Scheduling Problems for Robotics in Precision Agriculture

Stefano Carpin

Abstract—**Robotics is playing an increasingly important role in precision agriculture and agricultural technology because it allows to tackle some important problems at scale at a time when the agricultural workforce is declining. Robots can collect data that can better inform farmers on the best course of actions for their crops. Robots can also perform tasks that are too labor intensive for workers. Despite increased availability, however, these technologies will not become as so pervasive that each problem instance will be taken care of by robots and it is therefore important to carefully select which ones should be addressed and which ones can be deferred. Starting from these premises, in this overview paper we discuss a series of scheduling problems we developed that is pervasive in these applications.**

Index Terms—**Precision agriculture, robotics, optimization.**

I. Introduction

The National Research Council defines *precision agriculture* as “the application of modern information technologies to provide, process and analyze multisource data of high spatial and temporal resolution for decision making and operations in the management of crop production” [8]. This definition is related to the emerging term ”AgTech” (agricultural technology), which is loosely defined as the use of (sophisticated) technologies in agriculture, with objectives such as higher crop yield, increased efficiency and reduced environmental impact [15]. In this space, robotics and artificial intelligence (AI) can play a crucial role in addressing some of the most pressing societal needs. Robotic technologies, in particular, can help in mitigating the problem of a dwindling workforce in agriculture which is emerging in countries such as the US and has been further exacerbated by the COVID-19 pandemic [5]. The steady growth of scientific publications in this domain is a testament to the growing awareness of the importance of robotics in AgTech. As pointed out in Section II, robots are being developed for a variety of tasks including both actuation (e.g., harvesting and weeding), as well as more efficient sensing (see e.g., figure 1)

An interesting area of application is related to water efficiency. Draughts are a problem that has gained much attention in recent years, especially because of their increased recurrence in the American southwest [9]. With agriculture consuming an estimate of 70% of managed freshwater in the US and abroad [17], there is a strong interest in developing AgTech systems that can mitigate this problem. Contributions may include improved approaches for more accurate soil moisture measurements or stem water potential, as well as systems for implementing irrigation adjustments matching local conditions. In response to these problems we have are developing different systems combining robotics and AI to improve water delivery efficiency and soil moisture assessment.

Stefano Carpin is with the School of Engineering, University of California, Merced. This work is partially supported by USDA-NIFA under award # 2021-67022-33452 (National Robotics Initiative). Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the view of the U.S. Department of Agriculture.

Fig. 1. An autonomous robot equipped with a soil moisture sensor can autonomously perform measurements at a set of preassigned locations.

A recurrent problem in numerous robotic applications in AgTech is resource optimization and scheduling. Indeed, while in the future one can anticipate that autonomous vehicles will be deployed in fleets, the sheer scale of operations in commercial agriculture is such that the problem of selecting the best subset of operations to perform will remain relevant for a long time because brute force approaches where all tasks are executed would be inefficient. For example, in sampling applications one is often given a large set of locations where samples could be collected. Based on prior available information, it is often the case that not all tentative locations have the same anticipated value. Therefore, given that an autonomous robot can only visit a subset of these places (for example, because the limited power autonomy limits the distance it can travel), which sampling locations should be prioritized? And in which order? Another complicating feature is that due to the semistructured nature of orchards, the effort to perform certain operations (e.g., reaching a certain tree and collect a leaf) may be subject to substantial variability and therefore iterative replanning is often necessary. This is one of the numerous venues outlining how practical problems driven by farmers need are associated with hard computational problems requiring advanced algorithmic solutions. The rest of this overview paper is organized as follows. In section II we provide selected pointers to relevant literature. In section III we discuss routing problems relevant for robots operating in agtech, and provide pointers to our recent works in this domain (algorithmic details are omitted for brevity). Finally, conclusions are offered in section IV.

II. Related Work

Robot use in agriculture is quickly expanding and has a very promising future, with a multitude of heterogeneous applications gaining in popularity [27]. Remote sensing for information gathering is a typical utilization, where Unmanned Aerial Vehicles gather images from the sky of grow sites providing unique insights not visible from the ground [3], [14]. Another growing trend is the deployment of robots (usually on the ground) to capture images of fruit on plants for use in yield estimation [2], [16]. Fruit harvesting [18] and plant pruning [7] are other utilizations of robots that interact directly with the plants themselves. Still, some other robots are specifically built to help with logistic problems, such as moving fruit bins to and from human harvesters within fruit orchards to promote optimal labor time usage [28]. Other robotic applications are built on top of the already mechanized farm processes, such as spraying of pesticides and fertilizer from tractors in minimal travel distance and working time [6]. Regarding irrigation optimization, literature review did not uncover any preceding work related to robotic tools in this domain, except for our few selected works related ot the USDA funded *RAPID* project [13], [20], [21], [23]–[25].

III. Optimization of Routing and Scheduling

In this section we provide an overview of a class of optimization problems that is pervasive in AgTech applications, i.e., scheduling and routing. In numerous instances, precision agriculture require the implementation of spatially distributed activities over large areas. The dynamic nature of the natural phenomena being tracked combined with the necessity of refining on the fly the action strategy often prevents the deployment of stationary structures to perform these tasks (e.g., a set of sensors permanently placed at certain locations). In addition, some activities like specimen collection and retrieval requires robots to physically visit certain locations to complete their tasks. Figure 2 shows an aerial view of a commercial vineyard located in Firebaugh, CA. This block alone, part of a much larger ranch, has more than 50,000 vines. Red crosses indicate locations where soil moisture samples should be collected. Because of the existing irrigation lines under the trellises, it is not possible to move along straight lines between any two points, and to switch tree row it is necessary to exit from either end of the vineyard. These motion constraints mean that a ground robot[1] moving in the vineyard to complete the task will have to spend additional energy to move around because of the motion constraints and will be unable to visit all desired locations.

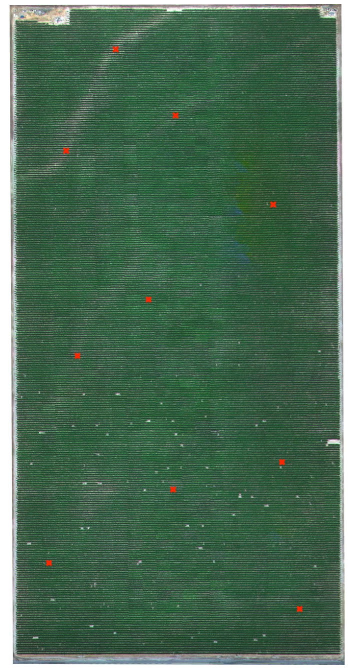

Fig. 2. Aerial view of a vineyard located in Firebaugh, CA.

From an abstract standpoint, this general optimization problem can be formulated as follows. Let $\mathcal{S}$ be the finite set of assigned target locations, i.e., the places where activities should take place (e.g., collecting a sample.). In general each location $s_i \in \mathcal{S}$ has a different utility or reward $r(s_i)$ representing how important it is to perform a task at that location. This reward or importance is assumed to be known and part of the input. For a subset of activities $S_c \subset \mathcal{S}$, let $\pi(S_c)$ be a path that visits all the locations in S_c, and let $c(\pi)$ be the cost of path π. Typically, the cost of a path is its length or the energy spent to travel along the path. Let $r(\pi)$ be the sum of the rewards of the target locations visited by path π. The following constrained optimization problem can therefore be formulated:

$$S^* = \arg\max_{S_i \subseteq \mathcal{S}} r(\pi(S_i)) \quad \text{s.t. } c(\pi(S_i)) \leq B$$

where B is a preassigned budget representing the maximum distance a robot can travel or the maximum energy it can consume in between two charges. Essentially, the problem is to determine a subset of locations that can be visited with the assigned travel budget and that maximizes the sum of collected rewards. This formulation is related to a graph optimization problem known in literature as *orienteering* that is NP-hard [11] and APX-hard [4]. Because of its intrinsic computational complexity, exact solutions can be found only for instance of moderate size, e.g., graphs with less than 1,000 vertices [10], [26] and heuristic methods are therefore extensively used [12]. Problem instances associated with agricultural problems generate much larger instances. For example, one may asso-

[1]While an aerial platform like a drone would not be affected by these constraints, their use in collecting specimens is limited because of the limited payload and energy availability.

ciate to an orchard a graph where every tree is a vertex, thus easily having tens of thousands of vertices. This renders exact methods impractical. General purpose heuristics can be used, but they may be not competitive because they miss the specific motion constraints associated with navigating in a vineyard, i.e., certain vertices must be visited in a sequence, or long detours must be taken when moving between vertices.

The basic version of the problem described above can also be extended or refined to include other aspects emerging when working on the field. For example, multiple ground vehicles may be simultaneously deployed to expedite operations and their motion should be coordinated to avoid wasted efforts (e.g., two robots performing the same operation), or negative interference (e.g., robots getting on each other's way or colliding with each other). This problem, known as team orienteering problem has received some attention in the past but is less studied than the single agent version. Another class of problems that is instead significantly less investigated is the stochastic orienteering problem that emerges when one more realistically considers that many of the involved quantities are only known with uncertainty. For example, the reward $r(s_i)$ associated with executing a task at location s_i may be uncertain, and the cost $c(\pi)$ of a path may also be uncertain because of unforeseen circumstances emerged while the robot moves, like for example the necessity to take a detour because a pathway is blocked.

For the deterministic single agent version of the problem, we recently proposed [23] different heuristic algorithms that factor in the specific motion constraints associated with vineyards when selecting the subset S_i. Using the common greedy approach that iteratively adds new vertices to the solution without ever reconsidering past choices, the key insight is in considering the benefit/cost ratio of each possible added location where the cost is informed by the constraints associated with moving in a vineyard. Figure 3 shows how the structure of a vineyard can be abstracted into a graph, where every vertex is associated with a tree and could also be an element of the target set $\mathcal{S}$. In this case, once the robot enters a tree row, if it needs to move to a different row it has to consider from which end to exit and whether it makes sense to spend additional budget to collect more samples or perform more tasks in the same row before exiting (algorithmic details are skipped and the reader is referred to [23] for details). In our work we show that if the cost is re-formulated to consider the motion constraints imposed by the environment, the custom designed heuristic can be used to solve large problem instances with tens of thousands of vertices and it outperforms formerly proposed general purpose methods that do not consider the associated motion constraints. This solution raises an important point, namely that for robots performing AgTech tasks, it may be necessary to re-design existing heuristics or design new ones informed by domain specific knowledge associated with the task being solved.

The single robot solution can then be used as a building block to determine how multiple robots can coordinate their motions, as we discussed in [22], [24]. If N robots are given, a simplistic approach consists in splitting the working area into N subregions, each approximately having the same amount of reward to be collected. Each robot then operates exclusively inside the assigned area using the single robot solution described above (see figure 3).

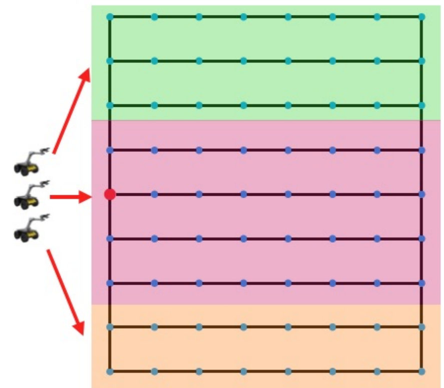

Fig. 3. With three robots given, a simple approach consists in splitting the work area in three subregions, and then let every robot run the single-robot algorithm on the assigned subregion. Subregions shall not necessarily have the same size, but should rather be sized based on the amount of available reward.

While this approach is simple, it has the drawback that if a robot is assigned to a sub-region with limited rewards, it cannot move to a different region and help other robots. A better strategy, instead consists in allowing each robot to move to any location of interest, provided that suitable coordination is added to prevent the aforementioned problems. In particular, given that tree rows are narrow, one should avoid having two robots traveling at the same time in opposite directions in the same tree row. This requires the implementation of so-called space/time coordination – a well known paradigm in robot motion planning. Two solutions embracing this approach are presented in our works [22], [24], which build upon the findings presented in [23]. The main difference is that in one case planning is done sequentially (i.e., the complete schedule for the first robot is completed before the scheduling for the second robot is comptued, and so on), while in the other case the planning is done in parallel. This second approach is generally better, but more complicated.[2] Extensive simulation results presented in the cited papers show that the speedup of these two last strategies is almost linear, i.e., the best one can hope for. One drawback, though, is that the planning approach is centralized and this is known to be problematic because if the central planner fails, then the entire system will not perform. Additionally, robots do not communicate with each other. While this may be advantageous as it does not require any communication infrastructure, the drawback is that replanning cannot be triggered on demand to address unforeseen circumstances.

[2]Note that even though planning may be done sequentially, at run time, all robots move at the same time and not one at the time.

A final class of scheduling problems that emerge in this domain and that has broad applicability is risk-averse planning and planning under chance constraints. With robots moving in semistructured environments such as vineyards and orchards, the cost of executing operations or moving between two locations is generally a stochastic variable. Henceforth, for a given subset $S_i \subset \mathcal{S}$, the associated cost to execute the path $c(\pi(S_i))$ is a random variable. The optimization problem can then be reformulated as follows:

$$S^* = \arg\max_{S_i \subseteq \mathcal{S}} \mathbb{E}[r(\pi(S_i))] \quad \text{s.t.} \;\; \Pr[c(\pi(S_i)) > B] < P_f$$

The constraint is now formulated in terms of a failure probability P_f and it imposes that the probability that the cost to execute the path exceeds the assigned budget B is less than the acceptable failure probability. This approach inherently requires replanning, because if a robot is running behind schedule then it will have to skip some locations to make sure it ends its mission before the budget B is spent. For this reason, the objective function is now formulated in terms of expectation, because the number of visited locations (and associated sum of rewards) is now a random variable, too. This problem has been scarcely studied in literature, and outlines how activities related to AgTech may require the solution of novel optimization problems. In [19]–[21] we took a first stab at the problem by introducing the idea of *path policy*. A path policy formalizes the intuition that if a robot is scheduled to visit in sequence locations $s_1, s_2, \ldots, s_k$ to perform certain tasks, it may have to adjust its route if it happens to be "running late" due to the stochastic nature of the motions between locations. To this end, it is necessary to track the amount of consumed budget (time, energy, etc.) and contrast it with the available budget B. In [19]–[21] the authors formulate a solution to this problem using a constrained Markov Decision Process (CMDP) [1]. The reader is referred to the cited papers for the numerous technical details necessary to solve this problem. An important finding, however, is that planning methods that specifically consider these failure constraints are much more robust in practice, and it appears that when robots are used for this type of problems an approach explicitly factoring in failure probabilities is a must.

IV. Conclusions

After having briefly motivated the use of robotic technologies in the area of precision agriculture and agricultural technology, we provided an overview of our proposed solutions for a class of routing and scheduling problems that is pervasive to numerous applications in this domain. Solving these problems at a scale relevant to commercial operations prevents the use of exact solutions and requires the development of heuristic approaches. Embedding domain specific knowledge driven by the agricultural domain helps in formulating solutions that are more performing than general purpose approaches. Finally, because vineyards and orchards are semistructured environments, planning methods explicitly featuring failure probabilities have been introduced and discussed.

References

[1] E. Altman. *Constrained Markov Decision Processes*. Stochastic modeling. Chapman & Hall/CRC, 1999.

[2] Suchet Bargoti and James P. Underwood. Image segmentation for fruit detection and yield estimation in apple orchards. *Journal of Field Robotics*, 34(6):1039–1060, 2017.

[3] Antonio Barrientos, Julian Colorado, Jaime del Cerro, Alexander Martinez, Claudio Rossi, David Sanz, and João Valente. Aerial remote sensing in agriculture: A practical approach to aera coverage and path planning for fleets of mini aerial robots. *Journal of Field Robotics*, 28(5):677–689, 2011.

[4] Avrim Blum, Shuchi Chawla, David R. Karger, Terran Lane, Adam Meyerson, and Maria Minkoff. Approximation algorithms for orienteering and discounted-reward tsp. *SIAM Journal on Computing*, 37(2):653–670, 2007.

[5] D. Bochtis, L. Benos, M. Lampridi, V. Marinoudi, S. Pearson, and C.G. Sørensen. Agricultural workforce crisis in light of the covid-19 pandemic. *Sustainability*, 12(19), 2020.

[6] D. Bochtis, H.W. Griepentrog, S. Vougioukas, P. Busato, R. Berruto, and K. Zhou. Route planning for orchard operations. *Computers and Electronics in Agriculture*, 113:51–60, 2015.

[7] Tom Botterill, Scott Paulin, Richard Green, Samuel Williams, Jessica Lin, Valerie Saxton, Steven Mills, XiaoQi Chen, and Sam Corbett-Davies. A robot system for pruning grape vines. *Journal of Field Robotics*, 34(6):1100–1122, 2017.

[8] National Research Council. *Precision agriculture in the 21st century: Geospatial and information technologies in crop management*. National Academy Press, 1997.

[9] Noah S Diffenbaugh, Daniel L Swain, and Danielle Touma. Anthropogenic warming has increased drought risk in california. *Proceedings of the National Academy of Sciences*, 112(13):3931–3936, 2015.

[10] Matteo Fischetti, Juan José Salazar González, and Paolo Toth. Solving the orienteering problem through branch-and-cut. *Journal on Computing*, 10(8):133–148, 1998.

[11] Bruce L. Golden, Larry Levy, and Rakesh Vohra. The orienteering problem. *Naval Research Logistics*, 34:307–318, 1987.

[12] Aldy Gunawan, Hoong Chuin Lau, and Pieter Vansteenwegen. Orienteering problem: A survey of recent variants, solution approaches, and applications. *European Journal of Operational Research*, 255(2):315–332, 2016.

[13] X. Kan, T. Thayer, S. Carpin, and K. Karydis. Task planning on stochastic aisle graphs for precision agriculture. *IEEE Robotics and Automation Letters*, 6(2):3287 – 3294, 2021.

[14] Zhengqi Li and Volkan Isler. Large scale image mosaic construction for agricultural applications. *IEEE Robotics and Automation Letters*, 1(1):295–302, January 2016.

[15] United States Department of Agriculture. Agriculture technology `https://nifa.usda.gov/topic/agriculture-technology`.

[16] Giuseppe Riggio, Cesare Fantuzzi, and Cristian Secchi. A low-cost navigation strategy for yield estimation in vineyards. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 2200–2205, 2018.

[17] G.D. Schaible and M.P. Aillery. *Challenges for US Irrigated Agriculture in the Face of Emerging Demands and Climate Change*, chapter Competition for Water Resources 2.1.1, pages 44–79. Elsevier, 2017.

[18] Abhisesh Silwal, Joseph R. Davidson, Manoj Karkee, Changki Mo, Qin Zhang, and Karen Lewis. Design, integration, and field evaluation of a robotic apple harvester. *Journal of Field Robotics*, 34(6), September 2017.

[19] T. Thayer and S. Carpin. An adaptive method for the stochastic orienteering problem. *IEEE Robotics and Automation Letters*, 6(2):4185–4192, 2021.

[20] T. Thayer and S. Carpin. A fast algorithm for stochastic orienteering with chance constraints. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 7398–7945, 2021.

[21] T. Thayer and S. Carpin. A resolution adaptive algorithm for the stochastic orienteering problem with chance constraints. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 6388–6395, 2021.

[22] T. Thayer, S. Vougioukas, K. Goldberg, and S. Carpin. Multi-robot routing algorithms for robots operating in vineyards. In *Proceedings of the IEEE Conference on Automation Science and Engineering*, pages 7–14, 2018.
[23] T. Thayer, S. Vougioukas, K. Goldberg, and S. Carpin. Routing algorithms for robot assisted precision irrigation. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 2221–2228, 2018.
[24] T. Thayer, S. Vougioukas, K. Goldberg, and S. Carpin. Multi-robot routing algorithms for robots operating in vineyards. *IEEE Transactions on Automation Science and Engineering*, 17(3):1184–1194, 2020.
[25] T.C. Thayer and S. Carpin. An adaptive method for the stochastic orienteering problem. *IEEE Robotics and Automation Letters*, 6(2):4185 – 4192, 2021.
[26] Pieter Vansteenwegen, Wouter Souffriau, and Dirk Van Oudheusden. The orienteering problem: A survey. *European Journal of Operational Research*, 209:1–10, 2011.
[27] S. Vougioukas. Agricultural robotics. *Annual review of control, robotics, and autonomous systems*, 2:339–364, 2019.
[28] Yawei Zhang, Yunxiang Ye, Zhaodong Wang, Matthew E. Taylor, Geoffrey A. Hollinger, and Qin Zhang. Intelligent in-orchard bin-managing system for tree fruit production. In *IEEE International Conference on Robotics and Automation Workshop on Robotics in Agriculture*, 2015.