

Smart Radiation Sensor Management

Radiation Search and Mapping using Mobile Robots

R.A. Cortez, X. Papageorgiou, H.G. Tanner, A.V. Klimenko, K.N. Borozdin,
R. Lumia, and W.C. Priedhorsky

Index Terms—Nuclear search, radiation mapping

MOTIVATION

The current geopolitical situation requires automated tools for quick and effective assessment of threats. Modern threats are subtle and ephemeral, and can be hidden across large areas. Classical information extraction methods, where data is randomly collected and then subsequently filtered and analyzed by human operators in search of particular signatures, are no longer effective against today's modern threats. Data collection must be *guided* by querying world models that afford the span and resolution needed for multi-scale problems.

Currently, searching for radiation sources is usually done manually, by operators waving radiation counters in front of them as they walk. This method does not provide any visual or statistical data map of the area in question. To quickly characterize the severity of the situation, an efficient way of obtaining this radiation map is needed. When searching for a weak radiation source, a speck of uranium for example, manual methods are unlikely to yield results.

In nuclear search, the strength of the signal relative to noise (SNR) falls with the *square* of the distance R to the source, as the latter increases. The relation between SNR and distance motivates bringing the sensor as close to the source as possible [1]. Mobile robots can carry sensors close to the source, and position them accurately for required measurement collection. Using traditional sequential testing theory we can only confirm the presence of a source of a particular strength at a given location. For locations where these specific nuclear signatures are not detected, no information is given regarding the local radiation levels. A different approach is therefore needed if the objective is to map the radiation intensity over a certain area.

In this article we suggest two different motion planning strategies for radiation map building. The first, named the *gradient-based Bayesian method*, an uncertainty metric is used to define a potential function with which to bias the search towards particular areas of the map where uncertainty regarding radiation levels is highest. The second strategy, named the *sequential-based Bayesian method*, the robot visits every area cell along a pre-determined path, and the time it spends at each cell depends on the local uncertainty levels.

The sequential-based Bayesian method ensures that each cell is only visited once, and thus it is time-optimal. However, due to the motion plan of the sensor being pre-determined, parts of the map that could be potentially the most interesting could be revealed last. In addition, this method is

not suitable when the prior is time-varying, that is, in the case of dynamic environments, since areas explored once are not revisited. The gradient-based Bayesian method, on the other hand, offers an approximate map of varying confidence at every time step, but it requires longer time for the completion of the map. However, the method outperforms the sequential-based Bayesian mapping *in the initial stages* of the area scanning, suggesting that when time constraints are imposed that will not allow the sequential-based method to terminate, a better map can be obtained with the gradient-based method. In addition, the gradient-based Bayesian method can accommodate real-time changes in the environment, through an on-line adaptation of the function that generates potential field.

The methods described in this article are not only suited to applications of “nuclear forensics,” where we need to determine in the least possible time, and at a given probability of a false positive, whether fissile material has been processed in a given area; they are also applicable to the problem of assessing the contamination due to accidental or malicious release of radioactive isotopes. As a result of a radiation map, decision makers can single out safe from unsafe regions and quantify contamination as a first step towards containment and cleanup.

WHAT DOES A MAP SHOW ?

There are significant differences between mapping walls and door locations in an office environment, and mapping the temperature distribution in the same space. First, robot localization is typically linked to map building in the first case, where in the second is not. Knowing where a measurement is taken is of paramount importance, but the two problems (localization and temperature mapping) are not linked. The second difference has to do with the underlying statistics: range measurements are typically associated with Gaussian distributions; measuring distributed quantities such as temperature, pressure, or radiation level, may obey different statistical laws, that result from the type of sensors used as well as the nature of the underlying physical process. For this reason we divide this section in three parts. The first refers to existing approaches to robot exploration, and is conceptually related to our mapping problem because our robots essentially “explore” radiation distributions. The second focuses on expressing the spatial distribution of physical quantities, and the third part specializes the discussion on radiation mapping.

Exploring the world

Robot exploration typically involves creating a map of the known workspace which depicts the location of ob-

stacles and landmarks. The process is often based on an occupancy grid, which discretize the area of interest into a large number of cells. The notion of grid maps or occupancy maps was first introduced to the area of mobile robotics in [2] and subsequently used in [3], [4]. In most cases the cells of a grid map contain a probability value of whether that cell is occupied. Yamauchi [3] uses occupancy grids to define a new *frontier* for the robot to investigate to expand the knowledge of the environment. Romero *et al.* [4] uses occupancy grids to minimize the cost to travel to an unoccupied cell for further investigation of the area. In this article, we use each cell of the grid map to hold a metric of the uncertainty regarding the radiation levels in that region. Our metric of uncertainty is the variance of a particular distribution over the expected radiation level at each spatial location.

Linking robot motion to uncertainty is not an entirely new concept, but is lately gaining momentum in robotic exploration, localization and mapping [5], [6]. Moorehead [5] uses the entropy, among other utility measures, to evaluate the benefit of visiting different locations. The entropy in [5] describes the uncertainty over a certain location being reachable; it is not directly associated with the quality of the model nor is it linked with the statistics of the measurements. In [6] the problem is to facilitate mobile robots in localizing target features in their environment, and mutual information is used as a metric of significance of different discrete locations in terms of sensing. Of course, the statistics for the detection problems addressed there are Gaussian, but the concept is nevertheless similar.

Mapping spatial distributions

Work on mapping the distribution of physical quantities over a region (e.g. gas concentration, temperature) is related to our approach in the sense that the map constructed is not related to the topology of the environment, nor does it include landmarks or other location identifiers. In [7], maps of gas concentrations are constructed, by maneuvering a robot using a predefined path that covers the entire area. An approach to search for ocean features is found in [8], in which multiple robots follow gradients to locate and track ocean features such as fronts and eddies. For the radiation mapping problem addressed here, sensor measurements at given locations are in theory random samples drawn from a Poisson distribution, and therefore vary widely making gradient calculations meaningless. Our approach is to follow gradients of uncertainty, rather than of measured radiation, and to steer the robot to locations where measurements make the most difference in situational awareness.

In [9] and [10] the focus is on efficiency. Kim and Hespanha [9] address the problem of determining paths for a group of unmanned combat air vehicles that cooperate in their use of jamming resources, such that the risk of being tracked and destroyed by surface-to-air missiles is minimized. The minimum risk path planning is reduced to a weighted anisotropic shortest-path problem. Bertucelli and How [10] propose an approach to calculate the minimum number of observations needed to achieve a given level of confidence for target existence in an uncertain grid-like environment.

The authors use a Beta distribution to model the imprecise knowledge of the prior probabilities in the individual cells. As in this article, [10] brings into play the variance of the distribution in order to strike a balance between the speed of observation and uncertainty on the existence of a target.

Mapping radiation

Sequential nuclear search allows us to quickly verify the existence of microscopic specks of radioactive material. When a gamma ray emitted from a source reaches a sensor, the latter is said to register a *count*. Radiation intensity is measured in counts per second, assuming that all emitted rays are detected and registered. Low-rate counting of radiation from nuclear decay is described by the Poisson statistics. Classical sequential testing theory [11] suggests the stopping rules, that is, when does one know with certain confidence that a source exists at a given point. These rules allow for rejection of certain sequences of observations at early stages. Either positive or negative identification can be made based on the likelihood ratio of the probability of observing a certain number of counts within some time period given that there exists a source emitting an average number of counts per unit time μ_s , over the probability of these counts corresponding to background radiation. The stopping rule is determined from the desired false negative and false alarm rates, P_{FN} and P_{FA} respectively,

$$C = \frac{P_{FN}}{1 - P_{FA}}, \quad A = \frac{1 - P_{FN}}{P_{FA}}. \quad (1)$$

The condition $\kappa_k \leq C$ rejects the hypothesis that the source is present at location κ_k , while the condition $\kappa_k \geq A$ confirms the presence of the source. When $C < \kappa_k < A$, longer exposure is required to make a decision. An example of the stopping rules is illustrated in Fig. 1, where the straight lines represent the limits of confidence intervals for certain conclusions: (i) when the radiation counts collected within a certain time period are in the upper region, the presence of a source is verified with a given confidence; (ii) if the counts fall in the lower region, then they are most likely due to background radiation; (iii) in between, no conclusion can be confidently drawn until more measurements are collected.

This strategy, however, is a detection strategy, not a mapping technique. To construct a contour radiation map using traditional sequential testing theory, one would have to scan the area for every single contour level. This approach is far from time-optimal, and the required completion time increases rapidly with the resolution of the desired map. There is also an important caveat here: the algorithm is sensitive to the assumed value of the strength of the source that is to be detected. If the source intensity is underestimated, the method will give a false negative by default, since the registered number of counts (triangles in Fig. 1) remains within the threshold boundaries for the whole time interval of 2.4 seconds.

A MODEL FOR THE DISTRIBUTION OF RADIATION

Natural gamma ray background radiation has a cosmic ray component, and a component from naturally occurring radioactive isotopes. Small detectors (such as a one cubic

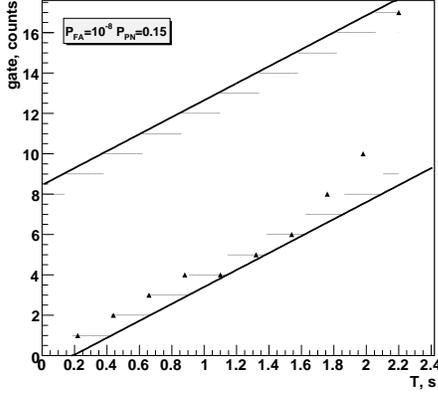


Figure 1. Applying sequential testing theory to a nuclear detection problem, involves calculating thresholds for a positive confirmation or rejection of the source hypothesis (from (1)). In the example depicted, the set of thin horizontal lines define count thresholds for positive (top) and negative (bottom) identification of a 10cts/s source within 1cts/s background. The bold solid lines are the linear fit to the conservative boundaries of the thresholds.

inch La_2Br scintillator of Fig. 5) typically record low count rates, and the probability of observing k counts, given a mean expected count rate λ , is well described by the Poisson distribution

$$P(X = k|\lambda) = \frac{\lambda^k e^{-\lambda}}{k!}. \quad (2)$$

We use a Gamma distribution for the initial estimate of the expected mean count rate λ

$$\pi(\lambda) = \beta^\gamma \lambda^{\gamma-1} e^{-\beta\lambda} \times \frac{1}{\Gamma(\gamma)},$$

where γ is the shape parameter, β is scale parameter and $\Gamma(\gamma) = \int_0^\infty t^{\gamma-1} e^{-t} dt$. The expected value and variance of the gamma distribution can be expressed in terms of its shape and scale parameters

$$\mathbb{E}(\Gamma) = \frac{\gamma}{\beta}, \quad V(\Gamma) = \frac{\gamma}{\beta^2}. \quad (3)$$

As new measurements are collected, the probability distribution of λ is updated using Bayes rule, in the form of the recursive formula

$$\pi(\lambda|X) = \frac{P(X = k|\lambda)\pi(\lambda)}{\int P(X = k|\lambda)\pi(\lambda)d\lambda}.$$

WHERE TO MOVE THE SENSOR

Area cell decomposition

Our main goal is to create a radiation field map of the area, at a given uncertainty level (variance at most V_0). We decompose the workspace to $m \times n$ cells, arranged in a two dimensional array, and indexed by i and j . In each cell we assume an *a-priori* radiation level, expressed in the form of an expected mean count rate λ_{ij} (Fig. 2). Being uncertain about this estimate, we assume that this mean count rate follows a Gamma distribution with mean $\lambda_{ij} = \frac{\gamma_{ij}}{\beta_{ij}}$, and variance $V_{ij} = \frac{\gamma_{ij}}{\beta_{ij}^2}$, according to (3).

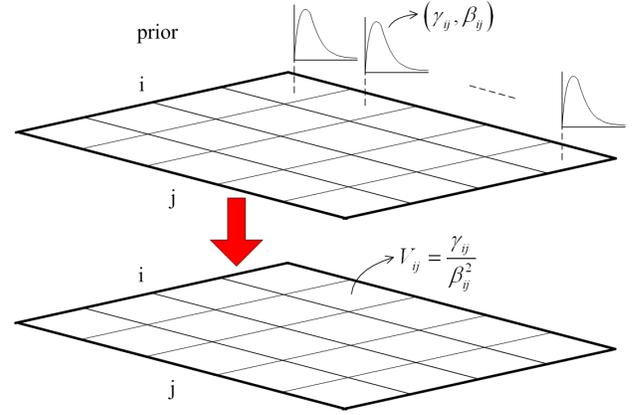


Figure 2. The robot's workspace is decomposed in $m \times n$ cells. The prior knowledge of this area is a Gamma distribution, with parameters γ_{ij} and β_{ij} in each cell. The mean value of the emission is λ_{ij} , and the variance is V_{ij} in any cell (i, j) of the grid.

Mapping sequentially

One method to move the robot and collect measurements is along the lines of sequential search: move from cell to cell when a statistically definitive conclusion can be drawn. The method we describe here, however, is not an instance of the traditional sequential search because we do not verify a hypothesis about the existence of a certain source, but rather we position the sensor at a given location for sufficient time to reduce the uncertainty over our radiation level estimate below a certain threshold. The similarities stop at motion planning. In this “hybrid” approach for implementing Bayesian-based radiation mapping, motion planning is done sequentially but the decision on next movement is based on Bayes rule. We call this type of strategy *sequential-based Bayesian search*.

In this method the robot stays in a cell and takes radiation measurements until the desired variance threshold is reached. Then, the robot moves to a neighboring cell along a certain direction and continues with mapping. The condition that enables the transition from cell (i, j) to, say $(i + 1, j)$, is $0 \leq V_{ij} \leq V_0$, where $V_0 \geq 0$ is the maximum acceptable variance. The mapping is completed when the robot has scanned every cell in the workspace.

Mapping using uncertainty gradients

Searching sequentially is an open loop strategy in the sense that the motion plan is predetermined and the radiation map is ready only after the whole area is scanned. Instead, we can close the loop on-line and drive the robot where measurements are more critical for reducing uncertainty. These places change as the robot moves around and more measurements are collected. We thus use measurements as feedback to determine motion, by means of an artificial potential field that is dynamically updated through Bayes rule. We construct a potential function of the form $\varphi = \frac{\gamma_d + 1}{e^{\beta_d^2/k}}$, where γ_d is the distance to the point where we want the robot to terminate the search, function β_o relates to the variance of each cell, and k is a positive tuning parameter.

The robot starts at cell (i, j) , where $i \in \{1, \dots, m\}$, $j \in \{1, \dots, n\}$. A prior map of distribution of the average

count rate over the area, in terms of parameters γ_{ij}, β_{ij} is assumed given. From this distribution we estimate the variance of the Gamma distribution that expresses λ_{ij} (Fig. 2) as $V_{ij} = \frac{\gamma_{ij}}{\beta_{ij}^2}$. Function β_o is then constructed as a strictly decreasing function of V_{ij} ; although many choices are possible, for simplicity we choose to set $\beta_{o_{ij}} = \frac{1}{V_{ij}}$. From cell (i, j) the robot moves to an adjacent cell after comparing the values of φ at neighboring cells (Fig. 3). To accelerate the search, instead of directly comparing variances in neighboring cells, we fit a smooth surface over the variance values of cells over the entire area. As a result, differences between adjacent cells are not skewed toward extreme values (zero or large positive and negative numbers), but are more uniformly distributed across the range of values. Using this method we accelerate the mapping process by 10 to 15%. If the robot chases the global

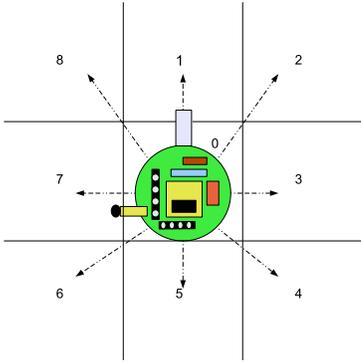


Figure 3. The proposed area is decomposed into cells creating an $m \times n$ grid. The Khepera II mobile robot is allowed to move to the eight neighboring cells or stay in its current cell depending on the calculated variance of the map, based on the measurements obtained from the radiation sensor.

variance maximum efficiency suffers. During both simulation and experimental tests of the gradient-based Bayesian mapping method, we observed a sharp increase in completion time due to the fact that such maxima can appear in very distant regions of the area to be mapped. The robot travels back and forth, and the increased length of the robot's path translates directly to increased completion time.

The indices of the neighboring cell which the robot moves to are given as the solution to

$$\max \arg \varphi_{pq}, \quad i-1 \leq p \leq i+1, \quad j-1 \leq q \leq j+1. \quad (4)$$

If the radiation counts registered within the following time step are $x \in \mathbb{N}$, then at the end of this time step the cell parameters are updated as $\gamma_{ij}^+ = \gamma_{ij} + x$, $\beta_{ij}^+ = \beta_{ij} + 1$, by applying the Bayesian rule on the Gamma distribution. The variance is updated to $V_{ij}^+ = \frac{\gamma_{ij}^+}{\beta_{ij}^+{}^2}$. The loop is repeated until every cell on the radiation map has a variance below the predefined threshold. Note that the cell's variance is not necessarily decreased with any new measurement; counts significantly different from the expected mean temporarily increase V_{ij} . Over time, however, the variance of revisited cells is decreased under a predefined threshold V_0 . We refer to this strategy as *gradient-based Bayesian search*.

ALGORITHM IMPLEMENTATION

Hardware description

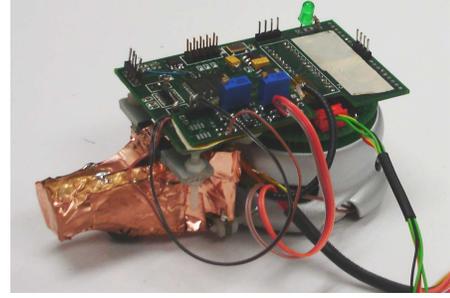


Figure 4. The Khepera II mobile robot, interfaced to a miniature radiation sensor (Fig. 5) detecting gamma-rays. The robot communicates with a central computer via RS232, for real-time radiation map building and motion control.

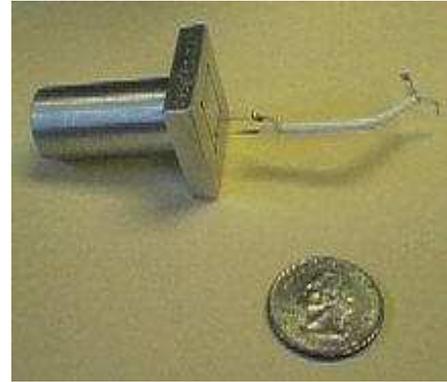


Figure 5. The miniature radiation sensor. It is interfaced to the Khepera II mobile robot as shown in Fig. 4. In our implementation the sensor is shielded with copper to reduce ambient noise.

Experimental tests are conducted using a Khepera II desktop mobile robot (Fig. 4). The robot is equipped with a custom-built turret interfacing the CsI (Cesium Iodide) radiation sensor (Fig. 5) with the robot's microprocessor, a Motorola 68331, running at 25MHZ. This processor executes the embedded C code that interprets and realizes motion commands coming from a desktop PC. On this PC the search is planned and collected data is visualized in real time. Three analog inputs available through the I/O robot interface are being used for sensor-robot communication, while collected sensor data are sent for visualization to a desktop computer through either a wireless or a cable RS232 link.

Gamma-rays passing through the CsI crystal may deposit some or all of their energy. This energy excites electrons into higher energy levels, which decay emitting visible light. The 4 cm long, 1.2 cm in diameter cylindrical CsI crystal is encapsulated into the Aluminum casing with the Hamamatsu S3509 pin photodiode mounted on it to detect light induced in the crystal by passing photons. The sensor is assembled for us by Alphaspectra, Inc. Pulses generated by the diode are weak and are amplified using an Amptek A250 preamplifier with external FET. The amplified pulse is then shaped through a

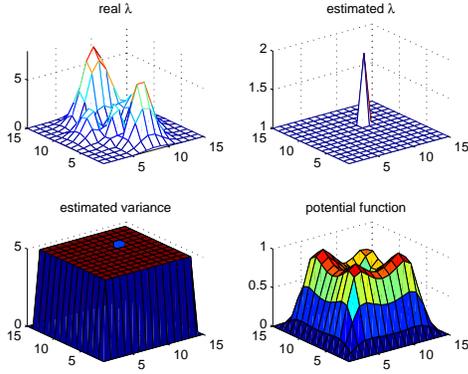


Figure 6. Initial configuration of the area. Upper left: real distribution of λ . Upper right: uniform prior information of the distribution of λ . Lower left: Uniform prior information of variance. Lower right: potential function based on surface fitting of variance data.

four-stage shaping amplifier. The shaping amplifier outputs an almost Gaussian waveform, the height of which corresponds to the energy deposited by the gamma-ray in the active region of the detector. This signal is then processed using a digital board consisting of a low power, high speed, 8bit National Semiconductor ADC08200 and an Altera Cyclone 2910 FPGA. The FPGA is programmed to perform peak finding and pulse counting. We estimate the total power consumption of the electronics to be below 200 mAh at 6V, which allows us to power them for several hours with four rechargeable digital camera batteries. The pin photodiode is in reverse bias and consumes negligible amount of power (nAh at 25V).

The added weight of the sensor, digital board, and power supplies, represents a challenge to the robot's motors. To reduce friction, a stainless steel ball caster wheel is added at the base of the sensor. Without external measurements to be used for localization, odometry errors build up and cause the robot to deviate from the reference path connecting one cell to the next. To address this issue, we manually issue corrective motion commands to keep the localization error bounded.

Experiment design and implementation

The Khepera II robot is programmed to accept high level, motion correcting commands from a controller implemented on a laptop computer running MATLAB, and interfaced with the serial port of the Khepera II. When sensor data is received from the robot, the central controller integrates them into the radiation map in real time.

The area to be mapped is a 60 cm \times 60 cm surface, decomposed into a 15 \times 15 grid. Initially, the robot is positioned at cell (10, 10). We assume a distribution of radiation levels over this area, λ , represented in upper left of Fig. 6. This distribution is unknown to the system, and the goal of the experiment is to reconstruct it up to a certain confidence level, using measurement data. Assuming no initial information about the λ distribution is available, we start with a uniform distribution for both λ_{ij} and V_{ij} , as shown in the

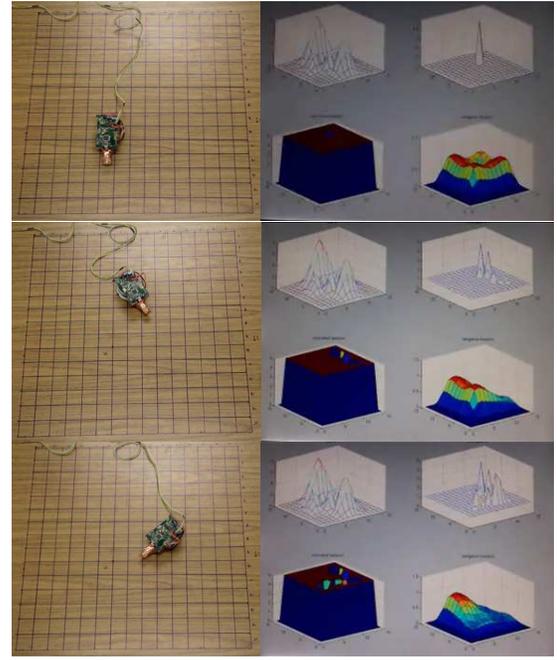


Figure 7. Implementation of the gradient-based Bayesian mapping. Left: Snapshots of the real execution of the experiment. Right: Snapshots of the updated data during the experiment. Upper left: real distribution of λ . Upper right: updated values of the distribution of λ , as the robot is moving and the sensor is collecting radiation data. Lower left: updated values of variance, during the experiment. Lower right: potential function based on surface fitting of variance data, during the experiment.

lower left and upper right of Fig. 6, respectively. The desired variance threshold for the constructed map is set at $V_0 = 0.5$.

Snapshots of the experimental test are shown in Fig. 7. The pictures on the left show the robot in different configurations on the grid, while the pictures on the right are screen captures showing how the map evolves and drives further measurements. On the upper left of each one of the screen captures, the real λ distribution of the area is shown. On the upper right is the radiation map, updated in real time as the robot is moving around taking measurements. On the lower left is the updated variance distribution, and on the lower right is the potential function that steers the robot through its gradient field.

Fig. 8 shows two-dimensional representations of the variance distribution at different time instances during the experiment. All variance values in Fig. 8(d) are below the threshold value of V_0 , and thus the search terminates.

Fig. 9 reveals how the completion time of the gradient-based and sequential based Bayesian mapping methods is affected when the variance threshold (quantifying map uncertainty) is reduced. We observe an exponential increase in the time required for map completion.

Comparison of different navigation strategies

We compare the proposed gradient-based Bayesian mapping algorithm, with the uniform and sequential-based Bayesian mapping techniques. The uniform mapping consists

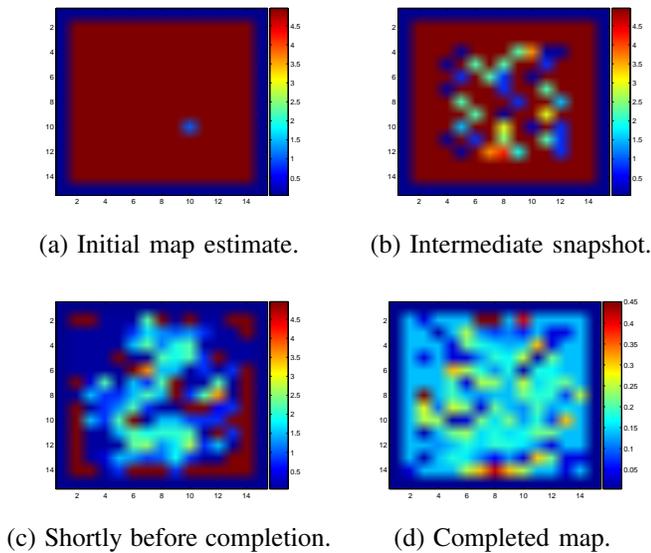


Figure 8. Successive snapshots of the radiation map real-time construction using the gradient-based Bayesian algorithm, from the initial map 8(a) to the final map 8(d) after the completion of the algorithm. Horizontal and (left) vertical axes on each matrix denote cell indices, and the right column labels the levels of variance of the Gamma distribution for radiation intensity at each cell.

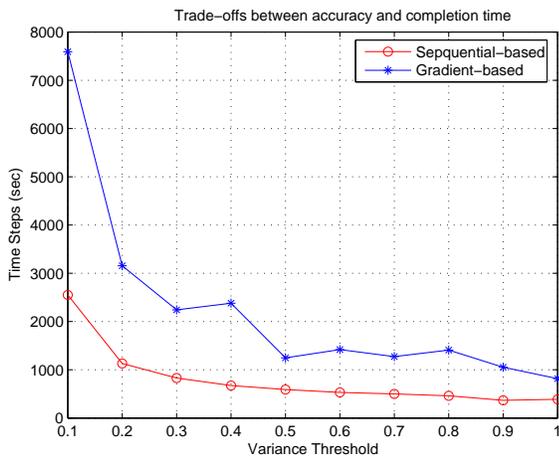


Figure 9. The effect of the variance threshold choice on the completion time of the gradient-based, and sequential-based Bayesian mapping algorithms. As the required accuracy increases, completion time appears to increase exponentially.

of scanning the area cell by cell along each row, spending a constant fraction of time at each cell. At the end of the scan, the maximum value for the variance over all cells is compared to the threshold value, and if found larger, the scan is repeated. Fig. 10 shows experimental results for the uniform search, tested in the same scenario as for Bayesian mapping algorithm.

In the sequential-based Bayesian mapping the time spent in each cell is adjusted to allow sufficient integration time for the sensor, and enough measurements to be collected, so that variance drops below the threshold before leaving the cell. Each cell is visited once. Fig. 11 shows experimental results

for the sequential search, for the same mapping scenario.

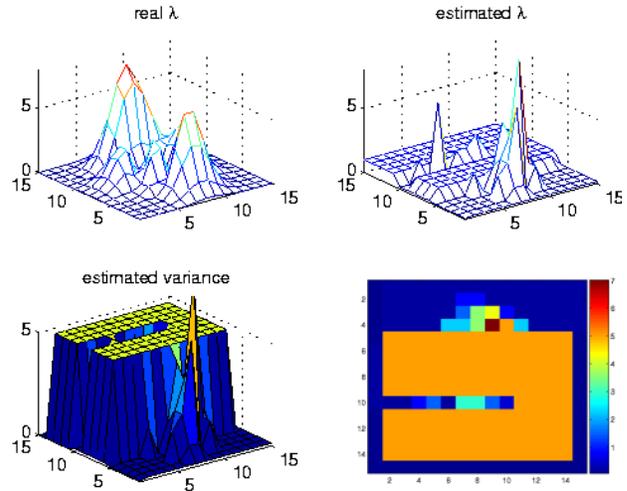


Figure 10. Experimental implementation of the uniform mapping algorithm. The left picture shows the real distribution of λ (upper left corner), the updated values of the distribution of λ as the robot is moving and the sensor is collecting radiation data (upper right corner), and the updated values of variance (Lower left corner). The right picture shows a 2D representation of the variance distribution at an intermediate time step. More than one scan of the whole area will be needed to achieve the required confidence level.

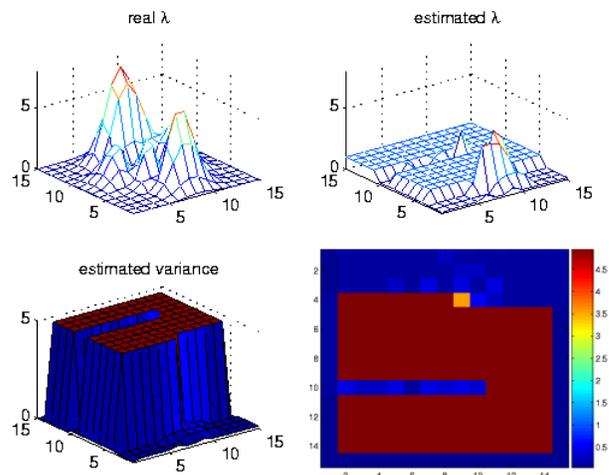


Figure 11. Experimental implementation of the sequential-based Bayesian mapping algorithm. The left picture shows the real distribution of λ (upper left corner), the updated values of the distribution of λ as the robot is moving and the sensor is collecting radiation data (upper right corner), and the updated values of variance (lower left corner). The picture on the bottom right shows a snapshot of the 2D distribution of radiation level variance over the area of interest.

The three mapping algorithms are compared in terms of completion time. Results indicate that the gradient-based Bayesian mapping is faster than uniform, but slower than the sequential-based Bayesian. In a typical run, the gradient-based Bayesian mapping requires approximately 1300 time steps (simulation seconds), the uniform mapping roughly 3000 time steps, and the sequential-based Bayesian close to 600 steps, to

complete the map at the same level of confidence. Thus, the sequential-based Bayesian mapping algorithm outperforms the gradient-based Bayesian mapping algorithm, but the latter has the advantage over both other techniques that at each time step, there is an available map constructed with a confidence that is improved with time. This is particularly important if there are severe time constraints for the completion of the mapping task: no alternative mapping technique can adequately address the problem of producing a reasonably accurate map of the *most interesting* portions of the search area, within a certain time interval. The information collected within 300 seconds by means of the gradient-based Bayesian algorithm may offer more clues for the distribution of radiation over the area than the half-built map resulting from the sequential-based Bayesian search.

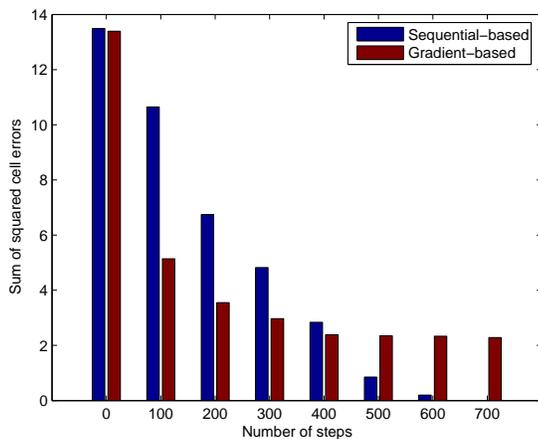


Figure 12. Comparison of the sequential-based and the gradient-based Bayesian methods in terms of accuracy for uniform, constant prior. The gradient-based algorithm reduces the error rapidly in the initial stages, but leaves a residual error after completion which is proportional to the given variance threshold. The sequential-based algorithm terminates faster, yields a more accurate map after completion, but map confidence increases almost linearly with time. The choice of method thus depends on the time constraints, the initial prior, and the dynamics of the environment.

The two algorithms that take advantage of the Bayesian update (the sequential-based Bayesian and the gradient-based Bayesian) offer different tradeoffs between optimality and flexibility (Fig. 12). The gradient-based Bayesian algorithm offers the ability to partially investigate areas that are of more interest first. The price to pay is that these areas may have to be visited again. On the other hand, the sequential-based Bayesian algorithm is time optimal when the prior variance distribution is uniform. The gradient-based Bayesian method cannot outperform the sequential-based in terms of map accuracy, the latter measured in terms of the residual sum of squared errors between the true radiation levels and the estimated radiation levels. This is because, and strange as it may sound, in nuclear measurement $1 + 1 \neq 2!$ The Poisson statistics of nuclear measurement imply that visiting the same cell twice and spending a total of two seconds there is not the same as getting there once and have the sensor integrate for two seconds continuously. When the

prior variance distribution is not uniform, however, a partially constructed map may be of limited value, depending on the initial position of the robot prior to exploration.

CONCLUSIONS

We develop two radiation mapping algorithms that can handle different situations based on prior information of the search area. The algorithms are developed in the framework of model-driven measurement, where a world model is used to drive measurement collection, and measurements are used to update the world model. We develop and experimentally test a robotic implementation of two Bayesian-based radiation mapping strategies in two dimensions, using a commercially available desktop mobile robot, fitted with a CsI radiation sensor. Our approach to implementing the Bayesian radiation mapping algorithms is to drive the robot over each segment of the search area, in real-time, according to the radiation counts collected by the sensor. Future research directions include extensions to three-dimensional mapping, exploring and characterizing the tradeoffs between time efficiency, map confidence level, and utilization of prior knowledge information, as well as the implementation of Bayesian statistics for the on-line update of the world model.

ACKNOWLEDGMENT

Xanthi Papageorgiou was supported by the Los Alamos National Laboratory Award No. STB-UC:06-36. Andres Cortez was supported in part by the aforementioned award, and in part by DoE URPR grant DE-FG52-04NA25590. The latter grant also supported Herbert Tanner and Ron Lumia. The authors wish to thank Nick Hengarter and Chuck Alexander of the Los Alamos National Laboratory for helping them in nuclear statistical modeling and experimental system integration, respectively.

REFERENCES

- [1] W. C. Priedhorsky, "Physical and economic tradeoffs between remote sensing and deployed sensor networks," *Journal of Nuclear Materials Management*, 2005.
- [2] A. Elfes and H. P. Moravec, "High resolution maps from wide angle sonar," in *IEEE International Conference on Robotics and Automation*, 1985, pp. 116–121.
- [3] B. Yamauchi, "Frontier-based exploration using multiple robots," in *Second International Conference on Autonomous Agents*, 1998, pp. 47–53.
- [4] L. Romero, E. Morales, and E. Sucar, "An exploration and navigation approach for indoor mobile robots considering sensor's perceptual limitations," in *IEEE International Conference on Robotics and Automation*, May 2001, pp. 3092–3097.
- [5] S. Moorehead, "Autonomous surface exploration for mobile robots," Ph.D. dissertation, Carnegie Mellon University, August 2001.
- [6] B. Grocholsky, J. Keller, V. Kumar, and G. Pappas, "Co-operative air and ground surveillance," *IEEE Robotics and Automation Magazine*, pp. 16–26, 2006.

- [7] A. Lilienthal and T. Duckett, "Building gas concentration gridmaps with a mobile robot," *Robotics and Autonomous Systems*, vol. 48, no. 1, pp. 3–16, 2004.
- [8] P. Orgen, E. Fiorelli, and N. E. Leonard, "Cooperative control of mobile sensor networks: Adaptive gradient climbing in a distributed environment," *IEEE Transactions on Automatic Control*, vol. 49, no. 8, pp. 1292–1302, August 2004.
- [9] J. Kim and J. P. Hespanha, "Cooperative radar jamming for groups of unmanned air vehicles," in *IEEE Conference on Decision and Control*, December 2004, pp. 632–637.
- [10] L. Bertuccelli and J. How, "Robust uav search for environments with imprecise probability maps," in *IEEE Conference on Decision and Control*, December 2005, pp. 5680–5685.
- [11] A. Wald, "Sequential tests of statistical hypotheses," *Annals of Mathematical Statistics*, vol. 16, pp. 117–186, 1945.

Andres R. Cortez received his B.S. degree in mathematics from New Mexico Highlands University, Las Vegas NM, in 2005. He is currently working towards the M.S. degree in mechanical engineering at the University of New Mexico under the supervision of Prof. Tanner.

Xanthi Papageorgiou received her Eng. Diploma in mechanical engineering from the National Technical University of Athens, (NTUA) Greece, in June 2003. In the Fall of 2006 she was a visiting research scholar with the Mechanical Engineering Department at the University of New Mexico. She is currently working toward the Ph.D. in mechanical engineering at NTUA. She is a student member of IEEE.

Herbert Tanner received his Eng. Diploma and Ph.D. in mechanical engineering from the National Technical University of Athens, Greece, in 1996 and 2001, respectively. From 2001 to 2003 he was a post doctoral fellow with the Department of Electrical and Systems Engineering at the University of Pennsylvania. In 2003, he joined the faculty of the Department of Mechanical Engineering at the University of New Mexico where he is currently an assistant professor. In 2007 he was awarded the University of New Mexico School of Engineering junior faculty research excellence award. He is also a recipient of the National Science Foundation Career award.

Alexei V. Klimenko received the B.S. and M.S. Degrees in electrical engineering from the Moscow Institute of Steel and Alloys (later Moscow University of Technology), Moscow, Russia, in 1998 and 1999, respectively, and the M.S. and Ph.D. degrees in nuclear physics from the Old Dominion University, Norfolk, VA, in 2001 and 2004, respectively. From 2004 to 2006, he was a Postdoctoral Research Associate at Los Alamos National Laboratory, Los Alamos, NM. He is currently a Staff Scientist at Passport Systems, Inc., Acton, MA, specializing in modeling and threat detection algorithm development for nonintrusive cargo screening systems.

Ron Lumia has been a Professor in the Mechanical Engineering Department since 1994. He received a B.S. from Cornell University in 1972 and the M.S. and Ph.D. degrees from the University of Virginia in 1977 and 1979, respectively. He has worked for industry, government, and academia. Re-

cently elected Fellow of the IEEE for leadership in the development of open architecture control systems for applications in robotics and automation, he has managed a variety of robotics, automation, and sensory processing research projects.

Konstantin N. Borozdin received the M.S. degree in experimental nuclear physics from the Moscow Engineering Physics Institute (Moscow, Russia) in 1988, and the Ph.D. degree in astrophysics and radioastronomy from the Moscow Space Research Institute in 1995. He is currently a technical staff member with the Space Science and Applications Group, Los Alamos National Laboratory, Los Alamos, NM. He has been at Los Alamos since 1998, working in the fields of astrophysics, physical modeling, information technologies, radiation detection and national security applications. He received an Order of Merit medal for his involvement in the Mir-Kvant space experiments in 1997.

William C. Priedhorsky received the B.A. summa cum laude with honors degree in physics from the Whitman College, Walla Walla, WA, in 1973, and the Ph.D. degree in physics from the California Institute of Technology, Pasadena, CA, in 1978. He has been at Los Alamos National Laboratory since 1978, most recently as a Program Director at the Laboratory Directed Research and Development Office. He also spent from several months to one year as a visiting scientist at the Max Planck Institute for Extraterrestrial Physics, Garching, West Germany, in 1985-1986, at the Danish Space Research Institute, Lyngby, Denmark, in 1994, and at the University of Melbourne, Australia, in 1995. His main research interests include high-energy astrophysics, X-ray and optical instrumentation and various national security applications. He is a recipient of numerous awards and honors.