Target-Motion Prediction for Robotic Search and Rescue in Wilderness Environments

Ashish Macwan, Goldie Nejat, *Member, IEEE*, and Beno Benhabib

***Abstract*—This paper presents a novel modular methodology for predicting a lost person's (motion) behavior for autonomous coordinated multirobot wilderness search and rescue. The new concept of isoprobability curves is introduced and developed, which represents a unique mechanism for identifying the target's probable location at any given time within the search area while accounting for influences such as terrain topology, target physiology and psychology, clues found, etc. The isoprobability curves are propagated over time and space. The significant tangible benefit of the proposed target-motion prediction methodology is demonstrated through a comparison to a nonprobabilistic approach, as well as through a simulated realistic wilderness search scenario.**

***Index Terms*—Lost-person motion prediction, multirobot coordination (MRC), wilderness search and rescue (SAR) (WiSAR).**

I. Introduction

SEARCH AND RESCUE (SAR) has been commonly classified according to the environment within which it takes place and the nature of the target's motion. Urban SAR (USAR) refers to activities amid collapsed structures and is concerned with locating stationary survivors within a bounded environment [1], [2]. Marine applications address search in potentially boundless environments, typically precluding consideration of terrain-difficulty issues [3], [4]. Wilderness SAR (WiSAR) refers to locating lost persons moving in unbounded inland environments with varying and often complex terrains [5], [6]. This presents rescuers with a significantly different problem than that for the two aforementioned applications. The main challenge is to locate a moving unobservable (i.e., nontrackable) target, whose state at any given time is unknown, but can be predicted through the use of probabilistic information. As part of our continuing research efforts on multirobot coordination (MRC) for autonomous WiSAR [7], this paper focuses on the development of a novel target (motion) behavior prediction methodology, using probability theory.

Manuscript received May 13, 2010; revised September 16, 2010 and February 22, 2011; accepted March 9, 2011. Date of publication May 12, 2011; date of current version September 16, 2011. The work of A. Macwan was supported in part by the Natural Sciences and Engineering Research Council of Canada through a PGS-D scholarship. This paper was recommended by Associate Editor E. Tunstel.

A. Macwan and B. Benhabib are with the Computer Integrated Manufacturing Laboratory, Department of Mechanical and Industrial Engineering, University of Toronto, Toronto, ON M5S 3G8, Canada (e-mail: a.macwan@utoronto.ca; beno@utoronto.ca).

G. Nejat is with the Autonomous Systems and Biomechatronics Laboratory, Department of Mechanical and Industrial Engineering, University of Toronto, Toronto, ON M5S 3G8, Canada (e-mail: nejat@mie.utoronto.ca).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TSMCB.2011.2132716

The literature on target-behavior prediction for WiSAR is sparse. Earlier methods have commonly dealt with only a limited number of aspects, assuming that both the initial *a priori* probability distribution for the target's location and the probabilistic description of how the target moves are given inputs (e.g., [3]–[5], [8], and [9]). Our proposed generic methodology, on the other hand, utilizes the new innovative concept of isoprobability curves to represent the probable location of the target at any given time. This representation mechanism explicitly and dynamically incorporates the influences of critical aspects that need to be considered, such as a growing search area, terrain topology, target physiology and psychology, etc.

A. Target Prediction in Search Theory

Numerous methods have been developed for the optimal allocation of search effort to locate a stationary or moving target [3], [8], [10]–[12]. An element common to these methods is the use of probability theory in the prediction of the mobile target's location, which is represented through a continuous or discretized probability distribution, propagated over time based on a stochastic process [3], [4].

Search-theory works have mainly utilized theoretical generalized search optimization (GSO) techniques for the optimal allocation of the search effort [10], [11]. Target prediction is addressed via a stochastic process, assumed to be a given input to the adopted GSO technique. However, no guidance is provided as to which stochastic process is applicable for a given search scenario. As part of the research on human–robot team formation in WiSAR scenarios in [5], a random-walk model is presented for deriving the initial probability distribution for the location of the missing person. It involves random perturbations to the speed and direction of the target to model drift in the target's motion.

Examination of missing-person incidents in the wilderness areas of Alberta, Canada, was carried out in [6]. The Wakeby distribution [13] was used to characterize the data for different categories of wilderness user. This distribution provides a function for computing percentiles of crow's flight distances in planning the boundaries of a search.

In [3], a marine SAR scenario is addressed, involving a floating target being carried by winds and waves. Starting with an initial given target-location probability density function (PDF), a prediction stage is applied to propagate the PDF using a Markov motion model. This work was extended in [8] to include a means of reconfiguring the search boundaries. In [9], the Bayesian probabilistic framework in [3] is used to also incorporate the effect of imperfect sensors.

A search planning system used by the U.S. Coast Guard is discussed in [4]. Given the estimates on the location and time of the distress incident, the system uses a Monte Carlo simulation, combined with oceanographic models for target drift computation, to build an initial PDF for target location and propagate it. Bayesian updating is also used to reflect "negative information" obtained from unsuccessful searches. The output is a sequence of time-phased PDFs giving the likely target positions over the time period of interest.

In [14], a method is proposed for constructing the *a priori* target-location PDF, as well as the transition matrix (assuming a first-order Markov process). The method requires that expert opinion be initially available to help in generating the PDF but allows for incorporating the influence of terrain features.

B. Robotic SAR

In the USAR domain, the design of specialized robots to address the challenges of this particular application is prevalent (e.g., [15], [16]). Research works focusing on teleoperation of robots and formation of effective mixed human–robot teams (e.g., [17] and [18]), as well as centralized (e.g., [19]–[21]) versus decentralized (e.g., [22]–[24]) control issues, have also been reported. Some of our recent work has also contributed to the design of sensing and control systems for real-time 3-D mapping and landmark identification (e.g., [25]–[27]).

For WiSAR, decentralized approaches for coordinating teams of robots have been reported. For example, the research in [5] investigates the use of multiple unmanned aerial vehicles for aerial imagery of the search area to aid human searchers. In [28], the focus is on searching for stationary targets within a bounded, albeit unknown, 2-D environment. The work in [29] compares various MRC mechanisms for automated SAR, but also for stationary targets.

In these and other works (e.g., [30] and [31]), although current interest and efforts toward the use, as well as the increased autonomy, of robots are apparent, search strategies are not specifically addressed. Those that do are often incompatible with WiSAR. For example, the search method in [32] addresses a visibility-based version of the pursuit–evasion problem, assuming an unlimited searcher-vision range in a bounded environment with a moderate distribution of obstacles. It was modified in [33] to accommodate multiple searchers and a searcher field-of-view model that more closely resembles that of a physical robot but still maintained the other limiting assumptions in [32]. Namely, current MRC methods do not utilize target prediction based on probabilistic information. However, they indicate clear recognition of the advantages that an autonomous multirobot solution, with target-behavior prediction, could bring to SAR.

II. Problem Statement

This paper addresses the problem of online prediction of a lost subject's behavior for the search component of a ground WiSAR scenario. It is based on representing the target's location in the form of a PDF that needs to be updated online over time and space. As the primary constraint, it is assumed that, during the search, the target is not trackable, as this is typically the case in real-life scenarios.

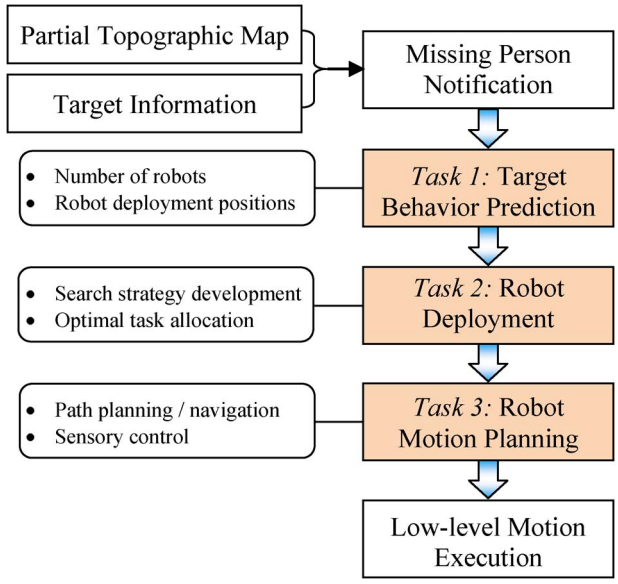

Fig. 1. Overall WiSAR problem.

A typical WiSAR scenario would proceed as follows: At a certain point in time, a notification of a missing person arrives, who is hereby referred to as the *target*—his/her last known position (LKP) is given, along with the time at which he/she was at this position, $t = 0$. The search agents (robots) are then transported to the region and arrive at time $t = T_{\rm hs}$ (i.e., the target's "head start")—the optimal number of robots, and their initial deployment, must be determined *a priori*. Basic topographic information of the region and elementary knowledge of the personal characteristics of the target, such as age, clothing, provisions, purpose of visit, health, experience, and familiarity with the area, are assumed to be known.

Fig. 1 shows the three major tasks that comprise the overall WiSAR problem with centralized control, which we wish to automate. Foremost is obtaining accurate probabilistic information about the target for behavior prediction. This task primarily involves generating a probability distribution for the location of the target within the search area and updating this distribution over time and space. The prediction must incorporate all the relevant sources of influences on the target's motion. The four influences, detailed hereinafter, that particularly need to be addressed are terrain, target physiology, target psychology, and clues left behind.

Robot deployment and robot-motion planning, which follow, depend on target-behavior prediction. Robot deployment must be conducted in accordance with the probability distribution for target location at the start of the search, so that search effort is proportioned according to the relative likelihoods of target presence in the search area at that time. Robot-motion planning, on the other hand, is a dynamic process that must keep updating the search task assignment based on new information. It requires the target-behavior-prediction module to propagate the probabilistic target-location information over time and space. Upon allocation, the robots autonomously execute their tasks—path planning and real-time navigation. The aforementioned process

is cyclic and continues until either the target is sighted or a maximum search time is reached.

A. Terrain

Terrains in WiSAR scenarios may vary from rugged mountainous regions to dense forests and smoother rolling hills. Thus, the terrain topography would influence the target's motion, and this effect must be reflected in the probabilistic prediction of the target's location. The search terrain could also include nontraversable obstacles, such as large rocks or boulders, deep bodies of water, swamps, ravines, etc., which may be *a priori* known or unknown to the search team.

B. Psychology

Pertinent literature on human psychology [34], [35], as well as data on past search incidents [6], [36], indicates that lost persons often employ predictable strategies to reorient themselves. For example, a lost hiker, although relatively unfamiliar with the terrain topology, may at least be aware of the existence and general direction of a popular destination (e.g., a well-known camp site or a parking lot) and could move toward it. Other reorientation strategies include searching high ground, wandering, trail running, signaling, etc. [34].

C. Physiology

Target-behavior prediction would also be influenced by physiological issues, such as fatigue setting in over time. For example, a particular PDF assumed for the target's speed may need to be scaled down in a controlled manner over time that is consistent with the decrease in forward speed experienced by the target as he/she gets tired. Changes in weather conditions may also cause predictable physiological responses.

D. Clues

An evolving source of information critical to the search process is the clues left behind by the target. Examples include articles of clothing belonging to the target, the time and location of a target sighting by an eyewitness, etc.

III. Proposed Solution: A Modular Target-Behavior-Prediction Methodology

Developed within the framework of our research is the novel key construct of isoprobability curves to predict target (motion) behavior. These curves are generated and propagated based on information about the target's behavioral descriptors, the topography of the search terrain, and other factors.

A. Isoprobability Curves

A lost target can move in any direction $\theta \in [0^\circ, 360^\circ]$ from his/her last LKP (x_0, y_0), achieving a trajectory that is difficult to predict. Instead of trying to predict the exact target trajectories, we choose to establish bounds on the target's location using conservative estimates of his/her motion. In particular, as the worst case motion scenario, we assume that the target

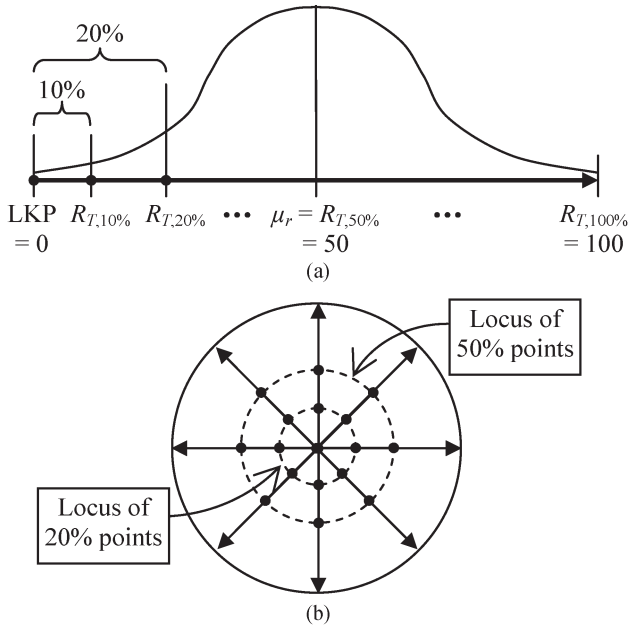

Fig. 2. Isoprobability curves. (a) Side view: a single ray. (b) Top view: locus of points.

may travel in a straight line radially outward from the LKP at its average speed, yielding a maximum distance from the LKP. Such an estimate ensures that the predicted and utilized PDF for the target's location would have the best chance of bounding the target.

Since the exact data about the lost individual may not be available at the time of the search, probabilistic information about his/her peer group must be utilized. These data could be in the form of a probability distribution $p(\nu, \theta)$, $\nu \in \Re$, $\nu \geq 0$, for all the mean speeds ν that individuals within the particular peer group can have. For a WiSAR scenario, one particular type of data gathered includes the percentiles of distances from the LKP where targets from past search incidents were found [6], [36], [37]. These data can be used to generate a mean-target-speed PDF, referred to hereinafter as the *"nominal" mean-target-speed PDF*. For the purposes of coordinating a search, however, a target-location PDF $p(r, \theta, t)$ is required, for the radial distance r, from the LKP and the given point in time t. This PDF can be derived by multiplying the nominal mean-target-speed PDF by the given time t.

In order to construct the probabilistic representation of target location, we first establish multiple *rays*, which are possible directions of *conservative* target motion emanating from the LKP. A target-location PDF, which is a 1-D distribution at this stage, can be calculated for any given point in time and then overlaid onto each established ray individually. Herein, a bounded normal distribution is assumed to define the nominal mean-target-speed PDF and, hence, the target-location PDF derived from it, along any given ray.

Fig. 2(a) shows the side view of one ray (bold arrow) where the position along the ray is given by the variable r, $r \in \Re$, $r \geq r_0$. The derived target-location PDF for the current time $t = T$ is overlaid onto this ray. The ray, and thus the target-location PDF, originates from the LKP [$r_0 = 0$ in Fig. 2(a)] and extends up to the upper limit of this PDF [$r = 100$ units

in Fig. 2(a)]. Thus, the $r = 100$ units point represents the maximum distance that the target can travel in the given time T, as dictated by the target-location PDF for this time.

One may subdivide a ray at key cumulative-probability points $p\%$ for the probable locations of the target $r = R_{T,p\%}$ along this ray for time T, using the overlaid target-location PDF [Fig. 2(a)]. Given that the normal distribution is unimodal and symmetric about its mean μ_r, the mean coincides with the 50% cumulative-probability point $R_{T,50\%}$, and the upper limit for the given time T is taken to be $R_{T,100\%}$.

Based on the normal-distribution assumption for $p(\nu, \theta)$ with a mean $\mu_{\nu,\theta}$ and a variance $\sigma_{\nu,\theta}$ along any given ray, one can define $p(r, \theta, t)$ on that ray through a transformation of variable ν to r for $r = r_0 + \nu \cdot t$, where $r_0 = \|(x_0, y_0)\|$. Since ν is a random variable, r is simply a linear combination of random variables, so that $p(r, \theta, t)$ is also a normal distribution with mean and variance of

$$\mu_{r,\theta,t} = E(r) = E(r_0 + \nu \cdot t) = E(\nu) \cdot t + r_0 = \mu_{\nu,\theta} \cdot t + r_0 \tag{1}$$

$$\sigma^2_{r,\theta,t} = \sigma^2_r = \sigma^2_{(r0+\nu \cdot t)} = t^2 \cdot \sigma^2_{\nu,\theta} \tag{2}$$

respectively.

The cumulative-probability points established on each ray incorporate probabilistic target-location information for each separate scenario (direction) of target motion independently. To estimate this probabilistic information for all the intermediate directions and to form a continuous representation over the search area, the loci of all common points among the established rays are connected through interpolation. These form a set of continuous 2-D closed contours over the search area, one for each cumulative percentage value considered, and are referred to hereinafter as *isoprobability curves*. It is important to note that, once the isoprobability curves are constructed, the volume of probability bounded by each contour will not be the same as the cumulative probability of the corresponding contour. However, given our conservative target-motion estimate, the contour positions *must* remain as computed from the loci of the common points along the rays to ensure that the 100% contour fully bounds the appropriate search area for that time.

Fig. 2(b) shows a set of isoprobability curves for eight rays, where the PDFs along the rays are assumed to be the same, yielding circular curves. When the PDFs along rays differ, due to terrain conditions or other factors, equal cumulative-probability points along the rays could be combined into (noncircular) isoprobability curves by fitting (for example, cubic) spline segments. The cumulative-probability points are also referred to, herein, as control points.

Since, in general, the target is continually in motion, it is insufficient to search areas only once and move on. Rather, it would be necessary to reinvestigate previously searched areas, as well as to increase the search area with time. Thus, as search time passes, continued possible target motion outward from the LKP necessitates the propagation of the isoprobability curves based on the probable target motion along each ray. It entails multiplying the mean-target-speed PDF corresponding to each ray by the total time passed since the target was at the LKP. Namely, based on the aforementioned derivation, at time $t + \Delta t$, the new mean and variance would respectively be

$$\begin{aligned}\mu_{r,\theta,(t+\Delta t)} &= E(r) = E\left[r_0 + \nu \cdot (t + \Delta t)\right] \\ &= E(\nu) \cdot (t + \Delta t) + r_0 \\ &= \mu_{\nu,\theta} \cdot (t + \Delta t) + r_0, \end{aligned} \tag{3}$$

$$\sigma^2_{r,\theta,(t+\Delta t)} = \sigma^2_r = \sigma^2_{[r0+\nu \cdot (t+\Delta t)]} = (t + \Delta t)^2 \cdot \sigma^2_{\nu,\theta}. \tag{4}$$

The more time that passes between curve updates, the more outdated and inaccurate the predictions become. It can be expected that this too may negatively impact the chances of finding the target. Available computational resources also influence how frequently and quickly curves can be updated.

B. Incorporating the Effect of Terrain

Varying terrain topography would influence target motion and thus must be accounted for in the determination of the isoprobability curves. Two factors are considered here: 1) difficulty of terrain and 2) *a priori* known obstacles. The former considers the impact of the difficulty in traversing relatively larger scale land formations. The latter considers smaller scale obstructions, such as large rocks or boulders, which the target may navigate around.

1) Terrain Difficulty: In order to consider terrain difficulty in the determination of isoprobability curves, the nominal mean-target-speed PDF along each ray needs to be scaled according to the instantaneous surface slope. The calculations begin with two basic assumptions: A topographic map of the area is available, and the LKP of the lost person is also known.

The terrain is first discretized by overlying a grid of cells onto the search area. Each cell is assigned its representative average terrain height value, yielding a *height map*. Different terrains were generated for our simulations using the Terragen Classic scenery-generation software [38]. Next, the terrain slope along each ray, for the portion of the terrain from the LKP up to any given position along the ray, is determined using a linear least squares fit to ground-elevation measurements.

Since the position of each control point along a ray depends on the corresponding mean target speed and time, an iterative process is needed to determine its exact location. Namely, although, at a given time, one could calculate the control point along the ray using the nominal mean target speed, the terrain traversed due to this distance traveled must be considered to scale the nominal mean speed to incorporate the terrain's influence on target motion. Scaling the mean speed would, in turn, change the distance traveled, and the terrain influence would have to be recomputed for this new travel distance. This process must be iterated until the difference between two successive computed travel distances is less than or equal to a user-specified maximum difference threshold value.

Step 1) Estimate the radial distance d_{k_est} of the kth control point $k \in [1, N_{cp}]$ computed along the ray under consideration, using the corresponding nominal mean target speed, and determine the coordinates of the control point $P_{k\_\mathrm{est}} = (x_{k\_\mathrm{est}}, y_{k\_\mathrm{est}})$.

Step 2) Determine the height measurements (from the height map) at regular distances along a straight line representing the distance traveled by the target from the LKP up to the current estimated control point.

Step 3) Calculate the average slope of the portion of the terrain lying along the direction vector from the LKP up to the current estimated control point using the sample terrain height measurements obtained in Step 2).

This average terrain slope is obtained through a linear-regression-line fit to the height data points, using the slope b of the resultant line to compute an angular measure of the average terrain slope sought after

$$\hat{h} = a + br \tag{5}$$

where $\hat{h}$ is the estimated terrain height value given by the fitted regression line; r is the radial distance r_i, $i \in [1, N_h]$, at which a given terrain height measurement $h_i, i \in [1, N_h]$, is taken; and $\{a, b\}$ denotes the estimated values for the regression coefficients, respectively.

The slope value is converted into an angular value $\gamma = \tan^{-1}(b)$, which represents the estimate of the average angular slope over which the target could have traveled in the direction of the given ray up to the estimated position of the control point.

Step 4) Use the computed average terrain slope to determine the speed scale factor corresponding to the control point being determined.

The determination of how much the mean speed of the target scales with different terrain slopes may be made by referring to empirical data [39], [40]. Herein, a linear relationship between average ground surface slope angle γ and speed scale factor q is assumed. This relationship is obtained by stipulating a maximum incline angle $\gamma_{\max,\text{inc}}$ and a maximum decline angle $\gamma_{\max,\text{dec}}$ that the target can handle, as well as the corresponding speed scale factors $q_{\max,\text{inc}}$ and $q_{\max,\text{dec}}$, respectively. In addition, an incline of $\gamma = 0^\circ$ is given a scale factor of $q = 1$

$$q(\gamma) = \begin{cases} q_{\text{dec}}(\gamma) = m_1\gamma + b_1, & \gamma_{\max,\text{dec}} \le \gamma < 0^\circ \\ q_{\text{inc}}(\gamma) = m_2\gamma + b_2, & 0^\circ \le \gamma \le \gamma_{\max,\text{inc}} \end{cases} \tag{6}$$

where m_1 and m_2 are parameter constants of the two portions of the linear relationship, respectively.

If the computed slope exceeds the stipulated bounds, the corresponding terrain is considered to be nontraversable by the target. The iterative process thus ceases, and the control point being computed remains at its current position P_{k_est}.

Step 5) Scale the corresponding nominal speed, using the computed speed scale factor q_k, and update the estimate of the radial distance d_{k_est} of the kth control point $k \in [1, N_{cp}]$. With speed having changed, the corresponding radial distance estimate of the control point must be changed as well. This relocates the control point along the ray and represents an improved estimate.

TABLE I
SUMMARY OF COMPUTATIONS FOR ITERATION 1 OF THE EXAMPLE PROBLEM

Step #	Results
1	$d_{k\ est} = (V_{10\%,nom}) \times (t) = (0.0796\text{ m/s}) \times (300\text{ s}) = 23.9$ m; $x_{k\ est} = (23.9\text{ m}) \times (\cos 40^\circ) = 18.3$ m; $y_{k\ est} = (23.9\text{ m}) \times (\sin 40^\circ) = 15.3$ m. The corresponding cell coordinates are [253 259].
2	Sample height data points from height-map: r_i = {105.1 m, 100.4 m, 106.7 m, 97.3 m, 94.1 m}
3	$b_k = -0.4292$; $\gamma_k = \tan^{-1}(b_k) = -23.2^\circ$
4	Equation (6) parameters used: $\gamma_{max,inc} = 60^\circ$, $\gamma_{max,dec} = -60^\circ$, $q_{max,inc} = 0.5$, and $q_{max,dec} = 1.5$; → ∴ $q_k = 1.1936$
5	$V_{k,scaled} = V_{10\%,nom} \times q_k = (0.0796\text{ m/s}) \times (1.1936) = 0.0950$ m/s; $d_{k\ est\_NEW} = V_{k,scaled} \times t = (0.0950\text{ m/s}) \times (300\text{ s}) = 28.5$ m
6	Using $\Delta d_{k,max} = 1\%$: $\Delta d_k = 100\% \times (\lvert d_{k\ est\_NEW} - d_{k\ est} \rvert / d_{k\ est}) = 100\% \times (\lvert 28.5 - 23.9 \rvert / 23.9) = 19.2\% > \Delta d_{k,max}$; ∴ since $\Delta d_k > \Delta d_{k,max}$, another iteration is necessary.

Step 6) Compare the updated control point radial distance estimate to the prior one used at the start of the current iteration, and determine whether another iteration is necessary

$$100\% \times (|d_{k\_\text{est\_NEW}} - d_{k\_\text{est}}|/d_{k\_\text{est}}) \le \Delta d_{k,\max}. \tag{7}$$

If the aforementioned difference is greater than the threshold $\Delta d_{k,\max}$, then the answer is "yes." The new radial distance estimate just obtained is set as the current radial distance estimate (i.e., set $d_{k\_\text{est}} = d_{k\_\text{est\_NEW}}$), and the process returns to Step 2). Otherwise, the iterative process is stopped, and the current estimate for the radial distance $d_{k\_\text{est\_NEW}}$ is accepted.

The aforementioned approach is illustrated through an example: A hiker is lost in a rugged region that does not allow for aerial surveying. The initial search region considered encompasses a 3000 m × 3000 m area. The LKP of the lost person is located at the center of this area (0, 0). The terrain is discretized by overlying a 513 × 513 grid array. Each cell is assigned the average height of the terrain enclosed within the cell, yielding a height map. The mean-target-speed PDF is a bounded normal distribution. Ten isoprobability curves are constructed for the elapsed search time of $t = 300$ s, utilizing 36 uniformly distributed rays with an interray spacing of 10°.

Table I summarizes the computations for the first iteration of the proposed process—a sample set of calculations carried out to determine the position of the control point corresponding to the 10% contour for the 40° ray.

For the second iteration, $d_{k\_\text{est}} = d_{k\_\text{est\_NEW}} = 28.5$ m, $P_{k\_\text{est}} = (21.8\text{ m}, 18.3\text{ m})$, b_k and γ_k remain -0.4292 and -23.2°, respectively, so that q_k and $V_{k,\text{scaled}}$ also remain 1.1936 and 0.0950 m/s, respectively; thus, $d_{k\_\text{est\_NEW}} = d_{k\_\text{est}} = 28.5$ m. Since the percentage difference between $d_{k\_\text{est\_NEW}}$ and d_{k_est} is less than the 1% threshold value, the stopping condition in Step 6) is satisfied, and no further iterations are required.

Fig. 3 shows the resulting set of isoprobability curves for the $t = 300$ s time point. In subsequent control-point computations,

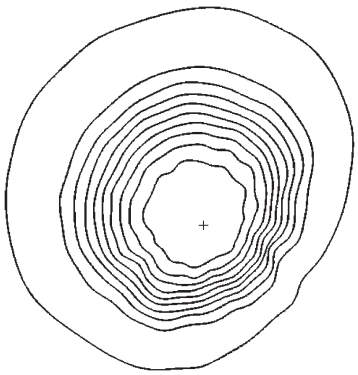

Fig. 3. Complete isoprobability curves for $t = 300$ s.

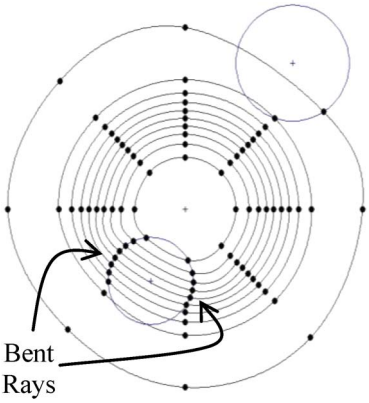

Fig. 4. Isoprobability curves modified to reflect the effect of obstacles.

when the isoprobability curves need to be updated, only the additional distance that would be traversed by each control point, between its previous and new positions, needs to be considered when computing the terrain slopes and scale factors in Steps 1) to 4). Thus, the "LKP" is replaced with the "current control-point position" in these steps.

2) Known Obstacles: The effect of *a priori* known obstacles is reflected through the changes in the shapes of the rays. The approach used herein is to "wrap" any ray that intersects an obstacle around its boundary, favoring the side leading to the shortest path around the obstacle as a conservative estimate. Neither the original path length of the ray nor the positions of the control points along the length of the ray are altered (i.e., the ray is not stretched). Hence, if the variable $s \in [0, s_{\max}]$ represents the curve length along the ray, starting from $s = 0$ at the LKP and going up to $s = s_{\max}$ at the control point on the 100% contour, then the length of the original straight ray $l = s_{\max}$ and the curve length values s_k corresponding to the positions of each of the control points all remain the same before and after wrapping the ray. The wrapping of rays mimics the slowing down of the outward progress of the target. Fig. 4 shows this approach for circular obstacles.

It is important to note that only the positions of the control points defined are changed to accommodate the *a priori* known obstacles. The interpolated curve segments between the control points are not influenced by the obstacles directly. As a result, these curve segments may pass through the obstacles (Fig. 4). However, such occurrences are not to be construed as a representation problem since it is the rays that model the different target-motion scenarios. The curve segments provide guidance for robot search movements, but only in those regions unobstructed by obstacles. When isoprobability curves intersect obstacles, the robots in those areas would inevitably need to engage in some type of obstacle-avoidance maneuver until the obstacle is circumvented.

C. Incorporating the Effect of Human Psychology

One aspect of psychology that may be accounted for is a preferred general direction of target travel [6]. The influence of such a travel direction can be addressed via the application of an appropriate scale factor to the PDF of a ray specifically chosen to align with the probable destination point. Namely, the target-location PDF along the ray in this direction would receive a higher scaling factor relative to other directions. Search commanders in real-life WiSAR scenarios often use their past experience and knowledge of the target at hand to assign relative likelihoods to different regions of the search area. Thus, a consensus among multiple such experienced personnel may be used to establish the relative likelihoods of different directions and, in turn, to obtain the scale factors.

For example, let us consider a linear relationship between directional likelihood and scale factor. In particular, two contrasting directions can be identified and assigned a relative likelihood: a "low" and a "high." The high-likelihood direction $\alpha_{\rm hi}$ can be designated to be Q_h times more likely than the low-likelihood direction $\alpha_{\rm lo}$. Thus, a probability distribution, having a linear probability model, can be established for the possible angular directions of travel by the target. This linear function would give the probability density value $p(\alpha)$ for a specified ray direction angle α and would range over all the possible directions, from -180° to $+180^\circ$.

Given the derived function $p(\alpha)$, the scaling factor $q(\alpha_i)$ for any given ray α_i can be obtained by taking the probability density ratio relative to that of the high-likelihood direction

$$q(\alpha_i) = p(\alpha_i)/p(\alpha_{\rm hi}). \tag{8}$$

The scale factor $q(\alpha_i)$ is applied to the nominal mean-target-speed PDF of ray α_i. Thus, all the nominal speeds corresponding to the contours are multiplied by this scale factor, which changes the radial distances of each of the corresponding control points accordingly. Scaling the rays in this way has the effect of proportionally "shrinking" the search area away from the lower likelihood directions, so that the searchers spend more time searching the regions closer to the high-likelihood direction. As an example, Fig. 5 shows a scenario where $\alpha_{\rm hi} = 0^\circ$, $\alpha_{\rm lo} = -180^\circ$, and $Q_h = 3$.

D. Incorporating the Effect of Clues Found

Clues can be abstracted into two general types: those that only give positional information about the target and those that give both a confirmed position and the time at which the target was at that location. The effect of clues is addressed by taking the coordinates of the newly found clue to be the new LKP. The isoprobability curves are then reconstructed based on the

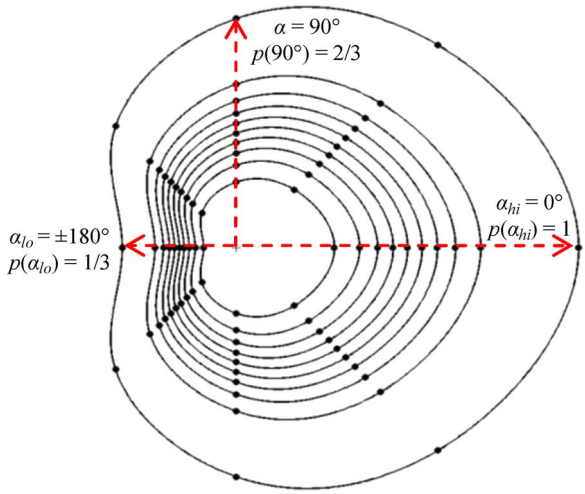

Fig. 5. Isoprobability curves modified to reflect the effect of a likely travel direction.

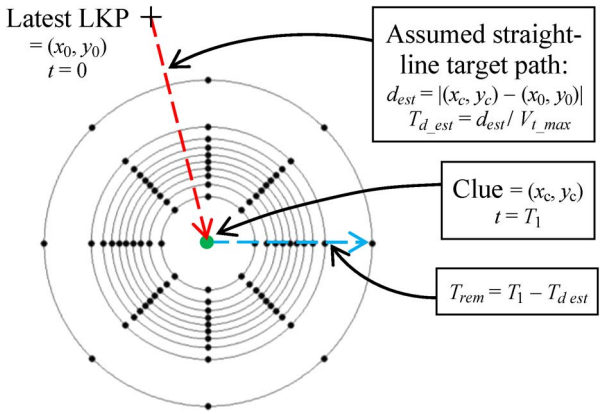

Fig. 6. Isoprobability curve relocation and reconstruction process when a clue is found.

elapsed time since the target had dropped the clue. For clues that only indicate position, conservative speed and path estimates can still be made to estimate the time at which the clue was dropped and thereby estimate the additional time for which the target would have been moving.

For example, if, at some time $t = T_1$, a clue is found without any information on its drop time $t = T_d$, an estimate T_{d_est} is calculated (Fig. 6): First, it is assumed that the target has moved on a straight-line path from the latest LKP position (x_0, y_0) to the newly found clue position (x_c, y_c). This constitutes a conservative path estimate, yielding a travel distance of $d_{\mathrm{est}} = |(x_c, y_c) - (x_0, y_0)|$. Next, a conservative target-speed estimate is made by assuming that he/she moved along this path at the maximum possible mean speed $V_{t\_\max}$, as defined by the upper bound of the nominal mean-target-speed PDF. Consequently, the isoprobability curves are reconstructed for the current time $t = T_1$, where the position of the new LKP becomes (x_c, y_c), and the curves propagate out from this point to a distance corresponding to a time increment of $t = T_{\mathrm{rem}} = T_1 - T_{d\_\mathrm{est}}$. Since the isoprobability curves stretch out further for clues found closer to the LKP (since T_{rem} will be larger), as opposed to clues found further out, the target is still ensured to be contained within the bounds of the 100% contour. This accounts for the inability to know the actual order in which clues without time information are dropped.

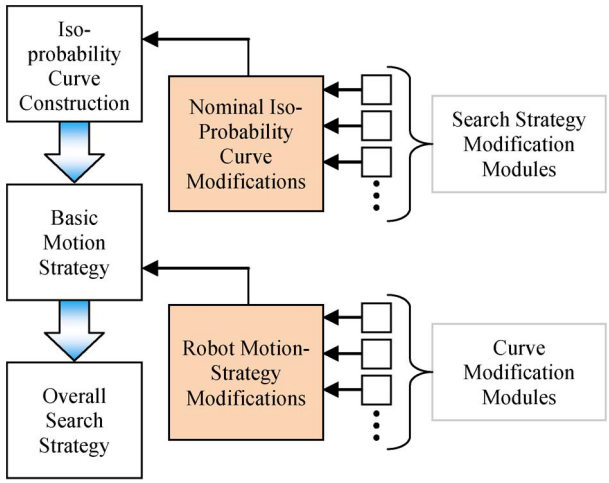

Fig. 7. Modularity of the proposed target-behavior prediction method.

There also exists the issue of uncertainty associated with any clue found as to its relevance to the target being sought. The conservative-approach philosophy applied so far in addressing other issues in our method may also be applied here. In particular, two sets of isoprobability curves can first be created: one for the case where the clue is certainly not associated with the target (i.e., maintain the current LKP and propagate the curves as usual) and one for the case where the clue is known for certain to belong to the target (i.e., update the curves as described previously and shown in Fig. 6). A truly conservative approach would then be to take the union of the curves for these two individual cases as the final set of isoprobability curves, so that both possibilities are considered.

E. Modularity of the Proposed Solution Method

The proposed methodology is modular in that the isoprobability curve concept and the search strategies devised can account for any influences (Fig. 7). Alterations to the nominal mean-target-speed and target-location PDFs, through the application of scale factors, account for influences at the isoprobability-curve-construction level. Modifications to the basic robot-motion strategy, in the form of specific detour motions conducted each time that a given set of conditions is satisfied, address influences at the search-strategy implementation level.

IV. Preliminary Search Strategy

A WiSAR strategy was devised to aid in the robustness analyses to follow, as well as to demonstrate our proposed target-motion prediction methodology. This strategy should be considered just as a means for utilizing the isoprobability curves and not as a proposal.

The devised search method guides a team of robots using a basic motion strategy augmented with specific detour maneuvers to account for obstacles, clues, and possible shelter-seeking psychological behavior displayed by the target.

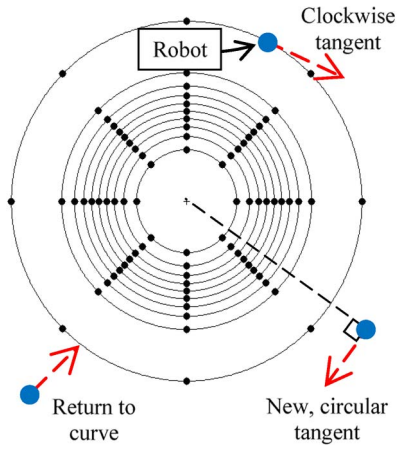

Fig. 8. Basic motion strategy.

A. Initial Robot Deployment

A lack of bias in the relative importance given to each isoprobability curve must be ensured. The relative curve density in different regions of the search area would then control the distribution of the search effort. The number of isoprobability curves can be determined based on the number of robots, ensuring that each curve is assigned with at least one robot.

B. Basic Robot Motion

A basic robot-motion strategy was devised to achieve a balanced search effort. It allows the robots to explore the intermediate regions between adjacent contours while still maintaining robot distribution in accordance with the isoprobability curves.

The basic motion strategy (Fig. 8) requires each robot to perform the following: 1) to start on its assigned curve; 2) to move in a direction tangent to it in a clockwise manner at a predetermined constant speed for a fixed short amount of time; 3) to change direction to move on a new line tangent to an imaginary circle centered at the current LKP, with a radius equal to the distance from the LKP to the current position of the robot; and 4) when the isoprobability curves are propagated outward, to return to its respective curve via the shortest path to restart the basic motion strategy.

C. Obstacle Avoidance

When a robot encounters an obstacle, *a priori* known or unknown, a detour strategy is implemented to circumnavigate it. For an *a priori* known obstacle, the robot traces the obstacle boundary on the side leading to the shortest path around it, which is determined in its entirety before implementing the circumnavigation motion.

For an *a priori* unknown obstacle, the robot extrapolates its current trajectory through the obstacle to determine the point of emergence on the opposite side and then traces the obstacle boundary until this point of emergence is reached. Once the obstacle has been circumnavigated, the robot continues to implement the basic motion strategy.

D. Clues

When a clue is found, the isoprobability curves are relocated and reconstructed. Due to this potentially significant shift, robots return to their respective curves via the shortest paths. This action returns the robots into the region bounded by the new 100% contour.

E. Target Psychology

In our example implementation, a detour strategy was devised to address the behavior of children between the ages of one to six years: When lost, they tend to passively seek out a place of shelter to lie down and sleep after an initial period of random motion [36]. Herein, it is assumed that the possible shelter locations within the search area are known *a priori*. Consequently, the overall motion strategy redirects robots, on the two nearest isoprobability curves within a certain radius from each shelter location, toward the shelter for a close-up investigation.

V. Robustness Analysis and Comparison to a Nonprobabilistic Approach

The effectiveness of isoprobability curves is analyzed in the following for target-behavior prediction in WiSAR scenarios. A measure of curve-fit error is devised to test robustness to the number of rays used. Success rates using a simple multirobot search approach based on the proposed methodology versus a nonprobabilistic technique are also presented.

A. Robustness

The accuracy of isoprobability curves can be improved by increasing the number of rays, i.e., by increasing the number of control points. One could thus conjecture that the "true" set of isoprobability curves can be constructed only by an infinite number of rays and that any other set would represent an approximation. Herein, for practical purposes, the true set of curves is defined as those resulting from the use of uniformly distributed 360 rays. The curve-fit-error metric is defined as the sum of the distances between corresponding points on the 360 rays intersecting the approximated and the true curves, respectively.

A study was conducted to determine how curve-fitting errors vary with increasing number of rays. Fig. 9 shows the plot of error ranges, as well as the average error, for the ray-set sizes tested. As can be noted, the errors in curve fits drop early and rapidly for increasing ray-set sizes. Furthermore, the decrease in the range widths and their increased overlap indicate a lack of significant difference in error-reduction capability for relatively larger number of rays.

In order to further analyze the effects of curve-fitting accuracy, a preliminary study was conducted to determine the impact on search success. These tests involved implementing the simple multirobot search strategy for a WiSAR scenario that was presented in Section IV. Since a theoretical success rate bound could not be determined, 500 different simulations

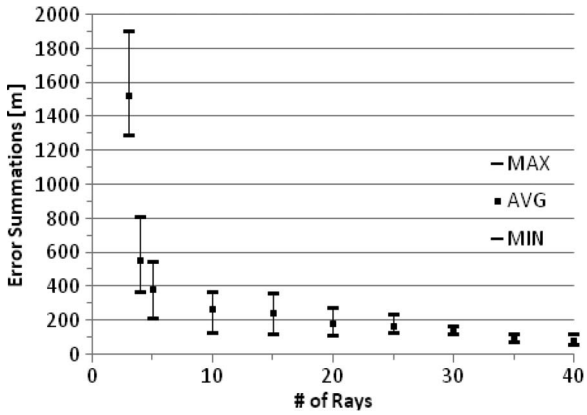

Fig. 9. Error ranges versus ray-set size.

TABLE II
SUMMARY OF OBSERVED SEARCH-SUCCESS RATES

# of Rays	3	5	10	15	25	40
Search Success Rate (%)	45	59	65	90	97	100

(10 target speeds for each of 50 different target directions) were conducted to establish empirical bounds. A search was considered to have *failed* if the target was not found by the 2-h search time limit (Table II). The trend observed from the study on isoprobability-curve-construction accuracy with respect to the number of rays, in Fig. 9, is noted here as well, namely, a rapid increase in success rates, followed by a gradual asymptotic increase.

The aforementioned findings raise the issue of whether optimal-ray placement, i.e., finding the minimum number of rays and their best locations, could be beneficial. In terms of numbers, one could conclude that computational time savings, during curve fitting, from using lesser rays, would be insignificant. Rather, it would be more beneficial to improve curve-fit accuracy through a simple increase in the number of rays. In terms of the locations of rays, similarly, one can note that any potential benefit that can be achieved (i.e., error difference between the worst and best sets of curves when comparing relatively larger and smaller set of rays, respectively) rapidly diminishes as the number of rays increases.

In our research, therefore, it was concluded that the more efficient and effective approach to constructing isoprobability curves would be to select a sufficiently large number of rays, separate them uniformly, and randomly select the angular position of the resulting ray set. However, for a given number of rays, the distance between rays will increase as the isoprobability curves propagate outward with time, thereby effectively decreasing the resolution of the curves and their accuracy. Hence, one may have to increase the number of rays over time to maintain a desired level of curve-fit error.

In order to determine how many rays are sufficient at any time, it would be necessary to estimate how curve-fit errors vary with time for different ray-set sizes. This may be done using intermediate curve-fit data obtained online throughout the search or through empirical data on typical curve-fit errors for the given search environment, if available. Interpolation may then be used to estimate how many rays are required at any given point in time to maintain curve-fit error below a maximum error threshold value.

Another source of isoprobability-curve-based target-behavior-representation inaccuracy is the erroneous guess of the nominal mean-target-speed PDF. If the assumed PDF is overestimated (i.e., it extends further outward from the LKP than would the true PDF), a theoretical maximum success rate of 100% can still be expected, since the searcher robots would still be distributed over the area containing the target. Thus, given enough time, the robots, moving in a sweeping search pattern as described earlier, would eventually find the target. With respect to underestimation, however, the searcher robots may fail to locate the target in some cases, so that the accuracy of the estimated mean-target-speed PDF is the second limitation to the isoprobability curve concept. This robustness though can only be quantified when a fully developed search strategy is available.

B. Comparison to a Nonprobabilistic Method

The purpose of explicitly accounting for the different aspects of target behavior in the target-location PDFs and representing this probabilistic information using isoprobability curves is to provide a means to guide the search process. Thus, in order to provide a basis for comparison, search-success rates were also studied for a nonprobabilistic approach to search (i.e., one in which isoprobability curves were not used).

In order to ensure fairness in the comparison, the same type of search strategy had to be established for the two methods. The simple strategy employed in the aforementioned robustness tests was used herein as well to guide the search robots. In the absence of isoprobability curves, the nonprobabilistic approach guides the search robots on ten random circles corresponding to the ten isoprobability curves used in the proposed method. The speed that each curve corresponds to is selected randomly from a uniform distribution with the upper bound equaling that of the nominal mean-target-speed PDF.

In these simulations, a large number of rays were used to construct the isoprobability curves, since the purpose was solely to compare the effects of using or not using probabilistic information about the target to conduct the search. For both the nonprobabilistic and the proposed methods tested, 500 simulations were run (10 target speeds and 50 target direction settings). The proposed method achieved the theoretical maximum success rate of 100%, while the nonprobabilistic approach was only able to obtain a 25% success rate. One may note that this rate is even lower than that achieved by the relatively inaccurate three-ray curve construction for the probabilistic method (Table II). In this comparison, the number of robots, their sensing ranges, and the maximum search time limit were all held the same so that the impact of the only differing factor, namely, the use of probabilistic target-location information, could be gauged. It should also be noted that, in these simulations, the particular settings used for key parameters, such as the number of robots and contours, as well as the maximum search time limit, were high enough to allow for a 100% success rate even in 500 simulations. However, in general, this is not guaranteed.

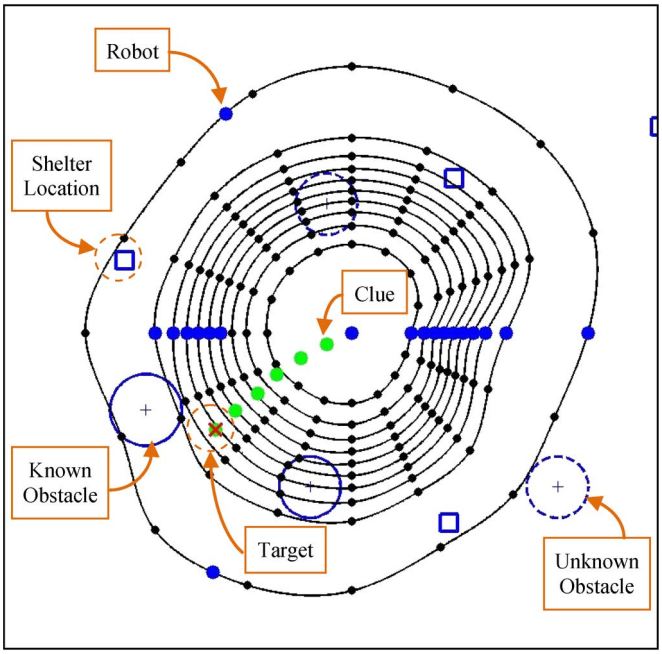

Fig. 10. System state at the start of search ($t = 1800$ s).

VI. Example Application—Autonomous MRC for WiSAR

A realistic multirobot WiSAR search scenario example is presented here in order to demonstrate the potential use of the proposed target-behavior prediction methodology. This example is one of the many that were conducted, using various parameter settings.

A. Simulated Experimental Procedure

A child target of age one to six is assumed to be lost in a forest-type terrain that only allows ground search. Sixteen uniformly distributed rays are used to determine ten isoprobability curves to model target-motion behavior (Fig. 10). The nominal mean-target-speed PDF is scaled online during the search according to the calculated instantaneous terrain slope along the 16 directions (rays).

The simulator randomly chooses a nominal target speed (from the assumed nominal mean-target-speed PDF) and a travel direction for the target. As the worst case scenario, the target is assumed to move outward from the LKP without stopping, with minor random variations to its direction and speed in order to mimic drift. The target is given a head start of 1800 s, representing the time taken for the searcher robots to be deployed. The target is simulated to leave behind clues every 300 s. After 900 s, the target starts to seek shelter and, upon finding a shelter location, moves straight toward it via the shortest path and remains there for the duration of the search.

A total of 19 robots are used in the search, 18 of which are assigned to the 10 isoprobability curves according to their proportional distances from the LKP (1-1-1-2-2-2-2-2-2-3), while the 19th robot is restricted to stay within the area bounded by the innermost (i.e., 10%) curve. The speed of the searcher robots is kept constant.

TABLE III
Nominal Isoprobability Curve Speeds Used for the Simulation

Contour (%)	Nominal Speed (m/s)
10	0.08
20	0.10
30	0.11
40	0.13
50	0.14
60	0.15
70	0.16
80	0.18
90	0.20
100	0.28

B. Results and Discussions

For the particular example simulation presented herein, the mean speed selected for the target was 0.17 m/s, and its initial direction was $-150°$. The target-speed variations were applied by randomly selecting the new speed from a normal distribution with a mean of 0.17 m/s and $\pm 3\sigma = \pm 10\%$ of the mean, giving a variance of 0.0056 m/s. The direction variations were simulated by applying a small randomly selected shift to the angular direction heading of the target using a normal distribution with a mean of $0°$ and $\pm 3\sigma = \pm 15°$. In this way, the target's speed always remained within a 0.15–0.18-m/s range, averaging at 0.17 m/s, while its heading direction would change based on the random (positive or negative) perturbations that would regularly be added on. These random variations to both the speed and the direction of the target were applied every 120 s.

The searcher robots maintained a constant speed of about 1.4 m/s throughout the search. Each robot also had a fixed detection radius of 3 m for clues and 10 m for the target. Furthermore, the robots changed their directions (i.e., recalculated trajectories and headings according to the basic motion strategy described in Section IV) every 30 s.

The assumed nominal mean-target-speed PDF used to construct the isoprobability curves was taken to be a normal distribution with a mean of about $\mu_\nu = 0.14$ m/s and $\pm 3\sigma$ ranging from 0 to 0.28 m/s, yielding a standard deviation of about $\sigma_\nu = 0.046$ m/s. The ten isoprobability curves used ranged from 10% to 100% in 10% increments. Table III summarizes the nominal speeds corresponding to each isoprobability curve. The isoprobability curves were propagated in time and space every 300 s.

The search was conducted using MATLAB V7, on an Intel Core 2 Duo E4600 2.40-GHz processor. Screenshots at four different points in time are shown in Figs. 10–13: The "×" represents the target, the large dark dots on the curves represent the robots, and the large lighter dots trailing the target represent the clues (that have not been found yet). The isoprobability curves are shown in black, where the smaller black dots indicate the cumulative-probability (control) points on the 16 rays; the obstacles are shown by dark hollow circles, where the dashed circles indicate *a priori* unknown obstacles; and the shelter locations are shown by squares.

The system state at the start of the search, i.e., at 1800 s, is shown in Fig. 10. Fig. 11 shows the modification of Robot #16's path (indicated by the arrow and dashed circle) to account for the shelter-seeking target psychological behavior at time 1890 s. Fig. 12 shows the isoprobability curve modification at time

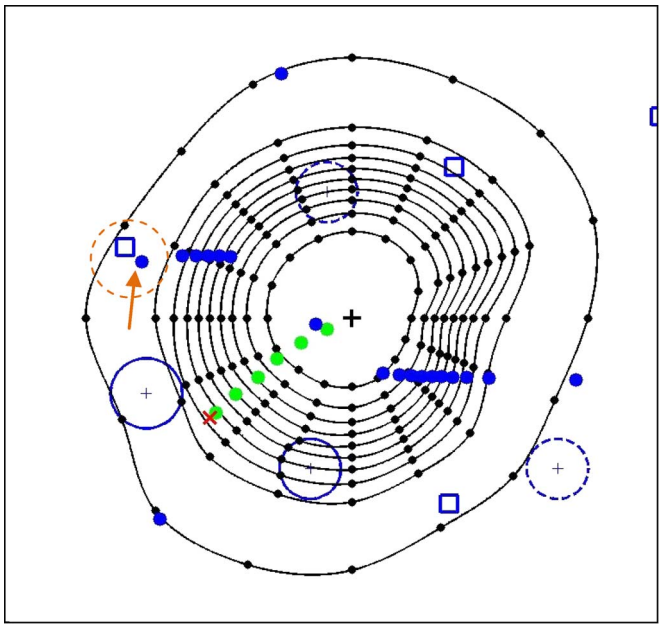

Fig. 11. Robot-search-strategy modification to address shelter-seeking psychological behavior of target ($t = 1890$ s).

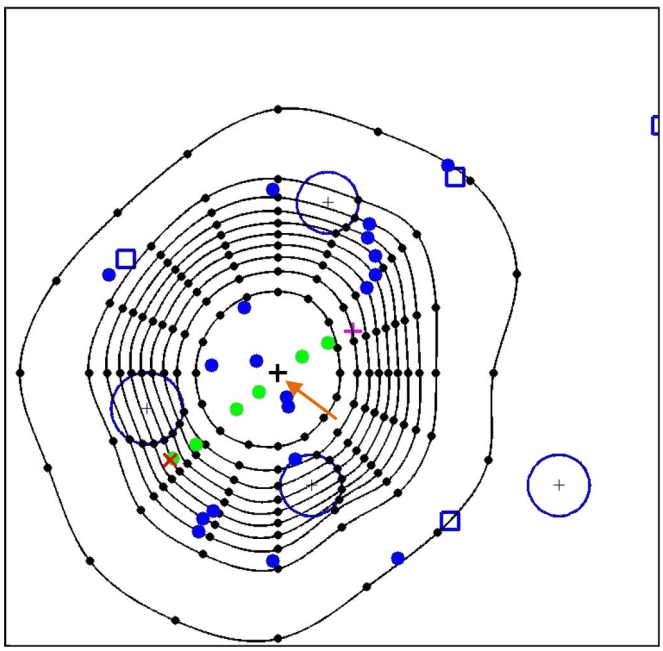

Fig. 12. Isoprobability curve modification due to a clue find ($t = 2440$ s).

2440 s, 30 s after a clue is found (the arrow indicates the new LKP). The screenshot was delayed by 30 s to allow time for the robot that found the clue to move away from the new LKP so that it could be seen. The target is eventually found at 3519 s.

Fig. 13 shows the complete set of paths followed by four of the searcher robots (four thin path lines connected to four large dots, respectively) and that by the target (thick path line connected to the "×"). The "jumps" on the paths are the result of isoprobability curve propagations at discrete time intervals, or isoprobability curve relocations due to clue finds, and robots trying to move to their new curves. An MPEG video of the complete simulation run has also been included with this paper.

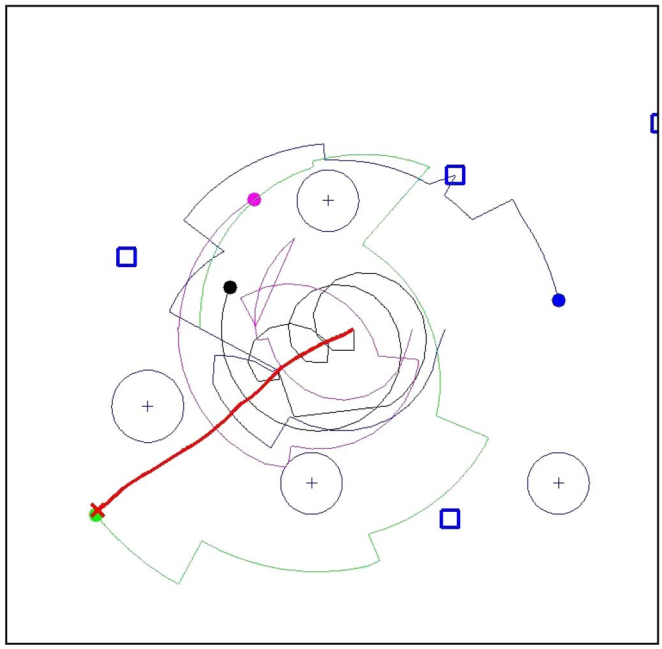

Fig. 13. Complete paths of the target and four of the searcher robots throughout the search process ($t = 3519$ s).

VII. Conclusion

In this paper, a novel methodology for predicting target-motion behavior for WiSAR scenarios has been presented, with a focus on autonomous coordinated search using a team of robots. The proposed modular methodology utilizes the unique key concept of isoprobability curves—an efficient means of representing the target's probable location. Modularity is achieved by independently accounting for different influences on target behavior: terrain topology, target physiology and psychology, clues found, etc. The isoprobability curves are propagated over time and space using the probabilistic target-motion information in the form of a nominal mean-target-speed PDF. Various simulations conducted, ranging from studying curve-fitting errors to search-success rates, verified that the proposed target-behavior prediction methodology effectively addresses the shortcomings of earlier limited-scope approaches reported in the search-theory literature.

References

[1] R. R. Murphy, S. Tadokoro, D. Nardi, A. Jacoff, P. Fiorini, H. Choset, and A. M. Erkmen, "Search and rescue robotics," in *Handbook of Robotics*, B. Siciliano and O. Khatib, Eds. Berlin, Germany: Springer-Verlag, 2008, pp. 1151–1173.

[2] R. R. Murphy and S. Stover, "Rescue robots for mudslides: A descriptive study of the 2005 La Conchita mudslide response," *J. Field Robot.*, vol. 25, no. 1/2, pp. 3–16, Jan. 2008.

[3] F. Bourgault, T. Furukawa, and H. F. Durrant-Whyte, "Optimal search for a lost target in a Bayesian world," in *STAR: Field and Service Robotics*, vol. 22, S. Yuta, Ed. Berlin, Germany: Springer-Verlag, 2006, pp. 209–222.

[4] H. R. Richardson, D. H. Wagner, and J. H. Discenza, "The United States Coast Guard computer-assisted search planning system," *Nav. Res. Logist. Q.*, vol. 27, no. 4, pp. 659–680, Dec. 1980.

[5] M. A. Goodrich, B. S. Morse, D. Gerhardt, J. L. Cooper, M. Quigley, J. A. Adams, and C. Humphrey, "Supporting wilderness search and rescue using a camera-equipped mini UAV," *J. Field Robot.*, vol. 25, no. 1/2, pp. 89–110, Jan. 2008.

[6] C. D. Heth and E. H. Cornell, "Characteristics of travel by persons lost in Albertan wilderness areas," *J. Environ. Psychol.*, vol. 18, no. 3, pp. 223–235, Sep. 1998.
[7] A. Macwan and B. Benhabib, "A multi-robot coordination methodology for autonomous search and rescue," in *Proc. IEEE TIC-STH*, Toronto, ON, Canada, Sep. 2009, pp. 675–680.
[8] B. Lavis, T. Furukawa, and H. F. Durrant-Whyte, "Dynamic space reconfiguration for Bayesian search and tracking with moving targets," *Auton. Robots*, vol. 24, no. 4, pp. 387–399, May 2008.
[9] G. Hollinger, J. Djugash, and S. Singh, "Coordinated search in cluttered environments using range from multiple robots," in *STAR: Field and Service Robotics*, vol. 42, C. Laugier and R. Siegwart, Eds. Berlin, Germany: Springer-Verlag, 2008, pp. 433–442.
[10] D. V. Chudnovsky and G. V. Chudnovsky, Eds., *Search Theory: Some Recent Developments*. New York: Marcel Dekker, 1989.
[11] L. D. Stone, "Generalized search optimization," in *Statistical Signal Processing*, vol. 53, E. J. Wegman and J. G. Smith, Eds. New York: Marcel Dekker, 1984, pp. 265–272.
[12] B. O. Koopman, *Search and Screening: General Principles With Historical Applications*. Toronto, ON, Canada: Pergamon, 1980.
[13] J. R. M. Hosking and J. R. Wallis, *Regional Frequency Analysis: An Approach Based on L-Moments*. New York: Cambridge Univ. Press, 1997.
[14] L. Lin and M. A. Goodrich, "A Bayesian approach to modeling lost person behaviors based on terrain features in wilderness search and rescue," in *Proc. Conf. Behav. Represent. Model. Simul.*, Sundance, UT, Mar./Apr. 2009, pp. 49–56.
[15] J. Borenstein, M. Hansen, and A. Borrell, "The Omni-Tread OT-4 serpentine robot—Design and performance," *J. Field Robot.*, vol. 24, no. 7, pp. 601–621, Jul. 2007.
[16] M. Arai, Y. Tanaka, S. Hirose, H. Kuwahara, and S. Tsukui, "Development of "Souryu-IV" and "Souryu-V": Serially connected crawler vehicles for in-rubble searching operations," *J. Field Robot.*, vol. 25, no. 1/2, pp. 31–65, Jan./Feb. 2008.
[17] J. Casper and R. R. Murphy, "Human–robot interactions during the robot-assisted urban search and rescue response at the World Trade Center," *IEEE Trans. Syst., Man, Cybern. B, Cybern.*, vol. 33, no. 3, pp. 367–385, Jun. 2003.
[18] J. L. Drury, J. Schholtz, and H. A. Yanco, "Awareness in human–robot interactions," in *Proc. IEEE Conf. Syst., Man Cybern.*, Washington, DC, Oct. 2003, vol. 1, pp. 912–918.
[19] Y. Mei, Y.-H. Lu, Y. C. Hu, and C. S. G. Lee, "Deployment of mobile robots with energy and timing constraints," *IEEE Trans. Robot.*, vol. 22, no. 3, pp. 507–522, Jun. 2006.
[20] G. Antonelli and S. Chiaverini, "Kinematic control of platoons of autonomous vehicles," *IEEE Trans. Robot.*, vol. 22, no. 6, pp. 1285–1292, Dec. 2006.
[21] R. Fierro, C. Branca, and J. R. Spletzer, "On-line optimization-based coordination of multi unmanned vehicles," in *Proc. IEEE Int. Conf. Netw., Sens. Control*, Tucson, AZ, Mar. 2005, pp. 716–721.
[22] R. Alur, R. Fierro, A. Das, J. Spletzer, J. Esposito, V. Kumar, J. P. Ostrowski, G. Pappas, C. J. Taylor, Y. Hur, I. Lee, G. Grudic, and J. Southall, "A framework and architecture for multi-robot coordination," *Int. J. Robot. Res.*, vol. 21, no. 10/11, pp. 977–995, Oct./Nov. 2002.
[23] M. B. Dias and A. Stentz, "Enhanced negotiation and opportunistic optimization for market-based multi-robot coordination," The Robotics Inst., Carnegie Mellon Univ., Pittsburgh, PA, Tech. Rep. CMU-RI -TR-02-18, 2002.
[24] Z. Cao, M. Tan, L. Li, N. Gu, and S. Wang, "Cooperative hunting by distributed mobile robots based on local interaction," *IEEE Trans. Rob.*, vol. 22, no. 2, pp. 403–407, Apr. 2006.
[25] Z. Zhang, H. Guo, G. Nejat, and P. Huang, "Finding disaster victims: A sensory system for robot assisted 3D mapping of urban search and rescue environments," in *IEEE Int. Conf. Robot. Automat.*, Rome, Italy, Apr. 2007, pp. 3889–3894.
[26] Z. Zhang and G. Nejat, "Robot-assisted 3D mapping of unknown cluttered USAR environments," *Disaster Robots Special Issue, Adv. Robot. J.*, vol. 23, no. 9, pp. 1179–1198, 2009.
[27] B. Doroodgar, M. Ficocelli, B. Mobedi, and G. Nejat, "The search for survivors: Cooperative human–robot interaction in search and rescue environments using semi-autonomous robots," in *Proc. IEEE Int. Conf. Robot. Automat.*, Anchorage, AK, May 2010, pp. 2858–2863.
[28] J. L. Baxter, E. K. Burke, J. M. Garibaldi, and M. Norman, "Multi-robot search and rescue: A potential field based approach," in *Studies in Computational Intelligence: Autonomous Robots and Agents*, vol. 76, S. Mukhopadhyay and G. Sen Gupta, Eds. Berlin, Germany: Springer-Verlag, 2007, pp. 9–16.
[29] Y. M. Chan, S. Wong, M. C. Foo, and R. Teo, "Engineering intuition for designing multi-robot search and rescue solutions," in *Proc. IEEE Conf. Cybern. Intell. Syst.*, Singapore, Dec. 2004, vol. 2, pp. 1238–1242.
[30] N. Ruangpayoongsak, H. Roth, and J. Chudoba, "Mobile robots for search and rescue," in *Proc. IEEE Int. Workshop Safety, Security Rescue Robot.*, Kobe, Japan, Jun. 2005, pp. 212–217.
[31] P. Lima, L. Custodio, M. I. Ribeiro, and J. Santos-Victor, "The RESCUE project—Cooperative navigation for rescue robots," in *Proc. Int. Workshop Adv. Serv. Robot.*, Bardolino, Italy, Mar. 2003.
[32] I. Suzuki and M. Yamashita, "Searching for a mobile intruder in a polygonal region," *SIAM J. Comput.*, vol. 21, no. 5, pp. 863–888, Oct. 1992.
[33] B. P. Gerkey, S. Thrun, and G. Gordon, "Visibility-based pursuit–evasion with limited field of view," *Int. J. Robot. Res.*, vol. 25, no. 4, pp. 299–315, Apr. 2006.
[34] K. A. Hill, *Lost Person Behavior*. Ottawa, ON, Canada: National SAR Secretariat, 1998.
[35] Y. Kato and Y. Takeuchi, "Individual differences in way-finding strategies," *J. Environ. Psychol.*, vol. 23, no. 2, pp. 171–188, Jun. 2003.
[36] D. Perkins, P. Roberts, and G. Feeney, "Missing person behaviour—An aid to the search manager," The Centre for Search Research, Northumberl, U.K., 2003.
[37] R. J. Koester, *Lost Person Behavior, a Search and Rescue Guide on Where to Look—For Land, Air, and Water*. Charlottesville, VA: dbS Productions, 2008.
[38] M. P. Fairclough, Terragen^{TM} Classic for Windows, Planetside Software, Cheshire, U.K., Sep. 7, 2005. [Online]. Available: www.planetside.co.uk/content/view
[39] R. W. Bohannon, "Comfortable and maximum walking speed of adults aged 20–79 years: Reference values and determinants," *Age and Ageing*, vol. 26, no. 1, pp. 15–19, Jan. 1997.
[40] M. Wojtyra, "Multibody simulation model of human walking," *Mech. Based Des. Struct. Mach.*, vol. 31, no. 3, pp. 357–379, Jan. 2003.

Ashish Macwan received the B.A.Sc. and M.A.Sc. degrees in mechanical engineering from the University of Toronto, Toronto, ON, Canada, in 2000 and 2002, respectively, where he is currently working toward the Ph.D. degree.

His current research interests include multirobot coordination and its applications to autonomous search and rescue.

Goldie Nejat (S'03–M'06) received the B.A.Sc. and Ph.D. degrees in mechanical engineering from the University of Toronto, Toronto, ON, Canada, in 2001 and 2005, respectively.

In 2005, she joined the Department of Mechanical Engineering, State University of New York, Stony Brook, as an Assistant Professor. In 2008, she became an Assistant Professor in the Department of Mechanical and Industrial Engineering, University of Toronto, where she is currently the Director of the Autonomous Systems and Biomechatronics Laboratory. Her current research interests include autonomous sensing, motion planning, control and intelligence of assistive/service robots for search and rescue, exploration, healthcare, and surveillance applications.

Dr. Nejat is a member of the IEEE Robotics and Automation Society and the American Society of Mechanical Engineers.

Beno Benhabib received the B.Sc. degree in mechanical engineering from Bogazici University, Istanbul, Turkey, in 1980, the M.Sc. degree in mechanical engineering from The Technion–Israel Institute of Technology, Haifa, Israel, in 1982, and the Ph.D. degree in mechanical engineering from the University of Toronto, Toronto, ON, Canada, in 1985.

Since 1986, he has been a Professor of mechanical and industrial engineering, electrical and computer engineering, and biomaterials and biomedical engineering at the University of Toronto. His current research interests focus on the development of autonomous robotic systems.