2013 IEEE/RSJ International Conference on
Intelligent Robots and Systems (IROS)
November 3-7, 2013. Tokyo, Japan

Navigation with Foraging

Michael Otte[1], Nikolaus Correll[2], and Emilio Frazzoli[1]

***Abstract*— We propose and study the *navigation with foraging* problem, where an agent with a limited sensor range must simultaneously: (1) navigate to a global goal and (2) forage en route as opportunities to forage are detected. Each foraging act causes a deviation from the shortest path to the long-term goal, with consequences for path length, mission duration, and fuel usage. We analytically calculate and/or bound the expected distance the robot actually travels, given the initial distance to the the global goal. In particular, for either of two non-trivial greedy strategies: (A) forage the point that minimizes goal-heading deviation. (B) forage the closest point ahead of the robot. Our results generalize to problems in higher dimensions.**

I. Introduction

We define *navigation with foraging* as the hybrid problem in which an agent with a limited sensor range must simultaneously: (1) navigate to a global goal and (2) forage en route as opportunities to forage become available. Each act of foraging increases the total distance that the robot must travel (with obvious consequences to path length, mission duration, and fuel usage). The problem is non-trivial, assuming that the agent simultaneously works to achieve both objectives. The cumulative acts of fulfilling many short-term foraging objectives must result in the fulfillment of one long-term navigational objective, and the two objective types occur at significantly different time scales.

To further motivate the problem, we now describe several scenarios in which navigation with foraging occurs:

- *Scientific Exploration*: a landing rover's mission involves visiting a distant crater while sampling interesting chemical/geological features that are detected along the way.
- *Combat*: an unmanned aircraft must visit a valuable target and then return to friendly territory, while also eliminating hostile targets that are detected along the way.
- *Search and Rescue*: a rescue vessel (boat or helicopter) must reach land before running out of fuel, but also desires to rescue people along the way.
- *Intelligence Gathering*: a spy plane's primary mission is to photograph a known enemy installation and then return to safety; however, it is also expected to photograph other unexpected enemy activity that it detects.
- *Salvage Operations*: A sinking ship has jettisoned cargo into the sea. The ship is the primary objective, but salvaging the floating cargo is also desirable.

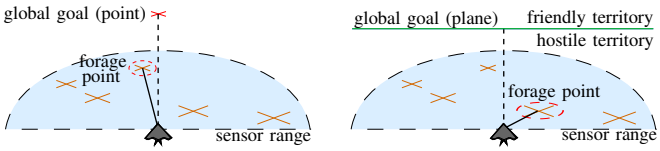

Fig. 1. Greedy navigation with foraging algorithms. Left: the agent forages the point with the heading nearest to that of the long-term global goal. Right: the agent forages the closest point. The global goal can either be a point (Left) or a plane (Right).

In this paper we focus on quantifying the cumulative effects of foraging on path length[1]. In particular, we investigate two special cases of navigation with foraging: (1) always forage the point that minimizes goal-heading deviation, (2) always forage the closest point that has a positive movement component vs. the goal. See Figure 1-Left and Right, respectively. These are the two *non-trivial*[2] pure-strategy extremes—i.e., they are respectively focused on reaching the long-term goal or foraging *as much as possible without foregoing the other mission constraint*.

We model this scenario as a first-order continuous-space Markov processes, and develop tools that allow us to solve for the expectation of the total distance traveled. All formulations are derived with respect to arbitrary dimensionality, and thus immediately generalize to higher dimensional spaces. To the best of our knowledge, we are the first to investigate navigation with foraging algorithms.

This paper is organized as follows: Section II contains a survey of related work. Section III contains the analysis of a single foraging act. Section IV considers foraging while moving toward a boundary (e.g., a geopolitical border, shoreline, etc.). Section V considers foraging while moving toward a point. Both IV and V contain theoretical and experimental results related to the particular scenarios they address. Conclusions are presented in Section VI.

II. Related Work

In [1] a predator-prey patch model is used to evaluate which tasks an agent should perform, assuming tasks are encountered randomly within a patch (a bounded subset of the environment), and the agent may move between patches. Tasks are defined broadly, and are arguably analogous to our forage points. The main theoretical contribution of [1]

[1]Michael Otte and Emilio Frazzoli are with the The Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA, ottemw@mit.edu

[2]Nikolaus Correll is with the Department of Computer Science, University of Colorado at Boulder, Boulder, Colorado, USA.

[1]Assuming constant speed, path length can be used to calculate the effects of foraging on mission duration and fuel usage.

[2]In contrast, *trivial* pure strategies involve ignoring one of the two mission requirements: (1) Move directly to the goal without foraging. (2) Always forage the closest point (which will reach the goal with probability 0).

is an analysis showing the subset of tasks that are expected to maximize an agent's long-term reward. There is no long-term navigational objective, and tasks are ranked based on a predefined (and static) expected return on foraging effort. [2] explores a similar multi-agent scenario. In contrast to [1] and [2], we consider a navigation with foraging scenario, study how path-length is affected by task density (i.e., using the language of [1]), and rank tasks based on their spatial location vs. the long-term navigational goal.

Other previous work on agent based foraging has primarily focused on communal and/or emergent foraging behavior in artificial colonies. Ant-like foraging is a canonical problem in the multi-agent robotic domain [3]. Common areas of focus include: statistics of the time required to find food and/or knowledge propagation within the colony as functions of ant speed, memory, life span, communication modality (e.g., pheromone trail and/or contact based knowledge exchange), and colony size. Early papers date back to the 1980s [4], [5], and there has been a steady stream of work to date, e.g., [6], [7], and [8]. Indeed, the field is so broad that it is impossible to do it justice here. The main differences between ant colony foraging and our work can be summarized as follows: Ant colony foraging is primarily concerned with the emergent behavior of a multi-agent system, the effects of communication, and a scenario where agents have the long-term goal of discovering resources and relocating them to a *nest*, i.e., a specific location to which the agents must also return. In contrast, our work focuses on a single nomadic agent (i.e., not assumed to return to the starting position), that must balance opportunistic local foraging with global path efficiency; we investigate how different algorithms, resource density, and dimensionality affect long-term path length.

Navigation, itself, is also an entire sub-field of artificial intelligence and another canonical robotics problem. However, from early works on bug algorithms [9], [10] and grid-based planning [11], [12] to more modern random graph techniques [13], [14], [15], nearly all previous navigational work in robotics and artificial intelligence has focused on navigation that *avoids* obstacles or other robots—usually by path planning or motion-planning, with an emphasis on *planning*. In contrast, our work is on navigation that *seeks* local points of interest, and does not produce a detailed long-term plan *a priori*.

Therefore, while it may be easy to misinterpret our work as belonging to the path planning sub-field of navigation, this is a critical error. Although the movement we investigate arguably constitutes a path, that path is not *planned* in the formal sense; rather, it emerges due to conflicting mission requirements. In this respect, our work shares similarities with reactive planning ideas [16], [17]. The two main differences between our work and other reactive planning ideas are: (1) We study a scenario in which local movement that *seeks* random opportunities is chosen, partially, based on the location of a static global objective, and (2) we are able to compute analytical expressions for the expected total distance the agent eventually travels.

Local navigation to successive waypoints along a predetermined sequence has also been studied (e.g., [18]); usually with a focus on calculating and/or learning a locally optimal planner or control policy for moving to the next waypoint. The origin of the sequence is largely irrelevant to the local plan or policy. In contrast, we investigate how conflicting local and global objectives influence the resulting emergent path.

$\ell_{i,*} = \mathbf{GreedyHeading}(r, n, x_i, \hat{\mathbf{g}}_i)$
1: $\mathbf{L}_i = \mathbf{Random}(r, n, x_i, \hat{\mathbf{g}}_i)$
2: **return** $\arg\min_{\ell_{i,j} \in \mathbf{L}_i} (\phi_{i,j})$

$\ell_{i,*} = \mathbf{GreedyProximity}(r, n, x_i, \hat{\mathbf{g}}_i)$
1: $\mathbf{L}_i = \mathbf{Random}(r, n, x_i, \hat{\mathbf{g}}_i)$
2: **return** $\arg\min_{\ell_{i,j} \in \mathbf{L}_i} (\|\ell_{i,j} - x_i\|)$

Fig. 2. Greedy foraging algorithms. r, n, x_i are the sensor radius, number of random local points considered at each iteration, and the agent's current location, respectively. $\hat{\mathbf{g}}_i$ is a unit vector that points at the goal from x_i. $\mathbf{L}_i$ is a set of random points found at iteration i, where $|\mathbf{L}_i| = n$. The subroutine $\mathbf{Random}(r, n, x_i, \hat{\mathbf{g}}_i)$ draws n points $\ell_{i,j}$ for $1 \leq j \leq n$, such that $\|\ell_{i,j} - x_i\| \leq r$ and $\|(\ell_{i,j} - x_i) \cdot \hat{\mathbf{g}}_i\| > 0$. The best member of $\mathbf{L}_i$, as defined by the greedy algorithm, is x_{i+1}. $\mathbf{GreedyHeading}()$ returns the member of $\mathbf{L}_i$ with the smallest angle $\phi_{i,j}$ away from the g_i-axis. $\mathbf{GreedyProximity}()$ returns the member of $\mathbf{L}_i$ that is closest to x_i.

III. Local movement

An agent desires to move to a global goal in a D dimensional Euclidean space $\Re^D$. However, it also desires to forage en route by visiting "locally interesting" points that it discovers along the way (i.e., "forage points"). The agent is equipped with a 180-degree sensor with range r, that it points at the long-term goal when searching for foragable points. Movement happens in a sequence of iterations. In each iteration the agent scans for n interesting points and then moves to the "best" one—as defined by one of the following two algorithms:

1) The agent moves to the point that requires the least amount of angular deviation from the long-term heading.
2) The agent moves to the point that is closest to its current location (and also in front of the robot).

Scenarios (1) and (2) are depicted in Figure 1-Left and -Right, respectively, and formalized in algorithms $\mathbf{GreedyHeading}()$ and $\mathbf{GreedyProximity}()$ in Figure 2-Top and -Bottom, respectively. By construction, the movement component along the desired heading is never negative.

We assume that the n points are independent and identically distributed (i.i.d.) uniformly at random at each iteration[3]. We assume that the agent moves directly to the appropriate forage point at each iteration[4]. We believe our

[3]This can happen in a number of interesting situations; e.g., whenever the time required to study/collect/process an interesting point is much greater than both the time required to move to it and the time in which the points shift location; or when the act of moving to a point causes the other points to redistribute randomly (for instance by "scaring them away").

[4]Although rotating in place is impossible for many vehicles (e.g., airplanes), our results still provide a reasonable approximation when the distance required to perform a rotation is small relative to the distance between forage points.

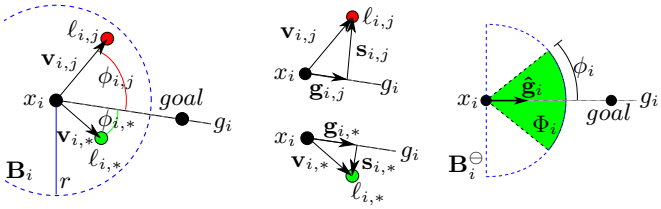

Fig. 3. $\ell_{i,*}$ (green/light gray) is the 'best' point discovered from x_i, another point is $\ell_{i,j}$ (red/dark gray). The unit vector $\hat{\mathbf{g}}_i$ points at the goal along the g_i-axis. Angles away from $\hat{\mathbf{g}}_i$ and vectors from X_i to $\ell_{i,*}$ and $\ell_{i,j}$ are also depicted, as is the D-ball $\mathbf{B}_i$ (dashed blue circle). Φ_i is the hypersector (green/light gray) defined by ϕ_i measured from the g axis, the goal is located along the g_i axis, and x_i is located at $x = 0$. $B_{1/2}$ is the portion of the D-Ball located in the positive g_i-direction.

technique can be modified to handle many non-Euclidean spaces, non-uniform distributions, and dynamics, but leave these for future work.

Let X_i denote the location of the robot at iteration i. The agent's movement can be modeled as a first-order continuous-space Markov process because the stochastic process of movement from one foraged point to the next depends only on the agent's location at the former point—and not the history by which it arrived there. Formally, if the agent starts at $X_1 = x_1$ and then visits a sequence of locations $x_i \in \Re^D$ for all $1 \leq i \leq k+1$ by making k moves, then

$$\mathbb{P}(X_i = x_i | X_{i-1} = x_{i-1}, \cdots, X_1 = x_1) \equiv \mathbb{P}(X_i = x_i | X_{i-1} = x_{i-1})$$

To avoid confusion with the navigational notion of time, we use the term "iteration" instead of "time" to describe the basic index of movement.

We use the subscripts '$\circ$' and '$\diamond$' to denote "best" with respect to $\mathbf{GreedyHeading}()$ and $\mathbf{GreedyProximity}()$, respectively. We use '$*$' as a proxy for '$\circ$' and '$\diamond$' in discussion/derivations that apply to both $\mathbf{GreedyHeading}()$ and $\mathbf{GreedyProximity}()$.

The rest of this section is devoted to understanding the effect that a single foraging act has on path length. In particular, we derive compact analytical expressions for the robot's expected movement, as a function of D, r, n, and the particular algorithm being used. A key insight is that the expected change in path length due to foraging at iteration i is related to the expectation of trigonometric functions of $\phi_{i,*}$, where $\phi_{i,*}$ is the angle between the optimal navigational heading and the optimal forage heading.

Let $\mathbf{L}_i$ be the set of foragable points that are available at iteration i (when the robot is at $X_i = x_i$), and let $\ell_{i,*}$ be the "best" member of $\mathbf{L}_i$. We shall consider the expectation of movement from x_i to $\ell_{i,*}$, assuming the long-term goal is further than r from x_i (in the more interesting scenarios in Sections IV and V we will drop this assumption). Let the vector $\mathbf{v}_{i,j} = \ell_{i,j} - x_i$, and let $\|\cdot\|$ denote magnitude. $\|\mathbf{v}_{i,j}\| = \|\ell_{i,j} - x_i\| < r$ for all $\ell_{i,j} \in \mathbf{L}_i$.

***Nomenclature note**: we use variations of the letter 'L' to represent forage points to emphasize their "local" effects. Similarly, we use variations of the letter 'G' for quantities associated with the long-term "global" navigational goal. Further, we use boldfaced uppercase to denote sets and boldfaced lowercase to denote vectors.*

Without loss of generality, assume that x_i is at the origin of a local coordinate system such that the g_i-axis contains the goal (see Figure 3-Left). By construction, we consider only movement in the positive g_i-direction. Let $\hat{\mathbf{g}}_i$ be a unit vector located at x_i that points along the g_i-axis. Let $\mathbf{g}_{i,j}$ be the projection of $\mathbf{v}_{i,j}$ onto the g_i-axis (Figure 3-Center).

$\mathbf{B}_i$ is defined as the D-ball of radius r centered at x_i. Let $\mathbf{B}_i^{\ominus}$ refer to the half of the D-ball that exists in the non-negative g_i-direction. $\mathbf{L}_i \subset \mathbf{B}_i^{\ominus}$. Let ϕ_i measure the angular distance from $\hat{\mathbf{g}}_i$, where $-\pi/2 \leq \phi_i \leq \pi/2$. The angle between $\hat{\mathbf{g}}_i$ and $\mathbf{v}_{i,j}$ is $\phi_{i,j}$, and the angle between $\hat{\mathbf{g}}_i$ and $\mathbf{v}_{i,*}$ is $\phi_{i,*}$.

Given a particular movement from x_i to $\ell_{i,j}$, the component of that movement toward the goal is $\mathbf{g}_{i,j}$. Let $\mathbf{g}_{i,j} + \mathbf{s}_{i,j} = \mathbf{v}_{i,j}$, where $\mathbf{s}_{i,j}$ is the component of $\mathbf{v}_{i,j}$ that is perpendicular to $\mathbf{g}_{i,j}$.

We shall use $\mathbb{P}_n(\cdot)$ and $\mathbb{E}_n(\cdot)$ to denote the probability density and the expectation of a quantity when $|\mathbf{L}_i| = n$, respectively. Given our i.i.d. uniform sampling assumptions, $\mathbb{P}(\|\ell_{i,j} - x_i\|)$ (proximity probability density) and $\mathbb{P}(\phi_{i,j})$ (heading probability density) are statistically independent; as a result, $\mathbb{P}_n(\|\ell_{i,j} - x_i\|)$ and $\mathbb{P}_n(\phi_{i,*})$ are also statistically independent. We now prove two lemmas based on this fact, regarding the effects of a single foraging act on the global path length.

Lemma 1: $\mathbb{E}_n(\|\mathbf{g}_{i,*}\|) = \mathbb{E}_n(\|\mathbf{v}_{i,*}\|)\mathbb{E}_n(\cos(\phi_{i,*}))$

Proof: $\mathbb{P}_n(\|\ell_{i,j} - x_i\|)$ and $\mathbb{P}_n(\phi_{i,*})$ are statistically independent; thus, so are $\mathbb{P}_n(\|\ell_{i,*} - x_i\|)$ and $\mathbb{P}_n(\cos(\phi_{i,*}))$. Note $\|\mathbf{v}_{i,*}\| = \|\ell_{i,*} - x_i\|$. The expectation operator supports multiplicativity between statistically independent variables; therefore, $\mathbb{E}_n(\|\mathbf{v}_{i,*}\|\cos(\phi_{i,*})) = \mathbb{E}_n(\|\mathbf{v}_{i,*}\|)\mathbb{E}_n(\cos(\phi_{i,*}))$, and $\|\mathbf{v}_{i,*}\|\cos(\phi_{i,*}) = \|\mathbf{g}_{i,*}\|$, since $-\pi/2 \leq \phi_{i,*} \leq \pi/2$. ■

Lemma 1 shows that $\mathbb{E}_n(\cos(\phi_{i,*}))$ relates the *expectations* of incremental movement toward the long-term goal vs. the actual movement required for local foraging at iteration i. This result is very intuitive—given that $\cos(\phi_{i,*})$ is the ratio between movement toward the goal vs. movement to point $\ell_{i,*}$—however, the functional non-invariance of the expectation operator requires that we prove it explicitly. The following corollary provides a similar result for $\mathbb{E}_n(\|\mathbf{s}_{i,*}\|)$.

Corollary 1: $\mathbb{E}_n(\|\mathbf{s}_{i,*}\|) = \mathbb{E}_n(\|\mathbf{g}_{i,*}\|)\mathbb{E}_n(\tan(|\phi_{i,*}|))$

Proof: $\mathbb{P}_n(\tan(|\phi_{i,*}|))$ and $\mathbb{P}_n(\|\mathbf{s}_{i,*}\|)$ are statistically independent. The rest of the proof is similar to Lemma 1, except that $\tan(|\phi_{i,*}|)$, $\mathbf{s}_{i,*}$, and $\mathbf{g}_{i,*}$ are used in place of $\cos(\phi_{i,*})$, $\mathbf{g}_{i,*}$, and $\mathbf{v}_{i,*}$, respectively. ■

Table I shows values of $\mathbb{E}_n(\cos(\phi_{i,\circ}))$ and $\mathbb{E}_n(\tan(|\phi_{i,\circ}|))$ for select D and n. Full derivations, including $\mathbb{E}_n(\cos(\phi_{i,\diamond}))$ and $\mathbb{E}_n(\tan(|\phi_{i,\diamond}|))$, are treated in the Appendix. However, we note that:

$$\mathbb{E}_n(\cos(\phi_{i,\diamond})) = \mathbb{E}_1(\cos(\phi_{i,\circ})) \quad (1)$$

$$\mathbb{E}_n(\tan(|\phi_{i,\diamond}|)) = \mathbb{E}_1(\tan(|\phi_{i,\circ}|)) \quad (2)$$

see the Appendix for more details.

TABLE I

SPECIAL CASES OF $\mathbb{E}_n(\cos(\phi_{i,\circ}))$ AND $\mathbb{E}_n(\tan(|\phi_{i,\circ}|))$

$\mathbb{E}_n(\cos(\phi_{i,\circ}))$

D	n				
	1	2	3	4	n
2	$\frac{2}{\pi}$	$\frac{8}{\pi^2}$	$\frac{24(\pi-2)}{\pi^3}$	$\frac{48(\pi^2-8)}{\pi^4}$	
3	$\frac{1}{2}$	$\frac{2}{3}$	$\frac{3}{4}$	$\frac{4}{5}$	$\frac{n}{n+1}$

$\mathbb{E}_n(\tan(|\phi_{i,\circ}|))$

D	n				
	1	2	3	4	n
2	∞	$\frac{\log(16)}{\pi}$	$\frac{3(\pi^2\log(4)-7\zeta(3))}{\pi^3}$	$\frac{4(\pi^2\log(4)-9\zeta(3))}{\pi^3}$	
3	$\frac{\pi}{2}$	1	$\frac{\pi}{4}$	$\frac{2}{3}$	$\frac{\sqrt{\pi}n\Gamma(\frac{n-1}{2})}{4\Gamma(1+n/2)}$

$\Gamma(\cdot)$ and $\zeta(\cdot)$ are the gamma and Riemann zeta functions, respectively

ForageHeadingA$(r, n, x_1, d_{goal}^{plane})$

$x_i = x_1$
while $x_i \notin \{x_{goal}\}$ **do**
  $x_{i+1} = \mathbf{Greedy}(r, n, x_i, \hat{\mathbf{g}}_1)$
  if $d_{goal}^{plane} - \|(x_{i+1} - x_1) \cdot \hat{\mathbf{g}}_1\| > 0$ **then**
    $x_i = x_{i+1}$
  else
    $x_i = x \in \overline{x_i x_{i+1}} \cap \{x_{goal}\}$

ForageHeadingB$(r, n, x_1, d_{goal}^{plane})$

$x_i = x_1$
while $d_{goal}^{plane} - \|(x_i - x_1) \cdot \hat{\mathbf{g}}_1\| > r$ **do**
  $x_{i+1} = \mathbf{Greedy}(r, n, x_i, \hat{\mathbf{g}}_1)$
  $x_i = x_{i+1}$
$x_i = x_i + \left(d_{goal}^{plane} - \|(x_i - x_1)\hat{\mathbf{g}}_1\|\right)\hat{\mathbf{g}}_1$

Fig. 4. Notation is similar to Figure 3. x_1 is the agent's initial position. The goal plane is at $g_1 = d_{goal}^{plane}$. At iteration i the agent's location is x_i, and x_{i+1} is the best point as defined by the greedy algorithm. $\overline{x_i x_{i+1}}$ is the line segment between x_i and x_{i+1}, while $\{x_{goal}\}$ is a set containing all points in the goal plane. In **ForageHeadingA**() the robot moves toward a locally interesting point on the last iteration, but stops at the goal while en route. In **ForageHeadingB**() the agent does not forage on the last iteration, but moves directly to the goal hyperplane.

IV. DIRECTIONAL NAVIGATION WITH FORAGING

In this section we consider the case where an agent's long-term goal is to reach a hyperplane located at a distance d_{goal}^{plane} from the agent. For example, a geopolitical boundary (such as country's border) or geographic boundary (such as shoreline). Without loss of generality, we assume the agent's initial global position x_1 is at the origin and the goal hyperplane is at $g_1 = d_{goal}^{plane}$.

The final movement to the goal can be addressed in two different ways, depending on if the agent favors the short- or long-term objective on the final move (see Figure 4). In algorithm **ForageHeadingA**() the agent never forages the last point that it detects, but moves toward it until reaching the goal. In algorithm **ForageHeadingB**() the agent moves directly to the goal whenever it is closer than r (See Figure 5). Both of these algorithms maintain the Markov property. However, $\mathbb{E}_n(\phi_{i,*})$ and $\mathbb{E}_n(\cos(\phi_{i,*}))$ remain unchanged

A simulated robot moving to a planer goal (20 trials)

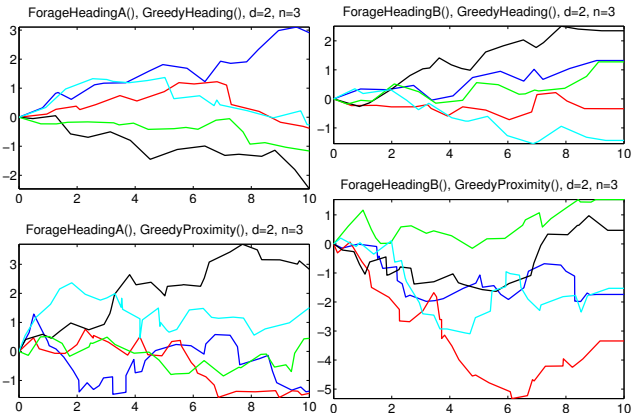

Fig. 5. A simulated robot navigates to a planer goal while also foraging. Each sub-figure contains five paths from $x_i = [0, 0]$ to the goal at $g_1 = d_{goal}^{plane} = 10$, given a particular combination of greedy algorithm, goal behavior, dimension D, and $n = |\mathbf{L}_i|$, and sensor radius $r = 1$.

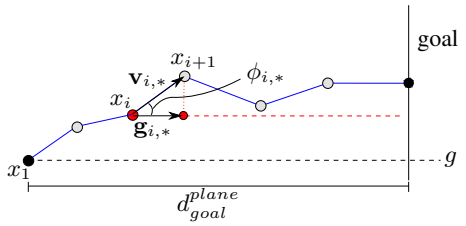

Fig. 6. Path from x_1 to a planar goal at $x = L$, and related quantities

for the final iteration of algorithm **ForageHeadingA**(), while they are not the same for the final iteration of algorithm **ForageHeadingB**(). That said, **ForageHeadingA**() and **ForageHeadingB**() are identical on all but the last move; thus, their differences vanish as $r/d_{goal}^{plane} \to 0$.

When the goal is a hyperplane at $g_1 = d_{goal}^{plane}$, all local g_i-axis are parallel due to symmetry. It is possible to define all local coordinate systems such that the transformation between them is a translation in the g_1-direction. As a result, statistics regarding the *incremental movement* between states i and $i+1$ are identical at all i such that $1 \le i < k$. For convenience, we define the global coordinate system to be the local coordinate system at $X_1 = x_1$. Let $\mathbf{P}_*$ be the path taken by the agent. Formally, $\{x_1, \cdots, x_{k+1}\} = \mathbf{P}_*$ such that $X_1 = x_1, \cdots, X_{k+1} = x_{k+1}$. Recall that $\mathbf{v}_{i,*}$ is the vector defined by the i-th movement along $\mathbf{P}_*$. Given our coordinate system, $\mathbf{v}_{i,*} = x_{i+1} - x_i$ (see Figure 6). Let ΔX_i denote the state transition at iteration i. Then for all i such that $1 \le i < k$ (note $i \ne k$)

$$\mathbb{P}(\Delta X_i = \mathbf{v}_{i,*} | X_{i-1} = x_{i-1}) \equiv \mathbb{P}(\Delta X_i = \mathbf{v}_{i,*}) \equiv \mathbb{P}(\Delta X_1 = \mathbf{v}_{1,*}).$$

Thus, understanding the behavior at $i = 1$ is equivalent to understanding the behavior at all other iterations except $i = k$. This simplifies the analysis. Further, we have already evaluated the relevant movement for when $i \ne k$ in the previous section, i.e., when $\mathbb{P}(\Delta X_1 = \mathbf{v}_{1,*}) \equiv \mathbb{P}(X_2 = x_2 | X_1 = x_1)$.

$\|\mathbf{P}_*\| = \sum_{i=1}^{k} \|\mathbf{v}_{i,*}\|$ is the cumulative length of $\mathbf{P}_*$. The rest of this section is devoted to calculating and/or bounding the expected path length $\mathbb{E}_n(\|\mathbf{P}_*\|)$ as simple functions of $\mathbb{E}_n(\cos(\phi_{i,*}))$. Casual readers wishing to skip the details of

the proofs should at least note the statements of Corollaries 2 and 3, as well as the favorable agreement between the theoretical and experimental results (Figure 7).

Recall that $\mathbf{g}_{i,*}$ is the projection of $\mathbf{v}_{i,*}$ onto the g_i-axis. $\|\mathbf{v}_{i,*}\|$ and $\|\mathbf{g}_{i,*}\|$ are the magnitudes of $\mathbf{v}_{i,*}$ and $\mathbf{g}_{i,*}$, respectively, and $\phi_{i,*}$ is the angle between $\mathbf{v}_{i,*}$ and $\mathbf{g}_{i,*}$, and $\cos(\phi_{i,*}) = \|\mathbf{g}_{i,*}\|/\|\mathbf{v}_{i,*}\|$. Although the value of k is random for any particular run of $\mathbf{ForageHeadingA}()$ or $\mathbf{ForageHeadingB}()$, our analysis only requires that iteration k be the final movement to the goal. $\mathbb{E}_n(\|\mathbf{v}_{k,*}\|)$ is the expected distance that the agent moves during the final iteration and $\mathbb{E}_n(\|\mathbf{g}_{k,*}\|)$ is the length of the projection of the final movement onto the g_1-axis.

Theorem 1: $\mathbb{E}_n(\|\mathbf{P}_*\|) = \frac{d_{goal}^{plane} - \mathbb{E}_n(\|\mathbf{g}_{k,*}\|)}{\mathbb{E}_n(\cos(\phi_{1,*}))} + \mathbb{E}_n(\|\mathbf{v}_{k,*}\|)$ for navigation to a planar goal with foraging.

Proof: By construction and the linearity of expectation $\mathbb{E}_n(\|\mathbf{P}_*\|) = \sum_{i=1}^{k} \mathbb{E}_n(\|\mathbf{v}_{i,*}\|)$. Substituting from Lemma 1 gives: $\mathbb{E}_n(\|\mathbf{P}_*\|) = \mathbb{E}_n(\|\mathbf{v}_{k,*}\|) + \sum_{i=1}^{k-1} \frac{\mathbb{E}_n(\|\mathbf{g}_{i,*}\|)}{\mathbb{E}_n(\cos(\phi_{i,*}))}$. We know that $\mathbb{P}(\Delta X_i = \mathbf{v}_{i,*}) \equiv \mathbb{P}(X_2 = x_2 | X_1 = x_1)$ for $1 \leq i < k$, and so $\mathbb{E}_n(\cos(\phi_{i,*})) = \mathbb{E}_n(\cos(\phi_{1,*}))$. Also by construction, $\sum_{i=1}^{k} \|\mathbf{g}_{i,*}\| = d_{goal}^{plane}$ and so by the linearity of expectation, and also that $\mathbb{E}_n(d_{goal}^{plane}) = d_{goal}^{plane}$, we have $d_{goal}^{plane} - \mathbb{E}_n(\|\mathbf{g}_{k,*}\|) = \sum_{i=1}^{k-1} \mathbb{E}_n(\|\mathbf{g}_{i,*}\|)$. Substitution finishes the proof. ■

Corollary 2: $\mathbb{E}_n(\|\mathbf{P}_*\|) = \frac{d_{goal}^{plane}}{\mathbb{E}_n(\cos(\phi_{1,*}))}$ for algorithm $\mathbf{ForageHeadingA}()$.

Proof: Lemma 1 is also valid when $i = k$ because stopping the agent at the global goal (i.e., somewhere between x_k and $\ell_{k,*}$) does not change $\mathbb{E}_n(\phi_{k,*})$ or $\mathbb{E}_n(\cos(\phi_{k,*}))$ ■

Corollary 3: $\frac{d_{goal}^{plane} - r}{\mathbb{E}_n(\cos(\phi_{1,*}))} + r \leq \mathbb{E}_n(\|\mathbf{P}_*\|) \leq \frac{d_{goal}^{plane}}{\mathbb{E}_n(\cos(\phi_{1,*}))}$ for algorithm $\mathbf{ForageHeadingB}()$.

Proof: By construction $0 \leq \mathbb{E}_n(\|\mathbf{g}_{k,*}\|) \leq r$ and $0 \leq \mathbb{E}_n(\|\mathbf{v}_{k,*}\|) \leq r$. ■

Figure 7 shows statistics from experiments with a **simulated robot** superimposed on the expected values predicted by our analytical results. 10000 experiments are performed per each algorithm combination. Various values of d and n are used, while $d_{goal}^{plane} = 10$ and $r = 1$. The expected values and bounds are within 0.005% and 0.02% of the average experimental path length, respectively. Note that we should expect the bounds to approach an equality as $d_{goal}^{plane}/r \to \infty$ and/or $n \to \infty$.

V. Point to Point Navigation with Foraging

We now consider the case where the long-term goal is a point. The agent uses the $\mathbf{ForageGoalPoint}()$ algorithm in Figure 8. Examples of paths taken by a simulated robot using $\mathbf{ForageGoalPoint}()$ are shown in Figure 9.

Without loss of generality, we define the local coordinate system at each $X_i = x_i$ such that x_i is at the origin and the goal point on the g_i-axis. Let d_{goal} be the distance between X_i and the goal. As in the previous section, we use the local coordinate system at $X_1 = x_1$ as our global coordinate system. However, it is important to note that

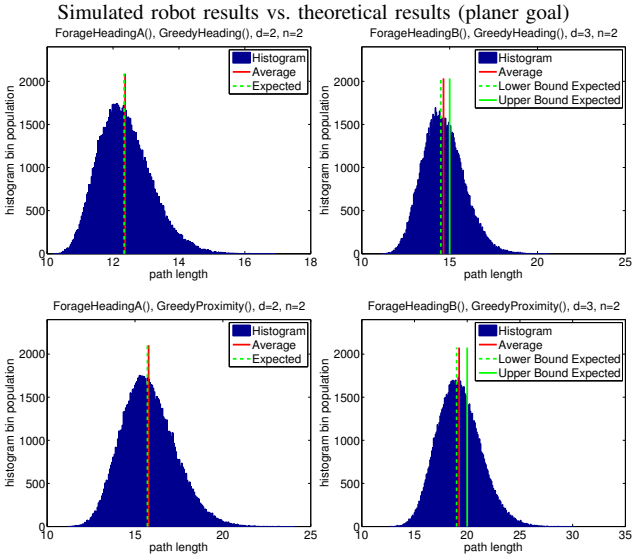

Fig. 7. Observed statistics over 100000 experimental trials with a simulated robot over four different combinations of algorithm, goal behavior, dimension d, and local forage points $n = |\mathbf{L}_i|$. In all experiments sensor radius $r = 1$ and $d_{goal}^{plane} = 10$. Histograms show the distributions of experimental path lengths and red-line depicts the mean value. Green line(s) show the theoretically predicted expectations based on our analytical results (exact value or upper and lower bounds, depending on algorithm).

$\mathbf{ForageGoalPoint}(r,n,x_1,x_{goal})$

1: $x_i = x_1$
2: **while** $\|x_{goal} - x_i\| > r$ **do**
3: $\quad x_{i+1} = \mathbf{Greedy}(r, n, x_i, \hat{\mathbf{g}}_i)$
4: $\quad x_i = x_{i+1}$
5: $x_i = x_{goal}$

Fig. 8. Notation similar to Figures 3 and 4. x_{goal} is the long-term goal.

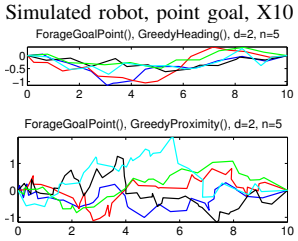

Fig. 9. Paths taken by a simulated robot navigating from $x_i = [0,0]$ to $x_{k+1} = [10,0]$. Five paths each for $\mathbf{GreedyHeading}()$ (Top) and $\mathbf{GreedyProximity}()$. $D = 2$, $n = |\mathbf{L}_i| = 5$, $r = 1$.

$\mathbf{g}_{i,*}$ is now defined as the projection of $\mathbf{v}_{i,*}$ onto the g_i-axis (and not the g_1 axis, in general). Unlike the previous section, g_i-axis and g_j-axis are not parallel for $i \neq j$ (with probability 1). Although the Markov property is maintained, the computation of an exact $\mathbb{E}_n(\|\mathbf{P}_*\|)$ becomes difficult, due to the loss of translational symmetry, and we must settle for computing bounds instead. We begin with a relatively tight lower bound, before moving on to calculate a loose upper bound.

The basic idea is to show that the point-goal problem can be transformed into the plane-goal problem of the previous section, but that the transformation increases the effective d_{goal}^{plane} of the resulting plane-goal problem vs. the d_{goal} of the original point-goal problem (see Figure 10). The transformation involves a rotation at each X_i of the rest of the problem. Thus, using d_{goal} instead of d_{goal}^{plane} leads to a lower bound on $\mathbb{E}_n(\|\mathbf{P}_*\|)$. Casual readers should note the

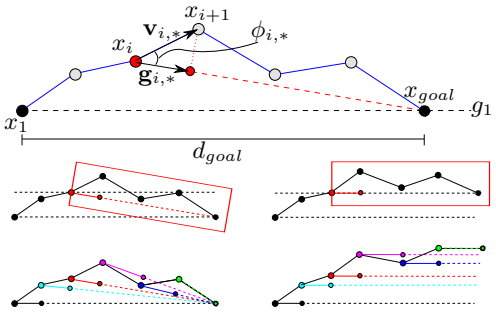

Fig. 10. Top: Path from x_1 to a point goal at x_{goal}, and related quantities. Middle: rotation of the sub-path $\{x_i, \cdots, x_{goal}\}$ does not increase path length. Bottom: performing rotation for all nodes such that $\mathbf{g}_{i,*}$ is parallel to the g_1 axis for all i.

statements of Theorems 2 and 3, as well as the comparison of theoretical to experimental results in Figure 12.

Lemma 2: *Given a path* $\mathbf{P}_* = \{x_1, \cdots, x_{k+1}\}$*, Rotating the sub-path* $\{x_i, \cdots, x_{k+1}\}$ *around* x_i *will not change* $\|\mathbf{P}_*\| = \sum_{i=1}^{k} \|x_{i+1} - x_i\| = \sum_{i=1}^{k} \|\mathbf{v}_{i,*}\|$.

Proof: $\sum_{i=1}^{k} \|\mathbf{v}_{i,*}\| = \sum_{j=1}^{i-1} \|\mathbf{v}_{j,*}\| + \sum_{j=i}^{k} \|\mathbf{v}_{j,*}\|$. It is obvious that $\sum_{j=1}^{i-1} \|\mathbf{v}_{j,*}\|$ remains unchanged because x_j for $1 \leq j \leq i$ are not affected by the rotation of $\{x_i, \cdots, x_{k+1}\}$ around x_i. Also, $\sum_{j=i}^{k} \|\mathbf{v}_{j,*}\|$ remains unchanged because no scaling occurs when $\{x_i, \cdots, x_{k+1}\}$ is rotated around x_i, and so $\|\mathbf{v}_{j,*}\|$ for all $i \leq j \leq k$ is the same before and after the rotation. (see Figure 10-middle) ■

Theorem 2: *For point-to-point long-term navigation with greedy foraging,* $\mathbb{E}_n(\|\mathbf{P}_*\|) \geq r + \frac{d_{goal} - r}{\mathbb{E}_n(\cos(\phi_{1,*}))}$.

Proof: This is a consequence of the triangle inequality and can be observed by examining a sequence of problems that have equal path length. Starting at $i = 2$, and then working forward for $i = \{2, \cdots, k\}$, each successive problem is obtained by rotating the sub-path $\{x_i, \cdots, x_{k+1}\}$ around x_i, such that $\mathbf{g}_{i,*}$ is parallel to the g_1-axis, and points in the positive direction (see Figure 10-bottom). Because rotations are performed around x_i, path length remains unchanged by Lemma 2. Further, we have warped the path such that the apparent location of the goal from every x_i "was" along the same heading during the calculation of $\ell_{i,*}$. However, each rotation moves the apparent location of a planar goal in a non-decreasing manner with respect to the g_1-axis (e.g., increases d_{goal}^{plane}), so $d_{goal}^{plane} = \sum_{i=1}^{k} \|\mathbf{g}_{i,*}\| \geq d_{goal}$ by construction. Substituting this into the lower bound of corollary 3 gives a slightly looser lower bound. We use corollary 3 because $\mathbf{ForageHeadingB}()$ handles movement at iteration k similarly to $\mathbf{ForageGoalPoint}()$. ■

We now calculate an upper bound on $\mathbb{E}_n(\|\mathbf{P}_*\|)$ by finding an upper bound on d_{goal}^{plane}. Let $\Delta_i x_{k+1}$ be the translation of x_k along the g_1-axis due to the rotation of $\{x_i, \cdots, x_k\}$ around x_i, and let ψ_i be the angle of that rotation (see Figure 11). $\sum_{i=1}^{k} \|\mathbf{g}_{i,*}\| = d_{goal}^{plane} = d_{goal} + \sum_{i=2}^{k} \Delta_i x_{k+1}$.

Lemma 3: $\Delta_i x_{k+1} \leq \|\mathbf{g}_{i-1,*}\| \tan(\phi_{i-1})$

Proof: It is easy to see (Figure 11) that $\Delta_i x_{k+1} \leq m_i$ when $|\psi_i| \leq \pi/2$. Using similar triangles we also know that $m_i \leq \|\mathbf{s}_{i-1,*}\| = \|\mathbf{g}_{i-1,*}\| \tan(\phi_{i-1})$ iff $|\psi_i| \leq \pi/4$, and for

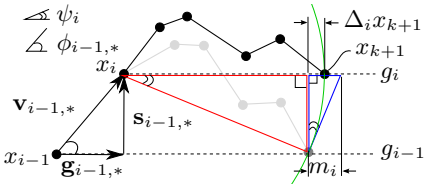

Fig. 11. The red (center) and blue (right-most) triangle are similar. ψ_i is the angle used to rotate $\{x_i, \cdots, x_{k+1}\}$ around x_i such that the g_i-axis is parallel to g_{i-1}-axis, and $\Delta_i x_{k+1}$ is the resulting shift in x_{k+1} along the g_{i-1} axis. $\Delta_i x_{k+1} \leq m_i$ by construction, and $m_i \leq \|\mathbf{s}_{i-1,*}\|$ when $|\psi_i| \leq \pi/4$ (which is always).

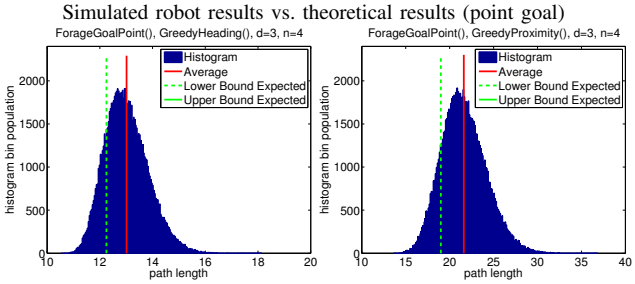

Fig. 12. Observed statistics over 100000 experimental trials with a simulated robot. Dimension $d = 3$, forage points $n = 4$, sensor radius $r = 1$, and $d_{goal} = 10$. Histograms show the distributions of experimental path lengths and red-line depicts the mean value. Green lines show the theoretically predicted bounds on expectation based on our analytical results. The upper bound is located at 58.2 in the left sub-figure, $\mathbf{GreedyHeading}()$, and is nonexistent in the right sub-figure, $\mathbf{GreedyProximity}()$.

the algorithms we are considering $|\psi_i| \leq \pi/4$ always. ■

Theorem 3: $\mathbb{E}_n(\|\mathbf{P}_*\|) \leq \frac{d_{goal}}{\mathbb{E}_n(\cos(\phi_{1,*}))(1 - \mathbb{E}_n(\tan(|\phi_{1,*}|)))}$, *given that* $\mathbb{E}_n(\tan(|\phi_{1,*}|)) \leq 1$.

Proof: Using Lemma 3 with the definition of d_{goal}^{plane} gives $d_{goal}^{plane} \leq d_{goal} + \sum_{i=2}^{k} \|\mathbf{g}_{i-1,*}\| \tan(\phi_{i-1})$. Taking the expectation and rearranging[5] gives: $\mathbb{E}_n(d_{goal}^{plane}) \leq \mathbb{E}_n(d_{goal}) + \sum_{i=1}^{k} \mathbb{E}_n(\|\mathbf{g}_{i-1,*}\|) \mathbb{E}_n(\tan(|\phi_{i,*}|))$. We observe that $\mathbb{E}_n(\tan(|\phi_{i,*}|)) = \mathbb{E}_n(\tan(|\phi_{1,*}|))$ and $\mathbb{E}_n(d_{goal}) = d_{goal}$ and $\sum_{i=2}^{k} \|\mathbf{g}_{i-1,*}\| \leq d_{goal}^{plane}$. Substituting, we get the slightly looser $\mathbb{E}_n(d_{goal}^{plane}) \leq d_{goal} + \mathbb{E}_n(d_{goal}^{plane}) \mathbb{E}_n(\tan(|\phi_{1,*}|))$. Rearranging gives: $\mathbb{E}_n(d_{goal}^{plane}) \leq \frac{d_{goal}}{1 - \mathbb{E}_n(\tan(|\phi_{1,*}|))}$, but comes at the price that we require $\mathbb{E}_n(\tan(|\phi_{1,*}|)) \leq 1$. Substituting this result into the upper bound in corollary 3 finishes the proof[6]. ■

A critical insight from Theorem 3 is that convergence is not guaranteed by our bound when $\mathbb{E}_n(\tan(|\phi_{i,*}|)) \geq 1$. Unfortunately, this means that the upper bound on $\mathbb{E}_n(\|\mathbf{P}_\diamond\|)$ is infinite for $\mathbf{GreedyProximity}()$. On the other hand, when using $\mathbf{GreedyHeading}()$ $\mathbb{E}_n(\tan(|\phi_{i,*}|)) \geq 1$ only happens when there are relatively few n vs. D; and because $\mathbb{E}_n(\tan(|\phi_{i,\circ}|)) \to 0$ and $\mathbb{E}_n(\cos(\phi_{i,\circ})) \to 1$ as $n \to \infty$, the upper bound on $\mathbb{E}_n(\|\mathbf{P}_\circ\|)$ shrinks as n increases.

Figure 7 contains statistics on a large number of experiments involving a **simulated robot** that are performed to

[5] recall the statistical independence of proximity and angle

[6] This is a slight abuse of notation, but recall that we made the substitution $d_{goal}^{plane} = \mathbb{E}_n(d_{goal}^{plane})$ in the derivation of Theorem 1, in the first place. So the factor d_{goal}^{plane} in Theorem 1 can be replaced by $\mathbb{E}_n(d_{goal}^{plane})$ before the final substitution is made in the current proof.

verify the accuracy/tightness of the results that we have obtained. 100000 experiments are performed per each algorithm combination and $d = 3$, $n = 4$, $d_{goal}^{plane} = 10$, and $r = 1$. The lower bound is less accurate than in the case of a planer goal, but still within 6% of the average experimental value observed with $\mathbf{GreedyHeading}()$ and 12% for $\mathbf{GreedyProximity}()$—it is expected to approach 0% as $n \to \infty$. On the other hand, the upper bound is very loose for $\mathbf{GreedyHeading}()$ and nonexistent for $\mathbf{GreedyProximity}()$. The latter happens because movement with a positive goal-wise component does not necessarily bring the agent closer to that goal; and provides an argument to avoid using $\mathbf{GreedyProximity}()$ in the point-goal scenario.

VI. Conclusions

We propose and study the *navigation with foraging* problem, where an agent must simultaneously (1) navigate to a global goal and (2) forage en route as opportunities to forage are detected. This problem has applications to combat, scientific exploration, search and rescue, intelligence gathering, and other areas. The problem is interesting because achieving a long-term objective must happen in parallel to achieving many small objectives. The latter each cause a small deviation from the former, and the two types of objectives occur at very different time scales.

We study two local foraging algorithms: (A) forage the point that minimizes deviation from the heading of the long-term goal, and (B) forage the closest point ahead of the agent. We consider both planar and point long-term goals.

Both analytical and experimental results show that the average length of the path decreases as global navigation becomes more important vs. local foraging, i.e., (A) vs. (B). This decrease is significantly more pronounced for scenarios with point goals than for scenarios with boundary goals, and is a consequence of the fact that boundary goals exist at many locations while point goals exist at a single location.

Our analytical bounds are tight vs. experimental results in the case of a planar goal, and for the lower bound in the case of a point goal. On the other hand, our upper bound for the point goal scenario is loose for (A) and nonexistent for (B). The latter is a consequence of the fact that movement with a positive component toward a point goal does not necessarily reduce the distance to that goal.

The navigation with foraging problem we study is unique from previous navigation and foraging problems. However, it may be possible to extend our work in either of the latter directions. For example, calculating the expected length of a planned-path in a random environment would undoubtedly be useful. Adding foraging as a secondary objective in the context of re-planning algorithms (or vice versa) also seems like a natural extension to this work.

Appendix

This appendix contains the derivation of $\mathbb{E}_n(\cos(\phi_{i,\circ}))$, $\mathbb{E}_n(\tan(|\phi_{i,\circ}|))$, $\mathbb{E}_n(\cos(\phi_{i,\diamond}))$, and $\mathbb{E}_n(\tan(|\phi_{i,\diamond}|))$.

In general, $\mathbb{E}_n(\cos(\phi_{i,\circ})) \neq \cos(\mathbb{E}_n(\phi_{i,\circ}))$ and $\mathbb{E}_n(\tan(|\phi_{i,\circ}|)) \neq \tan(\mathbb{E}_n(|\phi_{i,\circ}|))$. However, the expectations $\mathbb{E}_n(\cos(\phi_{i,\circ}))$, $\mathbb{E}_n(\tan(|\phi_{i,\circ}|))$, and $\mathbb{E}_n(\phi_{i,\circ})$ can be calculated from the probability density functions of $\cos(\phi_{i,\circ})$, $\tan(|\phi_{i,\circ}|)$, and $\phi_{i,\circ}$—which themselves can be calculated using order statistics given the probability density and distribution functions of $\cos(\phi_i)$, $\tan(|\phi_i|)$, and ϕ_i over $\mathbf{B}^{\ominus}$. *For ease of notation, we shall drop the subscript 'i' for the intermediate steps of these derivations in which we consider the quantities relevant to a single iteration of* $\mathbf{GreedyHeading}()$.

The distribution function F_ϕ of ϕ can be found using the problem's geometry. Given our assumptions, the probability a point is sampled from any particular region of space is proportional to the Lebesgue measure of that region. Let Φ denote the hypersector of $\mathbf{B}$ that is bounded by the revolution of ϕ around the g-axis, see Figure 3-Right (e.g., if $D = 2$ then Φ is a sector and if $D = 3$ then Φ is a spherical cone, etc.). Let $\lambda_{\mathbf{B}^{\ominus}}$ and λ_Φ represent the Lebesgue measure of $\mathbf{B}^{\ominus}$ and Φ, respectively. Thus, $F_\phi = \frac{\lambda_\Phi}{\lambda_{\mathbf{B}^{\ominus}}}$. From [20], we know $\lambda_{\mathbf{B}^{\ominus}} = \frac{r^D \pi^{D/2}}{2\Gamma(D/2+1)}$, where $\Gamma(\cdot)$ is the gamma function and $\lambda_B = 2\lambda_{\mathbf{B}^{\ominus}}$. From [21] we know $\lambda_\Phi = \lambda_{\mathbf{B}^{\ominus}} I_{\sin^2\phi}\left(\frac{D-1}{2}, \frac{1}{2}\right)$, where $I_{\sin^2\phi}\left(\frac{D-1}{2}, \frac{1}{2}\right)$ is the regularized incomplete beta function $I_z\left(\frac{D-1}{2}, \frac{1}{2}\right)$ evaluated at $z = \sin^2\phi$. Thus, $I_{\sin^2\phi}\left(\frac{D-1}{2}, \frac{1}{2}\right) = \frac{\mathsf{B}\left(\sin^2(\phi); \frac{D-1}{2}, \frac{1}{2}\right)}{\mathsf{B}\left(\frac{D-1}{2}, \frac{1}{2}\right)}$, where $\mathsf{B}\left(\frac{D-1}{2}, \frac{1}{2}\right)$ and $\mathsf{B}\left(\sin^2(\phi); \frac{D-1}{2}, \frac{1}{2}\right)$ are the corresponding beta function and incomplete beta function, respectively [22]. Substituting the integral form of the beta functions yields: $I_{\sin^2\phi}\left(\frac{D-1}{2}, \frac{1}{2}\right) = \frac{\int_0^{\sin^2(\phi)} t^{(D-3)/2}(1-t)^{-1/2}dt}{\int_0^1 t^{(D-3)/2}(1-t)^{-1/2}dt}$. Note that ϕ ranges from 0 to $\pi/2$ in this calculation, with the consequences that $|\sin(\phi)| = \sin(\phi)$ and $|\cos(\phi)| = \cos(\phi)$ and $\tan(|\phi|) = |\tan(\phi)| = \tan(\phi)$. Also note, $\cos(\phi)$ is monotonically decreasing vs. ϕ on the range $\phi = [0, \pi/2]$, while $\tan(\phi)$ is monotonically increasing on the range $\phi = [0, \pi/2)$. This means that while '$\circ$' is being used to denote the minimum value with respect to ϕ and $\tan(\phi)$, it will denote the maximum value with respect to $\cos(\phi)$. The decreasing $\cos(\phi)$ vs. ϕ also implies that its corresponding distribution function is $F_{\cos(\phi)} = 1 - F_\phi$, while the increasing $\tan(\phi)$ vs. ϕ means that $F_{\tan(\phi)} = F_\phi$.

Probability density functions of ϕ and $\cos(\phi)$ and $\tan(\phi)$ are $f_\phi = F'_\phi$ and $f_{\cos(\phi)} = F'_{\cos(\phi)}$ and $f_{\tan(\phi)} = F'_{\tan(\phi)}$, respectively, where F'_ϕ and $F'_{\cos(\phi)}$ and $F'_{\tan(\phi)}$ are the derivatives of F_ϕ and $F_{\cos(\phi)}$ and $F_{\tan(\phi)}$ with respect to ϕ.

It is now possible to use order statistics to find the probability density function of $\phi_\circ$ and $\cos(\phi_\circ)$. Since $\phi_\circ$ represents the minimum ϕ over a set of size n, we are interested in the first order statistic of ϕ. This is $f_{\phi_\circ} = n(1 - F_\phi)^{n-1} f_\phi$. Likewise, $f_{\cos(\phi)_\circ}$ is given by the n-th order statistic of $\cos(\phi)$ as follows: $f_{\cos(\phi_\circ)} = f_{\cos(\phi)_\circ} = n(F_{\cos(\phi)})^{n-1} f_{\cos(\phi)}$, where $f_{\cos(\phi_\circ)} = f_{\cos(\phi)_\circ}$ comes from the fact that $\cos(\phi)$ is non-increasing on $\phi = [0, \pi/2]$. Similarly, $f_{\tan(\phi)_\circ}$ is the first order statistic of $\tan(\phi)$ as follows:

$f_{\tan(\phi_\circ)} = f_{\tan(\phi)_\circ} = n(1 - F_{\tan(\phi)})^{n-1} f_{\tan(\phi)}$. Expected values for $\phi_\circ$, $\cos(\phi_\circ)$, and $\tan(\phi_\circ)$ can now be computed:

$$\mathbb{E}_n(\phi_{i,\circ}) = \int_0^{\pi/2} \phi n(1 - F_\phi)^{n-1} f_\phi \mathrm{d}\phi$$
$$\mathbb{E}_n(\cos(\phi_{i,\circ})) = \int_0^{\pi/2} \cos(\phi)\, n(F_{\cos(\phi)})^{n-1} f_{\cos(\phi)} \mathrm{d}\phi$$
$$\mathbb{E}_n(\tan(|\phi_{i,\circ}|)) = \int_0^{\pi/2} \tan(\phi)\, n(1 - F_{\tan(\phi)})^{n-1} f_{\tan(\phi)} \mathrm{d}\phi$$

Given D and n it is possible to solve for any $\mathbb{E}_n(\phi_{i,\circ})$, $\mathbb{E}_n(\cos(\phi_{i,\circ}))$, and $\mathbb{E}_n(\tan(|\phi_{i,\circ}|))$. A few cases are presented in Table I.

We now consider $\mathbb{E}_n(\cos(\phi_{i,\diamond}))$, and $\mathbb{E}_n(\tan(|\phi_{i,\diamond}|))$, which are relevant to $\mathbf{GreedyProximity}()$. *Note we resume our use of the subscript 'i'.* Although it is possible to derive $\mathbb{E}_n(\cos(\phi_{i,\diamond}))$ and $\mathbb{E}_n(\phi_{i,\diamond})$ in a similar fashion to $\mathbb{E}_n(\cos(\phi_{i,\circ}))$ and $\mathbb{E}_n(\phi_{i,\circ})$, respectively, there is a much easier way based on on the scale symmetry of spherical shells and the statistical independence of angle and proximity.

Moving any point $\ell_{i,j} \in \mathbf{B}_i^\ominus$ directly toward or away from x_i changes $\|\mathbf{v}_{i,j}\| = \|\ell_{i,j} - x_i\|$ but not $\phi_{i,j}$, $\cos(\phi_{i,j})$, or $\tan(\phi_{i,j})$. (we can ignore the measure 0 set where $\ell_{i,j} = x_i$ and $\phi_{i,j}$ is undefined, as well as the measure 0 set where $\phi_{i,j} = \pi/2$ and $\tan(\phi_{i,j})$ is undefined).

Let $\mathbf{B}_{i,\tilde{r}}^\ominus$ denote the level-set of $\mathbf{B}_i^\ominus$ that is the half-spherical shell located at radius $\tilde{r}$, where $0 < \tilde{r} \leq r$. That is, $\mathbf{B}_{i,\tilde{r}}^\ominus = \bigcup x$ for all x such that $\tilde{r} = \|x - x_i\|$. Let $\mathbf{B}_{i,\tilde{r}_\diamond}^\ominus$ be the particular level-set such that $\ell_{i,\diamond} \in \mathbf{B}_{i,\tilde{r}_\diamond}^\ominus$. Given our i.i.d uniform sampling assumptions, with probability 1 there is only one member of $\mathbf{L}_i \cap \mathbf{B}_{\tilde{r}_{i,\diamond}}^\ominus$, and this member is $\ell_{i,\diamond}$. Thus, the problems of calculating $\mathbb{E}_n(\phi_{i,\diamond})$ and $\mathbb{E}_n(\cos(\phi_{i,\diamond}))$ and $\mathbb{E}_n(\tan(|\phi_{i,\diamond}|))$ for $\mathbf{L}_i \subset \mathbf{B}_i^\ominus$ are reduced to calculating $\mathbb{E}_n(\phi_{i,\diamond})$ and $\mathbb{E}_n(\cos(\phi_{i,\diamond}))$ and $\mathbb{E}_n(\tan(|\phi_{i,\diamond}|))$ for a *single point* $\ell_{i,\diamond}$ drawn randomly from $\mathbf{B}_i^\ominus$ such that $\ell_{i,\diamond} \in \mathbf{B}_{i,\tilde{r}_\diamond}^\ominus$ where $\tilde{r}_\diamond = \|\ell_{i,\diamond} - x_i\|$. The scale symmetry of $\mathbf{B}_{i,\tilde{r}}^\ominus$ for $0 < \tilde{r} \leq r$ with respect to $\mathbb{P}_n(\phi_{i,j} = \phi_i | \ell_{i,j} \in \mathbf{B}_{i,\tilde{r}}^\ominus)$, gives:

$$\mathbb{P}_n(\phi_{i,\diamond}) = \mathbb{P}_n(\phi_{i,\diamond} | \ell_{i,\diamond} \in \mathbf{B}_{i,\tilde{r}_\diamond}^\ominus) = \mathbb{P}_n(\phi_{i,\diamond} | \ell_{i,\diamond} \in \mathbf{B}_{i,r}^\ominus)$$

In other words, replacing $\tilde{r}_\diamond$ with any other value $\tilde{r} \in (0, r]$ will yield identical results vs. $\mathbb{E}_n(\phi_{i,\diamond})$ and $\mathbb{E}_n(\cos(\phi_{i,\diamond}))$ and $\mathbb{E}_n(\tan(|\phi_{i,\diamond}|))$ because it will not change $\phi_{i,\diamond}$ or $\cos(\phi_{i,\diamond})$ or $\tan(|\phi_{i,\diamond}|)$. For convenience, we use $\tilde{r} = r$.

Again using scale symmetry and statistical independence, we realize that any point $\ell_{i,j} \in \mathbf{B}_i^\ominus$ can be projected directly away from x_i to the surface of $\mathbf{B}_i^\ominus$ without affecting $\phi_{i,j}$ or $\cos(\phi_{i,j})$ or $\tan(|\phi_{i,\diamond}|)$. Therefore, calculating $\mathbb{E}_n(\phi_{i,\diamond})$ and $\mathbb{E}_n(\cos(\phi_{i,\diamond}))$ for a *single point* $\ell_{i,\diamond}$ drawn randomly from $\mathbf{B}_i^\ominus$ such that $\ell_{i,\diamond} \in \mathbf{B}_{i,r}^\ominus$ is the dual of calculating $\mathbb{E}_n(\phi_{i,\diamond})$ and $\mathbb{E}_n(\cos(\phi_{i,\diamond}))$ for a *single point* $\ell_{i,\diamond}$ drawn randomly from $\mathbf{B}_i^\ominus$. Finally, when only one point is drawn from $\mathbf{B}_i^\ominus$ then, by construction, $\mathbf{GreedyHeading}()$ and $\mathbf{GreedyProximity}()$ must return the same result. Thus $\mathbb{E}_n(\phi_{i,\diamond}) = \mathbb{E}_1(\phi_{i,\circ})$, and $\mathbb{E}_n(\cos(\phi_{i,\diamond})) = \mathbb{E}_1(\cos(\phi_{i,\circ}))$, and $\mathbb{E}_n(\tan(|\phi_{i,\diamond}|)) = \mathbb{E}_1(\tan(|\phi_{i,\circ}|))$.

REFERENCES

[1] B. Andrews, K. Passino, and T. Waite, "Foraging theory for autonomous vehicle decision-making system design," *Journal of Intelligent and Robotic Systems*, vol. 49, pp. 39–65, 2007.

[2] D. Gillen and D. Jacques, "Cooperative behavior schemes for improving the effectiveness of autonomous wide area search munitions," in *Cooperative Control and Optimization*, ser. Applied Optimization, R. Murphey, P. M. Pardalos, P. M. Pardalos, and D. W. Hearn, Eds. Springer US, 2002, vol. 66, pp. 95–120.

[3] E. Bonabeau, M. Dorigo, and G. Theraulaz, *Swarm Intelligence: From Natural to Artificial Systems*. Oxford, 1999.

[4] L. Steels, "Cooperation between distributed agents through self organization," in *In Proceedings of the First European Workshop on Modelling Autonomous Agents in a Multi-Agelit World*, 1989.

[5] A. Drogoul, J. Ferber, and A. D. J. Ferber, "From tom thumb to the dockers: Some experiments with foraging robots," in *In Proceedings of the Second International Conference on Simulation of Adaptive Behavior*. MIT Press, 1992, pp. 451–459.

[6] J. de Lara and M. Alfonseca, "On some properties of artificial foraging ant communities," in *17th European Simulation Multiconference*, Nottingham, 2003.

[7] C. Fernandes, V. Ramos, and A. C. Rosa, "Varying the population size of artificial foraging swarms on time varying landscapes," in *Proceedings of the 15th international conference on Artificial Neural Networks: biological Inspirations - Volume Part I*, ser. ICANN'05. Berlin, Heidelberg: Springer-Verlag, 2005, pp. 311–316.

[8] K. Lerman and A. Galstyan, "Mathematical model of foraging in a group of robots: Effect of interference," *Autonomous Robots*, vol. 13, no. 2, pp. 127–141, 2002.

[9] V. J. Lumelski and A. A. Stepanov, "Dynamic path planning for a mobile automaton with limited information on the environment," *IEEE Transactions on Automatic Control*, vol. AC-31, no. 11, pp. 1057–1063, 1986.

[10] A. Sankaranarayanan and M. Vidyasagar, "A new path planning algorithm for moving a point object amidst unknown obstacles in a plane," in *Proc. IEEE International Conference on Robotics and Automation*, 1990, pp. 1930–1936.

[11] P. Hart, N. Nilsson, and B. Raphael, "A formal basis for the heuristic determination of minimum cost paths," in *Proc. IEEE Transactions On System Science and Cybernetics*, 1968, pp. 100–107.

[12] E. W. Dijkstra, "A note on two problems in connection with graphs," in *Numererical Mathematics*, vol. 1, 1959, pp. 269–271.

[13] M. Overmars and P. Svestka, "A probabilistic learning approach to motion planning," in *Algorithmic Foundations of Robotics (WAFR)*, 1995, pp. 19–37.

[14] A. M. Ladd and L. E. Kavraki, "Measure theoretic analysis of probabilistic path planning," *IEEE Transactions on Robotics and Automation*, vol. 20, no. 2, pp. 229–242, 2004.

[15] S. M. LaValle and J. J. Kuffner, "Rapidly-exploring random trees: Progress and prospects," in *Algorithmic and Computational Robotics: New Directions*, 2001, pp. 293–308.

[16] R. J. Firby, "An investigation into reactive planning in complex domains," in *Proceedings of The Sixth National Conference on Artificial Intelligence*, 1987.

[17] M. P. Georgeff and A. L. Lansky, "Reactive reasoning and planning," in *Proceedings of the Sixth National Conference on Artificial Intelligence*, pp. 677–682.

[18] A. Hornung, M. Bennewitz, and H. Strasdat, "Efficient vision-based navigation – Learning about the influence of motion blur," *Autonomous Robots*, vol. 29, no. 2, 2010.

[19] M. Otte, "Any-com multi-robot path planning," Ph.D. dissertation, University of Colorado at Boulder, 2011.

[20] G. Huber, "Gamma function derivation of n-sphere volumes," *The American Mathematical Monthly*, vol. 89, no. 5, pp. 301–302, 1982.

[21] S. Li, "Concise formulas for the area and volume of a hyperspherical cap," *Asian Journal of Mathematics and Statistics*, vol. 4, no. 1, pp. 66–70, 2011.

[22] R. B. Paris, "incomplete beta functions," in *NIST Handbook of Mathematical Functions*, F. W. J. Olver, D. M. Lozier, and R. F. Boisvert, Eds. Cambridge University Press, 2010.