterior < 0.5 that would have been eliminated as false alarms by the classifier.




Figure 3. Sample paths using constraint bound on risk approach.