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

Sensor Planning for a Symbiotic UAV and UGV system for Precision Agriculture

Pratap Tokekar, Joshua Vander Hook, David Mulla and Volkan Isler

Abstract— **We study the problem of coordinating an Unmanned Aerial Vehicle (UAV) and an Unmanned Ground Vehicle (UGV) for a precision agriculture application. In this application, the ground and aerial measurements are used for estimating nitrogen (N) levels on-demand across a farm. Our goal is to estimate the N map over a field and classify each point based on N deficiency levels. These estimates in turn guide fertilizer application. Applying the right amount of fertilizer at the right time can drastically reduce fertilizer usage.**

Towards building such a system, this paper makes the following contributions: First, we present a method to identify points whose probability of being misclassified is above a threshold. Second, we study the problem of maximizing the number of such points visited by an UAV subject to its energy budget. The novelty of our formulation is the capability of the UGV to mule the UAV to deployment points. This allows the system to conserve the short battery life of a typical UAV. Third, we introduce a new path planning problem in which the UGV must take a measurement within a disk centered at each point visited by the UAV. The goal is to minimize the total time spent in traveling and measuring. For both problems, we present constant-factor approximation algorithms. Finally, we demonstrate the utility of our system with simulations which use manually collected soil measurements from the field.

I. Introduction

Precision agriculture can improve crop productivity and farm profits through better management of farm inputs, leading to higher environmental quality [1]. By measuring soil nitrogen levels across a farm and applying the right level of nitrogen at the right time and place, it is possible to reduce fertilizer usage by 25% without affecting corn yield [2].

One of the key components of precision agriculture is data collection, typically done manually and through remote sensing. However, satellite and aerial remote sensing can be severely limited by cloud cover [3], and may not be available at desired times (update frequency can be from 3 to 26 days). Remote sensing from a manned or remotely-piloted aerial device is costly and difficult to plan against weather conditions. Further, soil moisture, crop height, and pest infestations cannot be measured remotely in a vegetated crop. Data can also be gathered manually or by guiding a vehicle equipped with sensors through the field [4]. This process can be tedious. Hence, we are building a robotic data collection system with small, low-cost Unmanned Aerial Vehicles (UAVs) and Unmanned Ground Vehicles (UGVs) working together. Our system will provide on-demand sensing capabilities, and combine the strengths of ground and aerial robots: ground robots are capable of traveling long distances, carrying large loads and measuring soil data but cannot obtain aerial imagery. Small aerial vehicles can take images from a low altitude but have limited battery life.

In this paper, we develop path planning algorithms for the application of estimating the nitrogen (N) levels in an agriculture plot. We start with a prior N level map (e.g., obtained from satellite imagery). The goal is to classify each point to identify regions with N deficiency. We first identify points in the prior map with a high probability of being misclassified (Section IV). The system is then charged with obtaining additional soil and aerial measurements near these points. Due to the UAV's limited battery life, it may not be feasible to visit all points. Therefore we seek to visit the maximum number of points using both the UAV and UGV.

Our proposed solution is to use the UGV to deploy the UAV at carefully selected locations. As the UAV is taking images, the UGV will take soil measurements nearby. The UAV can then land on the UGV which will take the UAV to the next deployment location. Landing and ascending consumes energy, leading to the problem of choosing how often to land. Our main contribution in Section V is to show how to compute a series of deployments which visit at least a constant factor of the points visited by an optimal algorithm.

Obtaining soil measurements with the UGV is likely to be time-consuming. We can reduce the total time by combining measurement locations of nearby points. This leads to a novel variant of the Traveling Salesperson Problem with Neighborhoods (TSPN) considering both travel and measurement time, which we call the Sampling TSPN Problem. We present a constant factor approximation algorithm for this problem. Finally, armed with these algorithms, we demonstrate the benefit of using the UAV+UGV system with simulations using real data collected from a corn field (Section VI).

We now review the related work in this area.

II. Related Work

The problem of designing informative sensor trajectories has recently received much attention. Low et al. [5] presented a control strategy to minimize the probability of misclassification with a Gaussian Process (GP). Zhang and Sukhatme [6] presented an adaptive sensing algorithm to estimate a scalar field. In these works, the measurements are assumed to take negligible time and the robots obtain measurements at each time-step. In this paper, we explicitly

P. Tokekar, J. Vander Hook, and V. Isler are with the Department of Computer Science & Engineering, University of Minnesota, U.S.A. {tokekar,jvander,isler} at cs.umn.edu

D. Mulla is with the Deparment of Soil, Water and Climate, University of Minnesota, U.S.A. mulla003 at umn.edu

This work is supported by NSF Awards #1111638, #0916209, #0917676, #0936710 and a fellowship from the Institute on the Environment at the University of Minnesota.

consider the measurement time and the goal is to find a path minimizing both travel and measurement time.

Singh et al. [7] presented an approximation algorithm to find a path maximizing mutual information between "cells" that are sampled and not sampled, subject to a distance budget based on the algorithm in [8]. However, the approximation guarantees depend on the size of the cells and the running time is exponential in the log of number of cells.

Finding a tour visiting most points (or collecting most reward) subject to a budget is known as the *orienteering problem*. Blum et al. [9] presented a 4-approximation for this problem on metric graphs. In Section V, we show how to model the problem of visiting most points with a symbiotic UAV and UGV as an orienteering problem.

Once the subset of points are found, the goal is to find a minimum cost UGV tour to visit them. The classical problem of finding the shortest length tour of a set of sites is known as the Traveling Salesperson Problem (TSP). Relevant to our application is the variant known as TSP with Neighborhoods (TSPN), where the tour must visit a point in each site's neighborhood. Dumitrescu and Mitchell [10] presented an 11.15-approximation algorithm for possibly overlapping unit disk neighborhoods. In our problem, the cost is the total time taken for traveling and obtaining soil measurements. As we discuss in Section V, a minimum length tour is not necessarily a minimum time tour.

Bhadauria et al. [11] studied the problem of computing a minimum time tour for k robots to wirelessly download data from sensors by visiting a point in the sensor's range. The robots have to separately query each sensor, unlike our case, where the soil measurements for multiple points can be combined by sampling in the intersection of their neighborhoods. Alt et al. [12] studied the problem of covering a given set of points with k radio antenna with circular ranges by choosing the center and radius r_i for each circle. The cost function is a weighted sum of the length of the tour and the sum of r_i^α for each disk (α models the radio transmission power). We do not require a fixed number of samples k, and instead penalize higher k in the cost.

Recently, there has been a significant interest in developing cooperative aerial and ground/surface/underwater robot systems. Tanner [13] and Grocholsky et al. [14] presented systems and algorithms for coordinating aerial and ground vehicles for detecting and locating targets. Sujit and Saripalli [15] studied the problem of exploration to detect targets using an UAV and inspection of the targets with Autonomous Underwater Vehicles. Unlike these works, we explicitly consider that the UAV can be carried between takeoff locations by the UGV in the sensor planning phase.

We begin by formalizing the problem. In this paper, we state the key lemmas leading to the main results. Full proofs are included in the accompanying technical report [16].

III. Problem Formulation

Our operating environment is a farm plot, discretized into a set of points $\mathcal{X} = \{x_1, x_2, \cdots, x_n\}$. A Gaussian Process is used to estimate N levels from prior measurements [17], [18]. For each point we associate the most-likely nitrogen estimate as $N(x)$, with variance $\sigma(x_i)$. A common task is to find regions in the plot with similar N levels. For example, the task can be to classify each point in the plot into three labels: low N, medium N, and high N. In general, we are given a set of labels, and each label L_i is specified by a minimum and maximum N level, L_i^-, L_i^+ respectively. Define $\mathrm{P}_{lj}(x_i)$ as the probability that the label j for point x_i is correct, i.e., $\mathrm{P}_{lj}(x_i) = \mathrm{P}(N(x_i) \in [L_j^-, L_j^+])$. We will use $\mathrm{P}_l(x_i)$ to denote the probability of the current (most likely) label.

We define Potentially Mislabeled (PML) points as all points in $\mathcal{X}$ for which the probability of assigning an incorrect label is greater than a desired value $\mathrm{P}_d \in (0, 1)$. Thus, $\mathcal{X}_{pml} = \{x_i \in \mathcal{X} : 1 - \mathrm{P}_l(x_i) \geq \mathrm{P}_d\}$. Our goal is to reduce the probability of mislabeling by taking additional soil and aerial measurements near the PML points.

The UAV spends some part of its energy budget (denoted by B) for each take-off and landing. We denote the average of these energy costs by C_a, so that a combined take-off and landing takes $2C_a$. We assume that the UAV and UGV travel at unit speed and the energy required to travel is proportional to the travel time. Hence, we use distance, time and energy interchangeably. If τ_a is a set of k deployments for the UAV, then the total cost of the tour is $\mathrm{len}(\tau_a) + 2kC_a$, where $\mathrm{len}(\tau_a)$ is the sum of lengths of paths in all deployments. Our problem can be concisely stated as,

Problem 1 (UAV coverage): Find a UAV tour τ_a consisting of one or more paths (each with a take-off and landing location), to sample the maximum number of PML points, such that the tour cost is not greater than the UAV budget.

Given the PML points visited by the UAV, our next objective is to design a UGV path that obtains soil measurements in the least time. The spatial correlation of soil properties means that nearby points are likely to have the same N level. Hence, as described in Section IV, the location of the measurement relative to any PML point should be within a specified radius. The UGV can thus combine measurements for multiple points if their radii overlap. The UGV is assumed to have sufficient fuel for the entire trip, but the time cost must be minimized. If τ_g is a UGV tour with k measurement locations, then the cost of the tour is given as $\mathrm{len}(\tau_g) + kC_g$, where C_g is the time cost for each measurement. Our second optimization problem is then,

Problem 2 (UGV cost): Given the set of points visited by the UAV and a radius associated with each point, find a UGV tour of minimum cost that obtains at least one measurement within the radius of each point.

IV. Finding Potentially Mislabeled Points

In this section we discuss how to find measurement locations for each point, such that the probability the point is mislabeled is less than a pre-defined threshold, P_d. Note P_d is an upper bound on the probability that $N(x_i)$ is outside the range $[L^-(x_i), L^+(x_i)]$. Let $\Phi(a)$ denote the Normal Gaussian CDF. Then for all points x_i, we require

$$\Phi\left(\frac{L^-(x_i) - N(x_i)}{\sigma(x_i)}\right) + 1 - \Phi\left(\frac{L^+(x_i) - N(x_i)}{\sigma(x_i)}\right) \leq \mathrm{P}_d$$

For any $N(x_i)$, there exists a corresponding $\sigma(x_i)$ which satisfies above equation. Define $\Delta(x_i) \triangleq \min(|L^+(x_i) - N(x_i)|, |L^-(x_i) - N(x_i)|)$. We have,

$$2 \cdot \Phi\left(\frac{-\Delta(x_i)}{\sigma(x_i)}\right) \leq \mathrm{P}_d$$
$$\therefore \frac{-\Delta(x_i)}{\Phi^{-1}\left(\frac{\mathrm{P}_d}{2}\right)} \geq \sigma(x_i) \quad (1)$$

$\Phi^{-1}(x)$ is negative-valued when $x < 0.5$. We refer to the left hand side of above equation by $\sigma_d(x_i)$.

Measurements taken near x_i reduce $\sigma(x_i)$ due to the spatial correlation of the N values. Let the measurement location be denoted by z, and the sensor noise variance be σ_s. Correlation between N levels at z and x_i is modeled by the GP equations [17]. The *new* variance at x_i, conditioned on the measurement z, satisfies,

$$\sigma^2(x_i|z) = \sigma^2(x_i) - K(x_i, z)[K(z, z) + \sigma_s^2]^{-1} K^T(x_i, z) \quad (2)$$

$K(\cdot, \cdot)$ is the *covariance* or *kernel* function of the GP [17]. We use the *squared exponential* function, which is commonly used in our intended application [18].

Recall $\sigma^2(x_i|z)$ should be no greater than σ_d^2 by Equation 1. Given z, Equation 2 simplifies as follows,

$$\sigma_d^2(x_i) - \sigma^2(x_i) \geq -\sigma_f^4(\sigma_f^2 + \sigma_s^2)^{-1} \exp(-\frac{1}{2l^2}||x_i - z||^2)$$

where σ_f and l are the *hyperparameters* of the covariance function, which are previously learned from the data [17].

Further rearrangement and taking the natural log of both sides, yields a constraint on the range of z from x_i.

$$||x_i - z||^2 \leq -2l^2 \log[(\sigma^2(x_i) - \sigma_d^2(x_i))(\sigma_f^2 + \sigma_s^2)\sigma_f^{-4}] \quad (3)$$

Thus, for every PML point $x_i \in \mathcal{X}_{pml}$ (i.e., points where N estimates do not satisfy Equation 1), we can find an upper bound on the distance of a measurement, z, from x_i using Equations 1 and 3. A sample obtained within this distance will yield small enough variance on $N(x_i)$ to determine the proper label with desired certainty for all PML points.

V. Path Optimization

In this section, we describe our algorithm for finding the UAV+UGV tours to visit the PML points. Our algorithm operates in two phases: we first find the largest subset of PML points to be visited by the UAV, subject to its energy budget (Problem 1). Then we find a UGV tour to obtain soil measurements near this subset of PML points (Problem 2).

A. Computing the UAV Tour

Visiting the largest subset of PML points with a UAV subject to an energy budget can be reduced to the orienteering problem. Given a graph $G(V, E, \pi, w)$ with edge weights $w(u, v)$, and vertex rewards $\pi(v)$, the orienteering problem is to find a tour of a subset of V collecting maximum reward, such that the sum of edge weights on the tour is less than the given budget. In the following, we show how to create such a graph for our problem.

First consider the simpler case of a UAV–only system with a camera whose footprint is a single point. The UAV tour will consist of a single path with one take-off and landing location. We build a complete Euclidean graph with PML points as the vertices. Each vertex is given a unit reward. The budget for the UAV equals $B - 2C_a$ to account for the single takeoff and landing. The solution for the orienteering problem will be a tour traversing a set of PML points. The reward equals the total number of PML points on this tour. Since the edge weights are Euclidean distances, this graph is a complete metric graph. Applying the 4-approximation algorithm in [9] will yield a UAV tour visiting at least $1/4^{\text{th}}$ of the PML points visited by the optimal algorithm.

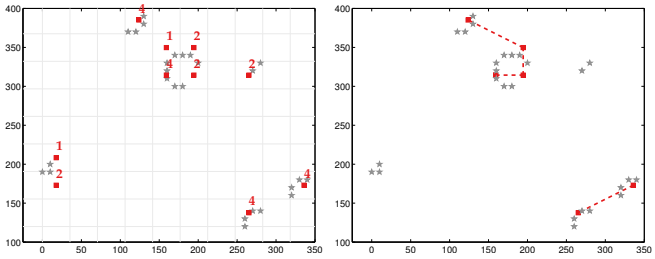

Fig. 1. Computing the UAV tour: (a) Grid of resolution $D/\sqrt{2}$. The reward for visiting a grid point (red square) is the number of PML points (gray stars) in each cell. (b) UAV tour found using orienteering with budget of 500 secs. 200 secs are spent for traveling and 240 secs for the 2 ascents/descents.

In an UAV+UGV system, the UGV can transport the UAV between two locations without affecting the UAV's energy budget. Additionally, since the UAV's camera has a footprint with diameter D, it can sample a point without flying directly over it. Thus, we modify the graph as follows:

(1) Create a square grid of resolution $D/\sqrt{2}$ (Figure 1). Store the number of $\mathcal{X}_{pml}$ within each grid cell (denoted by $\pi(v)$). Let V be the set of grid vertices with $\pi(v) > 0$.

(2) Build a complete undirected graph $G = \{V, E, \pi, w\}$. For each edge between $(u, v) \in V$, add a weight $w(u, v) = \min\{d(u, v), 2C_a\}$. This implies there are two types of edges between grid points: The UAV can either use the UGV to travel paying only for the ascent/descent $(2C_A)$ or travel directly between points paying the distance cost $(d(u, v))$.

We can verify that G is a metric graph [16]. Applying the algorithm in [9] yields a 4-approximation to Problem 1.

B. Sampling TSPN

Once we have the subset of PML points to be sampled, we must construct a minimum cost UGV tour to measure each point. Equation 3 gives a radius within which to obtain a sample for all points. This radius is not necessarily uniform. In practice, we have found the radii to be comparable and hence, use the minimum radius to simplify the algorithm.

We must now find a tour for the UGV which takes a measurement in all disks. In the standard TSPN with unit disks [10] only the travel time is considered. However, recall from Problem 2, the cost of the UGV tour equals the sum of time spent for traveling and for obtaining measurements. Let C be the time cost of obtaining a measurement. We are studying the following new variant of TSP.

Problem 3 (Sampling TSPN): Given a set of disks in the plane with uniform radius r, find a tour τ of k distinct sample

locations to minimize the cost $\text{len}(\tau) + kC$ such that each disk contains a sample location.

We present a brief description of the algorithm in [10] for the standard TSPN problem and show how to modify it to accommodate the additional cost of measurements. Dumitrescu and Mitchell in [10] first find a maximal independent set (MIS) of non-intersecting disks. Then they find a TSP tour visiting the centers of the disks in MIS (Figure 3(a)). This tour intersects the circumference of each disk in the MIS twice. The final tour that visits each disk is constructed as follows. Start from an arbitrary disk in the MIS. Follow the tour in the clockwise order up to the intersection point of the next disk. Follow the circumference of this disk in the clockwise order up to the second intersection point of the same disk. Continue until the starting point is encountered again. Repeat this process but now follow the untraveled portion of the circumference of the disks in the MIS.

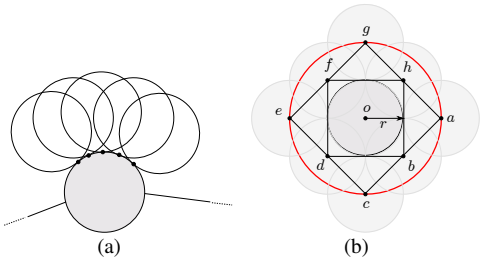

(a) (b)

Fig. 2. (a) The TSPN algorithm [10] visiting just the circumference of a disk in the MIS (shaded) will be forced to take a separate measurement for each outer disk yielding $O(n)$ measurement locations. The optimal algorithm can instead visit a small number of locations where the disks overlap. (b) We modify the tour to visit a fixed number of sites (shown as dots) around each disk in the MIS yielding a constant factor approximation.

When the UGV has to stop and obtain a measurement in each disk, restricting the motion to the circumference can be potentially costly in terms of the number of measurements. Figure 2(a) shows an instance where this algorithm is forced to take $O(n)$ measurements, where as an optimal algorithm can visit only a small number of intersection points.

We modify the local strategy (going around the circumference) in order to simultaneously find sampling locations for the tour. We will bound the additional length due to this new local strategy, and bound the number of samples obtained with respect to the optimal. Our algorithm is as follows:

(1) Find an MIS of non-overlapping disks from input disks.
(2) Starting from an arbitrary point, follow the TSP tour of the centers of disks in the MIS (Figure 3(a)).
(3) When a new center (say x) is reached on the TSP tour, before visiting the next center on the tour, first visit the eight sites as shown in Figure 2(b). Sites b, d, h, f lie on a square of side $2r$ centered at o and a, c, e, g lie on a square of side $2\sqrt{2}r$ rotated by $\frac{\pi}{4}$. Then continue towards the center of the next disk in MIS (Figure 3(b)).
(4) Restrict the candidate sampling locations to the set of centers of disks in MIS and the set of eight sites as described above. Denote this set of candidate sampling locations by S.

Algorithm 1: Sampling TSPN Algorithm.

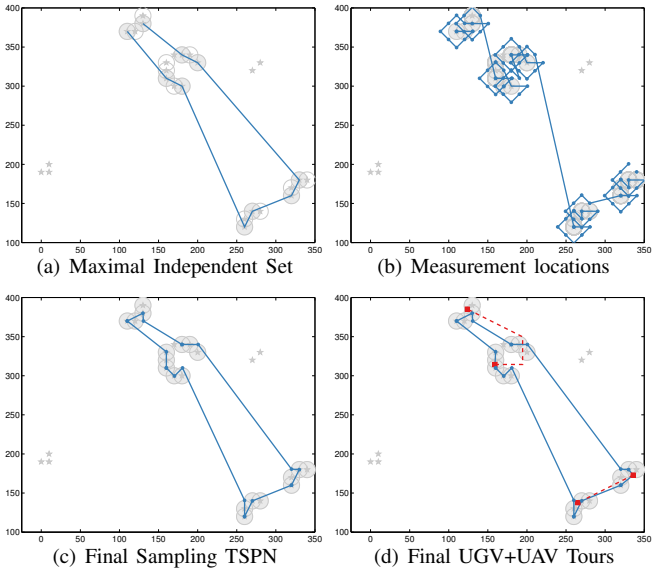

(a) Maximal Independent Set (b) Measurement locations

(c) Final Sampling TSPN (d) Final UGV+UAV Tours

Fig. 3. Steps in constructing a Sampling TSPN tour (Algorithm 1) and final UGV tour including UAV take-off locations (red squares) from Figure 1.

First we will show that the set S defined above intersects all disks in X. Then we will bound the size of S with respect to the size of the optimal number of samples N^*. Finally, we will bound the length of the tour and thus the total cost. The full proofs are presented in [16].

Lemma 1: Let S be the set of all centers of the disks in an MIS of non-overlapping disks of X. Let S also contain the eight sites as described in Algorithm 1. Then, for each disk in X, there exists a point in S lying in its interior.

Next, we bound the size of S with respect to the number of samples in an optimal Sampling TSPN algorithm. We also show that the total length of our tour is no more than a constant times the length of the tour of an optimal algorithm for the Sampling TSPN problem.

Lemma 2: If N^* is the number of samples by an optimal algorithm for the Sampling TSPN problem, then $|S| \leq 9N^*$.

Lemma 3: Let T_{ALG} be the tour constructed by the algorithm above, and T^* be the tour for the optimal Sampling TSPN algorithm. Then $\text{len}(T_{ALG}) \leq 58\ \text{len}(T^*)$.

Using above lemmas, we can show that the total cost of our algorithm is at most a constant times the optimal cost.

Theorem 1: Algorithm 1 is a valid Sampling TSPN tour with cost at most 58 times that of the optimal algorithm.

In practice, we do not have to sample at all sites in S. We can discard sample locations that do not intersect any disk or find a smaller subset (e.g. greedily) that samples all the disks (Figure 3(c)). We also add all the take-off locations to S, and find a TSP tour of the combined set of points (Figure 3(d)). Then, after visiting a take-off location, the UGV continues along its tour. When its time to reach the landing location from the tour becomes equal to the landing time, the UGV deviates from its pre-planned route, visits the landing location, and returns to its pre-planned route. The total distance overhead for the UGV is at most the budget of the UAV and in practice, not very significant.

VI. Simulations

In previous sections, we presented theoretical bounds on the number of PML points selected and the distance traveled by our algorithm with respect to optimal. We expect the UAV+UGV system to sample more PML points as compared to a UAV–only system with the same battery constraints. We explore this through simulations using actual system parameters and real data collected from an agricultural plot.

A. System Description

We present the details of the system we are developing to motivate the choice of our simulation parameters. Our UGV is a Husky A200 by Clearpath Robotics [19]. The UGV has a typical battery life of two hours and can be extended to over six hours with additional batteries. The UGV will measure soil organic matter as a proxy for soil N supply to the crop using a Minolta SPAD-502 Chlorophyll meter [20]. Our UAV is a Hexa XL by MikroKopter [21]. This UAV can operate for 25 mins on a single charge. An elevation of 100m yields a camera footprint of 50m diameter and takes about 2 minutes to ascend/descend. The images include multi-spectral information, such as near-infrared reflectance, used to estimate the crop N status [22].

B. Modeling

To generate realistic data, we constructed a generative model of N levels and determined sensor noise parameters using a remote sensing and soil sampling dataset [22]. The dataset consists of 1375 soil measurements taken manually in a 50×250m corn field, and 1m spatial resolution remote sensing images in the green (G), red (R), and near infrared (NIR) portions of the spectrum. Soil measurements were taken with dense uniform coverage giving soil organic matter (OM) levels at three different times during the season.

We performed GP regression over the set of sample points and OM measurement values using the GPML Toolbox [23]. This densely-sampled GP defined the hyperparameters which were used to generate new simulated ground-truth N maps. From the ground-truth GP regression, we can estimate the OM sample noise at each point from the data directly (σ_s in Equation 2). We used the OM variance directly as σ_g.

The UAV was modeled as measuring the Normalized Difference Vegetation Index (NDVI), which is a combination of NIR and R levels [24]. We assume the NDVI levels are corrupted by sensor noise σ_a. We calculated σ_a from the sample covariance between NDVI (from hand-measured R and NIR levels) and OM (directly measured), yielding the covariance matrix. We can then find the variance in OM given a measurement of NDVI as,

$$\sigma_a = \sigma^2_{\text{OM|NDVI}} = \sigma^2_{\text{NDVI}} - \sigma^2_{\text{OM,NDVI}}/\sigma^2_{\text{OM}}$$

From the dataset, we obtained $\sigma_a = 0.31$ and $\sigma_g = 0.05$.

We randomly generated 100 N maps for a 600×400 meter field. A prior estimate of each generated OM map was created by down sampling the ground-truth map by a factor of 20 and fitting a GP. We then found the PML points and ranges using $\text{P}_d = 0.4$ as described in Section IV.

C. Results

We first compare the number of PML points covered by the UAV+UGV system versus an UAV–only system, subject to the battery constraint of 25 mins. We used the implementation from the SFO Toolbox [25] for finding an orienteering tour, and the Concorde TSP solver [26] as a subroutine in the Sampling TSPN algorithm implementation.

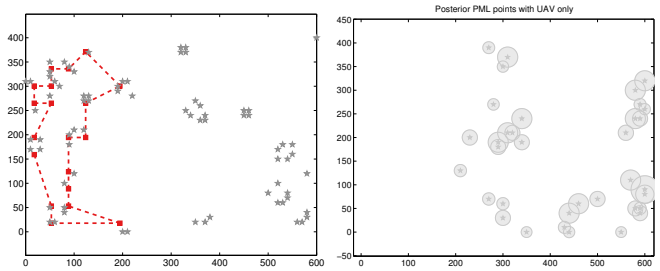

(a) UAV–only tour visiting 38 points. 35 posterior points in updated map.

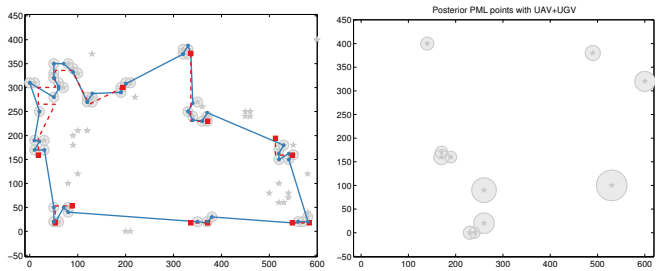

(b) UAV+UGV tour visiting 50 points. 11 posterior points in updated map.

Fig. 4. Sample simulation instance. The input consists of 75 PML points. The UAV+UGV tour consists of 6 subtours. The UGV allows the UAV to transport to farther locations in the plot which is reflected in fewer posterior PML points generated after incorporating aerial and ground measurements.

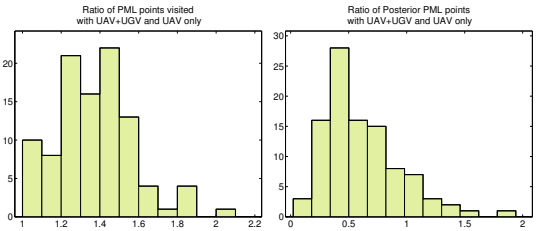

(a) (b)

Fig. 5. Histograms of the ratio of (a) number of PML points visited and (b) posterior PML points generated after updating N map with simulated measurements, for a UAV+UGV system and a UAV–only system for 100 random instances. Both systems are given an equal budget of 25 minutes.

Figure 4 shows a sample simulation run. We observe that the UAV–only tour is constrained to only one part of the field unlike the UAV+UGV system. This instance consisted of 75 PML points, the UAV–only tour visits 38 points whereas the UAV+UGV tour visits 50 points. Figure 5(a) shows a histogram of the ratio of the points covered by the UAV+UGV and the UAV–only tours for 100 random instances. The ratio is always greater than 1 as the UAV+UGV system is at least as good as a UAV–only system in terms of the number of points visited. Table I shows the effect of varying the budget on the percentage of input PML points visited.

The UAV+UGV system can cover points that are spread across the field. We expect the posterior map to have fewer

TABLE I

PERCENTAGE OF INPUT PML POINTS VISITED (AVG. OF 30 INSTANCES).

Budget (sec)	UAV–only	UAV+UGV
500	19	25
1000	36	49
1500	55	72

mislabeled points than if all measurements are in one part of the field. After calculating the tours, measurements for the sensors were sampled directly from OM values given the dense (ground truth) GP. We added noise to the measurements using the modeled variances $\sigma_a = 0.31$ and $\sigma_g = 0.5$. These values were then used to update the prior GP, and to find the posterior PML points (Figures 4(a) & 4(b)). For a fair comparison, we also added UGV measurements for each PML point visited by a UAV–only tour.

Figure 5(b) shows a histogram of the ratio of the posterior PML points with a UAV+UGV system and a UAV–only system. Since the number of PML points depend on both the variance, and the estimated $N(x)$ values, occasionally there are instances when the number of posterior PML points with UAV–only system are lesser than that of UAV+UGV system. However, as seen in Figure 5(b) the UAV+UGV system generally outperforms the UAV–only system in terms of number of posterior PML points.

VII. CONCLUSION

In this paper, we studied the problem of designing sensing strategies for obtaining aerial images and soil samples with a UAV+UGV system to estimate the nitrogen level in a plot. Since the battery life of the UAV is limited, the UAV and UGV can only sample a limited number of points. We studied the problem of maximizing the number of points visited by the UAV and UGV. Unlike traditional approaches, our algorithm takes into consideration the situation where the UAV can land on the UGV and thus be carried between points without expending energy. We also studied the problem of minimizing the time for sampling these points with a UGV. We presented a constant-factor approximation algorithm which finds a set of sampling locations and a tour of these locations, such that each point has a sampling location within its disk neighborhood.

We have started building the complete system using a Clearpath Husky A200 ground robot and a hexacopter from MikroKopter. In order to execute the algorithms presented in this paper, additional capabilities such as autonomous landing and soil sampling are necessary which is part of our ongoing work.

REFERENCES

[1] D. Mulla, "Mapping and managing spatial patterns in soil fertility and crop yield," in *Soil Specific Crop Management*. American Society of Agronomy, 1993, pp. 15–26.

[2] G. W. Randall and D. J. Mulla, "Nitrate nitrogen in surface waters as influenced by climatic conditions and agricultural practices," *Journal of Environmental Quality*, vol. 30, no. 2, pp. 337–344, 2001.

[3] M. S. Moran, Y. Inoue, and E. Barnes, "Opportunities and limitations for image-based remote sensing in precision crop management," *Remote Sensing of Environment*, vol. 61, no. 3, pp. 319–346, 1997.

[4] V. I. Adamchuk, J. Hummel, M. Morgan, and S. Upadhyaya, "On-the-go soil sensors for precision agriculture," *Computers and Electronics in Agriculture*, vol. 44, no. 1, pp. 71–91, 2004.

[5] K. H. Low, J. Chen, J. M. Dolan, S. Chien, and D. R. Thompson, "Decentralized active robotic exploration and mapping for probabilistic field classification in environmental sensing," in *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, 2012, pp. 105–112.

[6] B. Zhang and G. S. Sukhatme, "Adaptive Sampling for Estimating a Scalar Field using a Robotic Boat and a Sensor Network," *Proceedings 2007 IEEE International Conference on Robotics and Automation*, pp. 3673–3680, Apr. 2007.

[7] A. Singh, A. Krause, C. Guestrin, and W. J. Kaiser, "Efficient informative sensing using multiple robots," *Journal of Artificial Intelligence Research*, vol. 34, no. 2, p. 707, 2009.

[8] C. Chekuri, N. Korula, and M. Pál, "Improved algorithms for orienteering and related problems," *ACM Transactions on Algorithms (TALG)*, vol. 8, no. 3, p. 23, 2012.

[9] A. Blum, S. Chawla, D. R. Karger, T. Lane, A. Meyerson, and M. Minkoff, "Approximation algorithms for orienteering and discounted-reward tsp," *SIAM Journal on Computing*, vol. 37, no. 2, pp. 653–670, 2007.

[10] A. Dumitrescu and J. S. Mitchell, "Approximation algorithms for tsp with neighborhoods in the plane," *Journal of Algorithms*, vol. 48, no. 1, pp. 135–159, 2003.

[11] D. Bhadauria, O. Tekdas, and V. Isler, "Robotic data mules for collecting data over sparse sensor fields," *Journal of Field Robotics*, vol. 28, no. 3, pp. 388–404, 2011.

[12] H. Alt, E. M. Arkin, H. Brönnimann, J. Erickson, S. P. Fekete, C. Knauer, J. Lenchner, J. S. Mitchell, and K. Whittlesey, "Minimum-cost coverage of point sets by disks," in *Proceedings of the twenty-second annual symposium on Computational geometry*. ACM, 2006, pp. 449–458.

[13] H. G. Tanner, "Switched uav-ugv cooperation scheme for target detection," in *Robotics and Automation, 2007 IEEE International Conference on*. IEEE, 2007, pp. 3457–3462.

[14] B. Grocholsky, J. Keller, V. Kumar, and G. Pappas, "Cooperative air and ground surveillance," *Robotics & Automation Magazine, IEEE*, vol. 13, no. 3, pp. 16–25, 2006.

[15] P. Sujit and S. Saripalli, "An empirical evaluation of co-ordination strategies for an auv and uav," *Journal of Intelligent & Robotic Systems*, vol. 70, pp. 373–384, 2013.

[16] P. Tokekar, J. Vander Hook, D. Mulla, and V. Isler, "Sensor planning for a symbiotic uav and ugv system for precision agriculture," Dept. of Computer Science & Engineering, University of Minnesota 13-010, Tech. Rep., 2013. [Online]. Available: http://www.cs.umn.edu/research/technical_reports/view/13-010

[17] C. E. Rasmussen and C. K. Williams, *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006, vol. 1.

[18] D. E. Myers, "Estimation of linear combinations and co-kriging," *Mathematical Geology*, vol. 15, no. 5, pp. 633–637, 1983.

[19] "Clearpath robotics," 2013 (accessed Jan 2013). [Online]. Available: http://clearpathrobotics.com

[20] "Spectrum technologies inc." 2013 (accessed Mar 2013). [Online]. Available: http://www.specmeters.com/

[21] "Mikrokopter," 2013 (accessed Jan 2013). [Online]. Available: http://mikrokopter.de

[22] D. Mulla, M. Beatty, and A. Sekely, "Evaluation of remote sensing and targeted soil sampling for variable rate application of nitrogen," in *Proceedings of 5th International Conference on Precision Agriculture*. American Society of Agronomy, Crop Science Society of America, and Soil Science Society of America, 2001.

[23] C. E. Rasmussen and H. Nickisch, "Gaussian processes for machine learning (gpml) toolbox," *Journal of Machine Learning Research*, vol. 11, pp. 3011–3015, 2010.

[24] D. J. Mulla, "Twenty five years of remote sensing in precision agriculture: Key advances and remaining knowledge gaps," *Biosystems Engineering*, vol. 114, no. 4, pp. 358 – 371, 2013, special Issue: Sensing Technologies for Sustainable Agriculture.

[25] A. Krause, "Sfo: A toolbox for submodular function optimization," *Journal of Machine Learning Research*, vol. 11, pp. 1141–1144, 2010.

[26] D. Applegate, R. Bixby, V. Chvatal, and W. Cook, "Concorde tsp solver," *URL http://www. tsp. gatech. edu/concorde*, 2006.