Use of intelligent/autonomous systems in crop irrigation

Stefano Carpin[1*], Ken Goldberg[2], Stavros Vougioukas[3], Ron Berenstein[2], Josh Viers[4]
[1]Department of Computer Science and Engineering, University of California, Merced
[2]Department of Industrial Engineering and Operations Research, University of California, Berkeley
[3]Department of Biological and Agricultural University of California, Davis
[4]Department of Civil and Environmental Engineering, University of California, Merced

Abstract

In this chapter we discuss how robotic and artificial intelligence can be used to improve precision irrigation in vineyards. In particular, we discuss RAPID (Robot Assisted Precision Irrigation Delivery), a novel system being developed and tested at the University of California. RAPID uses remote sensing and machine learning to predict soil moisture content, and uses this information to determine how variable rate emitters should be adjusted by one or more robots moving through the vineyard. As part of this project we are developing machine learning algorithms for soil moisture prediction, routing algorithms for efficiently scheduling multiple robots through vineyards, and we are also designing novel grippers that can be used to adjust inexpensive variable rate emitters that can be retrofit existing irrigation systems. Early results demonstrate the promise of the proposed effort.

Keywords: robotics, irrigation, artificial intelligence.

1 Introduction

Agriculture is the major consumer of managed freshwater worldwide. Some studies have estimated that more than 80% of managed freshwater in the US is used for agriculture, and worldwide the percentage is around 70% (Schable et al., 2017). Climate change, combined with the need to feed an increasing population with decreasing arable land requires to radically re-think the way water is delivered to crops to increase efficiency and minimize wasted water (Wallace, 2000). Precision irrigation systems using drip emitters have been used for decades but only offer a partial solution. In particular, the problem of regulating the amount of water delivered to an individual plant is still outstanding, in particular for large-scale operations. This problem is evident in wine grapes and particularly acute in California.

In recent years, California has endured one of its most severe droughts (Wang et al., 2017). While these cyclic phenomena are known to periodically repeat, because of climate change they are in the future expected to increase in frequency, intensity, and length (Rahmstorf & Coumou, 2011). California produces more than 80% of the US wine (WineInstitute, 2018), and the wine sector is growing in quantity and quality. Vines are best grown with a strategy known as *stress irrigation* or *deficit irrigation* whereby the amount of water delivered to vines is limited to improve grape quality and curb vegetative growth. Given the intrinsic soil

* Corresponding author: scarpin@ucmerced.edu

heterogeneity found in large-scale farming, commonly used strategies using uniform irrigation over large blocks are suboptimal and lead to non-uniform growth. In particular, when in doubt, ranch managers tend to over-irrigate, with the doubly negative effect of consuming more water and possibly reducing the quality of the final product. One of the leading challenges in this domain is that if one looks at an irrigation infrastructure as a control system (McCarthy et al., 2011), we find an abundance of sensor data, but a scarcity of fine grain actuation mechanisms (see Fig. 1).

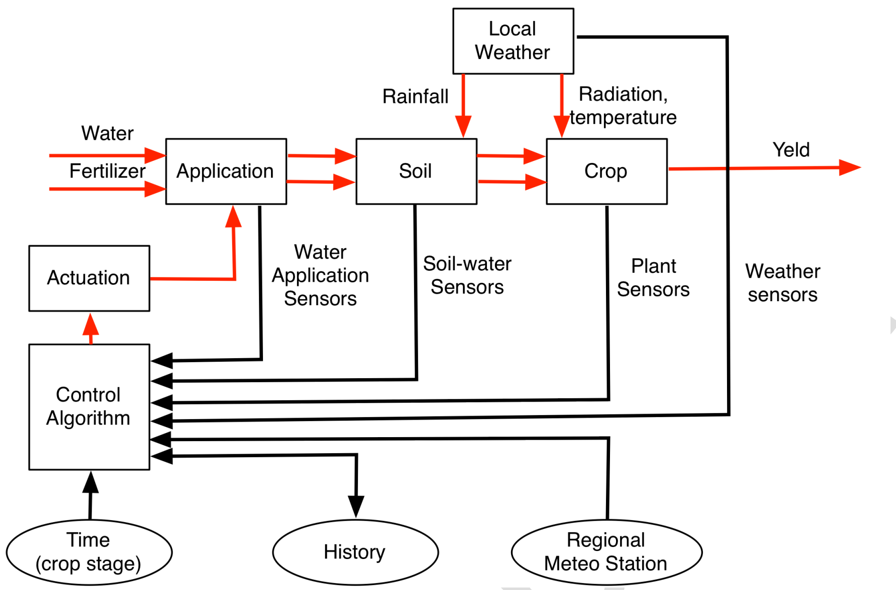

Figure 1: *Irrigation system as a control system. Figure adapted from (McCarthy et al., 2011).*

Funded by the US Department of Agriculture as part of the National Robotics Initiative led by the National Science Foundation, RAPID (Robot Assisted Precision Irrigation Delivery) is a multi-year, multi-campus research project aiming at overcoming these limitations by developing a novel irrigation system relying on a combination of remote sensing and mobile robots to carefully tune the amount of water delivered to each vine. The basic idea is to retrofit existing irrigation systems with passive, low-cost drip emitters that can be individually regulated to adjust the amount of water delivered to each vine. Due to the sheer scale of the ranches commonly found in intensive grape farming, manually adjusting them is an approach that does not scale. To this end, we are designing a fleet of robots that will periodically traverse the tree rows and make the necessary adjustments to increase or decrease water flow to match desired conditions. Decisions about the necessary adjustment will be made based on an inference system combining, though learning, weather data, remote sensing, and soil moisture measures. Note that electromechanical systems or IoT based solutions are currently not considered viable on a large scale by end users, because of their cost and the necessity to withstand for extensive periods of time extremely harsh environmental conditions.

The rest of this chapter is organized as follows. In section 2 we discuss selected works in the area of estimation and robotics. Section 3 provides an overview of the RAPID system and then in Section 4 we provide preliminary results for estimation and routing. We conclude the chapter with section 5, where conclusions and opportunities for future work are discussed.

2 Related Work

Soil and Field Modeling

The RAPID project incorporates modeling water levels available at a per-plant basis, and evaluating sensing modalities for updating those models. The problem of modeling how subsurface water propagates in space and over time has been extensively investigated and a variety of models and software packages are available, both commercially and as open source. These methods aim at solving a problem known as *soil water balance calculation* and compute the water balance as the difference between inputs (precipitation and irrigation) and outputs (evapotranspiration, runoff, and drainage). The problem of spatially varying moisture measurement and simulation has been extensively studied (Clement et al., 1994; Romano, 2014; Tian et al., 2011) using models based on finite differences, nonlinear differential equations and partial differential equations. Temporal variability has been considered as well, for example in (Choi & Jacobs, 2007; Ojha et al., 2014; Vereecken et al., 2014). Methods specifically aiming at modeling subsurface moisture with drip irrigation have been developed (Kandelous & Šimůnek, 2010) and experimentally shown to be quite accurate.

Building upon these models, several simulation packages are available for modeling surface, subsurface, and groundwater flow. Models for water flow through soil typically involve solving three-dimensional partial differential equations, e.g. the commonly used Richards equation for variably saturated soil water flow (Lafolie et al., 1989; Or & Coehlo, 1996; van Dam & Feddes, 2000; Tian et al., 2011). Software packages like HYDRUS 2D/3D (Hydrus, 2018) or LandLab (Adams et al., 2014) have been used for modeling flow and designing drip irrigation systems (Deb et al., 2013; Rahimi et al., 2004). However, such models have not been implemented to design irrigation systems at the per-plant scale, and have not been used for model-based closed-loop regulation. Furthermore, such models rely on many different parameters, including physical parameters related to soil or topographical properties, as well as being able to estimate initial and boundary conditions for moisture levels.

Environmental Sensing

Different sensing modalities for estimating plant water availability across vineyards have received a lot of attention, and there is a rich history regarding soil moisture measurement techniques (Dane & Topp, 2002; Robinson et al., 2008; Smith & Mullins, 2000). Recently, airborne thermal imagery has been used to assess the spatial variability of water stress, an indicator of soil moisture availability across vineyards (Baluja et al., 2012; Bellvert et al., 2014), as well as soil moisture probes based on soil electroconductivity. Wireless sensor networks (WSN) have been proposed for environmental monitoring (Pon et al., 2005; Rahimi et al., 2004), and applications in agriculture (Rehman et al., 2014). A WSN composed of 135 soil moisture and 27 temperature sensors was deployed in an apple tree orchard of about 5000 m^2 (Majone et al., 2013). The network is in charge of estimating soil moisture, but does not include an actuation subsystem capable of adjusting the application of water. A similar system was proposed in (Navarro-Hellín et al., 2015) where it was demonstrated that current sensing technology is mature to determine soil moisture levels. In (Gutierrez et al., 2015) an irrigation sensor was developed using a smartphone camera, and a water pump was then activated based on the results provided by the image-processing algorithm. This system assumes that a human will manually take a picture of its crops and therefore does not scale when large extensions are considered. Moreover, the pump activates a whole irrigation line and does not allow a per-plant adjustment.

Irrigation Scheduling

While irrigation schedules based on moisture measurements and water balance calculation are extensively used, in plant sciences there is increasing interest to switch towards plant-based methods, i.e., to determine the amount of water based on data related to plant physiology and water stress (Jones, 2004; Romero et al., 2010). However, to date, the question of which quantities should be measured and how these measures could be turned into a decision about how much water should be delivered to the plant is an active area of study. This is a research problem pertaining to the domain of plant science and orthogonal to the one tackled in this project. With RAPID we develop technology that enables individual adjustment of water emitters so that every plant can receive a different amount of water. In our control algorithm the amount to be delivered is based on soil water measurements and soil water balance calculations; our control algorithm is however amenable to systems where emitter states are controlled based on plant water stress levels.

Mobile Robotics

From the robotics and control theory communities there is a large body of literature concerning the use of robots and mobile sensor networks for data collection over stationary or time varying distributions (Bullo et al., 2009; Martinez et al., 2007; Schwager et al., 2009; Smith et al., 2012). Among the many approaches proposed in literature the Voronoi based coverage control algorithms introduced by Bullo and collaborators have gained notable traction (Bullo et al., 2009; Martinez et al., 2007). These approaches have in general only been studied in theory and have seen limited applications in the field (Pierson et al., 2015).

A related problem of interest is how to combine together multiple maps or models providing partial information about a domain of interest. This work is in the general area of map merging (Aragüés et al., 2011; Aragüés et al., 2012; Erinc & Carpin, 2014; Saeedi et al., 2016) and different solutions are available depending on the structure of the underlying model, i.e., occupancy grids, topological maps, appearance based maps, and more. Moisture maps essentially represent a scalar field, and the problem of combining them together can then be seen as a special instance of image stitching. Methods aiming at combining together multiple occupancy grid maps can be used, e.g., (Birk & Carpin, 2006; Carpin, 2008; Saeedi et al., 2012). Recently, these techniques have indeed been used to stitch together multiple images for agricultural applications (Li & Isler, 2016) where images were obtained from UAVs flying at low elevations.

Robot Task Allocation

Robot and multi-robot task allocation (MRTA) is a mature, active area of research drawing insights from multiple domains, like economy, operations research and more (Korsah et al., 2013; Khamis et al., 2015). When a cooperative system is considered with multiple tasks to be assigned and multiple heterogeneous agents available to execute them, the question "Which robot should execute which task?" (Gerkey, 2003) naturally emerges, and MRTA algorithms aim at answering such question optimizing some objective function accounting for the tasks being solved and the agents. Many MRTA algorithms can handle heterogeneous agents and could then in principle be used to solve task allocation problems where some of the agents are robots and some are humans.

With increased interest in the development of co-robot systems, the problem of scheduling human-robot teams became a question of independent value (Clare et al., 2012; Gombolay & Shah, 2014; Murphy, 2015). In this case the main focus is to leverage the peculiar strengths of humans and robots and to implement an allocation of tasks that does not alienate humans. Note that in general this is a different problem than human-robot interaction, because the scheduling algorithm may generate an allocation of tasks such that humans and robots to not interact at all. In the context of RAPID we face a situation where the set of tasks to be allocated is homogeneous, i.e., it consists of a set of spatially distributed emitters to be adjusted. Hence, the problem of picking a set of tasks engaging for humans (Gombolay et al., 2015) is not relevant because all tasks are the same. What is instead important is assigning tasks so that humans can improve their productivity without being overwhelmed.

3 RAPID Overview

Figure 2 shows how introducing RAPID would refine the control system presented in figure 1.

Rapid implements both a new estimation algorithm to determine how much water is to be delivered, as well as the actuation system. Both components rely on existing infrastructure and methods to gather data in the field, and consequently novel sensors or sensing systems are not being developed as part of this effort. Figure 3 illustrates in mode detail the overall architecture of the RAPID system we are developing, and evidences the feedback loop. As part of this project, we are currently developing three different efforts to develop three subsystems that will interact to implement the envisioned solution.

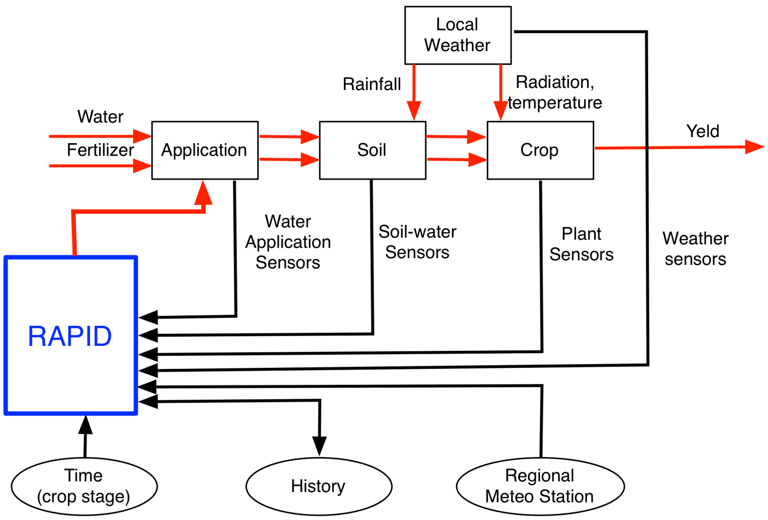

Figure 2: *revised irrigation as control system with the introduction of RAPID.*

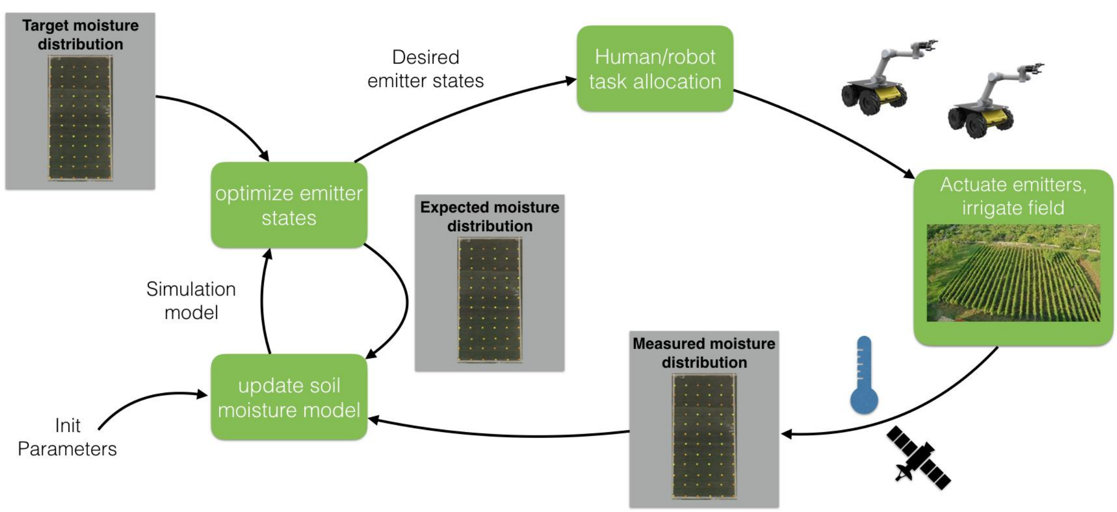

Figure 3: *RAPID architecture.*

The first effort consists in developing a learning based system to infer soil moisture levels from aerial images (Tseng et al., 2018). Based on deep convolutional neural networks (CNN), this module aims at learning a mapping between soil moisture data and aerial images.

The second thread consists in the co-design of a variable rate emitter and a robotic gripper that can perform the desired adjustments (Gealy et al., 2016; Berenstein et al., 2018). Mounted on a robotic arm placed on a mobile robot, or on a hand-held device carried by a human, the gripper can latch on a variable rate emitter and autonomously perform the desired adjustment. In fact, while we anticipate an extensive use of robots to address the scalability aspect, preliminary discussions with farmers and growers have evidenced that a human component will also be important, as the presence of a human expert in the vineyard has multiple benefits in terms of quality assessment and the overall process monitoring.

The last component consists in a suite of planning algorithms to distribute the load between robots operating in the vineyard (Thayer at al., 2018a; Thayer et al., 2018b). An essential part underlying these three efforts is the experimental validation in the field in collaboration with our commercial partners. In this chapter, due to space limitations, we mostly discuss the first and third task, i.e., inference and routing algorithms.

Our underlying assumption, corroborated by multiple interactions with vine growers, is that in most instances uniform growth is desired over a given block. Starting from this premise, the overall information flow of the RAPID system is as follows. Through the collection of aerial images, the trained CNN infers the moisture level in the soil of a given vineyard. This soil moisture map is then presented to a human expert (e.g., a ranch manager) who, based on domain knowledge, can then indicate a desired (*target)* soil moisture level. The mismatch between the inferred soil moisture level and the desired soil moisture translates into actionable information about how each variable rates emitter should be adjusted to increase or decrease the amount of water released and match the target distribution. This information is provided to the planner that then schedules one or more robots to make the desired adjustments throughout the vineyard. Once the mobile robot reaches an emitter to be adjusted, the robotic arm mounted on the robot performs the required change.

4 Preliminary Results

4.1 Soil Moisture Inference

Critical to the RAPID system is the knowledge of soil moisture to determine how emitters should be adjusted. Due to the sheer size of the vineyards managed by large-scale growers, measurements based on stationary sensors or mobile probes would provide only coarse scale data. To overcome these limitations, one possible alternative is using a data-driven approach to infer soil moisture conditions from aerial images. This approach is potentially advantageous because thanks to recent advances in UAV related technologies it has become relatively easy and inexpensive to acquire images covering large areas. Aerial images could be color images, hyperspectral, or thermal, just to name a few. In fact, there exist various previous works attempting to solve this very problem using a variety of machine learning algorithms, such support vector machines (SVM), Bayesian Neural Networks, and multilayer perceptrons (Hassan-Esfahani et al., 2014; Hassan-Esfahani et al., 2015; Hassan-Esfahani et al., 2017; (Zaman et al., 2012; Ahmad et al., 2010; Bachour et al., 2014). Estimation methods based on deep learning and related approaches (Goodfellow et al., 2016) have enormous potential but to the best of our knowledge have not yet been applied to this specific problem. One of the main limitations of these techniques is the necessity to acquire large-scale datasets for training. In a domain like the one considered, this could be prohibitively time consuming. To sidestep this entry-level obstacle and assess the potential of deep learning in this domain, we developed a simulator synthesizing artificial aerial images of a vineyard subject to different irrigation regimes (Tseng et al., 2018). The simulator is based on a discretized version of the Richards equation. In particular, the volumetric water content at time *t* is approximated as

$$s_t = s_{t-1} + \Delta s_t$$

where the variation Δs_t is defined as

$$\Delta s_t = -(r - a_t) - U(t)$$

In this last equation r is the soil drainage, a_t is the irrigation rate, and $U(t)$ is the plant uptake that is assumed constant in our simulation. The quantity $d = r + U$ is called *dissipation rate* and is what the algorithm tries to learn from the images generated by the simulator. More specifically, the simulator generates 320x200 RGB images where the growth of nm plants spaced in a regular grid with m row and n columns is determined based on a vector $\bar{d}$ of soil moisture dissipation rates and a vector $\bar{a}$ of applied irrigation rates (see figure 4).

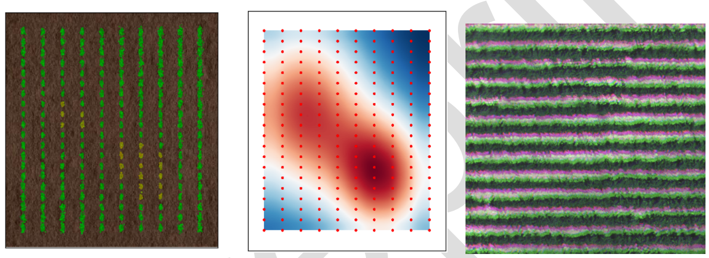

Figure 4: The left side shows one of the synthetic aerial images generated by the simulator, while the center figure shows the corresponding drainage image used for training. On the right, an actual cropped aerial image of a vineyard collected during our field experiments.

Both vectors have nm real components and are input parameters to the simulator. In all experiments presented in the following, $m = 20$ and $n = 10$. Starting from an initial uniform soil moisture distribution, the simulator propagates the distribution of moisture and adds leaves to a tree if the soil moisture value is positive. If the soil moisture is zero, then no leaves are added and the color of the existing leaves is altered by adding a shade of yellow. With this method, 1200 images were generated and then randomly split into 1000 training images and 200 validation images.

Starting from these images, the objective is to learn the dissipation rate. More formally, let x_i be the i-th synthetic image and $\overline{d_i}$ be the corresponding dissipation rate. The objective is to determine the parameters θ of a model H_θ that predicts the dissipation rate $\overline{d_i}$ associated with the image x_i. Given an aerial image x, for a given choice of parameters θ' the predicted soil moisture dissipation will be $H_{\theta'}(x)$. As common in machine learning literature, for a given choice of the vector of parameters θ we define a loss function L as

$$L(\bar{d}, H_\theta(x)) = \frac{1}{p}\sum_{j=1}^{p}(\bar{d}_j - H_\theta(x_j))^2$$

The optimal set of parameters would therefore be the vector θ^*solving the following minimization problem

$$\theta^* = argmin_\theta \sum_{j=1}^{N_{train}} L(\bar{d}_j, H_\theta(x_j))$$

To test the potential of deep convolutional neural networks in this domain, we consider two different architectures. The first, dubbed Deep Convolutional Neural Networks Uncorrelated Plants (CNNUP), is fed with a tree cropped out from the aerial image and produces the dissipation rate for the single tree (see figure 5).

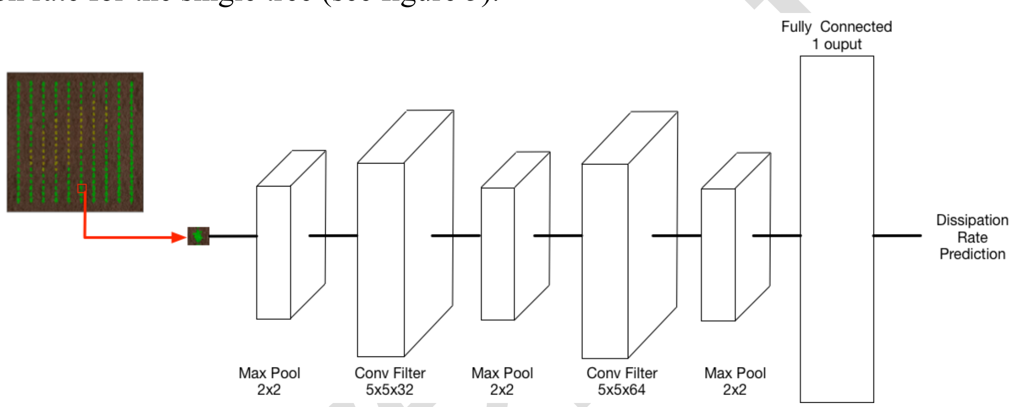

Figure 5: *The CNNUP architecture takes as input a single tree cropped from an image with 200 trees and predicts the dissipation rate for the input tree.*

The second, dubbed Deep Convolutional Neural Networks Correlated Field (CCNCF) instead receives as input the entire image and predicts $nm = 200$ dissipation rates, i.e., one for each tree in the input image (see figure 6).

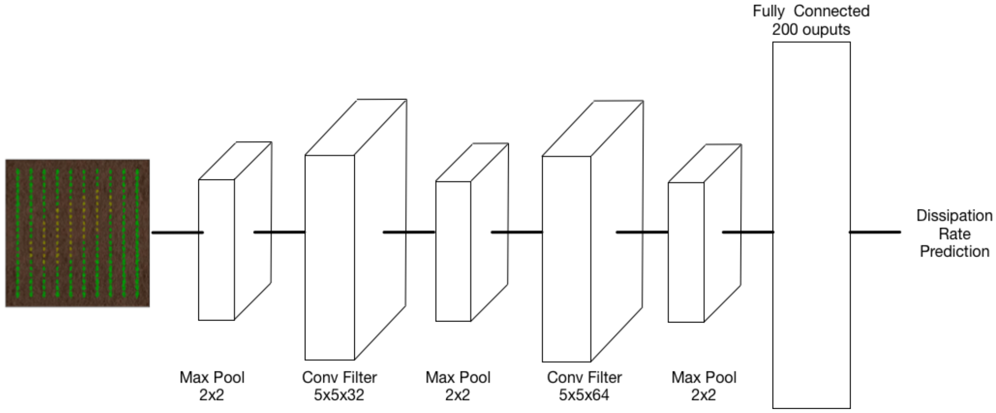

Figure 6: *The CCNCF architecture takes as input an entire image with 200 trees and produces as output the soil moisture prediction for the whole set of trees.*

These architectures were implemented using TensorFlow and the Adam algorithm was used for optimization (see (Tseng et al., 2018) for more implementation details).
To assess the potential of convolutional neural networks in this domain, we compared the two networks against five alternative approaches:

1. a constant baseline always producing a fixed value c for every input;
2. a set of 200 support vector machine (SVM) using a linear kernel, each predicting the moisture dissipation rate for one vine;
3. an uncorrelated method based on random forests using the RGB values to make predictions about moisture dissipation rate for one vine (RFUP);
4. a correlated method based on random forests where the whole image is used to make 200 predictions (RFCF);
5. A two-layer feed-forward neural network (NN) inspired by the work presented in (Hassan-Esfahani et al., 2017).

In all cases 1000 examples were used for training and 200 examples were in the validation set. To assess the merit of each approach and contrast them with each other, we consider the absolute error in the dissipation rate prediction $|\overline{d}_j - H(x)_j|$,i.e., the absolute difference between the predicted and actual dissipation rates. Table 1 shows mean, median, as well as the 25th and 75th percentile for this performance measure. Training times are also provided and refer to a system equipped with and Intel Core i7-6850k and three Nvidia Titan X Pascal GPUs.

Method	25th perc.	Mean	Median	75th perc.	Training Time (s)
Constant	0.082	0.177	0.162	0.246	--
SVM	0.015	0.038	0.030	0.053	4141

RFUP	0.030	0.071	0.060	0.101	578
RFCF	0.047	0.109	0.094	0.152	655
NN	0.037	0.086	0.075	0.119	243
CNNUP	0.021	0.054	0.044	0.077	538
CNNCF	0.012	0.034	0.027	0.047	603

Table 1: *results obtained comparing the proposed convolutional neural network architectures and five alternative machine-learning algorithms.*

From the table we can see that the CNNCF architecture outperforms all other methods in terms of the considered error metric. As expected, its training time exceeds the time needed to train the two-layer neural network, but the gap is moderate and the absolute time is contained. In particular, given that this refers to the offline training stage, this difference is acceptable from a practical standpoint. The overall objective of this preliminary study is to explore possible network architectures to identify promising approaches and to determine whether this solving method is practical at all. Results suggest that this is possible and efforts are ongoing to extend and adapt these deep architectures to deal with real data collected in the field.

4.2 Routing Algorithms in Vineyards

Once a uniform target moisture level for the block has been determined, either by a human expert or an algorithm, it will be necessary to adjust variable rate emitters to match it. The discrepancy between the current soil moisture level and the desired moisture is an indicator of how much the emitters should be adjusted. Therefore, for each emitter e we define a quantity $R(e) = |m(e) - T|$ as the difference between the moisture measured at the emitter $m(e)$ and the desired target moisture T. The quantity $R(e)$ indicates by how much the water flow at each emitter should be adjusted. In particular, if $R(e) = 0$ for a certain emitter, no adjustment is needed because the current soil moisture level matches the desired value. Large values of $R(e)$ indicate significant mismatches, either due to over-irrigation or under-irrigation.

In general, in large vineyards a mobile robot, or a person, will not be able to visit and adjust all emitters at once, due to the necessity to periodically recharge the batteries (or rest, in the case of a person). In the following, we indicate this limitation as *travel budget,* and we will indicate it with the letter B. Intuitively, B captures the constraint that a robot can only travel a finite distance between two successive charges. Therefore, an optimization problem emerges, i.e., selecting the subset of spatially distributed emitters that can be adjusted given the travel budget B. In particular, it appears intuitive to prioritize adjusting emitters with a high mismatch value $R(e)$.

This problem is related to a classic optimization problem known as *orienteering* (Golden et al., 1987). More precisely, the orienteering problem is defined as follows. Let B be a positive budget and G a graph where every edge has a non-negative weight and every vertex has a value. The orienteering problem asks to determine a path in G maximizing the sum of the values of the visited vertices and whose cost is no larger than B. It shall be noted that if a vertex is visited multiple times the reward is collected only once, but the cost of an edge is incurred every time it is traversed. In our problem setting this corresponds to the case where a robot visits more than once the same emitter. In such case the emitter would be adjusted only once, and accordingly the reward would be collected only once. It was shown in (Golden et al., 1987) that orienteering is an NP-hard optimization problem, and it was more recently shown that it is also APX-hard (Blum et al., 2007).

Numerous variations have been developed through the years. One that is particularly relevant to our discussion is the team orienteering problem (TOP) where one is given M agents, all subject to the same budget B, and the objective is to determine M paths maximizing the sum of the rewards collected by the M agents (Chao et al., 1996). Note that in TOP if a vertex is visited multiple times then its reward is collected only by the first robot. Because of its intrinsic complexity and the difficulty in developing suitable approximation algorithms, heuristic approaches have been the most common choice for solving large instances of the orienteering problem (Gunavan et al., 2016). General-purpose heuristics working for a large variety of graphs have however remained elusive.

The connection between orienteering and our vineyard routing problem is as follows. For a given vineyard, we build a graph where every emitter is associated with a vertex, and edges are added between vertices to model the motion constraints typical of a vineyard. One such graph is shown in Fig. 7. The connections between the vertices model the constraints that in a vineyard a robot can move along tree rows (horizontal edges), but row switching is allowed only at either end of the rows (see also Fig. 8). A graph of this type will be called *irrigation graph (IG)* in the following, and in particular we write *IG(m,n)* for an irrigation graph with *m* rows and *n* emitters along each row.

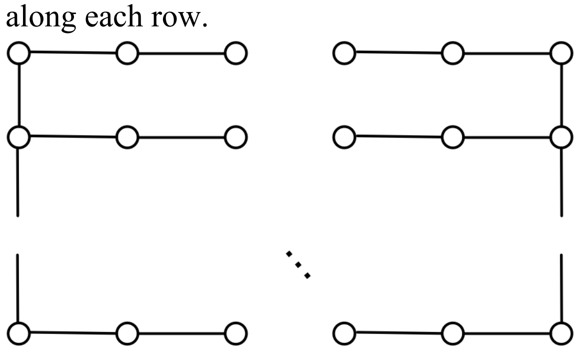

Figure 7: *An irrigation graph modeling the spatial layout of emitters in a vineyard.*

Figure 8: *A robot moving in a vineyard can change row only at either ends because of the irrigation lines running between vines.*

For each vertex the reward is then set to $R(e)$, as defined above. For each edge, the cost is set as the distance between the vertices it connects. Graphs like the one shown in Fig. 7 are a particular instance of *BP3,* i.e., bipartite, planar graphs with vertices of degree at most three. A special instance is obtained when the cost of each edge is constant, for example to model the fact that emitters are regularly spaced in a vineyard. In (Thayer et al., 2018) we defined the following problem.

> ***Constant Cost Orienteering on BP3 Problem (CCOBP3P)*:** let $G = (V, E)$ a *BP3* graph, $r: V \rightarrow \mathbb{R}$ a value function and $c: E \rightarrow \mathbb{R}$ a constant cost function defined over the edges. Moreover let v_1, v_n be two vertices in V and T_{MAX} a positive constant. Find a path from v_1 to v_n of maximum value with cost no larger than T_{MAX}.

In (Thayer et al., 2018a) we have shown that CCOBP3P is NP-hard and therefore we developed domain specific heuristics both for the single robot case and for the multi-robot case (Thayer et al., 2018b). We refer the reader to the aforementioned paper of the detailed algorithmic formulation and we here sketch the underlying ideas. In all cases we assume that $v_1 = v_n$, i.e., at the end the robot must return to the start location. For the single robot case, we have developed two heuristics called *greedy row* and *greedy partial-row.* The greedy row heuristic on an *IG(m,n)* works as follows. Preliminarily, for each of the *m* rows we compute the quantity

$R_i = \sum_{j=1}^{n} R(i, j)$

i.e., the overall reward collected if the robot fully traverses the i-th row in the graph[†]. Moreover, to each row we associate a *feasible* flag, indicating whether with the remaining budget the robot can move from the current position to the row, traverse it, and go to the final position. This flag is incrementally updated, as the robot route is updated and more of its allocated budget T_{MAX} is consumed. The overall route is then built as follows. Considering the current robot location v and remaining budget, the robot updates the feasible flag for all rows, and for each feasible row it computes the normalized cost

$$R_i' = \frac{R_i}{cost(v,i)}$$

i.e., the ratio between the reward of the *i-th* row and the cost to reach it, and traverse it. Among all feasible rows, the path is then extended by appending the row maximizing the normalized cost and the residual budget is updated subtracting the budget spent to extend the route. This process is iterated until no more feasible rows are left. At that point the robot returns to the starting location. Note that by construction this is always possible because a row is appended to the tour only if it is feasible, i.e., if there is enough residual budget to reach the final location. Therefore a valid solution is always found.

Observing that soil moisture displays a ‘‘local’’ structure, i.e., wet and dry locations tend to be clustered, we designed a different heuristic called *greedy partial-row* (GPR). The intuition is that if there is a region with high rewards at either end of the rows, it may be more advantageous to enter the row, adjust some of its emitters, and then come back from the same side from which it was entered, and then move to an adjacent row. The tradeoff is that the robot will spend some budget without getting any reward when traveling back along the row, but this may enable to move to an adjacent row with high rewards without having to first traverse an entire row. This would not be possible in the previous heuristic where once one row is entered it must be fully traversed. This idea can be implemented following an approach similar to the previous one, i.e., the route is iteratively extended by greedily selecting feasible rows or partial rows. The definition of a feasible row or partial row is exactly the same, i.e. there has to be sufficient residual budget to cover the partial row and then proceed to the final vertex v_n. However, given that in this case a row can be partially entered by either side, we compute two rewards, i.e., $L(i,j)$ and $R(i,j)$. $L(i,j)$ is the reward obtained if the robot enters the *i-th* row from the left, proceeds until the *j-th* emitter, and then goes back. The definition of $R(i,j)$ is the same but considers starting from the right. Such costs are then scaled by the travel cost as in the full row heuristic, and the path is then extended greedily (see (Thayer et al, 2018a) for more details). Both these heuristics have computational complexity polynomial in the number of rows *m* and columns *n*. In particular, the former is $O(m^2)$ while the latter is $O(m^2 n)$.

[†] Since emitters are placed in a grid layout, *R(i,j)* is the reward associated with the emitter placed at position *(i,j)* on the grid.

Besides the single robot heuristics, we also developed three heuristics for the TOP problem, as it is foreseeable that multiple robots will be developed in a vineyard to service it more efficiently (Thayer et al, 2018b). In the following we indicate with M the number of robots. The first approach, dubbed *sectioning*, splits the vineyard into M sections (each consisting of a certain number of rows) and assigns a robot to each section. Within each section the robot then runs the greedy partial row heuristic. Sections are defined aiming at having the same sum of rewards and not the same area, so that each robot can potentially make the same impact. With this approach, all robots stay separate and cannot collide with each other. The second heuristic, called *SeriesGPR*, runs M times the GRP heuristic in sequence. At every iteration, the path for a single robot is planned, and the map is updated at every iteration to zero out the rewards associated with the emitters already serviced. With the SeriesGPR algorithm, all robots can visit any area in the vineyard. Finally, the *ParallelGPR* approach plans the path for all robots at the same time. This means that at every iteration the algorithm determines for each robot what is the best next move (using the GPR heuristic) and the path is extended for that robot. The approach is iterated until all robots have expended their budget.

Extensive results are presented in (Thayer et al, 2018a; Thayer et al, 2018b) and we here report just a small representative set. The two single robot heuristic approaches we developed have been numerically compared with a general purpose heuristic for orienteering called the *S-algorithm* (Tsiligirides, 1984). For the TOP case, instead, we compared our findings with the guided local search (GLS) heuristic formerly proposed in literature (Vansteenwegen et al., 2009).

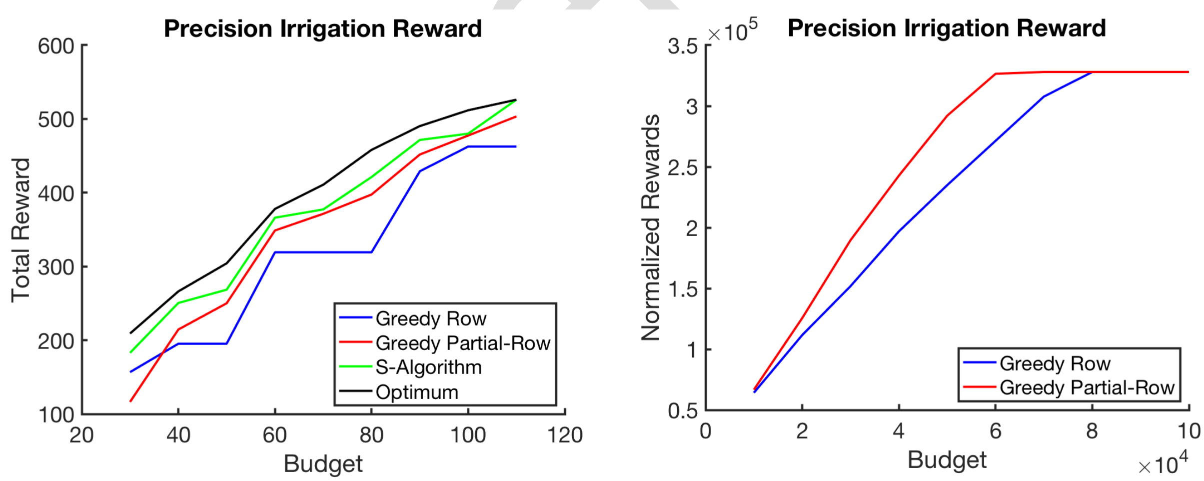

Figure 9: *performance of the single robot routing heuristics.*

Figure 9 shows the performance of the single robot heuristics. The left figure shows the results for a small vineyard where the optimal solution can be found with an integer programming approach. The figure shows that the greedy partial row heuristic and the S-algorithm are relatively close to the optimal solution and show comparable performance. However, the S-algorithm does not scale well and cannot efficiently solve instances with hundred of thousands of vines. In the right figure we then compare the performance of the two heuristics we proposed for the case of a full size vineyard. The chart confirms that the greedy partial row is the most competitive one. In figure 10, instead, we compare the performance of the various heuristics as a function of the budget for a case with 6 agents. In this case the ParallelGPR algorithm emerges as the most competitive and the result is representative of other cases, too.

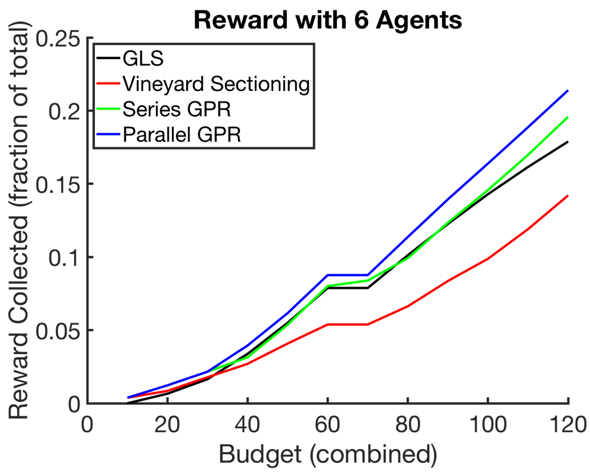

Figure 10: *performance of the heuristics for the multi-agent case.*

5 Conclusions and Future Work

We are currently in the second year of a project expected to last for four years. Therefore, while many of the individual components have been prototyped, system integration of the whole RAPID system is still ongoing. Due to limited space, we described just selected contributions but refer the reader to cited references for additional results, for example regarding the design and of the gripper.

Despite the fact we are still in an early stage of the project, there are some conclusions that can be already drawn. With regard to estimation, deep convolutional neural networks seem to be a promising method to infer soil moisture content from aerial images and they appear to have the potential to outperform methods proposed in the past. With regard to routing in vineyards, the peculiar structure of the environment calls for the development of domain-specific heuristics capable of scaling to very large instances. Our preliminary results have identified methods superior to the state of the art and scalable.

In terms of future work, besides the necessary integration of all components, we identify the following venues for further investigation. First, it will be necessary to test the CNN architecture on data collected in the field, and possibly later its structure to better model the data we gathered. For routing, we will incorporate uncertainty in the rewards, since those are obtained through an estimation process that is necessarily uncertain. Uncertainty shall therefore be considered when deciding where the robot should go next. Finally, extensive field evaluation will be needed to assess the viability of the whole system and ultimately quantify water savings.

Acknowledgments

This material is based upon work that is supported by the USDA-NIFA under award number 2017-67021-25925 (National Robotics Initiative). Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the authors and do not necessarily reflect the view of the U.S. Department of Agriculture or the US government.

We gratefully acknowledge Luis Sanchez and Nick Dokoozlian from E&J Gallo Winery for having granted access to their vineyards for data collection, and for the valuable information provided during this project. We also thank the many lab members and collaborators who contributed to this project: Thomas Thayer, Jose Manuel Gonzalez and Carlos Diaz Alvarenga, Christine Breckenridge, Jonathan Garache, David Tseng, David Wang, Carolyn Chen, Lauren Miller, William Song, and Juan Aparicio Ojea.

References

Adams, J.M., Nudurupati, S.S., Gasparini, N.M., Hobley, D., Hutton, Eric W.H., Tucker, G.E. and Istanbulluoglu, E. (2014). LandLab: sustainable software development in practice. In: Proceedings of the *Second Workshop on Sustainable Software for Science: Practice and Experiences.*

Ahmad, S., Kalra, A. and Stephen, H. (2010). Estimating soil moisture using remote sensing data: A machine learning approach. *Advances in Water Resources*, 33(1), pp. 69–80.

R. Aragüés, J. Cortés, and C. Sagüés (2011). Distributed consensus algorithms for merging feature- based maps with limited communication. *Robotics and Autonomous Systems*, 59(3-4), pp. 163–180.

R. Aragüés, J. Cortés, and C. Sagüés (2012). Distributed consensus on robot networks for dynami- cally merging feature-based maps. *IEEE Transactions on Robotics*, 28(4), pp. 840– 854.

Bachour, R., Walker, W.R., Ticlavilca, A.M., McKee, M. and Maslova, I. (2014). Estimation of spatially distributed evapotranspiration using remote sensing and a relevance vector machine, *Journal of Irrigation and Drainage Engineering*, 140(8), pp. 04014029.

Baluja, J., Diago, M.P., Balda, P., Zorer, R., Meggio, F., Morales, F. and J. Tardaguila (2012). Assessment of vineyard water status variability by thermal and multispectral imagery using an unmanned aerial vehicle (UAV). *Irrigation Science*, 30(6), pp. 511-522.

Bellvert, J., Zarco-Tejada, P.J., Girona, J. and Fereres, E. (2014). Mapping crop water stress index in a 'pinot-noir' vineyard: comparing ground measurements with thermal remote sensing imagery from an unmanned aerial vehicle. *Precision agriculture*, 15(4), pp. 361-376.

Berenstein, R., Fox, R., McKinley, S., Carpin, S. and Goldberg, K. (2018). Robustly Adjusting Indoor Drip Irrigation Emitters with the Toyota HSR Robot. In: *Proceedings of the IEEE International Conference on Robotics and Automation*, pp. 2236-2243.

Birk, A. and Carpin, S. (2006). Merging occupancy grids from multiple robots. *Proceedings of the IEEE*, 94(7), pp. 1384–1397.

Blum, A., Chawla, S., Karger, D.R., Lane, T., Meyerson, A. and M. Minkoff (2007). Approximation algorithms for orienteering and discounted-reward TSP. *SIAM Journal on Computing*, 37(2), pp. 653–670.

Bullo, F., Cortes, J. and Martinez, S. (2009). *Distributed Control of Robotic Networks*. Princeton.

Carpin, S. (2008). Fast and accurate map merging for multi-robot systems. *Autonomous Robots*, 25(3), pp. 305–316.

Chao, I., Golden, B. and Wasil, E. (1996) The team orienteering problem. *European Journal of Operations Research*, 88, pp. 464–474.

Choi, M. and Jacobs J.M. (2007). Soil moisture variability of root zone profiles within SMEX02 remote sensing footprints. *Advances in Water Resources*, 30(4), pp. 883-896.

Clare, A.S., Cummings, M.L., How, J. P., Whitten, A.K. and Toupet, O. (2012). Operator object function guidance for a real-time unmanned vehicle scheduling algorithm. *Journal of Aerospace Computing, Information, and Communication*, 9(4), pp. 161–173.

Clement, T.P., Wise, W.R. and Molz, F.J. (1994) A physically based, two-dimensional, finite-difference algorithm for modeling variably saturated flow. *Journal of Hydrology*, 161(1), pp. 71-90.

Dane, J.H. and Topp, G.C., editors (2002). *Methods of soil analysis. Part 4. Physical methods*. Soil Society of America.

Deb, S.K., Shukla, M.K., Šimunek, J. and Mexal J.G. (2013). Evaluation of spatial and temporal root water uptake patterns of a flood-irrigated pecan tree using the HYDRUS (2D/3D) model. *Journal of Irrigation and Drainage Engineering*, 139(8), pp. 599-611.

Erinc, G. and Carpin, S. (2014). Anytime merging of appearance-based maps. *Autonomous Robots*, 36(3), pp. 241–256.

Gealy, D.V., McKinley, S., Guo, G., Miller, L., Vougioukas, S., Viers, J., Carpin, S. and Goldberg, K. (2016). DATE: a hendheld co-robotic device for automated tuning of emitters to enable precision irrigation. In: *Proceedings of the IEEE International Conference on Automation Science and Engineering*, pp. 922-927.

Gerkey, B.P. (2003). *On multi-robot task allocation*. PhD. University of Southern California.

Golden, B.L., Levy, L. and Vohra, R. (1987). The orienteering problem. *Naval Research Logistics*, 34(3), pp. 307-318.

Gombolay, M. and Shah. J. (2014). Challenges in collaborative scheduling of human-robot teams. In: *Proceedings of the Association for the Advancement of Artificial Intelligence Symposium*, pp. 73-75.

Gombolay, M.C., Gutierrez, R.A., Clarke, S.G., Sturla, G.F. and Shah, J.A. (2015). Decision-making authority, team efficiency and human worker satisfaction in mixed human–robot teams. *Autonomous Robots*, 39(3), pp. 293–312.

Goodfellow, I., Bengio, Y. and Courville, A. (2016). *Deep learning.* MIT Press.

Gunavan, A., Chuin Lin, H., and Vansteenwegen, P. (2016). Orienteering problem: A survey of recent variants, solution approaches and applications. *European Journal of Operational Research*, 255(2), pp. 315–332.

Gutierrez Jaguey, J., Villa-Medina, J.F., Lopez-Guzman, A. and M.A. Porta-Gandara (2015). Smartphone irrigation sensor. *Sensors Journal*, 15(9), pp. 5122-5127.

Hassan-Esfahani, L., Torres-Rua, A., Ticlavilca, A.M., Jensen, A. and McKee, M. (2014). Topsoil moisture estimation for precision agriculture using unmmaned aerial vehicle multispectral imagery. In: *Proceedings of the IEEE International Geoscience and Remote Sensing Symposium*, pp. 3263–3266.

Hassan-Esfahani, L., Torres-Rua, A., Jensen, A. and McKee, M. (2015). Assessment of surface soil moisture using high-resolution multi-spectral imagery and artificial neural networks. *Remote Sensing*, 7(3), pp. 2627–2646.

Hassan-Esfahani, L., Torres-Rua, A., Jensen, A. and Mckee, M. (2017). Spatial root zone soil water content estimation in agricultural lands using bayesian-based artificial neural networks and high-resolution visual, NIR, and thermal imagery. *Irrigation and Drainage*, 66(2), pp. 273–288.

Hydrus (2018). Hydrus 2D/3D [online] Available at:

https://www.pc-progress.com/en/Default.aspx?hydrus-3d [Accessed October 2018]

Jones, H.G. (2004). Irrigation scheduling: advantages and pitfalls of plant-based methods. Journal of Experimental Botany, 55(407), pp. 2427–2436.

Kandelous, M.M. and Šimůnek, J. (2010). Comparison of numerical, analytical, and empirical models to estimate wetting patterns for surface and subsurface drip irrigation. *Irrigation Science*, 28(5), pp. 435-444.

Khamis, A., Hussein, A. and Elmogy, A. (2015). Multi-robot task allocation: A review of the state-of-the-art. In: Koubâa, A. and Martínez-de Dios, J. (editors), *Cooperative Robots and Sensor Networks*, Springer, pp. 31-51.

Korsah, G.A., Stentz, A. and Dias, M.B. (2013). A comprehensive taxonomy for multi-robot task allocation. *International Journal of Robotics Research*, 32(12), pp. 1495–1512.

Lafolie, F., Guennelon, R. and Van Genuchten, M.T. (1989). Analysis of water flow under trickle irrigation: I. Theory and numerical solution. *Soil Science Society of America Journal*, 53(5), pp. 1310-1318.

Li, Z. and Isler, V. (2016). Large scale image mosaic construction for agricultural applications. *Technical Report 16-002*, University of Minnesota.

Majone, B., Viani, F., Filippi, E., Bellin, A., Massa, A., Toller, G., Robol, F., and Salucci, M. (2013). Wireless sensor network deployment for monitoring soil moisture dynamics at the field scale. *Procedia Environmental Sciences*, 19, pp. 426-435.

Martinez, S., Cortes, J. and Bullo, F. (2007). Motion coordination with distributed information. *IEEE Control Systems Magazine*, 27(4), pp. 75–88.

McCarthy, A.C., Hancock, N.H. and Raine, S.R. (2011). Advanced process control of irrigation: the current state and an analysis to aid future development. *Irrigation Science*, 31(3), pp. 183-192.

Murphy, R.R. (2015). Meta-analysis of autonomy at the DARPA robotics challenge trials. *Journal of Field Robotics*, 32(2), pp. 189–191.

Navarro-Hellín, H., Torres-Sánchez, R., Soto-Valles, F., Albaladejo-Pérez, C., López-Riquelme, J.A. and R. Domingo-Miguel (2015). A wireless sensors architecture for efficient irrigation water management. *Agricultural Water Management*, 151, pp. 64-74.

Ojha, R., Morbidelli, R., Saltalippi, C., Flammini, A. and Govindaraju. R.S. (2014). Scaling of surface soil moisture over heterogeneous fields subjected to a single rainfall event. *Journal of Hydrology*, 516, pp. 21-36.

Or, D. and Coelho, F.E. (1996). Soil water dynamics under drip irrigation: transient flow and uptake models. *Transactions of the ASAE*, 39(6), pp. 2017-2025.

Pierson, A., Figueiredo, L.C., Pimenta, L.C.A. and Schwager, M. (2015). Adapting to performance variations in multi-robot coverage. In: *Proceedings of the IEEE International Conference on Robotics and Automation*, pp. 415–420.

Pon, R., Kansal, A., Liu, D., Rahimi, M., Kaiser, W.J., Pottie, G.J., Srivastava, M., Sukhatme, G. and D. Estrin (2005). Networked infomechanical systems (NIMS): next generation sensor networks for environmental monitoring. In: *IEEE MTT-S International Microwave Symposium Digest*, pp 373-376.

Provenzano, G. (2007). Using hydrus-2d simulation model to evaluate wetted soil volume in subsurface drip irrigation systems. *Journal of Irrigation and Drainage Engineering*, 133(4), pp. 342-349.

Rahimi, M., Pon, R., Kaiser, W.J., Sukhatme, G., Estrin, D. and M. Srivastava (2004). Adaptive sampling for environmental robotics. In: *Proceedings of the IEEE International Conference on Robotics and Automation*, pp. 3537-3544.

Rahmstorf, S. and Coumou, D. (2011). Increase of extreme events in a warming world. *Proceedings of the National Academy of Sciences*, 108(44), pp. 17905-17909.

Rehman, A., Zafar Abbasi, A., Islam, N. and Shaikh, Z. A. (2014). A review of wireless sensors and networks' applications in agriculture. *Computer Standards & Interfaces*, 36(2), pp. 263-270.

Robinson, D.A. , Campbell, C.S. , Hopmans, J.W., Hornbuckle, B.K., Jones, S.B., Knight, R., Ogden, F., Selker, J. and O. Wendroth (2008). Soil moisture measurement for ecological and hydrological watershed-scale observatories: A review. *Vadose Zone Journal*, 7(1), pp. 358-389.

Romano, N. (2014). Soil moisture at local scale: Measurements and simulations. *Journal of Hydrology*, 516(4), pp. 6-20.

Romero, P., Ferandez-Fernandez, J.I. and Martinez-Cutillas, A. (2010). Physiological thresholds for efficient regulated deficit-irrigation management in winegrapes grown under semiarid condi- tions. *American Journal of Enology and Viticulture*, 61(3), pp. 300–312.

Saeedi, S., Paull, L., Trentini, M., Seto, M. and Li, H. (2012). Map merging using Hough peak matching. In: *Proceedings of the IEEE/RSJ International Conference on Intelliegent Robots and Systems*, pp. 4683–4688.

Saeedi, S., Trentini, M., Seto, M. and Li, H. (2016). Multiple-robot simultaneous localization and mapping: A review. *Journal of Field Robotics*, 33(1), pp. 3–46.

Schaible, G.D. and Aillery, M.P. (2017). Challenges for US irrigated agriculture in the face of emerging demands and climate change. In: Ziolkowska, J. and Peterson, J.M., eds., *Competition for Water Resources*, Elsevier, 2017, pp. 44-79.

Schwager, M., Rus, D. and Slotine, J.J. (2009). Decentralized, adaptive coverage control for networked robots. *International Journal of Robotics Research*, 28(3), pp. 357–375.

Smith, K.A. and Mullins, C.E., editors (2000). *Soil and environmental analysis: physical methods*. Marcel Dekker.

Smith, S.L., Schwager, M. and Rus, D. (2012). Persistent robotic tasks: Monitoring and sweeping in changing environments. *IEEE Transactions on Robotics*, 28(2), pp. 410–426.

Thayer, T., Vougioukas, S., Goldberg, K. and Carpin, S. (2018a). Routing Algorithms for Robot Assisted Precision Irrigation. In: *Proceedings of the IEEE International Conference on Robotics and Automation*, pp. 2221-2228.

Thayer, T., Vougioukas, S., Goldberg, K. and Carpin, S. (2018b). Multi-Robot Routing Algorithms for Robots Operating in Vineyards. In: *Proceedings of the IEEE International Conference on Automation Science and Engineering*, 7-14.

Tian, F., Gao, L. and Hu, H. (2011) A two-dimensional Richards equation solver based on CVODE for variably saturated soil water movement. *Science China Technological Sciences*, 54(12), pp. 3251-3264.

Tseng, D., Wang, D., Chen, C., Miller, L., Song, E., Viers, J., Vougioukas, S., Carpin, S., Aparicio Ojea, J. and Goldberg, K. (2018). Towards Automating Precision Irrigation: Deep Learning to Infer Local Soil Moisture Conditions from Synthetic Aerial Agricultural Images. In: *Proceedings of the IEEE International Conference on Automation Science and Engineering,* pp. 284-291.

Tsiligirides, T. (1984). Heuristic methods applied to orienteering. *Journal of the Operational Research Society*, 35(9), pp. 797-809.

van Dam, J.C., and Feddes, R.A. (2000). Numerical simulation of infiltration, evaporation and shallow groundwater levels with the Richards equation. *Journal of Hydrology*, 233(1-4), pp. 72-85.

Vansteenwegen, P., Souffriau, W., Vanden Berghe, G., and Van Oudheusden, D. (2009). A guided local search metaheuristic for the team orienteering problem. *European Journal of Operations Research*, 196:118–127.

Vereecken, H., Huisman, J.A., Pachepsky, Y., Montzka, C., van der Kruk, J., Bogena, H., Weihermüller, L., Herbst, M., Martinez, G. and J. Vanderborght (2014) On the spatio-temporal dynamics of soil moisture at the field scale. *Journal of Hydrology*, 516(4), pp. 76-96.

Wallace, J. (2000). Increasing agricultural water use efficiency to meet future food production. *Agriculture, ecosystems & environment*, 82(1), pp. 105–119.

Wang, S.-Y., Yoon, J.-H., Gillies, R.R. and Hsu, H.-H. (2017). The California Drought. In: Wang, S.-Y., Yoon, J.-H., Funk, C.C. and Gillies, R.R. eds., *Climate Extremes: Patterns and Mechanisms*, Wiley, pp. 223-235.

Wineinstitute.org (2018) The Wine Institute Website [online], Available at: https://www.wineinstitute.org/resources/statistics [Accessed October 2018].

Zaman, B., McKee, M. and Neale, C.M. (2012). Fusion of remotely sensed data for soil moisture estimation using relevance vector and support vector machines. *International journal of remote sensing*, 33(20), pp. 6516–6552.